

---

# ESSAYS IN FINANCIAL AND BEHAVIOURAL ECONOMICS

---

A THESIS SUBMITTED TO THE UNIVERSITY OF DUBLIN, TRINITY COLLEGE  
IN APPLICATION FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

BY  
**SHANE BYRNE**

SUPERVISED BY  
PROF. MICHAEL KING

**FEBRUARY 2023**



**Trinity College Dublin**  
Coláiste na Tríonóide, Baile Átha Cliath  
The University of Dublin



# Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

Chapter 2 is based on joint work with Kenneth Devine, Michael King, Yvonne McCarthy, and Christopher Palmer. Chapter 3 is based on joint work with Channing Jang, Michael King, and Daniel Putman. I declare that I was the lead author for each of the co-authored chapters in the thesis. I agree to deposit this thesis in the University's open access institutional repository or allow the library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

---

Shane Byrne

# Summary

This thesis comprises three essays. The first two essays (Chapters 2 and 3) describe randomised controlled trials which experimentally test interventions to arm consumers with informational tools to navigate the promise and peril of the financial landscape more optimally. The third essay (Chapter 4) uses pan-European survey data to shed new light on an old puzzle in consumer finance: widespread non-participation in stock markets.

Chapter 2 tackles the well-documented failure to refinance puzzle, which can cost mortgage holders tens of thousands of euro over the lifetime of a loan, leave them with unnecessarily elevated debt-service burdens, inhibit the effective transmission of monetary policy, and reduce the incentive for competitive market entry. The chapter reports results from a large-scale field experiment with 12,000 borrowers with a large retail bank in Ireland. It tests the impact of smart enhancements to mandatory consumer disclosure in prompting take-up of refinancing opportunities. The best performing alternative increases the amount of refinancing by 76% compared against the pre-existing standard. With average savings of €1,209 achieved by refinancers in the first year, this represents no small achievement. To interpret results, the chapter extends and estimates a mixture model of inattentive financial decision-making to allow for disclosure treatment effects. The analysis shows that a simple reminder decreases the likelihood mortgage holders are inattentive by 15 percentage points. The results suggest that reminders could have larger effects on household refinancing than a 200 basis point rate cut and that reminders could strengthen the refinancing channel and stimulate local consumption even when policy rates are at the zero-lower bound or set in a monetary union.

Chapter 3 addresses the pernicious problem of fraud faced by small business owners using digital financial services (DFS) in low and middle-income countries, which can lead to severe financial cost, psychological damage, and frustrate the development of the strong and trusted financial ties critical for economic development. In a lab-in-the-field experiment in Northern Nigeria with 780 participants in the network of a partnering

digital market platform, this chapter tests the impact of a series of learning interventions, and a technical solution for the authentication of inbound communications, in improving the ability of small business owners to accurately discern fraudulent and genuine communications, as well as in building trust in DFS. We show the difficulty of improving discriminant ability between genuine and fraudulent communications, and the severity of the challenge faced by users in navigating the navigating the noisy landscape of competing communications. The results cast doubt on the utility of light-touch quick-fix remedial learning interventions to combat fraud susceptibility in DFS.

Chapter 4 examines a new dimension to another established puzzle in household finance: the role of private emergency financial safety nets in helping to explain widespread non-participation in stock markets. This topic has continued to exercise minds due to the considerable opportunity cost implied by non-participation, and its consequent implication for long-term patterns of wealth-inequality, as well as the implication of engendering divergent attitudes towards tax policy as well as risk-sharing and redistribution among those with and without financial assets. Consistent with a theory of insurance-induced consumption, the analysis shows that those households that enjoy the option of emergency financial support are more likely to participate in the stock market, and to report high financial risk appetite. In illustrating a mechanism of advantage compounding advantage, the results point towards the importance of effective and sustainable policies for financial inclusion, and highlight a potential use case for downside protection policies in replicating the insurance role provided by private safety nets in inducing stock-market participation.

# Acknowledgements

I would like to acknowledge the support of the Central Bank of Ireland in the completion of this thesis, in particular, my colleague and mentor Yvonne McCarthy, who has supported me unfailingly in my professional and academic journey.

I am grateful for the support of my supervisor, Michael King, whose wise counsel, constructive challenge, and steadfast encouragement made this thesis possible.

I would like to acknowledge the valuable contributions of my co-authors Christopher Palmer, Kenneth Devine, Daniel Putman, Channing Jang, and the research team at the Busara Center for Behavioural Economics, and Ahmaddo Bello University.

My research has benefited enormously from insightful suggestions and feedback received from John Fitzgerald, John Gathergood, Patrick Honohan, Tara McIndoe-Calder, Davide Romelli, and Michael Wycherley. Above all, I wish to thank my parents, Colbert and Sheila, for their unwavering support and encouragement. They inspired my curiosity and appetite for learning, opened a world of opportunity, and showed me the consolations of the outdoors, which have sustained me throughout my research journey. I thank my brother Lorcan, for his inspirational example, tireless humour, and solidarity.

To Eoin, Domhnall, Harry, Kevin, Mark, Neil, and Stephen, thank you for your lifetime friendship, it has been a precious and unparalleled gift. I look forward to nurturing it more faithfully than has been possible recently.

To Eimear, thank you for your love, belief, and your amazing spirit. You kept my head above water in the challenging moments. This thesis would not have been possible without you.

Finally, I would like to acknowledge the teachers I met along the way. I am indebted to them all.

*For my parents*

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>The Last Mile of Monetary Policy: Consumer Inattention, Disclosures, and the Refinancing Channel</b>	<b>7</b>
2.1	Introduction . . . . .	7
2.2	Context in literature . . . . .	10
2.2.1	The last-mile problem of monetary policy . . . . .	11
2.2.2	Central Bank communications . . . . .	12
2.2.3	Barriers to refinancing . . . . .	13
2.2.4	Consumer disclosures and reminders . . . . .	14
2.3	Institutional setting . . . . .	15
2.4	Experimental design . . . . .	17
2.4.1	Treatment arms . . . . .	18
2.4.2	Descriptive statistics . . . . .	19
2.5	Results . . . . .	21
2.5.1	Treatment effect heterogeneity . . . . .	24
2.5.2	Assessing potential Covid impacts . . . . .	26
2.5.3	General equilibrium considerations . . . . .	28
2.6	Inattention estimates . . . . .	29
2.6.1	Baseline refinancing model . . . . .	30
2.6.2	Allowing for inattention . . . . .	32
2.6.3	Comparison to interest-rate changes . . . . .	36
2.7	Conclusion . . . . .	39
2.8	Robustness to alternative refinancing parameters . . . . .	55
<b>3</b>	<b>Identifying Fraud and Building Trust in a Digital Market Platform: Evidence from a Lab Experiment in Nigeria</b>	<b>69</b>
3.1	Introduction . . . . .	69
3.2	Literature . . . . .	73
3.2.1	The impact of financial education interventions . . . . .	73
3.2.2	User-centered digital tools to combat fraud . . . . .	76
3.2.3	Confidence and performance . . . . .	77
3.3	Experimental design . . . . .	78
3.3.1	Experimental flow . . . . .	78
3.3.2	Experimental design . . . . .	79
3.3.3	Empirical strategy . . . . .	81
3.4	Results . . . . .	85
3.4.1	Effects of fraud education interventions . . . . .	85
3.4.2	Results of UCC intervention . . . . .	90



3.5	Conclusion . . . . .	91
3.6	Empirical methodology - extended . . . . .	105
3.6.1	Standard error adjustments . . . . .	105
3.6.2	Multiple hypothesis testing . . . . .	105
3.6.3	Statistical power . . . . .	105
3.7	Departures from pre-analysis plan . . . . .	107
3.7.1	Pre-specified but not included . . . . .	107
3.7.2	Included but not pre-specified . . . . .	107
3.8	Additional tables and figures . . . . .	108
<b>4</b>	<b>A Lender of Last Resort: Access to Emergency Financial Support and Household Financial Risk Taking</b>	<b>149</b>
4.1	Introduction . . . . .	149
4.2	Literature . . . . .	154
4.2.1	Stock market participation . . . . .	154
4.2.2	Impact of support mechanisms . . . . .	156
4.2.3	Contribution . . . . .	160
4.3	Mechanism . . . . .	160
4.4	Methodology . . . . .	162
4.5	Data and model . . . . .	164
4.5.1	Descriptive composition . . . . .	165
4.6	Empirical results . . . . .	167
4.6.1	Assessing model performance . . . . .	167
4.6.2	Estimated effects . . . . .	170
4.7	Discussion . . . . .	173
4.8	Conclusion . . . . .	175
4.9	Methodology and data - extended . . . . .	205
4.9.1	Weighting . . . . .	205
4.9.2	Model . . . . .	206
<b>5</b>	<b>Conclusion</b>	<b>208</b>

# List of Figures

2.1	US policy rates and new and outstanding mortgage interest rates . . . .	42
2.2	Ireland policy rates and new and outstanding mortgage interest rates .	43
2.3	Refinancing treatment effects . . . . .	44
2.4	Refinancing rates by treatment arm: no reminder sample . . . . .	45
2.5	Refinancing rates by treatment arm: reminder sample . . . . .	46
2.6	Refinancing rates by employment sector . . . . .	47
2.7	Market share of outstanding mortgages by product type . . . . .	57
2.8	External mortgage refinancing rates in Ireland . . . . .	58
2.9	Example control-group disclosure letter . . . . .	59
2.10	Example treatment-group disclosure letter . . . . .	60
2.11	Example reminder letter . . . . .	61
2.12	Distribution of refinancing incentives . . . . .	62
3.1	Direction of treatment effects in overall accuracy . . . . .	118
3.2	Accuracy in judgements through the sequence of scenarios . . . . .	119
3.3	Level of confidence reported through the sequence of scenarios . . . . .	119
3.4	Predictive margins in accuracy (fraud experience interaction) . . . . .	120
3.5	Difference in predicted values (no direct experience:1, direct experience:0)	120
3.6	Predictive margins in true positives (fraud experience interaction) . . .	121
3.7	Difference in predicted values (no direct experience:1, direct experience:0)	121
3.8	Predictive margins in true negatives (fraud experience interaction) . . .	122
3.9	Difference in predicted values (no direct experience:1, direct experience:0)	122
3.10	Predictive margins in accuracy (DFS experience interaction) . . . . .	123
3.11	Difference in predicted values (low DFS experience:1, high DFS experi- ence:0) . . . . .	123
3.12	Predictive margins in confidence (DFS experience interaction) . . . . .	124
3.13	Difference in predicted values (low DFS experience:1, high DFS experi- ence:0) . . . . .	124
3.14	Predictive margins in true positives (employment sector interaction) . .	125
3.15	Difference in predicted values (agriculture:1, non-agriculture:0) . . . . .	125
3.16	Predictive margins in true negatives (employment sector interaction) .	126
3.17	Difference in predicted values (agriculture:1, non-agriculture:0) . . . . .	126
3.18	Predictive margins in true negatives (age interaction) . . . . .	127
3.19	Difference in predicted values (above median age:1, median age or below:0)	127
3.20	Predictive margins in accuracy (trust interaction) . . . . .	128
3.21	Difference in predicted values (low trust:1, high trust:0) . . . . .	128
3.22	Predictive margins in true negatives (trust interaction) . . . . .	129
3.23	Difference in predicted values (low trust:1, high trust:0) . . . . .	129
3.24	Predictive margins in confidence (trust interaction) . . . . .	130
3.25	Difference in predicted values (low trust:1, high trust:0) . . . . .	130

3.26	Predictive margins in confidence (self-control interaction) . . . . .	131
3.27	Difference in predicted values (low self-control:1, high self-control:0) . .	131
3.28	Predictive margins in accuracy (gender interaction) . . . . .	132
3.29	Difference in predicted values (female:1, male:0) . . . . .	132
3.30	Predictive margins in confidence (gender interaction) . . . . .	133
3.31	Difference in predicted values (female:1, male:0) . . . . .	133
3.32	Predictive margins in confidence (risk appetite interaction) . . . . .	134
3.33	Difference in predicted values (low risk appetite:1, high risk appetite:0)	134
3.34	Accuracy and confidence: scenario averages . . . . .	137
3.35	Accuracy and confidence: scenario averages, separated by status . . . .	137
3.36	Example of a genuine scenario in the experimental task . . . . .	140
3.37	Example of a fraudulent scenario in the experimental task . . . . .	141
3.38	Elicitation of confidence in judgements following each presented scenario	142
4.1	Households benefiting from the option of emergency financial support, percentage by country . . . . .	177
4.2	Ownership of shares or mutual funds by country . . . . .	178
4.3	High financial risk appetite by country . . . . .	179
4.4	Value of investment holdings among stock market participants . . . . .	180
4.5	Probability of emergency financial support option (by wealth, income, age) . . . . .	181
4.6	Participation and risk appetite across households with and without the option of emergency financial support (wealth, income, age) . . . . .	182
4.7	Estimated effects on stock market participation: robust to alternative specifications . . . . .	186
4.8	Estimated effects on stock market participation: subgroup analysis . . .	187
4.9	Estimated effects on financial risk appetite: robust to alternative speci- fications . . . . .	188
4.10	Estimated effects on financial risk appetite: subgroup analysis . . . . .	189
4.11	Covariate balance: graphical representation (participation model) . . .	194
4.12	Density balancing plot (participation model) . . . . .	195
4.13	Cumulative distribution balancing plot (participation model) . . . . .	196
4.14	Covariate balance: graphical representation (financial risk appetite model)	199
4.15	Density balancing plot (financial risk appetite model) . . . . .	200
4.16	Cumulative distribution balancing plot (financial risk appetite model) .	201

# List of Tables

2.1	Treatment arms overview . . . . .	48
2.2	Descriptive statistics . . . . .	49
2.3	Internal refinancing treatment effects . . . . .	50
2.4	Internal refinancing treatment effects by treatment arm . . . . .	51
2.5	Refinancing treatment effect heterogeneity . . . . .	52
2.6	Parameter values used in optimal refinancing model . . . . .	53
2.7	Inattentive refinancing model maximum likelihood estimates . . . . .	54
2.8	Descriptive statistics across treatment cells . . . . .	63
2.9	Test of covariate balance by treatment arm . . . . .	64
2.10	External refinancing treatment effects . . . . .	65
2.11	Covid classification of employment sectors . . . . .	66
2.12	Internal refinancing treatment effects by employment sector . . . . .	67
2.13	Mixture ML estimates: robustness to alternative prepayment assumptions	68
3.1	Anti-Fraud Campaign Interventions . . . . .	80
3.2	Variables of interest for heterogeneous treatment effects . . . . .	94
3.3	Research hypotheses for core research questions . . . . .	95
3.4	Descriptive statistics across treatment cells . . . . .	96
3.5	Covariate balance by treatment . . . . .	97
3.6	Overall effect . . . . .	98
3.7	Confidence effect . . . . .	99
3.8	Effect on trust (ICW trust index) . . . . .	100
3.9	Effect on likelihood of future use . . . . .	101
3.10	Learning by doing effects . . . . .	102
3.11	Impact of authentication on engagement probability . . . . .	103
3.12	Impact of treatment on subsequent quiz score . . . . .	104
3.13	Estimated statistical power under alternative design scenarios . . . . .	106
3.14	Overall effect: partialling out lasso approach . . . . .	108
3.15	Confidence effect: partialling out lasso approach . . . . .	109
3.16	Heterogeneous effects in accuracy: experience and behaviour . . . . .	110
3.17	Heterogeneous effects in accuracy: demographic . . . . .	111
3.18	Heterogeneous effects in true positives: experience and behaviour . . . . .	112
3.19	Heterogeneous effects in true positives: demographic . . . . .	113
3.20	Heterogeneous effects in true negatives: experience and behaviour . . . . .	114
3.21	Heterogeneous effects in true negatives: demographic . . . . .	115
3.22	Heterogeneous effects in confidence: experience and behaviour . . . . .	116
3.23	Heterogeneous effects in confidence: demographic . . . . .	117
3.24	Overall effect - pooled treatments . . . . .	135
3.25	Confidence effect - pooled treatments . . . . .	136
3.26	Key signs of fraud . . . . .	138

3.27	Overview of scenarios . . . . .	139
3.28	Definition of variables used . . . . .	143
4.1	Main effects and robustness to alternative specifications (participation model) . . . . .	183
4.2	Main effects and robustness to alternative specifications (financial risk appetite model) . . . . .	184
4.3	Robustness to alternative matching algorithms (participation model) . . . . .	185
4.4	Robustness to alternative matching algorithms (financial risk appetite model) . . . . .	185
4.6	Total sample breakdown - heterogeneous treatment effect categories . . . . .	189
4.5	Potential confounders with proxy solutions . . . . .	190
4.7	Descriptive correlates of key dependent and independent variables. . . . .	191
4.8	Covariate balance pre and post matching procedure: standardised differences (participation model) . . . . .	192
4.9	Covariate balance pre and post matching procedure: variance ratios (participation model) . . . . .	193
4.10	Covariate balance pre and post matching procedure: standardised differences (financial risk appetite model) . . . . .	197
4.11	Covariate balance pre and post matching procedure: variance ratios (financial risk appetite model) . . . . .	198
4.12	Additional diagnostic assessment of matching procedure . . . . .	202
4.13	Definition of variables used . . . . .	203



# Chapter 1

## Introduction

The promise of the contemporary consumer financial system is enormous. It holds out the potential to expand opportunities and transform lives for the better. But the contemporary financial landscape is also complex, sometimes overwhelmingly so, and is dotted with hazards. Consumers navigating this landscape face a formidable array of choices, opportunities, risks, and pitfalls. We are confronted every day with by a steady stream of financial data, product offerings, innovations, and urgent opportunities. Behind each door lies an uncertain set of risks and rewards, potential regrets and wins. We are far from perfectly adapted to navigate this landscape smoothly.

Quite apart from the intrinsic uncertainty in our financial, economic, and personal lives which can frustrate our ability to make confident choices, we also face a wide range of cognitive and behavioural biases which can intrude on our decision-making process and lead us into error. On top of this, we are the products of our family, social, and professional networks, subject to subtle influences, privileges or disadvantages, and cultural aversions and attachments. All this before even mentioning the innumerable competing demands on our attention outside of the financial domain.

The contextual daily reality in which we actually make the choices that shape our financial lives appears sometimes comically removed from the idealised conditions of the perfectly rational, complete information harnessing utility maximising economic agent.

With these baggage and limitations, we step into the financial landscape as intrepid travelers, and we do our best to ‘do it ourselves’. We try to avoid the snakes and find the ladders. We duck, dive, and improvise. It is possible that 10,000 years of additional evolution will equip us with the tools we need to navigate this landscape

seamlessly, knowledgeably, and less haltingly, without stress, regret, or predation. More likely within a near time horizon is that our economic models, policy instruments, and consumer protection frameworks will continue to slowly adapt themselves to the contextual realities faced by the agents they are designed to represent and serve. That is, to meet people where they are, and put humans at the centre of policy design.

In the meantime, and in furtherance of this latter objective, this thesis comprises three essays that explore some of the most important fault lines, pressure points, and puzzling patterns of behaviour, and offer some practical ideas to address these challenges. They offer evidence and insight that can help us to understand and remediate, and help consumers to more fully and safely exploit the vast opportunities that a modern financial system presents. In the broadest terms, the thesis asks: why do we not seize the advantageous opportunities that lie right in front of our noses, and fall victim to traps that are set for us? And what can we do to help to move the dial towards the promise and away from the peril?

The settings are diverse: we move from the puzzle of sluggish mortgage refinancing in Ireland, to the corrosive problem of digital fraud in Northern Nigeria, and finally to the widespread opting out of the stock-market by households across Europe. The unifying theme is nonetheless clear. In each case we have a puzzle or friction in the efficient operation of consumer financial markets holding back and frustrating the full realisation of the social value that contemporary markets can provide. But similarly in each case, solutions are possible. There is scope to push back and help to untie these knots, and reclaim some of the value lost to these deadweight problems.

## **Interrupting the failure to refinance**

Chapter 2 tackles the well-documented failure to refinance puzzle, which carries considerable importance along a number of distinct dimensions. First, despite mandatory disclosures of attractive refinancing opportunities, under-refinancing costs mortgage holders in many countries a significant fraction of income annually. Second, the refinancing channel of monetary policy transmission has been shown to be significant such that frictions impeding refinancing have first-order implications for effective monetary stimulus. Third, when borrowers stand to make substantial savings on mortgage repayments from updating their terms but fail to do so, they carry an elevated debt service ratio above what would be carried in a scenario where switching was more frequent, leaving such borrowers more vulnerable to mortgage distress arising from more modest income shocks. Finally, an observably low propensity of customers to switch mortgage providers could both diminish the incentive for providers to compete on the basis of



price and send a discouraging signal to potential entrants who might bring competition to the market.

The chapter tests the impact of a series of behaviourally-informed enhancements to mandatory consumer disclosure about mortgage refinancing options in prompting greater take-up, in a field experiment with 12,000 borrowers from a larger retail bank in Ireland. While the experiment finds only modest impacts from disclosure design improvements, a simple reminder letter increases the amount of refinancing by 76%, from 8.9% to 15.7% of mortgage holders, with the average savings achieved by refinancing mortgagors being €1,209. To interpret this reminder effect, the chapter extends and estimates a mixture model of inattentive financial decision-making to allow for disclosure treatment effects. The analysis shows that a reminder decreases the likelihood mortgage holders are inattentive by 15 percentage points from 76% to 61%. The results suggest that reminders could have larger effects on household refinancing than a 200 bp rate cut and that reminders could strengthen the refinancing channel and stimulate local consumption even when policy rates are at the zero-lower bound or set in a monetary union. This chapter contributes a policy lever that has the potential to significantly improve refinancing responsiveness, which should be of interest to Central Banks seeking to strengthen the impact of interest rate reductions on the real economy, fiscal authorities seeking to stimulate refinancing and consumption, competition authorities aiming to improve the competitiveness of the mortgage market, and consumer protection authorities focused on improving households debt service burdens.

This paper also contributes to the literature on consumer inattention, and additionally contribute to work demonstrating the potential of behaviorally-informed, modernized mandatory financial disclosures. The chapter presents the first mortgage refinancing field experiment outside of the United States, and the first large-scale refinancing experiment not targeted at distressed or low-income mortgage borrowers but instead at the wider population of outstanding mortgage holders. The study is the first to show statistically and economically meaningful impact from experimental treatment arms in the domain of mortgage refinancing.

## **Pushing back against digital fraud**

Chapter 3 addresses the pernicious problem of digital fraud and its discouraging influence on the adoption of digital financial services among small businesses in low- and middle-income countries (LMIC). Non-institutional fraud targeted at small businesses is pervasive across LMICs. Non-institutional fraud causes immediate and long-term damage. Immediately, it can lead to monetary loss, but also to psychological impacts

including anger, difficulties with trust, feelings of violation, stress, and social embarrassment. Additionally, because being defrauded is a violation of trust, it may reduce trust, a key driver of macroeconomic growth. Specifically, it can frustrate economic development through use of financial services and the exploitation of advantageous opportunities. Loss of institutional trust may lead to underuse of financial services.

There is limited knowledge on what mitigation strategies can be taken to reduce fraud, effectively signal authenticity and engender trust, and thereby facilitate the better realisation of the promise of digital financial services in developing countries. This project seeks to address this gap with a lab-in-the-field experiment in Northern Nigeria with 780 small business owners, which tests the impact of a series of learning interventions in improving the ability of small business owners to accurately discern fraudulent and genuine communications, as well as in building trust in DFS. We also test the potential for a technical solution for the authentication of inbound communications to establish confidence and engagement (a ‘unique communications code’, or UCC).

While we find some evidence of positive impacts on certain adjacent outcomes, we do not find evidence that these learning interventions significantly improve discriminant ability between genuine and fraudulent communications, nor do we find evidence that our UCC authentication solution acts as a sufficiently strong signal to increase user engagement with experimental outreach. We observe significant increases in the confidence that treated users report in their judgements, notwithstanding the absence of any corresponding improvement in actual underlying accuracy. In this, we highlight the potential for false confidence effects from learning interventions, which may engender the subjective feeling of competence, and unintentionally increase the susceptibility through complacency.

This chapter contributes to the literature by being the first of its kind initiative to undertake a targeted experimental evaluation of anti-fraud learning interventions among business owners in a LMIC. A preponderance of existing literature in this area is focused on high-income country contexts, or with one-size fits all interventions for universal consumption. Our interventions, by contrast, are specifically adjusted to resonate within the information-environment and digital financial landscape faced by small business owners in a LMIC. The chapter additionally contributes to our understanding of the subtle and sometimes unintended dynamics of light-touch learning interventions, as well as the interplay between trust, confidence, and ability in respect of deception detection. Finally, the laboratory setting allows us to study anti-fraud learning interventions and fraud detection in a controlled setting.

Our failure to deliver meaningful treatment effects from learning speaks to the severity of the challenge that small business owners are likely to face in reality in successfully

navigating the noisy landscape of competing communications, and cast doubt on the utility of relatively light-touch, quick-fix learning interventions as meaningful antidotes, even when delivered in a timely fashion. Our results also highlight the risk of false confidence effects from ineffectual learning interventions which may engender the feeling but not the reality of heightened competence. As such, they can be offered as a cautionary lesson for policy in this domain, which recommends re-thinking about the type of interventions tested here, and their content.

## **Private safety nets in stock market participation**

Chapter 4 studies the impact of private networks of emergency financial support on household financial risk appetite and stock market participation in Europe, and can be seen at the juncture of two distinct strands of the household financial literature. Firstly, the long-running research focus exploring the determinants and obstacles to stock market participation, and secondly, the literature which addresses the impact of financial and economic support mechanisms on household financial behaviour and outcomes.

Research in the under-participation of households in stock markets has been motivated by a diverse set of factors from political economy: including the opportunity cost in the failure to take account of the equity premium, and its consequent implication for long-term patterns of wealth-inequality, the under-diversification of household portfolios and the implied concentration risk in labour income and real estate assets, the elasticity of intertemporal substitution, and even the political economy of financial regulation, by engendering divergent attitudes towards corporate and investment income tax policy as well as risk-sharing and redistribution among those with and without financial assets.

Various explanations have been advanced to help account for the puzzle of non-participation, including the role of demographic and socio-economic factors, personality traits, social network and community effects, institutions and trust, social norms, and informational and entry costs. Collectively, factors such as these can go a long way empirically to accounting for asymmetrical patterns of stock market participation among households. This chapter examines a previously unexplored factor: how an emergency financial bailout facility from friends or family impacts upon household financial risk appetite and stock market participation.

The chapter harnesses data from the pan-European *Household Finance and Consumption Survey*, and propensity score matching to establish comparability between households with and without such a lender of last resort. It is one of few papers to address

the question of how the option of financial support, as distinct from the actual delivery of financial support influences the household financial behaviour, and it is the first paper to address how the option of financial support relates to the likelihood of stock market participation and to risk appetite. In doing so, the paper contributes to our understanding about advantage, portfolio choice, and access to finance which play an important role in driving patterns of widening income and wealth inequality. The research helps to illuminate an important aspect of the psychology of money, namely, how the peace of mind brought about by emergency financial buffers may relate to financial risk-taking.

Consistent with a theory of insurance-induced consumption, the analysis shows that those households that enjoy the option of emergency financial support are 6% more likely to participate in the stock market, and 2% more likely to report high financial risk appetite. In illustrating a mechanism of advantage compounding advantage, the results point towards the importance of effective and sustainable policies for financial inclusion. Our results additionally highlight a potential use case for downside protection policies in replicating the insurance role provided by private safety nets in inducing stock-market participation. In so doing, the results emphasise the importance of agile and effective consumer protection regimes surrounding investment and associated add-on products, to uphold standards of clarity and transparency for potential consumers in the retail market.

# Chapter 2

## The Last Mile of Monetary Policy: Consumer Inattention, Disclosures, and the Refinancing Channel

### 2.1 Introduction

Across many countries, researchers have documented a widespread “failure to refinance,” where substantial savings available to mortgage holders through refinancing remain unclaimed.<sup>1</sup> From a macroeconomic perspective, suboptimal refinancing may significantly limit the power of the refinancing channel of monetary policy transmission (Beraja et al., 2019; Di Maggio et al., 2020; Cloyne et al., 2020).<sup>2</sup> From a microeconomic perspective, suboptimal refinancing implies many households are overpaying mortgage interest and foregoing current or future consumption as a result (Financial Conduct Authority, 2019). The modulation of monetary policy transmission by refinancing frictions is an example of what we term the “last-mile problem” of monetary policy because it inhibits the delivery of accommodative policy to the real economy.<sup>3</sup>

---

<sup>1</sup>See evidence of mortgage borrowers’ low take-up of seemingly advantageous refinancing opportunities in the US (Campbell, 2006; Keys et al., 2016; Agarwal et al., 2016; Johnson et al., 2019), Italy (Bajo and Barbi, 2018), Denmark (Andersen et al., 2020), the UK (Financial Conduct Authority, 2019), Ireland (Byrne et al., 2020), and Australia (Australian Competition and Consumer Commission, 2018).

<sup>2</sup>Below, we illustrate this transmission friction by documenting how US outstanding mortgage interest rates react to monetary policy rates slowly and incompletely relative to interest rates on new mortgages.

<sup>3</sup>In supply-chain management and telecommunications the “last-mile problem” refers to the disproportionate difficulty of delivery to a final user in contrast to the relative ease of intermediate transmission. We mention research documenting what we see as other last-mile problems for monetary policy transmission in section 2.2.1 below.

While a growing body of work finds that consumers are often inattentive financial decision makers, in this paper we demonstrate experimentally how an inexpensive and simple communication can have substantial effects supporting attentive refinancing decision-making. We analyze a field experiment conducted by a large retail bank in Ireland testing whether behavioral design changes to mandatory consumer disclosures prompt borrowers to take-up beneficial internal refinancing opportunities. We see only modest improvements from most disclosure design enhancements, consistent with overall inattention to disclosure (similar to Adams et al., 2021). However, we also find that the best-performing treatment and reminder letter combination significantly increases the probability of internal mortgage refinancing by 79%, from 8.9% to 15.7%, substantially larger than effects found by the other two mortgage refinancing trials of which we are aware. A conservative back-of-the-envelope calculation suggests that our reminder treatment generates €42 of mortgage borrower consumption for every €1 spent on disclosure reminders (€605 per refinancing household), highlighting the value of improving the last mile of monetary policy delivery.

We interpret our treatment effects through the lens of the Andersen et al. (2020) behavioral model of inattentive refinancing. After extending their framework to allow for disclosure treatment effects on attention, maximum likelihood estimates of the model imply that reminders have large attention effects, increasing the share of attentive households from 24% to 39%. We use the estimated model to contrast the partial-equilibrium relative effectiveness of cutting interest rates with sending refinancing reminders.<sup>4</sup> Holding baseline inattention fixed, we find that the best performing communication combination increases refinancing by significantly more than would be achieved by a 100 bp decrease in mortgage rates. To demonstrate the representativeness of our setting, we estimate the model without experimental treatment effects on US data and also compare to estimates from Denmark, finding similar results. Heterogeneity exercises below demonstrate that this effect is unlikely to be an artifact of the contemporaneous Covid pandemic.

These findings contribute to a nascent literature exploring the potential effectiveness of direct central bank communication. The simplicity of reminder letters suggests that they could be a cost-effective communication tool to help policy rates reach the household sector directly. Moreover, refinancing reminders also have the potential to be effective even when conventional monetary tools are de facto constrained by a lower bound or a monetary union, both of which limit the flexibility of monetary policymakers to decrease policy rates. Finally, research also finds that conventional monetary policy

---

<sup>4</sup>As we discuss in section 2.5.3 below, these partial-equilibrium arguments are unlikely to be reversed by the typical magnitude of the general equilibrium effects estimated in the mortgage literature.

is less powerful in recessions than expansions, further motivating the development of alternative accommodative tools (e.g., Tenreyro and Thwaites, 2016).

The failure to refinance puzzle continues to attract considerable academic and policy attention for at least three reasons.<sup>5</sup> First, recent work shows the refinancing channel of monetary policy transmission to be quite significant such that frictions impeding refinancing have first-order implications for effective monetary stimulus. See Amromin et al. (2020) for a review and Altavilla et al. (2020) for recent evidence of subdued pass-through of European Central Bank (ECB) policy rates to retail interest rates. Second, there are financial stability implications which potentially arise from low mortgage switching rates. Failing to realize substantial potential savings on mortgage repayments from refinancing leaves borrowers with an elevated debt service ratio and more vulnerable to mortgage distress from income shocks (Gerardi et al., 2013; Giordana and Ziegelmeyer, 2020; European Commission, 2015, 2019). Third, a low propensity of customers to switch mortgage providers could both diminish the incentive for providers to compete on the basis of price and send a discouraging signal to potential entrants who might otherwise bring competition to the market (Aghion and Bolton, 1987; Farrell and Klemperer, 2007).

To illustrate the connection between this work on pass-through to refinancing frictions and the last-mile problem, we construct the time series of average interest rates on new and outstanding mortgages in the US using CRISM data (as in Di Maggio et al., 2020). Figure 2.1 plots these interest rates against US policy rates (the effective Federal Funds Rate). All three series are highly correlated in levels. However, the figure shows that while the interest rates on *new* originations follow policy rates reasonably closely, *outstanding* interest rates move slowly and incompletely with policy rates. At the quarterly frequency, while changes in the Federal Funds Rate have a pass-through coefficient to changes in new interest rates of 0.35 with an  $R^2$  of 0.24, FFR changes have a pass-through coefficient to outstanding mortgage rates of 0.03 with an  $R^2$  of 0.05. Figure 2.2 demonstrates similar dynamics between outstanding and new mortgages in Ireland. The combination of many mortgages without indexed rates and the last-mile friction of inattention leads low policy rates to have limited impact on outstanding mortgage rates in both the US and Ireland. Figure 2.3 summarizes these arguments by contrasting the magnitude of refinancing effects from the experimental treatments with the implied effect of 100 bp lower mortgage rates holding attention fixed in Ireland, the

---

<sup>5</sup>As we discuss below, even though there are several rational reasons why a household may not refinance in the face of interest-rate savings, the evidence—including our results—suggests that behavioral factors are important in explaining this behavior. Our paper is the first to demonstrate that some of these frictions are at least partially addressable.

US, and Denmark. Given these pass-through dynamics, our experimental results highlight the potential of non-monetary interventions by policymakers, including national central banks in a currency union, to stimulate the refinancing channel.

The experiment we study is the first large-scale refinancing experiment targeted at a wider population of outstanding mortgage holders instead of distressed or low-income mortgage borrowers. To our knowledge, only two previous papers undertake field experiments in the domain of mortgage refinancing. Johnson et al. (2019) carry out a series of field experiments to encourage uptake of preapproved refinance mortgages under the US Home Affordable Refinance Program, a 2009 federal program to help underwater and near-underwater homeowners refinance their mortgages. Keys et al. (2016), among other things, test for effects of mailed notices to 193 borrowers from lower-income communities in Chicago. Among these peer efforts, our trial is the first to show statistically and economically meaningful impacts on refinancing from experimental treatment arms and to contrast these effects with conventional monetary policy effectiveness.<sup>6</sup>

The paper proceeds as follows. In section 2.2, we contextualize our contribution in the relevant literatures. Section 2.3 provides a brief overview of the Irish mortgage market, including a description of the regulation on which the experiment is based. Section 2.4 presents the experimental design and summarizes the data. Section 2.5 reports and discusses experimental treatment effects. In section 2.6, we extend the inattentive refinancing model of Andersen et al. (2020) to allow for disclosure effects, estimate treatment effects on attention, and contrast the effectiveness of targeted communication with rate changes. Section 2.7 concludes with back-of-the-envelope consumption estimates and qualifications.

## 2.2 Context in literature

To explain the relevance of our contribution to several related literatures, we first situate our findings in the literature on monetary policy transmission to the real sector. The backdrop of generally imperfect and sluggish pass-through of monetary policy to the household sector heightens the importance of last-mile studies such as ours, which document policy interventions that could successfully strengthen monetary policy transmission. Next, we summarize the literature on central bank communication and argue that our results demonstrate that personalized messaging about refinancing

---

<sup>6</sup>For example, the statistically insignificant effects of the treatments studied by Johnson et al. (2019) on applying to refinance range from -0.6 to +0.1 percentage points (pp) in contrast to our estimated refinancing effect of +6.8 pp.



opportunities is a promising policy tool. We then overview the literature on demand- and supply-side factors that impact household engagement with interest rate changes. Finally, we consider the literature on the effectiveness of reminders and consumer-facing disclosures, which shows that the design of communications matters for their capacity to prompt consumer engagement.

### 2.2.1 The last-mile problem of monetary policy

Analogous to the high marginal cost of delivering a good or service to dispersed and possibly remote end-users, the last-mile problem of monetary policy is that frictions inhibit monetary stimulus from reaching its final destination in the real economy. A range of frictions could slow and reduce the pass-through of monetary policy shocks to retail interest rates or prevent final demand by households and firms from responding to interest rate changes even when they do adjust. In this section, we briefly review recent research documenting other frictions that we see as last-mile problems for monetary policy relative to the faster pass-through of policy rates to other yields in financial markets.

For example, the key mechanism behind the investment channel of monetary policy is that changes to benchmark interest rates change a firm's cost of capital and thereby affect its real investment decisions. However, many firms do not update their investment hurdle rates when nominal rates change, thus preventing their investment decisions from responding to monetary policy (Gormsen and Huber, 2023). Other work shows that many households find it costly to acquire information on the menu of available retail interest rates on car loans and mortgages, reducing their responsiveness to policy rate changes when passthrough is uneven (Argyle et al., 2022; Kim, 2022). Similarly, households with low financial literacy are less responsive to rate changes when they fail to appreciate the connection between rate changes and their budget constraints (Blinder et al., 2022).

The degree to which monetary policy transmits to retail interest rates for households has also been explored in detail over recent decades. The importance of the refinancing channel of monetary policy has focused particular attention on mortgage rates (Calza et al., 2013; Di Maggio et al., 2020; Amromin et al., 2020; Cloyne et al., 2020). In the US, Gertler and Karadi (2015), Gilchrist et al. (2015), and Wong et al. (2019) estimate contemporaneous pass-through coefficients from policy rates to mortgage rates of 0.17-0.68. However, the literature generally documents sluggish and heterogeneous policy interest rate pass-through across retail financial products and across countries, motivating research into policies that can successfully improve pass-through (Andries and Billon, 2016). Early research for the euro area found a typical pass-through rate of

about 30% from policy rate changes to retail interest rates in the first month following a change and nearly 100% within 3-10 months, with business loan rates converging faster than loans to households (De Bondt, 2002). More recently, Altavilla et al. (2020) find sluggish and incomplete interest-rate adjustments to policy rates, with a medium-run pass-through coefficient of 0.65.

We study the last-mile problem of monetary policy within the context of the refinancing channel, with implications for why policy rate changes sluggishly filter through to household consumption and investment decisions. Relative to the literature above, our work provides experimental evidence that monetary policy pass-through to final demand is often limited because of the last-mile problem of household inattention. In so doing, we document the first intervention that significantly improved refinancing responsiveness and thus could increase the pass-through of monetary policy to the household sector.

### 2.2.2 Central Bank communications

Our paper also contributes to a developing literature on central bank communication by demonstrating the effectiveness of a new form of messaging that central banks could use to provide non-monetary stimulus. We briefly review the frontier of this literature and explain its connection to our work.

Central bank communications affect economic outcomes through routine communications such as Federal Open Market Committee (FOMC) minutes and press releases, unconventional measures such as forward guidance (McKay et al., 2016), and informal leaks (Vissing-Jørgensen, 2019).<sup>7</sup> Recent research has explored how both standard and non-standard central bank communications are ultimately processed by the general public. Evaluating the status quo, Haldane and McMahon (2018) find central banks' main communications relatively inaccessible to a wide audience in part due to linguistic complexity; typical central bank publications have reading grade levels equivalent to college-level (Hernández-Murillo et al., 2014). Lamla and Vinogradov (2019) find that FOMC announcements have little impact on consumers' perceptions and expectations of inflation or interest rates, with 65% of consumers in their data unaware of a FOMC announcement during the average announcement week. Binder (2017) and Blinder et al. (2022) offer progress reports detailing the lack of both policy and research consensus about optimal central bank communications with the general public, despite overall calls for more and better communication (e.g., Ehrmann et al., 2021). Such emerging communications strategies range from social media and music videos

---

<sup>7</sup>See Blinder et al. (2008) for a survey of the literature on traditional central bank communication, which has often focused on the roles of communication to reduce financial market volatility and influence expectations.

to listening events to reach different stakeholders and reduce the cost of acquiring and processing information. Overall, households seem to have a low desire to be informed about monetary policy and are inattentive to news linked to it except when adverse conditions arise.

Relative to this developing literature on central bank communication aimed at the general public, our paper is the first to demonstrate the promise of targeted, direct communication to households. In particular, we illustrate how the credit registries available at many central banks can be leveraged to support more attentive financial decision-making by households through customized communication. While the typical style of informational disclosure mandated by banking regulation appears relatively ineffective on its own, we show how simple reminder letters can strengthen the transmission of accommodative monetary policy to final demand. Moreover, because such a tool could be deployed without control over interest rates or interest-rate expectations, it has the potential to be effective even when monetary policy is constrained by the zero-lower bound or because rates are set in a currency union.

### 2.2.3 Barriers to refinancing

As cited above, a growing literature has documented under-refinancing in multiple countries. In the US, for example, Keys et al. (2016) find that 20% of households for whom refinancing would be advantageous have not refinanced, with a foregone annual savings of almost \$2,000. Both supply- and demand-side factors can inhibit refinancing and thus the final delivery of monetary policy to the household sector—see Amromin et al. (2020) for a review. Supply-side barriers to refinancing can result from underwriting constraints that are binding for households most acutely during a recession (Beraja et al., 2019; Di Maggio et al., 2020; DeFusco and Mondragon, 2020).

On the demand side, inattention, low financial literacy, present bias, and distrust have each been implicated in sluggish refinancing.<sup>8</sup> Andersen et al. (2020) show that Danish borrowers with lower education, income, housing wealth, and financial wealth are less likely to consider refinancing and less attentive to the financial incentive to do so. Similarly, Bajo and Barbi (2018) find that take-up of an attractive no-cost refinancing program in Italy was positively correlated with loan size, remaining term, age, wealth, experience with financial products, and financial literacy. Johnson et al. (2019) document the role played by borrowers' trust in financial institutions and their present bias, which discourages households from incurring time costs today for lower interest

---

<sup>8</sup>The pattern of inertia and disengagement in mortgage markets echoes many other product markets, including retirement accounts, deposit accounts, energy, telephone, and internet broadband markets, which also have subdued levels of consumer switching against a market backdrop of meaningful price dispersion (Madrian and Shea, 2001; Yang, 2014; Yin and Matthews, 2016; Shcherbakov, 2016; Lunn and Lyons, 2018; Harold et al., 2019; Adams et al., 2021).

payments realized in the future. Agarwal et al. (2016) find a large role for financial sophistication in explaining suboptimal refinancing. Consistent with our findings, Keys et al. (2016) find that two-thirds of low-income survey respondents who received refinancing offers did not open the letters or “planned to call the loan officer but did not get around to it or were simply too busy to make the phone call.”

Building on the literature studying barriers to refinancing, our contribution is demonstrating that inattention in financial decision-making is partially addressable through direct communication to households. In contrast to the prior literature that generally documents predictors of less responsive refinancing without finding evidence of effective or actionable prescriptions, we provide experimental evidence that the refinancing channel of monetary policy can be strengthened. The small effects of disclosure design improvements and the strong effects of a simple reminder letter suggest that inattention in the form of absentmindedness or procrastination is a significant impediment to refinancing (Schacter, 1999). Our estimates imply that reducing monetary policy’s last-mile problem caused by the final demand sector’s inattention can result in cost-effective stimulus. We discuss the literature exploring potential general equilibrium implications of increased refinancing responsiveness in section 2.5.3.

#### **2.2.4 Consumer disclosures and reminders**

In many settings, the policy response to potentially suboptimal consumer choice has been to provide additional information, leading to a proliferation in mandatory disclosures and plausibly contributing to inattention (Ben-Shahar and Schneider, 2014; Kell, 2016). Our results contribute to a growing body of evidence that behaviorally-informed disclosures can—but do not always—deliver meaningful impacts on various public policy challenges (e.g., Lee and Hogarth, 2000; Bar-Gill, 2012; Adams et al., 2015; Wang and Burke, 2022). One approach to address inattention to information provision, tested by a growing number of studies, is to use personalized services to improve household engagement or the take-up of publicly provided services (e.g., Bergman et al., 2019; Finkelstein and Notowidigdo, 2019; Guryan et al., 2023; Evans et al., 2023). One advantage of the approach we test is that, while personalized, the reminder treatments we study are low-cost and readily scalable because they do not require personally provided services.

Reminders have a significant history in health sciences, with evidence that reminders can increase vaccination take up and cancer check-up rates (e.g., Mayer et al., 2000; Hirani, 2021; Milkman et al., 2022). In consumer finance, research finds some scope for well-timed reminder letters positively affecting financial behavior (Adams and Hunt, 2013; Adams et al., 2015, 2021; Karlan et al., 2016). These studies demonstrate the

significant effect that well-timed reminders can play in prompting financial action, pointing to the role of procrastination and inattention in shaping consequential financial decisions.<sup>9</sup> Nevertheless, looking across settings, the effects of reminders are mixed, with reminders in domains with straightforward choice architecture and salient benefits more likely to have significant effects.

Our work extends the literatures on behavioral interventions, disclosures, and reminders in several ways. First, despite mixed success of reminders in other settings, this paper is the first to test their use in the high-stakes mortgage refinancing context. Second, the reminder effects in our setting—including the contrast to the small effects of disclosure design—add the first experimental evidence indicating a role for limited attention in explaining both the failure to refinance and the typical ineffectiveness of disclosures more broadly. Third, we show that reminder letters can be a cost-effective tool for policymakers such as central banks that had previously not considered such communication in their toolkits.

## 2.3 Institutional setting

To provide context for our experimental setting, this section briefly overviews the Irish mortgage market and its relevant institutional features. By several metrics, the Irish mortgage market is fairly representative of mortgage markets in other advanced economies (Calza et al., 2013). Nearly one in three Irish households has an outstanding mortgage on their main residence (Central Statistics Office, 2020). Ireland’s mortgage debt-to-GDP ratio (50%), typical loan-to-value ratio (70%), typical mortgage duration (20 years), and score on the IMF Mortgage Market Index measuring market development (0.39) are all roughly average compared with mortgage markets in the US, Canada, Europe, Japan, Australia, and New Zealand.

There are three primary types of residential mortgages in Ireland: fixed-rate mortgages, variable-rate mortgages, and tracker mortgages, accounting for approximately 55%, 20%, and 25% of 2021 outstanding balances respectively.<sup>10</sup> Fixed-rate mortgages in Ireland are similar to those in the UK and to adjustable-rate mortgages in the US: a fixed interest rate for an initial term (usually 1-5 years) that converts to a variable rate thereafter. There is generally a prepayment penalty of approximately 2% of the outstanding balance if a borrower prepays their mortgage during the fixed-rate period. Variable-rate mortgage interest rates in Ireland adjust at the sole discretion of

<sup>9</sup>See also Banerjee et al. (2010), who find that small nudges can have large effects by helping people coordinate their attention to a task at a specific time instead of delaying it indefinitely.

<sup>10</sup>Appendix Figure 2.7 provides a time series of this breakdown, highlighting the growing prominence of fixed rate mortgages in recent years.

the lender (as opposed to floating debt in other markets that is usually indexed to an interest-rate benchmark). There is no penalty for prepayment of a variable-rate mortgage and refinancing internally (i.e., with the current lender) is allowed without a fee, unless the borrower wishes to pay for an appraisal to justify a lower loan-to-value ratio classification. Tracker mortgages in Ireland usually track the ECB refinancing rate plus a spread of around 100 bp. However, Irish lenders stopped originating new tracker mortgages in 2008, and their share of outstanding mortgages has steadily declined since then (Appendix Figure 2.7). In March 2020, the Irish government instituted a mortgage forbearance program, through which mortgage borrowers who had experienced financial hardship due to the Covid pandemic could apply for a temporary payment break. We observe and exploit heterogeneity in who received Covid payment forbearance in section 2.5.2.

Refinancing in Ireland generally maintains the original maturity at origination and does not extend a mortgage's term. Similar to other mortgage markets, subdued refinancing activity in the Irish mortgage market contrasts with widespread opportunities to realize substantial financial savings from refinancing and with policy and commercial advertising efforts to facilitate refinancing. Byrne et al. (2020) estimate that three in every five mortgages could save over €1,000 within a year of refinancing (over €10,000 over their remaining term) but that just 2.9% of mortgages switched provider during the second half of 2019.<sup>11</sup> As we discuss below, this level of potential savings is representative of the savings available to borrowers in the experimental sample. A 2016 survey found that while most surveyed borrowers would consider refinancing for interest-rate savings, over half were uncertain how much money they could save and many borrowers believed that the refinancing process would be too complex or time-consuming (Central Bank of Ireland, 2017b).

To provide mortgage borrowers with enough information to refinance their mortgages when advantageous to do so, Irish banking regulations require mortgage lenders to disclose the availability of such opportunities at least annually (Central Bank of Ireland, 2017a). Provision 6.5(g) of the Central Bank of Ireland's Consumer Protection Code requires regulated lenders to provide variable-rate mortgage holders a letter summarizing other mortgage products that could provide them with savings on their mortgage at that point in time. Because fixed-rate mortgages automatically convert to variable-rate mortgages after their fixation periods end, many borrowers are on variable rates at any given time and qualify for these disclosures. Notably, the regulations do not

---

<sup>11</sup>Appendix Figure 2.8 plots mortgage external refinancing rates over time, showing that there was no significant increase in the level of external refinancing during the pandemic in 2020. During the period of the experiment, external refinancing options included both mortgages with more attractive interest rates and mortgages with less attractive interest rates but with upfront, and highly advertised, cash-back bonuses (King et al., 2018).

stipulate how the disclosed information should be presented. As we discuss below, these mandatory annual disclosures form the basis of our field experiment.

## 2.4 Experimental design

We partnered with a large retail bank in the Irish mortgage market to analyze the results of a field experiment testing whether a series of behaviorally enhanced versions of the mandatory financial disclosures described above have an effect on refinancing behavior. The letters were delivered by mail to a total of 12,050 variable-rate mortgage holders between January 28 and February 3, 2020, randomly drawn from the population of variable-rate mortgage customers of the partnering bank. Figure 2.2 illustrates the interest-rate savings available to the average bank customer receiving the disclosure, from an average outstanding interest rate of 4.2% to the shortest fixed-rate mortgage offered on the disclosure of 2.9%.

Participants were randomly assigned to seven equally sized groups: six treatment arms and one control group. The modified version of the disclosure letter sent to each treatment arm was rigorously evaluated to ensure it provided at least the baseline information required by the Consumer Protection Code (i.e., no key information was removed, which could have led to a treated mortgage borrower having less information available than they would under the baseline scenario).<sup>12</sup>

Power analysis indicates that with a sample size of approximately 1,700 customers per group, the minimum detectable treatment effect on mortgage refinancing is a 1.6 pp improvement over the baseline rate of refinancing, an increase of 13%. Within each treatment arm, the sample was randomly divided in half, with one half receiving an additional follow-up reminder notification by mail between February 27 and March 6, 2020 (4-6 weeks after the original communication).

We gathered detailed baseline data on each trial participant from prior to the intervention and assess the impact of the intervention using data snapshots provided by the bank four and seven months after the disclosure distribution. The loan-level dataset recorded loan and borrower characteristics, including the interest rate prevailing on the loan, years to maturity, outstanding loan balance, current loan-to-value ratio, pre-trial available savings on the mortgage with respect to the best available alternative product option, borrower age, and indicators for Dublin residents, first-time homebuyers, and

---

<sup>12</sup>To avoid the potential for observer effects that could affect the integrity of the experimental design, the bank did not inform treatment-group participants that the version of the mandatory disclosure they received was experimental.

borrowers who had received Covid-related payment forbearance. Follow-up administrative data from the bank allow us to identify those loans that refinanced internally, refinanced externally, reached maturity, or otherwise exited the bank. We drop around 300 borrowers in arrears from the final estimation sample (less than 3%) to remove borrowers who may have perceived themselves ineligible for refinancing, although our results are robust to including them.

### 2.4.1 Treatment arms

The treatment arms' disclosure redesigns addressed informational, procedural, financial literacy, and behavioral obstacles to refinancing and showed promise in the encouragement of consumer engagement in other settings. Below, we explain the disclosure design elements featured by various treatment arms and provide citations for their rationale. Table 2.1 summarizes how each treatment arm incorporates a particular combination of the disclosure redesign elements described below. Appendix Figures 2.9 and 2.10 reproduce example control-group and treatment-group letters, respectively.

***Simplification:*** Each treatment communication included a box on the front page of the letter with key points highlighted, including the current interest rate and monthly repayment and the lowest alternative interest rate and associated monthly repayment available to the customer from refinancing internally. The box was designed to ensure that key information could be accessed quickly to target customer inattention and information overload, both of which have been found to affect the ability of consumers to make informed choices (e.g., Adams and Hunt, 2013; Lunn et al., 2016).

***Personalized Savings:*** The retail bank's standard communication (Appendix Figure 2.9) included a table of the interest rate associated with each alternative product option available to the customer, but there was no translation into the associated monthly savings. Each treatment supplemented the interest-rate table with the monthly payment amount associated with each option and the savings (where available) relative to the current monthly repayment. This personalization was designed to target financial illiteracy and present bias by making the immediate savings more salient (Financial Conduct Authority, 2016).

***Prominent Subject Line:*** The subject line in the control letter stated, "You may be able to save money on your mortgage." To increase the likelihood consumers would perceive the letter worth their attention, Treatment groups 3-6 tested the use of color, increased font size, and emboldened the subject-line text (Behavioural Insights Team, 2014).

***Framing:*** Three of the treatment arms varied the presentation of financial savings to be either in a gain frame or a loss frame to counteract loss aversion, the tendency



for people to prefer avoiding losses to acquiring equivalent gains (Genakos et al., 2015; Adams et al., 2015). For the one gain-framed treatment arm (#4), the language read “With a different rate, you could save up to €X a year on your mortgage,” and in the two loss-framed treatment arms (#5-6), the language used was “You could be missing out on savings of up to €X a year by not choosing a lower mortgage interest rate.” The remaining treatment arms’ letters adopted a more neutral tone.

**Color:** Treatment group 2 received letters that used color at key junctures to draw attention to salient information (Behavioural Insights Team, 2014).

**Process Clarification:** To reduce potential ambiguity aversion over the potentially unknown complexity of the refinancing process, Treatment group 6 included a clarified process box, which clearly delineated the steps required for a mortgage holder to take action and move onto a lower cost interest rate option (Behavioural Insights Team, 2014). Testing whether a series of disclosure redesigns had any effect on savings account holders switching to a higher-paying savings product, Adams et al. (2021) found the strongest treatment effects by simplifying the switching process.

**Reminder:** As reviewed above, reminders can effectively target customer inattention, procrastination, and forgetfulness (Adams et al., 2015). Half of each of the six treatment arms received a follow-up reminder letter 4-6 weeks after the initial treatment disclosure. See Appendix Figure 2.11 for an example.

## 2.4.2 Descriptive statistics

The trial sample of loans consists of a random subsample of the outstanding variable rate mortgages held by the partnering institution because this is the cohort eligible for receipt of the mandatory disclosure from which we build our experimental treatment arms. Our total sample of 12,050 letter recipients reduces to an estimation sample of 11,200 following the attrition of 850 observations which exited the loan book, reached maturity during the trial period, or were excluded from estimation due to mortgage arrears history. Of our estimation sample, 1,345 (12%) go on to refinance internally, and 373 (3%) refinance externally with a different provider.

Table 2.2 reports summary statistics for several mortgage and borrower characteristics in our data, pooling treatment arms together with the same reminder status.<sup>13</sup> Around 20% of borrowers in our sample live in Dublin and 40% are currently living in their first purchased home. The average borrower in our data is 50 years old and has 13 years left on their €83,000 mortgage. At baseline, the average interest rate among trial borrowers is 4.2% with a standard deviation of 20-30 basis points. Calculating

<sup>13</sup>See Appendix Table 2.8 for descriptive statistics by treatment arm.

the interest savings that borrowers could realize in the first year if they refinanced to the shortest fixed-rate mortgage available (2.9%), the average savings is €1,044 in the control group and similar for the treatment groups.

To formally test for random assignment and the balance of treatment status across observables, Appendix Table 2.9 shows a regression of an indicator for each treatment arm (in a sample restricted to observations from that treatment arm and the control group) on a vector of covariates. We find a high degree of statistical balance and low  $R^2$  values ranging from 0.001-0.005. Although years to maturity shows marginally statistically significant differences across treatment and control, a coefficient of 0.003 years corresponds to an economically meaningless difference in remaining maturity of approximately one day. In every column, we fail to reject that all of the slope coefficients are jointly equal to zero.

**External Validity** To assess the representativeness of our experimental sample, we compare its summary statistics to the near-universe of outstanding Irish mortgages. Columns 4-5 of Table 2.2 report summary statistics representing about 90% of outstanding mortgages in Ireland using a database updated every six months by the largest banks in Ireland. Column 4 describes outstanding variable rate mortgages and column 5 describes residential mortgages in Ireland. Overall, the mortgages held by our partnering bank have similar characteristics to the nationwide sample, helping to address external validity concerns. Exceptions include that mortgages in the experimental sample are 7-8 pp less likely to be in Dublin and have €20,000 smaller balances on average.

The average variable-rate mortgage at the partner bank has a 0.6 pp higher interest rate than the average variable-rate mortgage across all providers, consistent with Figure 2.2. A priori, the high potential interest-rate savings of our sample has theoretically ambiguous effects on the effectiveness of the treatments we study. On the one hand, borrowers with high savings might be the most responsive to refinancing nudges. On the other hand, borrowers who have not yet refinanced in the face of high available savings may be particularly inattentive or constrained for unobservable reasons and therefore the least responsive to the treatments. However, when compared to the sample of all outstanding variable-rate mortgages in column 4 of Table 2.2, the amount of savings available to the experimental sample looks typical, suggesting that our experimental sample is not particularly unique within Ireland.<sup>14</sup> Moreover, the

---

<sup>14</sup>In the near-population of outstanding mortgages in column 5, the average mortgage has no savings available from refinancing, mostly owing to tracker mortgages pegged to the ECB policy rate that were significantly below market-available refinancing rates at the time. Excluding tracker mortgages, the average outstanding mortgage in Ireland has €834 of first-year interest payment savings available.

research cited in section 2.2.3 also demonstrates that mortgage borrowers worldwide are frequently found to be ignoring in-the-money refinancing options. Finally, the lack of substantial treatment-effect heterogeneity found in section 2.5.1—especially for the reminder treatment effects—further indicates that our results do not seem to be driven by some particular, unique characteristic of our sample.

## 2.5 Results

In this section, we estimate the causal effects of the experimental treatments on the observed rate of internal mortgage refinancing compared to the baseline standard represented by the control group. Our impact analysis is based on administrative bank data captured in June 2020, four months after the distribution of disclosure letters, although our results are robust to using outcomes measured as of September 2020. After measuring treatment effects and testing for heterogeneity in effectiveness across treatment arms, we test for differences in treatment effects across observable demographic and financial differences in borrowers. We also examine whether our results are likely to be an artifact of the extra time and motivation to refinance some borrowers may have had during the Covid pandemic. Finally, we conclude this section by discussing general equilibrium considerations in the degree to which our findings would potentially change in a large-scale implementation.

Figure 2.3 summarizes our core results, which we examine in more detail below. We find that without a follow-up reminder letter, the average disclosure redesign treatment increases refinancing 20% (1.8 pp) from a base of 8.9 pp (the control-group internal refinancing rate) to 10.7 pp. An accompanying follow-up reminder letter 4-6 weeks after the initial disclosure increases the refinancing rate by an additional 3.6 pp for a total communication effect for the average treatment arm of 5.6 pp. The best-performing treatment arm with a reminder letter (V2) increases internal refinancing by a total of 76% (6.8 pp). The average 12-month savings realized by refinancing mortgage borrowers in our data is €1,209.

These results—particularly the large proportional effects of treatments with reminders—contrast with much smaller effects found in two preceding mortgage refinancing experiments. Keys et al. (2016) found no statistical differences in refinancing across three treatment arms that drew attention to the amount of savings that mortgage holders could achieve in different ways.<sup>15</sup> Similarly, Johnson et al. (2019) found that none of the

---

<sup>15</sup>However, a much smaller sample size (N=193) meant that the authors were underpowered and unable to reject the possibility of economically meaningful results.

experimental interventions they tested had a statistically or economically meaningful impact on refinancing take-up rates.

Table 2.3 reports the magnitude and formally tests for the significance of these treatment effects, with and without controls. We estimate

$$Refinance_i = \beta_0 + \beta_1 Treatment_i + \beta_2 Treatment_i \times Reminder_i + X_i' \gamma + \varepsilon_i, \quad (2.1)$$

The indicator  $Refinance_i$  equals one if borrower  $i$  internally refinanced within four months of receiving the legally required refinancing opportunities disclosure. Without controls,  $\beta_0$  estimates the refinancing rate of the control group. The treatment effects  $\beta_1$  and  $\beta_2$  capture the increase in refinancing by borrowers randomly assigned to a redesigned disclosure treatment arm. The coefficient  $\beta_2$  on the interaction term estimates the differential refinancing by treated borrowers who also received the follow-up reminder letter (no borrowers received the reminder letter without also receiving a redesigned disclosure). In specifications that include them to improve precision, the individual- and loan-level controls  $X_i$  are the covariates listed in Table 2.2. All of our estimates use heteroskedasticity-robust standard errors.

Across all columns in Table 2.3, effects are statistically significant at at least a 95% confidence level and most often at the 99% confidence level. Column 1 reports a treatment effect on refinancing pooling all six treatment arms of 3.6 pp without conditioning on reminder status, thus conflating the effects of the treatment redesigns alone and the reminders. Adding the borrower and loan controls from Table 2.2 as covariates in column 2 increases this estimate slightly. In column 3, we additionally control for whether each borrower also received a follow-up reminder communication. The main pooled treatment effect decreases to 1.8 pp, indicating that mortgage borrowers who received a redesigned disclosure but *not* a reminder were only 1.8 pp more likely to internally refinance than the control group, of which 8.9% refinanced internally (the constant term when there are no other controls). Adding the pooled reminder effect of 3.6 pp to the pooled treatment effect of 1.8 pp yields a total refinancing effect of the average treatment and reminder of 5.4 pp. Column 4 again adds individual-level controls, with the treatment and reminder effects changing only slightly.

Given the modest overall effectiveness of the treatment redesigns without an accompanying reminder, is any one of the treatments particularly effective? Is a reminder letter more effective when combined with some treatments than others? To test for the relative performance of the various redesigns elements, we next estimate treatment effects separately by treatment arm. We first plot internal refinancing rates by treatment arm for subsamples without and with reminders in Figures 2.4 and 2.5, respectively. In both

figures, there are economically small differences in refinancing rates across treatment arms. In the no-reminder sample of Figure 2.4, only some treatments have refinancing rates that are individually significantly different from the control group. Moreover, across treatment arms, the sizes of disclosure redesign effects and reminder effects seem negatively correlated; the strongest treatment effects without reminders are not from the treatment arms with the strongest reminder effects.

Table 2.4 reports the magnitude of the refinancing treatment effect differences across treatment arms, including reporting a formal test of joint equality across treatment arms. Column 1 pools treatments and reminders and finds small differences across treatment arms. The  $p$ -value of 0.99 for a joint  $F$ -test of the null hypothesis that all of the treatment arm coefficients are equal to each other fails to reject that all of the treatments had the same effect on refinancing. Adding controls in column 2 increases the estimates slightly compared with column 1, with the joint test again failing to reject equality of the effects across treatment arms. In column 3, we add controls for the interaction of each treatment arm indicator with the reminder indicator for whether each mortgage borrower received a reminder letter; column 4 adds controls. In both columns 3 and 4, only some treatment arms have treatment effects or reminder effects that are individually statistically different from zero. As was apparent in Figures 2.4 and 2.5, the treatment arms with the largest and most statistically significant treatment effects are not the treatment arms with the largest or most statistically significant reminder effects. However, consistent with Figure 2.5, the *total* treatment and reminder effect is statistically significant for each treatment arm. Testing for equality of the treatment without reminder effects across treatment arms and the equality of the reminder effects across treatment arms in columns 3 and 4, we again fail to find statistically significant evidence of differential effects across disclosure redesign versions.

The results above indicate that the reminder communication had particularly strong effects stimulating internal refinancing. Do redesigned disclosures or the reminder letters affect *external* refinancing, where borrowers refinance and switch providers? We evaluate whether any treatment effect can be observed in terms of this secondary channel in Table 2.10. In a series of regression specifications that mirror our main regression analysis in Table 2.3, we find no evidence for treatment or reminder effects, estimating economically small, precise, and statistically insignificant effects on external refinancing. This contrast between internal and external refinancing effects could be one reason for the strength of our estimated treatment effects. Whereas most other studies focus on the drivers of external refinancing and document a reticence among borrowers to switch providers, our results suggest that inertia is much weaker when

borrowers have the opportunity to refinance while staying with their current provider.<sup>16</sup> Policymakers seeking to support active refinancing could consider efforts to facilitate internal refinancing, with the caveat that success improving refinancing responsiveness could be partially offset in general equilibrium (see discussion in section 2.5.3).

### 2.5.1 Treatment effect heterogeneity

Next, we test whether some types of borrowers responded to the treatments more strongly than others. This exercise tests whether the overall effectiveness of the communication treatments is driven by strong effects for a particular subset of borrowers, is an input into questions of cross-subsidization (Fisher et al., 2022; Zhang, 2022), informs external validity assessments, and guides estimates of welfare effects (Finkelstein and Notowidigdo, 2019).<sup>17</sup> We estimate heterogeneous treatment effects by augmenting (2.1) with interaction terms between the treatment variables and a given borrower or loan characteristic  $x$ :

$$\begin{aligned} Refinance_i = & \beta_0 + \beta_1 Treatment_i + \beta_2 Treatment_i \times Reminder_i + \beta_3 x_i \\ & + \beta_4 Treatment_i \times x_i + \beta_5 Treatment_i \times Reminder_i \times x_i + \varepsilon_i. \end{aligned} \quad (2.2)$$

To ease interpretation, we discretize the controls in Table 2.2 into indicator variables. For example, instead of the variable age measured in years, we calculate an indicator for age greater than 50, which is the mean age in our sample. When  $x = 1(\text{Age} > 50)$ ,  $\beta_0$  and  $\beta_3$  correspond to the refinancing rate of younger and older control-group borrowers, respectively, and  $\beta_1$  and  $\beta_2$  correspond, respectively, to the disclosure redesign and reminder treatment effects for younger borrowers. The interaction terms coefficients  $\beta_4$  and  $\beta_5$  measure the differential treatment effects for older borrowers, with the  $t$ -test on each of these coefficients testing the hypothesis of no heterogeneity in treatment effects along the age dimension.

Table 2.5 estimates (2.2), with each column reporting results of a separate regression with a different characteristic standing in as  $x$ , as indicated by each column header.<sup>18</sup> Panel I estimates heterogeneous treatment effects along borrower characteristics (indicators for Dublin residence, age over 50, first-time homebuyers, and borrowers with

<sup>16</sup>We note that despite also studying internal refinancing opportunities with their small-scale field experiment, Keys et al. (2016) found insignificant effects, as discussed above.

<sup>17</sup>Recent work by Gerardi et al. (2023) documents disparities between groups in their responsiveness to refinancing incentives. In this section, we test whether such disparities are compounded by potential differential responsiveness to disclosures and reminders.

<sup>18</sup>To test whether we can reject that all of the treatment-covariate interaction terms are jointly equal to zero, we also estimate a version of (2.2) controlling for all of the treatment-covariate interaction terms and covariate main effects simultaneously. However, because of the conceptual similarity between some of these variables (e.g., baseline interest rates and potential savings), interpreting the jointly-estimated coefficients is less intuitive than the one-at-a-time version presented in Table 2.5.

Covid forbearance), and panel II tests for heterogeneity by loan characteristics (outstanding loan balances above €75,000, baseline interest rates above 4.2%, more than 13 years left in the mortgage, and first-year potential refinancing savings exceeding €1,000). As expected given the literature reviewed in section 2.2.3, we find evidence that baseline refinancing rates differ across consumer types. Examining estimates of  $\beta_3$  in the rows labeled Covariate  $x$ , control-group borrowers who are younger, with high balances, have longer until loan maturity, or stand to save more after refinancing have higher incentives to refinance and are 3, 9, 5, and 9 percentage points more likely to refinance, respectively.

Turning to treatment effect heterogeneity, overall we find modest but statistically insignificant heterogeneity in the disclosure redesign’s effectiveness and small and insignificant heterogeneity in the reminder’s effectiveness.<sup>19</sup> While the estimated main effect  $\hat{\beta}_1$  of the disclosure redesigns continues to be small, several borrower types have disclosure redesign treatment effects estimated to be more than 2 percentage points higher. Unfortunately, these tests also appear to be underpowered, and we cannot reject that a Treatment  $\times x$  coefficient of, for example, 2.5 pp is statistically different from zero. The exception is for years to maturity, where we estimate that the disclosure redesign treatments had a de minimis effect for borrowers with 13 or fewer years until maturity and a statistically significant 4.6 pp treatment effect for borrowers with more than 13 years until maturity. The pattern for reminder effects is more uniform. While the estimated reminder effect  $\hat{\beta}_1 + \hat{\beta}_2$  is consistently large and statistically significant across columns, we estimate the reminder treatment effect heterogeneity  $\beta_5$  to be economically small in magnitude across all characteristics and in every case less than 2 pp.

Because of power and multiple hypothesis testing concerns, we hesitate to draw strong conclusions from the generally marginal evidence in Table 2.5 for treatment effect heterogeneity. Moreover, data limitations prevent us from testing for heterogeneity by income, race, or financial sophistication. However, the contrast between a) significant heterogeneity in baseline refinancing rates along the dimensions we observe, b) the imprecise and modest disclosure redesign effect heterogeneity, and c) the much smaller heterogeneity in reminder effects suggests that reminder communication stimulates refinancing for a majority of borrowers and that such messages do not favor one group over another. Furthermore, the similarity in refinancing effects supports applying a representative mortgage borrower MPC when estimating aggregate mortgage borrower consumption effects, as opposed to a lower (higher) MPC if primarily unconstrained (constrained) borrowers responded to reminders.

<sup>19</sup>Joint tests that all of the Treatment  $\times x$  or Treatment  $\times$  Reminder  $\times x$  interaction terms are simultaneously equal to zero fail to reject at the 0.05 significance level.

## 2.5.2 Assessing potential Covid impacts

Finally, we address the possibility that our treatment effects are driven by borrowers with an atypical surplus of time or refinancing motivation due to Covid lockdown measures during our estimation window. Such circumstances might plausibly facilitate, for certain households, a greater degree of attention to administrative mail communications than they would ordinarily allocate. The potential for our results to be driven by the unique circumstances induced by the pandemic is particularly important from an external validity standpoint. If our results are driven by something about the specific state of the world at the time of our experiment then that lessens the likelihood that direct communication to households about refinancing activities in other times or in other markets would also have strong effects.

We have three main strategies to test whether our results are driven by pandemic effects. First, our use of a randomized controlled trial helps us not misattribute contemporaneous time-series variation in refinancing to the experimental treatments. If refinancing during Covid were extraordinarily high or low for reasons unrelated to our treatments, such aggregate effects would impact the control group as well as our treatment groups. Because both groups would be affected, our treatment effects estimates of the differential refinancing by treated borrowers relative to control-group borrowers would not be biased. Two findings we discuss above suggest that this channel is not a large concern. First, we do not observe a significant increase in external refinancing rates in Ireland from 2019-2020 (Appendix Figure 2.8). Second, control-group borrowers with and without Covid forbearance had similar internal refinancing rates ( $t$ -statistic of 0.6 in Table 2.5). Similarly, some types of borrowers may be more likely to respond to the pandemic by refinancing than other types of borrowers. However, the large sample and balance of both observable and unobservable borrower types across treatment and control ensured by randomization—including the balance of the Covid forbearance flag—prohibits such heterogeneity from affecting our results. In specifications with covariate controls, we control for the Covid forbearance flag and find it to have a positive but relatively small and statistically insignificant coefficient.

Second, we address potential interplay between the pandemic and the treatments by interacting our treatment indicator with the Covid forbearance flag. The concern addressed here is that the use of a randomized control-group does not help if the treatment effects themselves are driven by Covid. For example, imagine hypothetically that the reminder treatment is only effective because a subset of treated borrowers had ample time to respond actively to a reminder follow-up (or because they were distracted with background stress and were thus particularly in need of a reminder). Such a mechanism could lead to our estimating large reminder effects that would be unlikely to



replicate in other settings where the Covid mechanism would not be present. Around 9% of borrowers in our data received mortgage payment forbearance by documenting a Covid-induced financial hardship limiting their ability to make their mortgage payments. If the strong reminder effects are only because of Covid, then the most Covid-affected borrowers in our data should show the strongest treatment effects. In Table 2.5, we indeed find that the Treatment  $\times$  Covid Forbearance coefficient is positive, suggesting that the treatment may be more effective for people who have time or particular motivation to seek payment savings. However, as in the case of the Covid Forbearance main effect, the Treatment  $\times$  Covid Forbearance interaction term coefficient is modest in magnitude and statistically insignificant, with a  $t$ -statistic of 0.5-0.6. We conclude that strength of the reminder effects is not driven by any heightened responsiveness by the most Covid-affected borrowers.

Finally, we test for other forms of differential treatment effects due to Covid that might not line up with whether a borrower received forbearance. Even among workers without sufficient Covid-related financial distress to qualify for forbearance, the pandemic might have been a much busier time or a time with much more slack depending on a borrower's employment situation. To test for Covid-related drivers of our treatment effects that are not captured by the forbearance flag, we estimate specifications that allow for heterogeneous effects across employment sectors. This approach allows us to examine whether the treatment effects are driven by an interaction between the treatments and pandemic-specific employment situations that would be unlikely to be present in future implementations of the communication treatments we study here.

We group mortgage borrowers into employment sectors that differ in the likelihood they face Covid-related disruptions using data on the employment industry of each borrower. Beginning in June 2020, the Central Bank of Ireland's administrative loan-level data contains a field collected at origination by lenders recording a borrower's employment industry using Eurostat's Statistical Classification of Economic Activities in the European Community. For loans in the June 2020 data, we merge their employment industry to our estimation sample using unique loan identifiers. Of our original 11,200 observations in our estimation sample, we obtain employment sector information on 10,260 loans, a 92% match rate. We then split our estimation sample into employment sectors that are more likely to be working from home (WFH), experiencing business as usual (BAU), and at home but not working (AHNW), using information on the borrower sector of employment at point of loan origination.<sup>20</sup>

---

<sup>20</sup>See Appendix Table 2.11 for this classification of employment industries. The WFH category includes those employed in industries more likely to have flexibility to work from home: information and communication, financial and insurance, professional, scientific and technical, public administration, and other service activities sectors. The BAU category includes those industries likely to have

Figure 2.6 summarizes this subgroup analysis by plotting the refinancing rates of each employment sector by treatment status. The left-hand side of the figure reports refinancing rates for borrowers that were treated with a redesigned disclosure letter but not a follow-up reminder letter. The right-hand side reports refinancing rates for borrowers that received both redesigned disclosure letters and follow-up reminder letters. If treated borrowers have similar refinancing rates across employment sectors that have very different levels of exposure to pandemic disruptions, this suggests that the treatment is not particularly effective or ineffective *because* of the pandemic. With the caveat that we are underpowered for this heterogeneity analysis—with wide confidence intervals especially for borrowers in the relatively small business-as-usual sector—refinancing rates are quite similar across employment sectors, inconsistent with our results being driven by the uniqueness of life during the pandemic.<sup>21</sup>

### 2.5.3 General equilibrium considerations

An important and complementary series of recent papers explores how equilibrium mortgage rates might change if inertial refinancing were reduced at scale.<sup>22</sup> The core idea behind this line of inquiry is that prevailing pricing anticipates status quo refinancing behavior. If refinancing were to become more responsive for a meaningful fraction of borrowers, in equilibrium firms may raise their interest rates to recoup lost interest income from formerly sluggish refinancers. We also note that the resulting competitive forces may alternatively motivate firms to lower their interest rates—either at origination or by offering internal refinancing opportunities—to eliminate the incentive to switch providers. Similarly, the economic efficiency of monetary policy might improve if origination interest rates were higher but more closely passed through policy rate changes. These papers generally consider hypothetical successful and yet-to-be demonstrated policy interventions. One contribution of our paper is providing such a policy that we show—experimentally in the field and in a partial equilibrium attention model—can stimulate refinancing.

---

continued requiring in-person work: agriculture, forestry, and fishing, electricity and gas supply, and transport and storage sectors. The AHNW category includes those more likely to be laid off or furloughed, who were employed in manufacturing, construction, wholesale and retail trade, and vehicle repair, and accommodation and food services sectors.

<sup>21</sup>Appendix Table 2.12 reports corresponding treatment and reminder effects by subgroup relative to the control group. We fail to reject equality of the treatment effects across employment sectors, although the standard errors are large for this subgroup analysis. Moreover, although the business-as-usual reminder effect is relatively large, it is also the most imprecise and has the noisiest constant term such that the total refinancing rate looks more similar across sectors, as seen in Figure 2.6.

<sup>22</sup>More broadly, Campbell (2006) notes the theoretically ambiguous welfare impact of financial product innovations or interventions designed to eliminate the cross subsidization of financially sophisticated households by naïve ones. Outside of finance, Grubb (2014) and Grubb and Osborne (2015) develop and estimate a model of inattentive consumers in which disclosure regulation that improves consumer attention leads to a pricing response by firms.

To understand the degree to which general equilibrium forces might offset some of the stimulative benefits from a targeted communication policy, we briefly review the literature on the theoretical general equilibrium effects of refinancing interventions. Zhang (2022) uses a structural model of the US mortgage market to show that automatically refinancing mortgage contracts could simultaneously reduce inequality in the market and improve total consumer welfare. Similarly, Fisher et al. (2022) develop a structural model of the UK mortgage market in which the elimination of cross-subsidies from slow to fast refinancers “democratizes” the mortgage market, with a potential for increased mortgage uptake by relatively poorer households. On the other hand, Berger et al. (2022) use an equilibrium pricing model for the US mortgage market to demonstrate that mortgage reforms can have negative distributional consequences when they increase equilibrium mortgage rates and reduce credit access. In particular, if a reform increases equilibrium origination rates because it causes more households to actively refinance, the households who still do not respond are worse off. Complementing these efforts, our results below provide attention effects plausibly generated by actually tested policies.<sup>23</sup> These magnitudes could be incorporated into general equilibrium models to consider what total effects could be expected from scaled versions of the treatments we study.

Finally, while these results suggest that research and policy efforts to reduce optimal refinancing barriers should consider equilibrium effects, we also note that the countervailing effects estimated by the papers above are generally modest. For example, the estimated average present value of the cost of the status quo cross-subsidization from slow to quick refinancers relative to no cross-subsidization in Zhang (2022) is equivalent to 26 bp higher origination interest rates. Similarly, Fisher et al. (2022) find that average UK outstanding rates would be 20 bp higher in a counterfactual world with no cross-subsidization. Berger et al. (2022) estimate that a counterfactual in which attention increases by 12 pp (roughly similar to the estimated effects on attention below) would increase origination rates 50 bp, followed by a steeper decline in rates over time as households refinance more frequently.

## 2.6 Inattention estimates

In this section, we interpret our treatment effects through the lens of the Andersen et al. (2020) model of inattentive refinancing. We adapt and build upon this existing

---

<sup>23</sup>In a similar spirit, an exercise in Berger et al. (2022) uses observational data on the correlation between borrowing from non-banks and subsequent refinancing sensitivity to imagine a world where only non-banks originate mortgages. Considering how our attention effects might affect equilibrium rates would directly connect general equilibrium model results to a feasibly scalable and already-demonstrated policy.

work to introduce a model of refinancing behavior adjusted to the Irish context and our experimental setting. First, we assess the degree to which inattention helps a model of refinancing fit the data with realistic and reasonable fixed cost parameters. Second, because attention is unobservable in our setting, we use the model to estimate the degree of inattention and the extent to which the effects of the disclosure treatments and reminder letters are consistent with a mechanism that operates through reducing inattention. This exercise also allows us to contrast estimates of attention treatment effects with changes to refinancing induced by conventional monetary stimulus holding the level of household attentiveness fixed.

### 2.6.1 Baseline refinancing model

The baseline model builds on the optimal refinancing model of Agarwal et al. (2013), which assumes that households are risk neutral and fully attentive, refinancing their mortgages if the expected net benefits of refinancing are positive. The model captures several reasons why not refinancing a mortgage might be a perfectly rational financial decision for mortgage holders. First, mortgage holders might deem the available savings insufficient to justify actual or psychological switching costs. We allow for and estimate the threshold of savings that is sufficient for attentive borrowers to consider the benefits of refinancing net such costs sufficient. Second, mortgage holders might be ineligible to switch as a consequence of their loan-to-value positions or their repayment history. Throughout the paper, we drop borrowers in arrears from our sample such that all remaining borrowers are eligible for an internal refinance. Furthermore, we calculate our interest rate savings conservatively assuming that borrowers do not qualify for a lower LTV category. Third, if mortgage borrowers are ex-ante likely to move in the near-term, they might decide not to switch or refinance because they will not be in the home long enough to recoup the fixed costs of switching or refinancing. The Agarwal et al. (2013) optimal refinancing model allows for the risk of exogenous mortgage prepayment, which we estimate from data on about 90% of outstanding Irish mortgages. Fourth, borrowers may expect rates to fall further soon and prefer to take the chance that an even more advantageous refinancing opportunity will soon arise. The optimal refinancing decision incorporates such forward-looking behavior with an expected interest-rate process calibrated to the historical volatility of Irish interest rates.

There are two components in the model to the net benefits of refinancing: the incentive to refinance  $I(x_i, \theta)$  that depends on observable mortgage characteristics  $x_i$  through parameter vector  $\theta$  and an idiosyncratic random shock  $\epsilon_i$  to the net benefits of refinancing. The unobserved component  $\epsilon_i$  of the decision allows for borrowers to differ in the private benefits or costs they derive from refinancing. The incentive to refinance

$I(x_i, \theta)$  is a function capturing a household's incentive to refinance in interest-rate points

$$I(x_i, \theta) = (r_i^{old} - r_i^{new}) - O_i(x_i, \theta) \quad (2.3)$$

where  $r^{old}$  is the household's current mortgage rate and  $r^{new}$  is the best prevailing mortgage rate available to the household. Each household in the model has a minimum decrease in interest rates  $O_i$  they require to be willing to refinance, and  $I(\cdot, \cdot)$  measures how far above that threshold they are currently. The household's optimal refinancing threshold  $O_i$  is calculated using the Agarwal et al. (2013) closed-form solution to optimal refinancing option exercise:

$$O_i = \frac{1}{\psi_i} [\phi_i + W(-\exp(-\phi_i))] \quad (2.4)$$

$$\psi_i = \frac{\sqrt{2(\rho + \lambda_i)}}{\sigma} \quad (2.5)$$

$$\phi_i = 1 + \psi_i(\rho + \lambda_i) \frac{\kappa(m_i)}{m_i(1 - \tau)} \quad (2.6)$$

where  $W(\cdot)$  is the Lambert W -function,  $\rho$  is the fixed household discount rate,  $\sigma$  is the volatility of  $r$ ,  $\tau$  is marginal tax rate (for the tax deductability of mortgage interest),  $m_i$  is the outstanding mortgage balance, and  $\kappa(m_i)$  is refinancing costs. In practice, we will allow for an additive term  $\exp(\gamma)$  in  $\kappa(\cdot)$  that will capture any non-monetary cost of refinancing that borrowers face, such as time or hassle costs. The expected rate of decline in real principal  $\lambda_i$  is the sum of expected inflation  $\pi$ , the exogenous probability of early termination  $\mu$ , and the amortization rate of the borrower's mortgage.<sup>24</sup>

Table 2.6 reports our calibration of these parameters, and the appendix discusses alternative formulations of the model to account for differences between US and Irish mortgage products. Appendix Figure 2.12 plots the distribution of refinancing incentives  $I(x_i, \theta)$  defined by (2.3) using the parameters in Table 2.6. The median and modal incentive to refinance is around 100 bp, reflecting substantial unclaimed refinancing opportunities in the experimental sample. Indicating that modeled  $I(x_i, \theta)$  relates to actual refinancing incentives, the refinancing share of each histogram bin (pooling treatment and control) plotted in Appendix Figure 2.12 is strongly increasing in the refinancing incentive. Further, the refinancing share is essentially zero for mortgage borrowers with negative refinancing incentives. Still, the absolute level of refinancing is small even for borrowers with substantial refinancing incentives, pointing to frictions such as inattention that limit borrower responsiveness.

<sup>24</sup>Following Agarwal et al. (2013), we approximate a borrower's time-varying amortization rate using the current amortization rate, measured as the difference between the borrower's annual mortgage payment to current mortgage balance ratio and the current interest rate:  $payment_i/m_i - r_i^{old}$ .

In the baseline full attention model, the household refinances if

$$e^\beta I(x_i, \theta) + \epsilon_i > 0, \quad (2.7)$$

where  $\beta$  measures the household's responsiveness to the incentive. For estimation,  $\epsilon$  is assumed to be distributed logistic, in which case the probability a borrower refinances is

$$\Pr(\text{Refinance}_i = 1 | x_i; \beta, \theta) = \Pr(e^\beta I(x_i, \theta) + \epsilon_i > 0) = \Lambda(e^\beta I(x_i, \theta)),$$

where  $\Lambda(\cdot)$  is the inverse logistic function  $\Lambda(x) = e^x / (1 + e^x)$ . We can then estimate  $\beta$  and  $\theta$  by maximum likelihood, finding the parameters that maximize the likelihood that we would observe the vector of refinancing decisions in the data.

### 2.6.2 Allowing for inattention

Inattention to a refinancing opportunity can take a number of forms. Inattention may be rational for the stressed consumer with a high current cost of processing all available information relative to low expected returns to doing so. A consumer may be distracted and simply overlook the potential savings in the moment they receive the information. Following an appreciation of the contents of a letter or other communication, inattention may occur as absent-mindedness, described by Schacter (1999) as shallow processing contributing to weak memories of key information and a related to-do action. Related to this third form of inattention is procrastination, often defined as postponing, delaying, or putting off a task or a decision in a way that is problematic rather than strategic.<sup>25</sup>

To allow for the possibility that a household is inattentive and thus not paying any attention to their refinancing incentive, we follow Andersen et al. (2020) and estimate a mixture model with each household inattentive with some probability. Inattentive households do not refinance regardless of their incentive to do so. We model households as inattentive if

$$\delta' w_i + \eta_i > 0$$

where  $\eta$  is a random shock to a household's attention and

$$\delta' w_i = \delta_0 + \delta_1 \text{Treatment}_i + \delta_2 \text{Reminder}_i.$$

This specification integrates experimental variation in treatment assignment into the probability that a household is attentive. The attention intercept  $\delta_0$  facilitates estimating the baseline attention level of the control group, and the disclosure redesign

<sup>25</sup>As we mention above, many survey respondents in Keys et al. (2016) cite procrastination as a reason for their inaction. Studies suggest that procrastination chronically affects 15-20% of adults, and that approximately 25% of adults consider procrastination to be a defining personality trait for them (Steel and König, 2006; Nguyen et al., 2013).

treatment and reminder treatment effects  $\delta_1$  and  $\delta_2$  allow for the attentiveness of each consumer to be impacted by the communication they receive.

**Identification** Intuitively, the sensitivity  $\beta$  to refinancing incentives is identified by cross-sectional variation in refinancing incentives  $I(x_i, \theta)$ . In practice, because  $r_i^{new}$  in (2.3) is constant in our experiment, variation in  $I(x_i, \theta)$  is driven by cross-sectional differences in initial interest rates  $r_i^{old}$ , mortgage balances  $m_i$ , and loan maturities (the remaining determinant of monthly payments). Heterogeneity in these variables leads to the distribution of incentives shown in Appendix Figure 2.12;  $\beta$  is related to the slope of the refinancing share line, indicating how refinancing propensities vary with refinancing incentives. Whereas  $\beta$  helps the model *scale* the refinancing incentive appropriately, the refinancing cost parameters in  $\theta$ —including the extension below to allow for an unobserved hassle cost of refinancing  $\gamma$ —help the model help *locate* the refinancing incentive. By imposing the condition in (2.3) of refinancing only when facing a positive incentive to do so, the functional form of  $I(\cdot, \theta)$  identifies the unknown terms in  $\theta$  by essentially shifting the refinancing incentive bins in Appendix Figure 2.12 to satisfy (2.3).

Extending the model to allow for treatment effects on attention helps with an identification challenge driven by the unobservability of attention in many empirical settings. Specifically, the mapping of  $\delta_0$  to purely attention is not necessarily identified given that  $\delta_0$  will capture any reason for a given borrower to not act on positive financial incentives to refinance. In many contexts, what this model estimates as under-refinancing due to inattention could also be driven by other state variables. For example, if all borrowers are fully attentive but some are missing required documentation to be approved by a bank’s underwriting department, the model described here will attribute these constraints to inattention, loading such unexplained refinancing failures onto  $\delta_0$ .

Our approach addresses this empirical challenge in four ways. First, in contrast to prior work estimating inattention, we exploit the random assignment of treatment and control to ensure a balance of borrower unobservables across treatment variables. If some unobserved constraint besides inattention leads some borrowers to not refinance, such a refinancing barrier would be present among both treatment-group and control-group borrowers. This balance allows us to identify  $\delta_1$  and  $\delta_2$  even if the interpretation of  $\delta_0$  is confounded by unobserved heterogeneity. Second, studying internal refinancing opportunities for which borrowers in our sample are eligible allows us to abstract away from settings where many borrowers face underwriting constraints unobservable to the econometrician. Again, by virtue of random assignment, any *misperceptions* about refinancing eligibility should be balanced across treatment and control. Third, economic intuition supports our interpretation of  $\delta_1$  and  $\delta_2$  as causal effects on inattention, given

that, for example, a reminder letter more plausibly affects attention than overcomes unmeasured refinancing constraints. Finally, we note that Andersen et al. (2020) address potential unobserved heterogeneity by estimating an extended version of their model that allows for random coefficients and unobserved borrower heterogeneity.<sup>26</sup> Their estimates with and without borrower heterogeneity are quite similar, suggesting that unobservable differences across borrowers are not a main driver of inattention estimates in their setting.

**Estimation** If the inattention shock  $\eta$  is also distributed logistic, then the probability that a given household is inattentive in any given period can be written as

$$\Pr(\delta_0 + \delta_1 Treatment_i + \delta_2 Reminder_i + \eta_i > 0) = \Lambda(\delta'w_i). \quad (2.8)$$

To refinance, households need to both be attentive (probability  $1 - \Lambda(\delta'w_i)$ ) and have positive net benefits of refinancing (probability  $\Lambda(e^\beta I(x_i, \theta))$ ). Households that do not refinance are either inattentive or attentive but do not have sufficient incentive to refinance. The likelihood that a household refinances at time  $t$  is then  $(1 - \Lambda(\delta'w_i))\Lambda(e^\beta I(x_i, \theta))$ . The overall likelihood  $\mathcal{L}(\cdot|\cdot)$  of observing a sample of refinancing given covariates  $x$  is the product of the relevant probabilities for the refinancers and the non-refinancers

$$\mathcal{L}(\beta, \delta, \theta|x, w) = \prod_{ref_i=1} (1 - \Lambda(\delta'w_i))\Lambda(e^\beta I(x_i, \theta)) \prod_{ref_i=0} \Lambda(\delta'w_i) + (1 - \Lambda(\delta'w_i))\Lambda(-e^\beta I(x_i, \theta)).$$

where the first and second products are taken over all borrowers  $i$  that did and did not refinance, respectively.

To estimate the model, we first set certain parameters in  $\theta$  that govern mortgage contracts and market expectations to fit the Irish context. Using a variety of data sources, we estimate expected Irish inflation as of 2020, household discount rates, nominal interest-rate volatility, mortgage-interest tax deductability, the likelihood of exogenous early mortgage termination, and the fixed cost of refinancing—see Table 2.6 for details. The maximum likelihood estimates  $(\hat{\beta}, \hat{\delta}, \hat{\theta})$  then maximize the log of the likelihood function above. These parameters estimate the importance  $\beta$  of the refinancing incentive, the importance  $\delta$  of the covariates in shifting attention, and the importance  $\theta$  of the covariates in determining private refinancing costs. Estimating this model in our setting with exogenous treatment variables allows us to characterize how valuable a given treatment is at focusing consumer attention on refinancing.

---

<sup>26</sup> Andersen et al. (2020) also demonstrate several predictions of a model where refinancing thresholds  $\gamma_{it}$  vary arbitrarily across borrowers and time that are inconsistent with panel data on refinancing behavior.



Table 2.7 reports estimates of this model using Maximum Likelihood along with robust standard errors. In column 1, we essentially constrain the model to follow only the Agarwal et al. (2013) model of refinancing without any fixed cost of refinancing or possibility of borrower inattention. In this specification, we estimate a strongly negative  $\beta$  such that the estimated coefficient  $\exp(\beta)$  on the incentive to refinance in (2.7) is essentially 0. Without allowing for fixed costs of refinancing or inattention, it would appear as if the model is a poor fit to actual behavior and that borrowers are completely insensitive to the incentive to refinance.

Starting in column 2, we allow for there to be an unobserved fixed cost of refinancing in the refinancing cost function  $\kappa(\cdot)$  in equation (2.6). Specifically, we let the total cost of refinancing be  $\kappa(m_i) = \kappa_0 + \exp(\gamma)$ .<sup>27</sup> As before, borrowers refinance when their expected gain from refinancing (including their logit private shock to refinancing costs) exceeds their optimal threshold, which—starting in column 2—also depends on  $\gamma$ . Once we model these unobserved refinancing costs with  $\gamma$ , estimates of  $\beta$  increase significantly. The estimate of  $\beta$  in columns 2 implies that a 10 bp decrease in rates increases refinancing conditional on being attentive by approximately 50 bp.<sup>28</sup>

However, the implied estimate of fixed costs in the specification of column 2, which does not yet allow for attention effects, is implausibly high ( $\exp(\hat{\gamma}) \approx \text{€}514,000$ ). Even allowing for the interpretation of this fixed cost to include the mental, time, and hassle costs of refinancing, the large estimates are perhaps more consistent with mortgage borrower inattention, which the specification in column 2 is constrained to attribute to borrowers behaving as if their costs of refinancing were incredibly high. When we allow for attention effects in column 3, the fixed cost parameter is reduced substantially from 13.2 to 6.4, demonstrating how allowing for a certain fraction of mortgage borrowers to be inattentive to refinancing improves the model’s fit of the data. The estimate of  $\gamma$  in column 3 implies a cost of refinancing of approximately  $\text{€}620$ .

The estimate of the probability of being inattentive is  $\Lambda(\hat{\delta}_0) \approx 78\%$  in column 3. Although consistent with a substantial likelihood of being inattentive, this estimate pools the control group and the treatment group. Columns 4 and 5 allow mortgage borrowers who received disclosure letters with design improvements and those that additionally received follow-up reminder letters 4-6 weeks later to have different levels of attention. The estimate of  $\delta_0$  in column 4 measures the control group’s average

<sup>27</sup>In our internal refinancing context, no fee is required paid to the lender as in other settings. We set a small nominal  $\kappa_0 = \text{€}100$  for estimation purposes to bound the refinancing cost function away from zero.

<sup>28</sup>To interpret refinancing magnitudes, we consider effects in the neighborhood of a refinancing incentive of 100 bp, which is roughly the median incentive in Appendix Figure 2.12. A 10 bp increase in the refinancing incentive then increases attentive refinancing by  $\Lambda(1.1e^\beta) - \Lambda(e^\beta)$ .

probability of inattention to be  $\Lambda(1.13) \approx 76\%$ . The treatment effects estimates in column 4 imply that the combination of redesigned disclosures and follow-up reminders decreased inattention by 16 pp in total: 6 pp from redesigned disclosure letters and 10 pp from the reminders.

The fixed-cost estimate increases when we allow for treatment effects on inattention, with the estimate of  $\gamma$  in column 4 of Table 2.7 implying a €6,000 cost of refinancing. This higher cost of refinancing in column 4 than column 3 suggests that the specification in column 3 was misattributing some of the more responsive refinancing of the treatment groups to having a lower cost of refinancing. Once allowing for the treatment groups to have lower inattention in column 5, it is clear that the control group still behaves as if they have a high cost of refinancing, consistent with overall pessimistic beliefs about the time and effort required to refinance a mortgage (Central Bank of Ireland, 2017b). Column 5 adds controls that allow for heterogeneity in refinancing costs along observable dimensions to test whether certain groups have stronger inertia, with refinancing inertia increasing in age and first-time homebuyer status and decreasing in Covid forbearance. The estimates of the treatment effects on attention and the fixed cost estimates are similar to column 4. Overall, the redesigned disclosure treatment and subsequent follow-up reminder decrease the probability of being inattentive using the column 5 estimates by 20 pp from 76% to 56%.

The estimates are consistent with the reminders having a large effect on refinancing by increasing the probability that a given borrower is attentive. Reconciling the nontrivial effects of the treatments without reminders on inattention in Table 2.7 with the more modest effects in Figure 2.4, recall that the total effect of the treatment on refinancing is the increase in the probability of attending to the task of refinancing times the probability of refinancing for a given refinancing incentive conditional on paying attention. Because this second term is low, the total effect of improving attention by a few percentage points through redesigned disclosures is still somewhat muted, consistent with the material implied fixed cost of refinancing  $\gamma$  in columns 2-4.

### 2.6.3 Comparison to interest-rate changes

We use our model estimates to measure the relative partial-equilibrium effectiveness of cutting interest rates (which increases the refinancing incentive  $I$  by lowering  $r^{new}$ ) versus sending a reminder as effective as our field experiment reminders that decrease the probability of inattention  $\Lambda(\delta'w_i)$ . This exercise is particularly policy relevant when monetary policy is de facto constrained by a zero lower bound, complicating efforts to decrease interest rates through conventional monetary policy, or when pass-through from policy rates to mortgage rates is otherwise impaired. Similarly, when rates are

set in a monetary union as in the euro zone or in the United States, the optimal policy rate may differ across regions, in which case non-monetary measures available to individual regions to stimulate demand may be valuable.

The estimates suggest that there is significant scope for direct-to-household communication from the central bank or other policymakers in the form of reminder notices to provide monetary stimulus by spurring refinancing. Even when the incentive to refinance is approximately 0, the estimates in columns 4-5 of Table 2.7 predict that reminders will increase refinancing by 8-9 percentage points, which is within the 95% confidence interval of half of the treatment-reminder combinations in Figure 2.5. Reminders and lower interest rates are also complementary. When the average incentive to refinance is 100 bp, reminders increase refinancing by about 9% (an additional percentage point), with the modesty of the complementarity driven by the presence of inattention and our small estimates of  $\beta$ .

We can further use the model to contrast the effectiveness of targeted communication with the implied effect of the more conventional approach stimulating refinancing through decreasing mortgage interest rates. Figure 2.3 contrasts estimates of our experimental treatment effects from section 2.5 with the implied effect of decreasing mortgage rates by 100 bp in Ireland, the US, and Denmark. The estimates of column 4 in Table 2.7 imply that if mortgage interest rates fell by 100 basis points, refinancing would only increase by 1.2 pp, comparable to the small effect of the average redesigned disclosure treatment without reminders.<sup>29</sup> This effect of even a large interest-rate change is small both because baseline inattention is so high (76% in column 4) and because the coefficient  $e^{\hat{\beta}}$  on refinancing incentives is small even when accounting for inattention. The latter reason for large equivalent effects could arise from the limited amount of cross-sectional variation in refinancing incentives in our data (Appendix Figure 2.12). However, even using the Andersen et al. (2020) estimate for Denmark of  $\hat{\beta} \approx 0.7$ , the effect of a 100 bp interest rate change on refinancing is still 2.4 pp—much smaller than the best performing treatment and reminder combination (6.8 pp in Figure 2.3).

To validate our refinancing model estimates and test whether they are representative of other contexts, we estimate the specification of column 3 of Table 2.7 on US data and calculate the implied effects on refinancing in Denmark using estimates from Andersen et al. (2020). While we cannot estimate experimental treatment effects on attention in other contexts, we can use observational data and the functional form of the mixture model to describe the sensitivity of US and Danish borrowers to refinancing incentives. This exercise also illustrates the degree to which Irish borrowers in our 2020 experiment

<sup>29</sup>The model-implied change in refinancing from increasing the incentive to refinance from 100 to 200 bp while holding inattention fixed at its baseline level is  $(1 - \Lambda(\hat{\delta}_0)) (\Lambda(2e^{\hat{\beta}}) - \Lambda(e^{\hat{\beta}}))$ .

are particularly unique in their refinancing attention or elasticity. Using the CRISM data used in Figure 2.1, we estimate the model on a similarly sized random sample of US mortgages from May to September of 2019, a period when US interest rates were relatively stable.<sup>30</sup> Using the results that  $\hat{\beta} = -.36$  and  $\hat{\delta}_0 = 2.15$ , we estimate how much decreasing mortgage interest rates by 100 bp in the US would increase mortgage prepayment rates over a four-month period. Given the interest-rate sensitivity among borrowers in the US data and holding their inattention fixed, the model predicts that increasing refinancing incentives from 100 to 200 bp would increase mortgage prepayment by 1.4 pp.

We can also use the estimates of Andersen et al. (2020) to calculate the implied effect on external refinancing in Denmark of lowering mortgage rates by 100 bp. The estimates of  $\beta$  and  $\delta_0(\chi)$  from Model 4 of Table 2 in Andersen et al. (2020) imply that increasing refinancing incentives from 100 to 200 bp in Denmark would increase refinancing by 1.5 pp.<sup>31</sup> The small size and similarity of these effects (1.2, 1.4, and 1.5 pp in Ireland, the US, and Denmark, respectively) are consistent with reminders being potentially more powerful than monetary stimulus at stimulating refinancing and further supports the representativeness of the Irish setting to study attention effects and mortgage refinancing.

Several qualifications apply to this exercise. First, an increase in refinancing incentives from an average of 100 bp to 200 bp is outside the data for the majority of borrowers in our sample, suggesting that caution should be exercised when extrapolating to larger rate swings (although we also note that estimated refinancing incentives are more dispersed in the US and Danish data). Relatedly, general equilibrium considerations loom when predicting the effects of large interest-rate changes from our partial equilibrium model using cross-sectional parameter estimates. Many features of the economy could change if rates were to fall by a large amount (Ascari and Haber, 2021). Particularly relevant to our setting is the possibility that aggregate attention to refinancing could increase substantially in response to a large rate cut given non-linearities in refinancing incentives (Berger et al., 2021; Eichenbaum et al., 2022). However, despite these caveats, the strong qualitative conclusion from this exercise is that the reminder effects are significantly more powerful than a typical change in policy rates. Importantly, we also note that the comparison above is to a *mortgage* rate decrease of 100 bp, which

---

<sup>30</sup>We treat  $r_i^{new}$  from equation (2.3) as 4.15%, the average of the weekly Freddie Mac new mortgage conventional interest-rate series over the four-month window we consider. For comparability with our Irish sample, we restrict the sample to conventional 30-year fixed-rate first mortgages with balances between \$50,000 and \$1,000,000 that were outstanding, current, and not scheduled to mature before September 2019. In the US data, we cannot observe the difference between refinancing and non-distressed prepayment (e.g., from a borrower selling her home).

<sup>31</sup>Figure 2.3 lacks confidence intervals for the Denmark effect because the variance-covariance matrix of the Andersen et al. (2020) ML estimates is unavailable.

would generally require aggressive or extraordinary measures to achieve given the limited pass-through from ECB policy rates to mortgage interest rates and the apparent lower bound on nominal policy rates.

## 2.7 Conclusion

In this paper, we study an intervention targeted at improving the last-mile delivery of monetary policy to the real economy. We use a field experiment combined with a mixture model of inattentive financial decision-making to demonstrate that targeted communication can help overcome the attention frictions that inhibit the refinancing channel of monetary policy transmission. The best performing combination of redesigned disclosures and follow-up reminders implemented by our partner bank increases the take-up of in-the-money refinancing opportunities by 76%—substantially more than any prior effort in the literature.

What impact might the refinancing effects we document have on borrower consumption? The average 12-month savings realized by refinancing mortgage borrowers was €1,209. Using the MPC out of interest rate changes among UK mortgage holders estimated by the Bank of England of 0.5 (Anderson et al., 2014), we estimate that refinancing households increased their consumption by €605.<sup>32</sup> Averaged across all households receiving a reminder letter, this suggests that the best redesigned disclosure letter and accompanying reminder increased borrower consumption by an expected €42 per household. Conservatively assuming that the redesigned disclosure and reminder letters cost €1 each to produce and deliver, this implies a borrower consumption multiplier of €42 for every €1 spent on communication to households about the opportunity to refinance.<sup>33</sup>

Estimates of an extended version of the Andersen et al. (2020) model of inattentive refinancing suggest that the reminder disclosures had large effects precisely because they increased the probability that a given consumer was attentive to the task of refinancing. Using our model estimates, we find that communication reminding mortgage borrowers of refinancing opportunities has significant potential to be an effective monetary policy tool to complement or substitute for lowering rates. We estimate that mortgage interest rates falling by 100 bp in Ireland, the US, or Denmark would have much smaller

---

<sup>32</sup>We caveat that the Bank of England estimate relates to the MPC of borrowers out of higher mortgage interest payments instead of the mortgage repayment savings we study here. However, this MPC choice is conservative relative to the Di Maggio et al. (2017) estimated MPC out of interest savings in the US of 0.75.

<sup>33</sup>The expected effects on *aggregate* consumption may be less than the effect on borrower consumption after taking into account a potential offsetting loss in income by the bank's domestic shareholders.

effects than the reminders we study. Moreover, given limited pass-through of policy rates to retail interest rates, a large decrease in mortgage interest rates would likely require unconventional monetary stimulus to achieve.

Several caveats apply to our estimates. Repeated reminders may be more or less effective than the one-shot reminder we studied here. Repeated reminders may lose their salience if households learn to rely on them instead of proactively acquiring their own information on refinancing activities (Ericson, 2017). Similarly, as the households with the largest incentive to refinance or the lowest cost of attention to refinancing attrit from the sample of mortgage borrowers with large refinancing incentives, the effect of an additional reminder may decrease. However, it's also possible that as consumers become attuned to reminder letters, they would trust them more with potential spillovers through social learning. We also note that reminders are more effective when rates have fallen and may not be as impactful in a rising rate environment. However, broadly speaking, policymakers are generally not keen to stimulate refinancing in such an environment anyway.


The treatment effects we study here are also likely to be more effective when the status-quo disclosure letter is less transparent to begin with. Streamlining, personalizing, simplifying, and highlighting are more valuable in the context of confusing, onerous, and overly detailed disclosures. The success of the communication also depends on the trust households place in the disclosing entity (Johnson et al., 2019). It may be advantageous for the communication to be sent directly by a government agency or central bank than from a for-profit bank, although emphasizing that the letter itself is mandated could help. Finally, as discussed above, some of the benefits of increased refinancing may be eroded by higher initial interest rates in general equilibrium. However, it's also possible that the need for reminders would decrease in equilibrium if more attentive refinancing led banks to decrease the spread between their offered variable rates and policy rates in the first place.

Communication with mortgage holders about refinancing opportunities may interest policymakers at different points in the business cycle. First, during an expansionary monetary phase as we study, offer letters and reminders can improve interest rate pass-through and strengthen the impact of interest rate reductions on the real economy. Second, during contractionary monetary policy phases, central banks may seek to improve household financial stability by encouraging households into fixed rate products as interest rates rise. In addition, several categories of governmental entities could be interested in the policy lever we evaluate here, including fiscal authorities seeking to stimulate refinancing and consumption, competition authorities aiming to improve the

competitiveness of the mortgage market, and consumer protection authorities focused on improving households debt service burdens.

Finally, our results contribute to a growing body of evidence that demonstrates the value of behaviorally informed approaches in delivering effective consumer protection in essential product markets. In particular, the results we document here are the first to demonstrate statistically and economically meaningful improvements in the stubbornly persistent puzzle of low take up of advantageous mortgage refinancing opportunities.

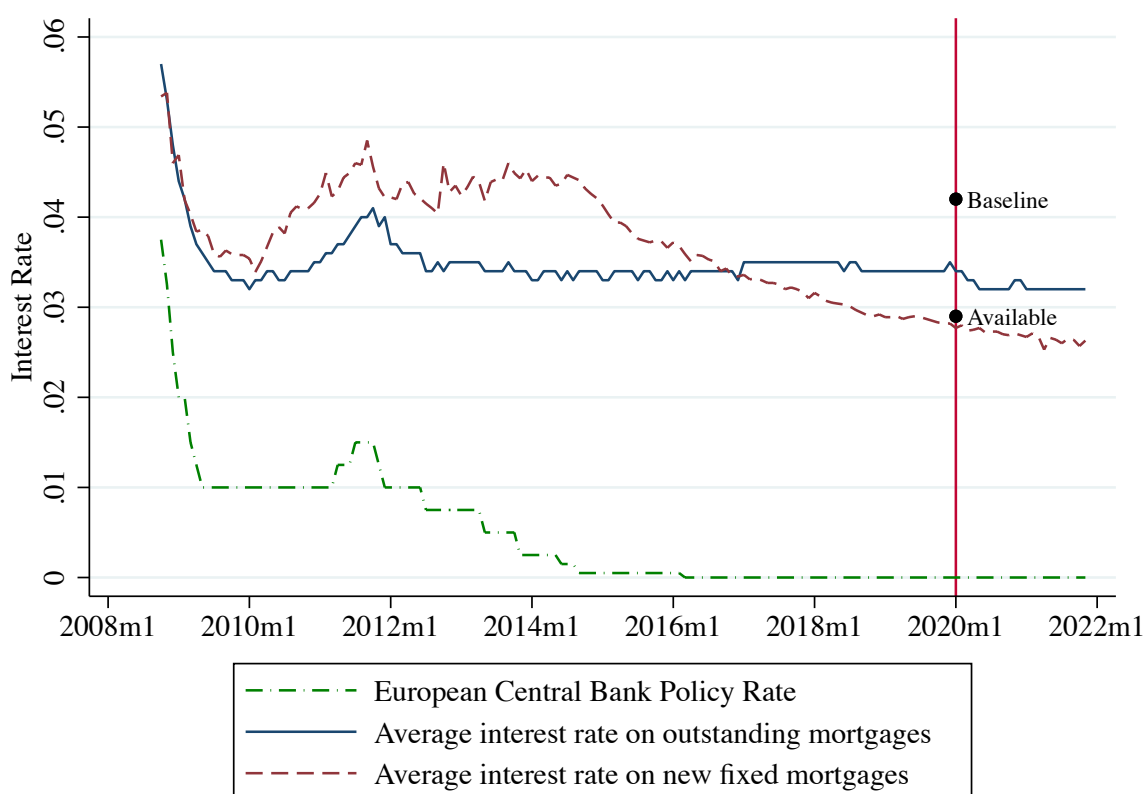
Figure 2.1: US policy rates and new and outstanding mortgage interest rates



Note: Figure plots average mortgage interest rates by month for outstanding mortgages (solid blue line), newly originated mortgages (dashed red line), and the effective Federal Funds Rate (dashed-dotted green line). Outstanding and new mortgage interest rates are calculated from CRISM (see Di Maggio et al. (2020) for details). Effective federal funds rate is from the Board of Governors of the Federal Reserve System.




Figure 2.2: Ireland policy rates and new and outstanding mortgage interest rates




Note: Figure plots average mortgage interest rates by month for outstanding mortgages (solid yellow line), newly originated variable rate mortgages (dashed red line), newly originated fixed rate mortgages (dashed green line), and the European Central Bank Main Refinancing Operations Rate (solid blue line). Outstanding and new mortgage interest rates are calculated from Central Bank of Ireland Retail Interest Rate data. The European Central Bank rate is from the ECB Statistical Data Warehouse. The baseline blue dot represents the average interest rate at the outset for those cohorts within the experiment. The opportunity blue dot is the interest rate achievable by taking up the refinancing offer.

Figure 2.3: Refinancing treatment effects




Notes: Figure plots treatment effects on refinancing rates in percentage points for observed and hypothetical events. Blue bars on the left report treatment effects estimated using our experimental data. Red bars on the right report the implied effect of a 100 bp decrease in mortgage rates using estimates of the Andersen et al. (2020) model using data from the indicated country. Treated no reminder is the increase in the internal refinancing rate relative to the control group for borrowers across all treatment arms who did not receive a follow-up reminder letter. Treated with reminder is the increase in the internal refinancing rate relative to the control group for borrowers randomly assigned to any treatment arm who did receive a follow-up reminder letter. Best treatment with reminder is the increase in the internal refinancing rate relative to the control group for borrowers randomly assigned to the best performing treatment arm with a follow-up reminder letter (V2). Effect of 100 bp lower mortgage rates in Ireland is the implied increase in the internal refinancing rate for the control group from an increase in the refinancing incentive from 100 to 200 bp holding attention fixed, calculated using the estimates in column 4 of Table 2.7 as  $(1 - \Lambda(\hat{\delta}_0))(\Lambda(2e^{\hat{\beta}}) - \Lambda(e^{\hat{\beta}}))$ . Effect of 100 bp lower mortgage rates in the US similarly uses  $\hat{\beta} = -.36$  and  $\hat{\delta}_0 = 2.15$  from estimating the specification of column 3 of Table 2.7 on US data on mortgage prepayment from May-September 2019. See section 2.6.3 for details. Effect of 100 bp lower mortgage rates on external refinancing in Denmark uses estimates from Model 4 of Table 2 in Andersen et al. (2020) to estimate the effect of increasing the refinancing incentive from 100 to 200 bp holding attention fixed. Internal refinancing is defined as a borrower switching mortgage products with the partner bank within four months of initial treatment. Error bars denote robust 95% confidence intervals. For the model-implied refinancing effects, confidence intervals are calculated by the Delta method using the robust covariance matrix of the model estimates, which is unavailable for the Denmark point estimates.

Figure 2.4: Refinancing rates by treatment arm: no reminder sample




Notes: Figure plots internal refinancing rates by treatment arm for the subset of the sample that did not receive a reminder letter. Brackets denote 95% confidence intervals based on robust standard errors for the difference between the control group and each treatment arm.

Figure 2.5: Refinancing rates by treatment arm: reminder sample



Notes: Figure plots internal refinancing rates by treatment arm for the subset of the sample that did receive a reminder letter along with the control group. Brackets denote 95% confidence intervals based on robust standard errors for the difference between the control group and each treatment arm.

Figure 2.6: Refinancing rates by employment sector



Note: Figure plots internal refinancing rates by employment sector and treatment category. WFH, BAU, and AHNW denote mortgagor employment industries more likely to be working from home, experiencing business as usual, and being at home but not working, respectively, during the estimation window. See Appendix Table 2.11 for our employment sector classification scheme. The left three bars plot refinancing rates for borrowers that received the redesigned disclosure treatment but not a reminder. The right three bars plot refinancing rates for borrowers that received the redesigned disclosure treatment and a follow-up reminder letter. Brackets denote 95% confidence intervals based on robust standard errors for the within-sector difference between the control group and each treatment arm.

Table 2.1: Treatment arms overview

Disclosure Redesign Element	Treatment Group						
	Control	V1	V2	V3	V4	V5	V6
Simplification		✓	✓	✓	✓	✓	✓
Personalized savings estimate		✓	✓	✓	✓	✓	✓
Neutral frame		✓	✓	✓			
Color			✓				
Prominent subject line				✓	✓	✓	✓
Gain frame					✓		
Loss frame						✓	✓
Clarified process box							✓
Follow-up reminder letter		12	12	12	12	12	12

Notes: Chart overviews the additional design elements incorporated into each treatment arm. See section 2.4.1 for a description of each element. The control group column indicates that the control group received the existing standard disclosure without any additional design elements. The reminder row indicates that a randomly assigned half of each of the six treatment arms received a follow-up reminder letter 4-6 weeks after the initial treatment.

Table 2.2: Descriptive statistics

Group	(1) Control	(2) Treated no reminder	(3) Treated with reminder	(4) Market (variable rate)	(5) Market (all)
Dublin	0.20 (0.40)	0.19 (0.39)	0.20 (0.40)	0.27 (0.44)	0.28 (0.45)
Borrower age	49.74 (9.26)	50.10 (9.41)	49.99 (9.31)	48.99 (9.90)	48.24 (9.63)
First Time Buyer	0.41 (0.49)	0.41 (0.49)	0.39 (0.49)	0.39 (0.49)	0.38 (0.49)
Mortgage balance	83,503 (84,125)	80,617 (87,748)	82,027 (92,103)	102,688 (95,037)	128,238 (111,522)
Interest rate	0.042 (0.003)	0.042 (0.002)	0.042 (0.002)	0.036 (0.006)	0.026 (0.01)
Years to maturity	13.87 (8.54)	13.22 (8.47)	13.29 (8.49)	14.63 (8.85)	15.90 (8.64)
1-Year savings	1,044 (1,010)	1,019 (1,115)	1,028 (1,093)	968 (1,120)	-60.88 (1,827)
Covid forbearance	0.09 (0.28)	0.08 (0.27)	0.08 (0.28)	0.11 (0.32)	0.13 (0.34)
Observations	1,613	4,796	4,791	206,083	538,956

Notes: Table reports means and standard deviations in parentheses of mortgage borrower characteristics for the control group in column 1, loans treated with redesigned disclosures but without a reminder in column 2, and loans treated with redesigned disclosures and with a follow-up reminder letter in column 3. In columns 4-5, we report descriptive statistics from the Loan Level Data of the Central Bank of Ireland covering about 90% of outstanding mortgages, regardless of lender. Column 4 reports statistics on all outstanding variable-rate mortgages and column 5 reports on all outstanding residential mortgages. Dublin is an indicator for whether the mortgaged property is located in Dublin. Borrower age of the oldest borrower on the mortgage. First Time Buyer indicates whether the borrower is a first time-buyer. Mortgage balance is amount outstanding on loan at the time of experiment in euros. Interest rate is the interest rate applicable on the loan at the outset of the experiment. 1-year savings is the amount in euros of savings available to the borrower in the first year after refinancing to the best available rate. Covid forbearance indicates whether the borrower was using Covid payment break (introduced in Ireland in March 2020 to alleviate short-term liquidity constraints faced by borrowers experiencing financial difficulties due to the impact of the pandemic). Covid forbearance shares for the market comparisons are measured from loan-level data collected by the Central Bank of Ireland as at June 2021, while all other variables are measured at the outset of the field trial.

Table 2.3: Internal refinancing treatment effects

	(1)	(2)	(3)	(4)
Disclosure Redesign Treatment	0.036*** (0.008)	0.040*** (0.008)	0.018** (0.008)	0.022*** (0.008)
Treatment $\times$ Reminder			0.036*** (0.007)	0.035*** (0.007)
Constant	0.089*** (0.007)	-0.311*** (0.067)	0.089*** (0.007)	-0.307*** (0.066)
Controls		✓		✓
Observations	11,200	11,200	11,200	11,200
R-squared	0.002	0.042	0.004	0.044

Notes: Table reports treatment effects on an indicator variable equal to one if the borrower internally refinanced, defined as a borrower changing their mortgage product with the partner bank within three months of treatment, and zero otherwise. Disclosure Redesign Treatment is an indicator that the borrower was randomly assigned to one of the six treatment arms. Reminder is an indicator for whether that borrower received a follow-up reminder letter 4-6 weeks after the initial treatment as in Appendix Figure A5. Control variables in columns 2 and 4 are listed in Table 2.2. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .



Table 2.4: Internal refinancing treatment effects by treatment arm

	(1)	(2)	(3)	(4)
Personalized Treatment (V1)	0.036*** (0.011)	0.040*** (0.011)	0.017 (0.013)	0.021* (0.013)
Color Treatment (V2)	0.038*** (0.011)	0.042*** (0.011)	0.007 (0.013)	0.011 (0.012)
Headline Treatment (V3)	0.034*** (0.011)	0.038*** (0.011)	0.026* (0.013)	0.029** (0.013)
Gains Treatment (V4)	0.041*** (0.011)	0.044*** (0.011)	0.030** (0.013)	0.032** (0.013)
Losses Treatment (V5)	0.032*** (0.011)	0.036*** (0.011)	0.020 (0.013)	0.025* (0.013)
Process Treatment (V6)	0.035*** (0.011)	0.038*** (0.011)	0.009 (0.013)	0.012 (0.013)
<u>Reminder ×</u>				
Personalized Treatment (V1)			0.040** (0.017)	0.037** (0.016)
Color Treatment (V2)			0.061*** (0.016)	0.060*** (0.016)
Headline Treatment (V3)			0.016 (0.016)	0.018 (0.016)
Gains Treatment (V4)			0.021 (0.017)	0.024 (0.016)
Losses Treatment (V5)			0.026 (0.016)	0.023 (0.016)
Process Treatment (V6)			0.052*** (0.017)	0.050*** (0.016)
Constant	0.089*** (0.007)	-0.312*** (0.067)	0.089*** (0.007)	-0.309*** (0.066)
Controls		✓		✓
Treatment effects equality $p$ -value	0.986	0.989	0.609	0.685
Reminder effects equality $p$ -value			0.310	0.362
Observations	11,200	11,200	11,200	11,200
R-squared	0.002	0.042	0.005	0.045

Notes: Table reports treatment effects by treatment arm on internal refinancing, defined as a borrower changing their mortgage product with our partner bank within four months of initial treatment. See Table 2.1 for summary of treatment arm features. Reminder is an indicator for whether that borrower received a follow-up reminder letter 4-6 weeks after the initial treatment as in Appendix Figure 2.11. Control variables in columns 2 and 4 are listed in Table 2.2. Treatment effects equality  $p$ -values are from a joint  $F$ -test that the six treatment coefficients are equal to each other. Reminder effects equality  $p$ -values are from a joint  $F$ -test that the six reminder × treatment arm coefficients are equal to each other. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 2.5: Refinancing treatment effect heterogeneity

	(1)	(2)	(3)	(4)
		<i>I. Individual Characteristics</i>		
Covariate $x \rightarrow$	Dublin	Age > 50	FTB	Covid
Treatment	0.014 (0.009)	0.028** (0.012)	0.009 (0.010)	0.016* (0.009)
Treatment $\times$ Reminder	0.036*** (0.007)	0.044*** (0.010)	0.030*** (0.008)	0.036*** (0.007)
Treatment $\times x$	0.020 (0.021)	-0.020 (0.017)	0.024 (0.017)	0.025 (0.033)
Treatment $\times$ Reminder $\times x$	-0.001 (0.017)	-0.019 (0.013)	0.016 (0.014)	-0.0002 (0.027)
Covariate $x$	-0.007 (0.017)	-0.031** (0.014)	0.016 (0.015)	0.017 (0.027)
Constant	0.091*** (0.008)	0.103*** (0.010)	0.083*** (0.009)	0.088*** (0.007)
Observations	11,200	11,200	11,200	11,200
R-squared	0.004	0.012	0.009	0.005
		<i>II. Loan Characteristics</i>		
Covariate $x \rightarrow$	High Balance	$r > 4.2\%$	YTM > 13	High Savings
Treatment	0.013 (0.009)	0.018 (0.013)	-0.003 (0.010)	0.011 (0.008)
Treatment $\times$ Reminder	0.036*** (0.007)	0.025** (0.010)	0.038*** (0.008)	0.035*** (0.007)
Treatment $\times x$	0.018 (0.018)	-0.0004 (0.017)	0.046*** (0.017)	0.027 (0.019)
Treatment $\times$ Reminder $\times x$	-0.003 (0.015)	0.017 (0.014)	-0.004 (0.014)	0.001 (0.015)
Covariate $x$	0.089*** (0.015)	0.015 (0.014)	0.053*** (0.014)	0.091*** (0.016)
Constant	0.052*** (0.007)	0.080*** (0.011)	0.063*** (0.009)	0.052*** (0.007)
Observations	11,200	11,200	11,200	11,200
R-squared	0.028	0.005	0.024	0.033

Notes: Table estimates treatment-effect heterogeneity by interacting the disclosure redesign treatment variable and the reminder treatment variable with borrower and loan characteristics in panels I and II, respectively. Each column estimates a different regression replacing the covariate  $x$  with the binary measure of heterogeneity indicated in that column's header. FTB stands for first-time buyer. Covid stands for Covid mortgage-payment forbearance. High Balance is an indicator for outstanding principal in excess of €75,000. The indicator  $r > 4.2\%$  refers to the baseline prevailing interest rate on each borrower's mortgage. The indicator YTM > 13 denotes there are at least 13 years remaining until a mortgage matures. High savings denotes borrowers who stand to save more than €1,000 in their first year after refinancing. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 2.6: Parameter values used in optimal refinancing model

Parameter	Name	Value	Source
Inflation	$\pi$	0.02	Average IE inflation
Real discount rate	$\rho$	0.05	Standard
Nominal interest rate volatility	$\sigma$	0.002	CBI monthly interest rate series
Marginal tax rate for interest deduction	$\tau$	0	Eliminated in Ireland in 2019
Exogenous Pr(termination)	$\mu$	0.11	Microdata from partner bank
Perceived fixed costs of refinancing (€)	$\kappa$	100	Usual cost is zero

Notes: Table reports parameter values used in the Agarwal et al. (2013) model of optimal refinancing discussed in section 2.6 adapted to the Irish mortgage market context.

Table 2.7: Inattentive refinancing model maximum likelihood estimates

Parameter	(1)	(2)	(3)	(4)	(5)
Incentive Sensitivity ( $\beta$ )	-125.48*** (1.12)	-1.61*** (0.01)	-0.23 (0.51)	-1.58*** (0.05)	-1.65*** (0.05)
Fixed Cost of Refinancing ( $\gamma_0$ )		13.15*** (0.70)	6.43*** (0.49)	8.71*** (0.03)	8.71*** (0.20)
Inattention Constant ( $\delta_0$ )			1.28*** (0.19)	1.13*** (0.11)	1.02*** (0.12)
Treatment on Inattention ( $\delta_1$ )				-0.31** (0.12)	-0.33** (0.13)
Reminder on Inattention ( $\delta_2$ )				-0.43*** (0.08)	-0.44*** (0.09)
Fixed Cost Controls					✓
Observations	11,200	11,200	11,200	11,200	11,200
Log likelihood	-7,763	-4,111	-3,977	-3,912	3,907

Notes: Table reports maximum likelihood estimates of the mixture model of inattentive refinancing described in the text. Incentive Sensitivity is the coefficient on the Agarwal et al. (2013) refinancing incentive described in section 2.6 using the parameters defined by Table 2.6, with coefficient  $\exp(\beta)$ . The fixed cost of refinancing constant  $\gamma$  estimates an average fixed cost term to rationalize observed refinancing variable. The fixed-cost controls allow for differences across groups in the estimated fixed cost of refinancing. The inattention constant  $\delta_0$  allows the inattention index in (2.8) to have a constant term. The inattention treatment effects allow borrowers who treated with redesigned disclosures ( $\delta_1$ ) and disclosure reminders ( $\delta_2$ ) to have different levels of attention. Age is demeaned. Covid indicates whether the borrower was approved for mortgage-payment forbearance with a Covid hardship. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

## Appendix

### 2.8 Robustness to alternative refinancing parameters

A concern with our use of the Agarwal et al. (2013) model of optimal refinancing is that their model considers only US-style fixed-rate mortgages. Specifically, Agarwal et al. (2013) study the optimal exercise of the option to refinance a US-style fixed-rate mortgage into another fixed-rate mortgage, resetting the term of the new mortgage back to 30 years. In contrast, Irish fixed-rate mortgages do not have fixed interest rates for their entire duration, instead converting to variable-rate mortgages by default after an initial fixation period of usually 1-5 years. Furthermore, when mortgage borrowers in Ireland refinance, they generally keep their remaining term constant instead of restarting at 30 years or switching to an entirely different duration. In this appendix, we consider alternative formulations of the incentive to refinance that account for these differences in mortgage product design. Before proceeding, we note that despite the shorter fixation periods in Ireland relative to the US, Irish mortgage borrowers still behave similarly in terms of duration, with a typical mortgage lasting for around 12 years despite rolling over to a variable rate.


One approach to tweak the Agarwal et al. (2013) model to accommodate differences between mortgage systems is used by Fisher et al. (2022). They set the likelihood of prepayment for exogenous reasons to  $\mu = 0.5$ , which makes the actual duration of a typical mortgage approximately two years. By making borrowers expect the need to go back to the market for a new market-rate mortgage with such a high probability, this mimics the effect of having a fixation period end with the mortgage rolling over to a variable rate. Strictly speaking, this is not what happens in the data in Ireland—typical borrowers hold their mortgages much longer. However, we also adopt this approach in a robustness check of setting  $\mu = 0.5$  to demonstrate that our core estimates are relatively insensitive to the particulars of the optimal mortgage refinancing model parameterization.

In Appendix Table 2.13, we report estimates from reestimating the maximum likelihood specification of section 2.6, formulating  $I(x_i, \theta)$  with  $\mu = 0.5$  and the other parameters the same as in Table 2.6. On the whole, the estimates are similar across the two tables. The fixed cost estimates are generally bigger in Appendix Table 2.13, with  $\exp(\hat{\gamma}) \approx \text{€}3,133$  in column 3, for example, but also more stable across specifications. The biggest change is a decrease in the baseline estimated rate of inattention  $\Lambda(\hat{\delta}_0) \approx 64\%$  in column 3, down from 78% in Table 2.7. While a majority of borrowers are inattentive to the opportunity to refinance in either parameterization of the refinancing decision, it is intuitive that the model would find fewer households inattentive when borrower horizons are short because exogenous prepayment is high. In this case, which approximates a world where mortgages are not fixed rate for their entire duration, households may optimally fail to refinance because they will likely have to refinance soon anyway, reducing the length of time over which they should expect to have enjoyed the benefits of refinancing. This force serves to alleviate some of the pressure for inattention to explain low refinancing levels, reducing the estimated baseline inattention rate. However, even when allowing for this possible force to be stronger in the model

than it seems in the data given slow Irish refinancing, inattention is still high. Moreover, even in Appendix Table 2.13, the combined treatment effect of receiving a redesigned disclosure and follow-up reminder letter is still large and of a similar magnitude to the original maximum-likelihood results in Table 2.7.


A final approach, also explored by Fisher et al. (2022) that abstracts away from the Agarwal et al. (2013) model is to remove the optimal threshold term  $O(x_i, \theta)$  from the specification of the incentive to refinance given in equation (2.3). This interest-rate gap definition of the refinancing incentive is also popular in the mortgage refinancing literature (for recent uses, see, e.g., Berger et al., 2021; Eichenbaum et al., 2022). Doing so remains agnostic about the precise threshold for optimal options exercise and instead lets the incentive to refinance just be proportional to the interest-rate gap, defined as the difference between a borrower's current interest rate  $r_i^{old}$  and their potential rate if refinancing  $r_i^{new}$ . Again, we find that our core results are unchanged, further emphasizing that our conclusions are not driven by the particular form or parameterization of the Agarwal et al. (2013) model.

Figure 2.7: Market share of outstanding mortgages by product type



Notes: Figure plots the share of total balances of outstanding residential mortgages in Ireland that are fixed rate, variable rate, or tracker rate. Source is the Central Bank of Ireland Retail Interest Rate Statistics series.

Figure 2.8: External mortgage refinancing rates in Ireland



Notes: Figure plots the share of mortgages that switched lenders each half year (left axis) and the volume of outstanding balances in millions of euros that switched lenders each half year (right axis). Source is Central Bank of Ireland Loan Level Data.



Figure 2.9: Example control-group disclosure letter

Mortgage Account Number: 1234567

**You may be able to save money on your mortgage**

Dear John,

This letter supplements the information we sent with your annual mortgage loan statement in the leaflet called “Information about your mortgage (You may be able to save money on your mortgage)”.

The standard variable interest rate we currently charge you on your mortgage loan is 4.34%. However, we want to make sure you are getting the best deal and we may have a lower interest rate for your mortgage.

**What rates are available?**

The lowest interest rate currently available to you is a one or two-year fixed rate of 2.9%. We also offer fixed rates for periods of three, five and ten years. The ten-year rate varies depending on your Loan to Value (LTV). We explain Loan to Value at the end of this letter.

**Explaining the tables below**

These tables show you the interest rates along with the Annual Percentage Rate of Charge (APRC). We explain APRC at the end of this letter.

**Fixed interest rates**

<b>Fixed interest rate options</b>	<b>Loan to Value Up to 60%</b>	<b>Loan to Value 61-80%</b>	<b>Loan to Value over 80%</b>
1-year	2.9% (3.9% APRC)	2.9% (4.2% APRC)	2.9% (4.4% APRC)
2-year	2.9% (3.8% APRC)	2.9% (4.0% APRC)	2.9% (4.3% APRC)
3-year	3% (3.7% APRC)	3% (3.9% APRC)	3% (4.1% APRC)
5-year	3.2% (3.7% APRC)	3.2% (3.8% APRC)	3.2% (4.0% APRC)
10-year	3.5% (3.7% APRC)	3.5% (3.8% APRC)	3.7% (4.0% APRC)

Notes: Figure shows page one of an example mandatory disclosure letter sent to the control group. Letterhead with customer and bank information is omitted.

Figure 2.10: Example treatment-group disclosure letter

Mortgage Account Number: 1234567

### You may be able to save money on your mortgage

Dear John,

Your current mortgage interest rate is a standard variable rate of 4.25%. We want to make sure you are getting the best deal and we may have a lower interest rate for your mortgage.

Current monthly repayment at 4.25%:	€717	<ul style="list-style-type: none"> <li>We have a range of interest rates that could save you money.</li> <li>Our lowest rate is a fixed rate of 2.9%, which could result in an immediate monthly saving to you of about €131. Over the course of a full year, that's approximately €1,572 in savings.</li> <li>Below, we outline the full range of interest rate options currently available, along with the next steps to take if you wish to choose one of these alternative options.</li> </ul>
Potential monthly repayment at 2.9% fixed:	€586	
Estimated difference in monthly repayments	-€131	
Potential difference over the year:	-€1,572	

### Explaining the tables below

These tables show you the interest rates along with the Annual Percentage Rate of Charge (APRC). We explain APRC at the end of this letter. The rates may vary by Loan to Value (LTV) ratio. We also explain LTV at the end of this letter.

### Fixed interest rates

Fixed interest rate options	Loan to Value Up to 60%	Loan to Value 61-80%	Loan to Value over 80%	Difference in monthly repayments	Difference over the year
1-year	2.9% (3.9% APRC)	2.9% (4.2% APRC)	2.9% (4.4% APRC)	-€131	-€1,572
2-year	2.9% (3.8% APRC)	2.9% (4.0% APRC)	2.9% (4.3% APRC)	-€131	-€1,572
3-year	3% (3.7% APRC)	3% (3.9% APRC)	3% (4.1% APRC)	-€123	-€1,476
5-year	3.2% (3.7% APRC)	3.2% (3.8% APRC)	3.2% (4.0% APRC)	-€108	-€1,296
10-year	3.5% (3.7% APRC)	3.5% (3.8% APRC)		-€84	-€1,008
10-year			3.7% (4.0% APRC)	-€67	-€804

Notes: Figure shows page one of an example redesigned mandatory disclosure letter sent to Treatment group 2. Letterhead with customer and bank information is omitted.

Figure 2.11: Example reminder letter

Mortgage Account Number: 1234567

**REMINDER: You may be able to save money on your mortgage**

Dear X,

We recently wrote to you about the availability of lower mortgage interest rate options and the potential for savings on your monthly mortgage repayments.

This is a reminder to take action to avail of one of these options.

If you wish to take up a lower interest rate for which you are eligible, you can go online at [websiteaddress.com/mortgages](http://websiteaddress.com/mortgages), call us on 01 XXX XXXX, or visit a branch.

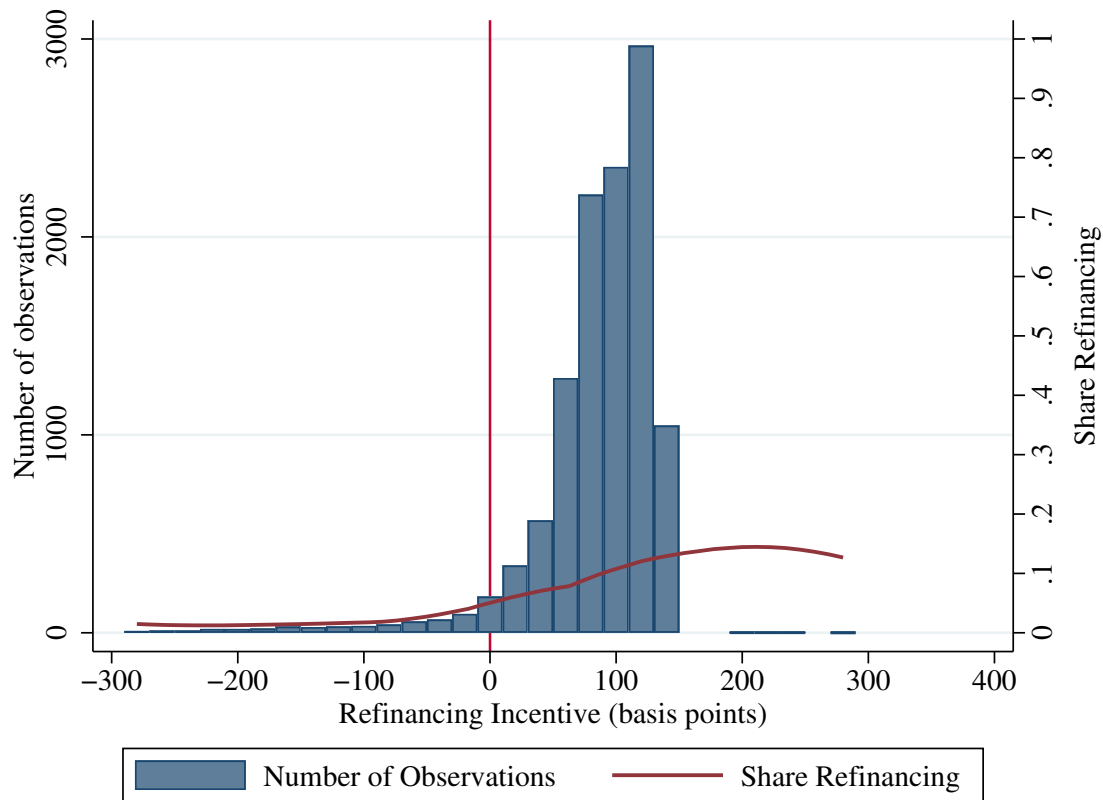
Yours sincerely,

Firstname Secondname

Head of Mortgages

Notes: Figure shows an example reminder letter sent to half of the treated borrowers in the experimental sample. Letterhead with customer and bank information is omitted.

Figure 2.12: Distribution of refinancing incentives



Notes: Figure plots a histogram (left axis) of the refinancing incentives calculated in the experimental data using the model of Agarwal et al. (2013) along with the share of each histogram bin that refinanced within four months of initial treatment (right axis).

Table 2.8: Descriptive statistics across treatment cells

Treatment group	Control	V1	V2	V3	V4	V5	V6
Dublin	0.20 (0.40)	0.21 (0.40)	0.20 (0.40)	0.19 (0.39)	0.21 (0.41)	0.19 (0.39)	0.19 (0.39)
Borrower age	49.74 (9.26)	50.29 (9.37)	49.80 (9.22)	50.08 (9.26)	50.13 (9.61)	50.10 (9.30)	49.87 (9.40)
First Time Buyer	0.41 (0.49)	0.40 (0.49)	0.40 (0.49)	0.40 (0.49)	0.40 (0.49)	0.38 (0.49)	0.40 (0.49)
Mortgage balance	83,503 (84,125)	81,425 (89,826)	80,098 (80,088)	81,530 (90,834)	81,020 (91,867)	81,351 (98,831)	82,548 (87,424)
Interest rate	0.042 (0.003)	0.042 (0.002)	0.042 (0.002)	0.042 (0.002)	0.042 (0.002)	0.042 (0.003)	0.042 (0.002)
Years to maturity	13.87 (8.54)	13.21 (8.54)	13.21 (8.47)	13.25 (8.48)	13.36 (8.50)	13.16 (8.41)	13.38 (8.50)
1-Year savings	1,044 (1,010)	1,037 (1,155)	1,007 (980)	1,021 (1,137)	1,022 (1,101)	1,018 (1,178)	1,037 (1,065)
Covid forbearance	0.09 (0.28)	0.07 (0.25)	0.08 (0.28)	0.07 (0.25)	0.09 (0.28)	0.09 (0.29)	0.08 (0.28)
Observations	1,613	1,587	1,616	1,602	1,629	1,585	1,568

Notes: Table reports means and standard deviations in parentheses of mortgage borrower characteristics in each treatment and control group. Dublin is an indicator for whether the mortgaged property is located in Dublin. Borrower age of the oldest borrower on the mortgage. First-time buyer indicates whether the borrower is a first time-buyer. Mortgage balance is amount outstanding on loan at the time of experiment in euros. Interest rate is the interest rate applicable on the loan at the outset of the experiment. 1-year savings is the amount in euros of savings available to the borrower in the first year after refinancing to the best available rate. Covid forbearance indicates whether the borrower was using Covid payment break (introduced in Ireland in March 2020 to alleviate short-term liquidity constraints faced by borrowers experiencing financial difficulties due to the impact of the pandemic).

Table 2.9: Test of covariate balance by treatment arm

Treatment group	V1	V2	V3	V4	V5	V6
Dublin	0.011 (0.023)	0.002 (0.023)	-0.018 (0.023)	0.022 (0.023)	-0.009 (0.023)	-0.017 (0.023)
Borrower age	0.001 (0.001)	-0.002 (0.001)	-0.000 (0.001)	0.000 (0.001)	-0.001 (0.001)	-0.001 (0.001)
First Time Buyer	0.007 (0.020)	0.000 (0.020)	0.001 (0.020)	0.003 (0.020)	-0.026 (0.020)	-0.001 (0.020)
Mortgage balance	-0.000** (0.000)	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)
Interest rate	-6.971 (4.784)	-0.800 (4.894)	2.058 (4.825)	0.263 (4.916)	2.649 (4.510)	-3.091 (4.922)
Years to maturity	-0.003* (0.002)	-0.004** (0.002)	-0.003* (0.002)	-0.002 (0.002)	-0.003** (0.002)	-0.003* (0.002)
1-year savings	0.063** (0.028)	0.011 (0.030)	-0.004 (0.027)	0.013 (0.030)	-0.002 (0.023)	0.026 (0.031)
Covid forbearance	-0.072** (0.033)	-0.013 (0.031)	-0.071** (0.033)	-0.002 (0.031)	0.010 (0.031)	-0.017 (0.032)
Constant	0.790*** (0.214)	0.684*** (0.222)	0.470** (0.218)	0.510** (0.221)	0.491** (0.207)	0.710*** (0.222)
Equality $p$ -value	0.054	0.453	0.263	0.835	0.408	0.740
Observations	3,200	3,229	3,215	3,242	3,198	3,181
R-squared	0.005	0.002	0.003	0.001	0.003	0.002

Notes: Table reports estimates of a regression of treatment status (an indicator for the treatment heading each column) on a vector of covariates. Each column's sample consists of participants assigned to the control group and the indicated treatment group. Equality  $p$ -value is from the  $F$ -test for joint equality of all of the slope coefficients in a given column. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 2.10: External refinancing treatment effects

	(1)	(2)	(3)	(4)
Disclosure Redesign Treatment	-0.005 (0.005)	-0.005 (0.005)	-0.003 (0.005)	-0.003 (0.005)
Treatment $\times$ Reminder			-0.003 (0.004)	-0.003 (0.004)
Constant	0.037*** (0.005)	0.177*** (0.047)	0.037*** (0.005)	0.177*** (0.047)
Controls		✓		✓
Observations	11,200	11,200	11,200	11,200
R-squared	0.000	0.007	0.000	0.007

Notes: Table reports treatment effects on external refinancing, defined as a borrower prepaying their mortgage with our partner bank and taking out a mortgage with another provider. Control variables in columns 2 and 4 are listed in Table 2.2. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 2.11: Covid classification of employment sectors

Sector Covid Status Category	Employment Sector
1. Working from home (WFH)	J: Information and communication K: Financial and insurance M: Professional, scientific, technical O: Public administration S: Other service activities
2. Business as usual (BAU)	A: Agriculture, forestry, fishing D: Electricity, gas supply H: Transport and storage
3. At home not working (AHNW)	C: Manufacturing F: Construction G: Wholesale and retail trade, vehicle repair I: Accommodation and food services

Notes: Table reports the classification of employment sectors into groups more likely to be working from home (WFH), experiencing business as usual (BAU), and at home not working (AHNW). Prefix letters represent Eurostat Statistical Classification of Economic Activities in the European Community, Rev. 2 (2008) available at Eurostat.



Table 2.12: Internal refinancing treatment effects by employment sector

Employment sector	WFH	BAU	AHNW
Disclosure Redesign Treatment	0.015 (0.011)	0.028 (0.040)	0.026 (0.017)
Treatment $\times$ Reminder	0.046*** (0.011)	0.080* (0.043)	0.066*** (0.017)
Constant	0.100*** (0.0009)	0.079* (0.034)	0.075*** (0.014)
Observations	7,218	494	2,548
R-squared	0.003	0.008	0.006

Notes: Table reports internal refinancing treatment effects within the employment sector subgroups more likely to be working from home (WFH), experiencing business as usual (BAU), and being at home but not working (AHNW) during the estimation window. See Appendix Table 2.11 for our employment sector classification scheme. Dependent variable is an indicator variable equal to one if the borrower internally refinanced, defined as a borrower changing their mortgage product with the partner bank within three months of treatment, and zero otherwise. Disclosure Redesign Treatment is an indicator that the borrower was randomly assigned to one of the six treatment arms. Reminder is an indicator for whether that borrower received a follow-up reminder letter 4-6 weeks after the initial treatment as in Appendix Figure 2.11. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 2.13: Mixture ML estimates: robustness to alternative prepayment assumptions

Parameter	(1)	(2)	(3)	(4)	(5)
Incentive Sensitivity ( $\beta$ )	-125.48*** (1.56)	-1.44*** (0.02)	-1.63*** (0.05)	-1.63*** (0.05)	-1.72*** (0.05)
Fixed Cost of Refinancing ( $\gamma_0$ )		9.66*** (0.03)	8.05*** (0.04)	8.04*** (0.04)	9.01*** (0.15)
Inattention Constant ( $\delta_0$ )			0.57*** (0.05)	1.05*** (0.11)	0.78*** (0.13)
Treatment on Inattention ( $\delta_1$ )				-0.32** (0.13)	-0.35** (0.14)
Reminder on Inattention ( $\delta_2$ )				-0.45*** (0.09)	-0.51*** (0.10)
Fixed Cost Controls					✓
Observations	11,200	11,200	11,200	11,200	11,200

Notes: Table reports maximum likelihood estimates of the mixture model of inattentive refinancing described in the text, but where we adjust the model parameters to take account of the typically short fixation periods which predominate in Irish (and UK) mortgage markets, as distinct from the long-term fixation periods to which the Agarwal et al. (2013) model is originally attuned. Incentive Sensitivity is the coefficient on the Agarwal et al. (2013) refinancing incentive described in section 2.6 using the parameters defined by Table 2.6, with coefficient  $\exp(\beta)$ . The fixed cost of refinancing constant  $\gamma_0$  estimates an average fixed cost term to rationalize observed refinancing variable. The fixed-cost controls allow for differences across groups in the estimated fixed cost of refinancing. The inattention constant  $\delta_0$  allows the inattention index in (2.8) to have a constant term. The inattention treatment effects allow borrowers who treated with redesigned disclosures ( $\delta_1$ ) and disclosure reminders ( $\delta_2$ ) to have different levels of attention. Age is demeaned. Covid indicates whether the borrower was approved for mortgage-payment forbearance with a Covid hardship. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

# Chapter 3

## Identifying Fraud and Building Trust in a Digital Market Platform: Evidence from a Lab Experiment in Nigeria

### 3.1 Introduction

Non-institutional fraud targeted at micro and small enterprises (MSEs) is pervasive across low- and middle-income countries (LMIC) and has risen in the wake of the COVID-19 pandemic (Tade, 2021; Kabanda et al., 2018; Fu and Mishra, 2022). Non-institutional fraud can include phishing,<sup>1</sup> scams to access passwords and log-ins, impersonating a formal institution, offering fake products or services and absconding with payments, and using psychological manipulation to persuade victims to part with money (Garz et al., 2021; Titus et al., 1995).<sup>2</sup> In surveys of digital financial service users, 56% of Kenyan respondents, 33% of Ugandan respondents, and 42% of Nigerian respondents had faced phishing scams in the months after the start of the pandemic (Innovations for Poverty Action, 2021a). The presence of these scams was the most prevalent cited challenge faced in consumer engagement with digital financial services (DFS) in Kenya and Uganda, and the third most prevalent in Nigeria.

---

<sup>1</sup>The fraudulent practice of sending emails purporting to be from reputable companies in order to induce individuals to reveal personal information, such as passwords and credit card numbers

<sup>2</sup>At the core, non-institutional fraud is carried out by individuals or groups who are not affiliated with a formal institution who seek to trick victims into directly sending money, or sending sensitive information that can be used to defraud the victim.

Non-institutional fraud causes immediate and long-term damage. Immediately, it can lead to monetary loss, but also to psychological impacts including anger, difficulties with trust, feelings of violation, stress, and social embarrassment (Vohs et al., 2007; DeLiema et al., 2017). Additionally, because being defrauded is a violation of trust, it may erode trust in counterparties and institutions, a key driver of growth (Francois and Zabojsnik, 2005). Specifically, it can frustrate economic development through the suppressed use of financial services and under-exploitation of advantageous opportunities (Caribou Digital, 2019; Banerjee et al., 2019).<sup>3</sup> Likewise, when fraud operates through interpersonal networks, this might encourage the pursuit of opportunities only within closely trusted networks (Nash et al., 2013).<sup>4</sup>

There is limited knowledge on what mitigation strategies can be taken to reduce fraud, effectively signal authenticity and engender trust, and thereby facilitate the better realisation of the promise of digital financial services to MSEs in LMICs. The economics of fraud are closely related those of adverse selection. Fraud takes place in settings where there is asymmetric information over the quality (and in many cases, existence) of goods, services, or financial rewards (Akerlof, 1970; Miles and Pyne, 2017).<sup>5</sup> While some can persist longer, the great majority of non-institutional fraud is designed for single interactions with unsuspecting victims (Titus et al., 1995; Herley, 2012).<sup>6</sup> Therefore, we test two interventions designed to address the information asymmetries that lead to fraud victimization by improving detection of fraudulent messaging. Similar interventions, particularly related to financial literacy, have been widely used.<sup>7</sup>

Existing efforts to combat the threat of fraud have centered variously on educational interventions which seek to arm users with the means to recognise and sidestep attempted

---

<sup>3</sup>Notwithstanding regulatory differences, the lower rate of digital payments in Nigeria compared to Kenya may reflect lower levels of trust in digital financial service (World Bank Group, 2019).

<sup>4</sup>If this limits the number of suppliers one purchases from, for example, it may lead to monopolistic competition (i.e., differentiation along the dimension of trust) and could lead to higher prices.

<sup>5</sup>Some fraud and scams are related to adverse selection in credit markets since they depend on promises of the delivery of goods in the future (Stiglitz and Weiss, 1981). This includes schemes like advanced fee fraud.

<sup>6</sup>Examples of longer term frauds are interesting, (Darby and Karni, 1973) give the example of joint diagnosis and provision of services, as in auto-mechanic work. Other examples include the credence goods such as antivirus software (Stone-Gross et al., 2013).

<sup>7</sup>A vast array of generic, general audience public information campaigns have developed around the world to encourage awareness in relation to digital fraud. For example, the US Federal Trade Commission has produced a list of actions to avoid fraud (<https://consumer.ftc.gov/articles/how-avoid-scam>) while also publishing a series of fotonovelas as part of its programme for combatting fraud in African American and Latino Communities (<https://consumer.ftc.gov/features/fotonovelas>). UK Finance leads a similar campaign ‘take five to stop fraud’ which has produced toolkits in various formats as well as scenario-based quizzes (<https://www.takefive-stopfraud.org.uk/toolkit/>) and spot the difference videos encouraging consumers to ‘stop, challenge, and protect’ against financial fraud (<https://quiz.takefive-stopfraud.org.uk/>). Similarly, Banco de Portugal has produced a brochure for school-age children of five tips for staying safe while using digital channels to access banking products and services (<https://clientebancario.bportugal.pt/en/material/5-tips-staying-safer-online-toptip>).

deception (Burgoon, 2015), technological solutions to verify counterparty identities and authenticate communications (Conroy, 2017), and centralised algorithmic tools used to detect and flag anomalous patterns of behaviour (Hilal et al., 2022).

In this paper, we analyse results from a behavioural laboratory experiment conducted in partnership with Amana Market, a digital market platform for agricultural trading, with a sample of 780 participants from the partner’s existing network in Nigeria. We test the impact of a series of learning interventions in improving the ability of small business owners to accurately discern fraudulent and genuine communications, as well as in building trust in DFS.<sup>8,9</sup> We also test the potential for a technical solution for the authentication of inbound communications to establish confidence and engagement (a ‘unique communications code’, or UCC). Our learning interventions are just-in-time, range in intensity, and are designed variously to forewarn and encourage vigilance in respect of fraud, up to arming business owners with key signs to watch out for with applied illustrative examples. Our most intensive intervention (Treatment 3) lasts a total of 25 minutes. As part of an experimental task, participants are required to evaluate 20 fictionalised communications scenarios for their authenticity, and report their subjective feeling of confidence in their judgements. To complement our core experimental evaluation, we use data from baseline and endline surveys to explore treatment effect heterogeneity across relevant demographic, behavioural, and experiential factors. We conduct a follow-up knowledge retention quiz four weeks after participants have left the lab.

We do not find evidence that these learning interventions significantly improve discriminant ability between genuine and fraudulent communications, nor do we find evidence that our UCC authentication solution acts as a sufficiently strong signal to increase user engagement with good faith customer outreach. We observe significant increases in the confidence that treated users report in their judgements, notwithstanding the absence of any corresponding improvement in actual underlying accuracy. In this, we highlight the potential for false confidence effects from ineffectual learning interventions, which may engender the subjective feeling of competence, and unintentionally increase susceptibility to fraud victimization through complacency.<sup>10</sup> We do, however, find positive treatment effects on other associated outcomes. We find a significant impact from treatment on trust in DFS, potentially reflecting a heightened confidence on

---

<sup>8</sup>This experiment was conducted in partnership with Amana Market, a digital platform in Nigeria that offers access to market information and financial service to MSEs.

<sup>9</sup>The experiment was pre-specified (Byrne et al., 2022) and pre-registered with the AEA RCT Registry (RCT ID: AEARCTR-0009470).

<sup>10</sup>While some theoretical work has dealt with the unintended consequences of fraud deterrence, our research places this firmly within a the more common everyday context of financial literacy interventions as opposed to detecting and arresting fraudsters (Miles and Pyne, 2017).

the part of treated participants in their ability to discern fraud, successfully navigate the digital financial landscape, and as such engage with confidence with legitimate digital financial counterparties. In addition, we find evidence of increased likelihood of future use of banks and mobile banking, as well as improved knowledge regarding the signs of fraud.

Our paper's contribution to the literature is threefold. Ours is the first paper of its kind to undertake a targeted experimental evaluation of anti-fraud learning interventions among business owners in a LMIC. Gathering evidence in LMIC contexts is particularly important given the heterogeneity in impact of broader financial literacy interventions by country income, as well as by socio-economic status within high income countries (Fernandes et al., 2014; Kaiser and Menkhoff, 2017). A preponderance of existing literature on anti-fraud interventions is focused on high-income country contexts, or with one-size fits all interventions for universal consumption. Our interventions, by contrast, are specifically adjusted to resonate within the information-environment and digital financial landscape faced by small business owners in a LMIC. In contrast to results from the broader financial literacy literature, where stronger effects are found in lower income country contexts, we do not find a strong effect in our population of small business owners in Nigeria.

Second, our paper contributes to our understanding of the subtle and sometimes unintended dynamics of light-touch learning interventions, as well as the interplay between trust, confidence, and ability with respect to deception detection. In this regard, our paper contributes depth and a layered approach to the problem of susceptibility and attempted inoculation to fraud, with several important and pre-specified and pre-registered dimensions. In particular, while in most cases confidence is strongly correlated with performance, we find evidence for the ways in which confidence regarding performance in a task can depart from actual performance (Woodman et al., 2010).

Third, the laboratory setting allows us to study anti-fraud learning interventions and fraud detection in a controlled setting. It is difficult to measure fraud detection ability and prevalence in field settings. While surveys are often subject to issues of recall, since fraudsters seek to deceive, this may add to the noise of these measures. Furthermore, losses from fraud as measured by surveys may rely on a minority of responses in the upper tail (Florêncio and Herley, 2013). Finally, administrative data of complaints may be biased by transaction costs to lodging a complaint (Ba, 2018), administrative capacity (Innovations for Poverty Action, 2021b), or other behavioral factors (Raval, 2020). In contrast, the laboratory setting allows us to circumvent the noise and focus directly on fraud detection ability.<sup>11</sup>

---

<sup>11</sup>It is important to note however that the lab setting includes certain disadvantages which may

Our failure to achieve meaningful treatment effects from learning interventions in respect of our primary detection outcome is in keeping with the relatively underwhelming pattern of results found in the literature on anti-fraud learning interventions, where no or modest effects have been frequently observed (Fernandes et al., 2014). It is, nonetheless, surprising in view of the direct nature of the instruction, the contextually-adjusted and engaging nature of the interventions, and the immediacy of the experimental task that followed. Our results speak to the severity of the challenge that small business owners are likely to face in successfully navigating the noisy landscape of competing communications, and cast doubt on the utility of relatively light-touch, quick-fix learning interventions as meaningful antidotes, even when delivered in a timely fashion.<sup>12</sup> Equipping LMIC small business owners to successfully navigate the contemporary torrent of digital fraud to safely exploit the promise of digital financial markets is evidently a steep challenge. Our results also highlight the risk of false confidence effects from ineffectual learning interventions which may engender the feeling but not the reality of heightened competence. As such, they can be offered as a cautionary lesson for policy in this domain, which recommends introspection with respect to the type of interventions tested here, and their content and intensity. Our results also highlight the critical importance of rigorous pre-testing of planned interventions in this domain to establish what works, and just as importantly, what doesn't.

The paper proceeds as follows: Section 3.2 provides a review of relevant literature, Section 3.3 describes the experimental design and empirical strategy, Section 3.4 reports our empirical results, and Section 3.5 concludes.

## 3.2 Literature

### 3.2.1 The impact of financial education interventions

Considerable evidence has been gathered on the impacts of financial education programmes on financial knowledge and behavior, showing a positive effect on financial knowledge and behaviors. Here we draw on four comprehensive meta analyses of these impacts.<sup>13</sup> Relying on RCT evidence, Kaiser and Menkhoff (2017) and Kaiser et al.

---

limit the reliability of our evaluation of learning interventions, chiefly the lack of ecological validity. In our case, participants do not face real financial or other exposure from erroneous judgement of communications, and this may limit the seriousness with which they treat the presented task. This is one respect in which field experimentation may offer a more reliable assessment of the learning interventions under evaluation.

<sup>12</sup>We cannot exclude the possibility that true effects from our learning intervention may fall below our minimum detectable effect. However, the economic significance of treatment effects in this range would be limited.

<sup>13</sup>Fernandes et al. (2014), aggregating 168 cases; Miller et al. (2015), aggregating across 188 cases; Kaiser and Menkhoff (2017), reviewing 126 studies; and Kaiser et al. (2021), reviewing 76 studies. While the latter two draw in fewer studies, they do so to put a higher weight on experimental evidence.

(2021) find that there are positive impacts on financial behavior from financial education initiatives, estimating effect sizes of 0.08 SD and 0.1 SD respectively.<sup>14</sup> In observational studies, evidence is more mixed with some behaviors unchanged or unexplained (Miller et al., 2015; Fernandes et al., 2014). While Fernandes et al. (2014) found weaker effects from initiatives in low-income samples within country, countries with higher average income (and education) have weaker effects (Kaiser and Menkhoff, 2017).<sup>15</sup>

Financial education may differ in its impacts as the intensity changes, but does not seem to change much as the mode of delivery changes. Both Fernandes et al. (2014) and Kaiser and Menkhoff (2017) find evidence of ‘returns to intensity’ of financial education, with more hours of education producing larger effects on behaviour, while Miller et al. (2015) find mixed evidence.<sup>16</sup> Other programme characteristics which could reasonably be expected to influence the effectiveness of interventions (e.g. the age and gender of participants, the delivery channel, duration exposed to the treatment, whether the intervention was staged at school, in the community, or in the workplace) were not shown to be significant (Miller et al., 2015; Kaiser and Menkhoff, 2017). Beyond heterogeneity in effects, other factors may also matter in understanding the size of effects: whether the outcome is difficult to change, if there are heterogeneous effects on subgroups of interest, the cost per participant, and ease of scalability (Kraft, 2020).

Finally, there is some evidence for the importance of the dynamics of financial education interventions, including the speed at which the effect of the education decays and the optimal timing of financial education. Fernandes et al. (2014) finds that effect size decays over time, with decay over time being stronger for larger interventions.<sup>17</sup> The authors envisage a reduced role for financial education that is not acted upon soon afterward, and suggest a role for ‘just-in-time’ financial education tied to specific behaviours it intends to help.<sup>18</sup> Similarly, Kaiser and Menkhoff (2017), find that positive effects are associated with providing financial education at a ‘teachable moment’ (i.e. when teaching is directly linked to decisions of immediate relevance to the target group). Kaiser et al. (2021) find a less rapid decay in treatment effects, though still little support for the long-run sustainability of effects.

---

<sup>14</sup>Interestingly, the authors note that this is similar in magnitude to effect sizes reported in meta-analyses of behaviour change interventions in other domains such as health or energy conservation.

<sup>15</sup>This may be attributable to diminishing marginal returns to additional financial education.

<sup>16</sup>In particular, they find that intensity was weakly significant in some specifications of the model it was not significant in the others.

<sup>17</sup>More specifically, the authors observe equal effects for 6 hours of intervention at no delay and 18 hours of intervention at 10 months of delay, and equal effects of 1 hour of instruction at no delay and 12 hours at 10 months of delay. Even large interventions with many hours of instruction have negligible effects on behaviour 20 months or more from the time of intervention.

<sup>18</sup>Although Lusardi (2015) notably argues that ‘just in time is too late’.



## Educational interventions to curb susceptibility to financial fraud

The high social cost associated with fraud, and the difficulty in deception detection has led to significant academic attention on the topic, including analysis of whether interventions can help and which are helpful (Burgoon, 2015). This has led to a diverse array of educational interventions including traditional training sessions, warning messages, as well as consumer advice and decision-making heuristics aimed at improving participants' performance in fraud detection in various digital settings. Many of these interventions have been experimentally tested, showing varying degrees of promise for improvement in detection deception. Several studies have demonstrated the potential for targeted fraud warning messages to reduce the susceptibility of consumers to fraud in the short term (up to five weeks post-intervention) (Anderson, 2003; Scheibe et al., 2014). Elsewhere, a weakening of the resistance offered by warning interventions to fraud susceptibility has been evidenced at intervals up to six months post-intervention (Burke et al., 2020). As in the broader financial education literature, a major limitation in durability assessments, however, is the time-horizon over which the possibility of decay is typically being evaluated, with little evidence available which demonstrates the longer-term durability of simple warning interventions beyond a number of months.

Work on deception detection has broadened the scope of outcomes beyond accuracy, to include 'hits' and 'false alarms', outcome variables borrowed from signal detection theory. While warnings, meant to make deception salient, are not always successful in improving overall accuracy in detection deception, their impact may be evident when we disaggregate by hits and false alarms. For example, while Biros et al. (2002) finds a positive impact on accuracy (driven by hits), both George et al. (2004) and Grazioli and Wang (2001) find null effects.<sup>19</sup> This might be because warnings increase skepticism generally. Xiao and Benbasat (2015) emphasize how carelessly designed warning messages may be effective at increasing hits, but at the cost of increasing false alarms.<sup>20</sup> In an online experiment, the authors test warnings with positive and negatively framed advice about bias in recommendations on an e-commerce website against a standard warning. While all interventions were effective in improving the 'hit' rate in detecting manipulations, only the negatively framed intervention was effective in improving overall detection precision (more hits and fewer false alarms), with the other interventions increasing false alarms as well.

<sup>19</sup>More specifically, George et al. (2004) find no impact of warnings on participants' accuracy in a hypothetical interview scenario. Likewise Grazioli and Wang (2001) find that the pre-issuance of priming material from the Federal Trade Commission on internet fraud did not improve the ability of subjects in a lab experiment to successfully detect fraudulent deceptive manipulations embedded in a commercial webpage.

<sup>20</sup>See also: Burgoon et al. (1994).

Results from training interventions are mixed. Biros et al. (2002) find that traditional training had no effect on detection success or the rate of false alarms, while warnings about data quality did increase detection success. Warnings combined with just-in-time training resulted in better detection success, but at the cost of a greater amount of false alarms. However, George et al. (2004) finds training in deception cue recognition one week in advance of the deception detection task does improve detection accuracy.

### 3.2.2 User-centered digital tools to combat fraud

Quite apart from educational initiatives designed to arm users with the know-how to discriminate effectively between genuine and fraudulent activity, a vast array of digital tools have been devised with the objective of detecting and preventing fraud in digital commerce. Rather than relying on training a market of discerning users with an adequate filter to separate fraudulent from genuine communications, digital tools offer the potential for variously ‘smart’ and automated solutions which efficiently perform that function on our behalf. These solutions can be broadly classified into two camps: detection (the ability to identify suspicious patterns indicative of fraudulent activity), and authentication (establishing the likelihood that a person is who they say they are). Here we focus on user centered authentication tools.<sup>21</sup>

The proliferation of such tools is driven not only by the need to stay ahead of fraudsters whose learning rate is such as to gradually erode the integrity of any given defensive protocol, but also by the virtue of the fact that such tools introduce frictions to the user experience, and a seamless digital customer experience has become a more important competitive factor (Herley, 2009). Platforms are therefore compelled to continually thread the needle between elegant user experiences on the one hand and robust fraud prevention through burdensome or intrusive layers of authentication on the other.

A number of user-centred authentication tools seek to establish that a product user is who they claim to be while interacting and transacting in digital space. In the industry lexicon, this is traditionally achieved by validating something the user has (possession), something they are (inherence), or something they know (knowledge) (Velásquez et al., 2018). A wide variety of tools have grown up around this basic principle as means of implementation at different times and in different contexts. These tools introduce varying degrees of friction and efficacy. These include personal log-in credentials (i.e., user name and password), two-factor authentication, physical biometrics, and know your customer (KYC) protocols. Each of these tools come with advantages and weaknesses,

---

<sup>21</sup>Beyond these, and of less immediate relevance to our purpose, is a set of other institutional tools used to detect financial fraud. Hilal et al. (2022) provides an instructive survey of machine-learning algorithmic anomaly detection methods in the field of financial fraud, where the majority of applications are found in insurance and credit card markets.

and can often be undermined by data breaches, creative social engineering on the part of fraudsters, or poor practices on the part of users. Frequently they are used in various overlapping layers and combinations as a means of reinforcing the reliability of gate-keeping.

Increasingly, vendors are seeking to flexibly match the level of friction imposed by authentication layers to the risk of the underlying transaction, and to the preferences of the user (Conroy, 2017).

### 3.2.3 Confidence and performance

A separate literature which our paper connects with is that which studies the relationship between confidence and task performance. This relationship has received considerable attention in diverse areas spanning from sports psychology, to education, to management science, where a strong positive correlation is typically deemed to hold. Confidence tends to be positively correlated with task performance. For example, in a meta-analysis of confidence effects in sports performance spanning 42 studies, 76% found a positive association, with a mean effect size (correlation coefficient) of 0.23 observed (Woodman and Hardy, 2003). Stankov and Crawford (1996) report correlation coefficients with performance in a series of perceptual tasks ranging from 0.32 in the judgement on the length of a line to 0.62 in a vocabulary test. This empirical relationship is supported theoretically by self-efficacy theory (Bandura, 1977), which stresses the importance of perceptions of personal capabilities as a central determinant of successful outcomes.<sup>22</sup> Some notable exceptions to the rule, however, have demonstrated how confidence may give rise to complacency, risk-taking, and reduced effort towards a task which can negatively impact upon performance (Vancouver and Kendall, 2006; Vancouver et al., 2001)<sup>23</sup> Finally, Woodman et al. (2010) demonstrate how the relationship between confidence and performance is not always linear because the introduction of doubt can result in increase confidence and focus. Testing the prediction that a reduction in self-confidence (i.e. an introduction of an element of self-doubt) in the context of a familiar task may aid performance by inducing additional effort in the context of a simple physical task, Woodman et al. (2010) find evidence in support,

<sup>22</sup>The proposed mechanism through which self-efficacy is said to operate via coping, persistence, and effort. See also Bandura et al. (1999).

<sup>23</sup>Vancouver and Kendall (2006) demonstrate how self-efficacy can negatively relate to motivation and performance in a learning context, specifically, college student exam study time and performance. The authors show that over the course of four examinations, as self-efficacy increased by a grade (i.e. self-efficacy magnitude increasing from an anticipated B to A), study time decreased by 15 minutes and exam performance decreased by nearly a quarter grade. Similarly, Vancouver et al. (2001) show how complacent self-assurance can undermine motivation to adversely affect a person's performance in across time in an experimental mastermind game. The authors assess what happens to a person's performance as their self-efficacy changes over time, and find a significant negative relationship, with suggestive evidence that when individuals had higher self-efficacy, they may have committed to their responses too early ("their self-efficacy encouraged them to act rather than think").

indicating how while confidence may indeed be positively associated with performance, the relationship is not necessarily positive and linear.

### 3.3 Experimental design

We investigate the effects of learning interventions on fraud detection, confidence, and trust, using a laboratory experiment. To conduct the experiment we worked with Amana Market, the Busara Center for Behavioral Economics and the experimental lab at Ahmadu Bello University (ABU) in Zaria, Nigeria. ABU is a public research university and the largest university in Nigeria. Zaria is a city of over 700,000 people in Kaduna State, which is based in Northern Nigeria. Study participants were recruited by Amana Market, and then lab sessions physically and virtually managed by the Busara and ABU staff from Kenya and Nigeria. A total of 45 respondents across the 3 treatment arms participated in the pilot session which was held on 21st July, 2022 while a total of 780 respondents across 52 lab sessions participated in the main experiment, which ran from July 25th - August 26th, 2022.

#### 3.3.1 Experimental flow

Amana Market recruited participants from communities in Zaria using their existing agent network. Participants were either Amana Market users or similar in profile to Amana Market users, thus predominantly those within the agriculture value chain. The profiles of the participants (name, phone number, gender, etc.) were uploaded into a database and then randomly assigned to a given lab session. Those assignments were then sent to the Amana Market agent who would bring participants to the experiment at the designated date and time.

A total of 15 participants were invited to the lab for each session. In the waiting room, participants verified their identity with a staff member to ensure they were participating in the correct session, and were randomly assigned a seat number in the lab. Upon entering the lab, informed consent was obtained from each participant. During the experiment, participants first underwent a baseline survey, followed by an educational intervention or control condition, followed by completing fraud detection experimental task and an endline survey. Upon finishing the endline survey, a UCC was created for participants to be used for the follow up process and survey. Finally, participants are paid a total of NGN 4,500 in cash as a transport reimbursement and attendance incentive before leaving the lab.

After each participant visited the lab, we conducted an SMS and phone survey. Three weeks after the lab session, we sent each participant an SMS in which they are requested

to respond confirming their month and year of birth. The SMS was randomly assigned to contain or not contain the UCC that was assigned to them at the end of the lab experiment. A week after we sent the SMS (and four weeks after the lab session), we conducted a follow up phone survey where we asked them about their experience with the SMS that was sent to them, and questions about key signs of fraud.

### 3.3.2 Experimental design

#### Education interventions

All participants are randomised into either a control group or one of three treatment groups. Each of the three treatment groups receives a variation of a learning intervention aiming at helping participants distinguish between genuine and fraudulent communication (either a general warning treatment, or one of two targeted educational interventions). These simple interventions are meant to be brief and replicate common approaches used in anti-fraud campaigns and training. The control group initially receives no additional warning about fraud, while the treatment arms receive some warning or education. The four experimental arms are presented in Table 3.1.<sup>24</sup>

#### Unique Customer Code (UCC) intervention

At the end of the lab experiment, all participants were assigned a UCC. Participants are randomly allocated into one of two equally weighted groups: the non-personalised UCC group (these subjects are assigned a randomly generated 5-digit UCC), and the personalised UCC group (these subjects are instructed to choose their own 5-digit UCC). The lab staff explain that this code will be used to verify the authenticity of future communications with participants. All codes are recorded centrally, and sent to subjects by SMS to keep as a record.

#### Survey and experimental task

A baseline survey collected demographic information, attitudes and experience with digital financial services, and previous experience with fraud. The primary outcome

---

<sup>24</sup>It could be argued that our control group is inherently primed to think about fraud by virtue of being confronted with the experimental task where the evaluation of authenticity and fraudulence is the clear objective, and as such that the control contains an element of implicit treatment which may serve to mute the additional treatment effect from the subsequent learning interventions. Such implicit treatment is inescapable in the lab experimental setting, and will be consistent across treatment arms, meaning that the additional benefit of the learning interventions under evaluation will still be identified. However, the possibility of such a placebo effect will mean that our estimated treatment effects cannot be treated as an ecologically valid estimate of a treatment effect that might be observed in the field, where the control benchmark condition involves no such implicit task-based priming.

Table 3.1: Anti-Fraud Campaign Interventions

<b>Control</b>	Control group receive the lab manager’s session introduction and undergo the consenting process, but receive no additional warning or educational information related to fraud.
<b>Treatment 1</b>	On top of the lab manager’s session introduction, and the consenting process, T1 subjects receive on-screen general warning messages stating, “Digital fraud represents a threat to small businesses in Nigeria. Fraudsters may contact you pretending to represent legitimate businesses or agencies, in an effort to take your information or your money. Be on the lookout for signs of potential fraudsters in the communications you receive – over the phone, by email, or in person”.
<b>Treatment 2</b>	On top of the lab manager’s session introduction, and the consenting process, T2 subjects receive an on-screen written list of 7 key signs of potential fraud which is narrated in an audio file (key signs are detailed in Table 3.26. This information is prefaced by a general warning message (see Treatment 1). To aid recall, subjects are prompted to write down the key signs upon completion, before replaying the 7 key signs and filling in the gaps in their answer sheets. Participants’ notes are collected before the remainder of the lab session.
<b>Treatment 3</b>	On top of the lab manager’s session introduction, and the consenting process, T3 subjects receive an on-screen written list of 7 key signs of potential fraud, complemented with applied illustrative examples which is narrated in an audio file. This information is prefaced by a general warning message (see Treatment 1). To aid recall, subjects are prompted to write down the key signs upon completion, before replaying the 7 key signs and filling in the gaps in their answer sheets. Participants’ notes were collected before the remainder of the lab session.

Notes: The table describes the four groups into which subjects are randomly assigned.

measure is performance on a task where participants attempt to identify fraudulent or legitimate communications from a fictitious sender. Participants are exposed to 20 fictitious scenarios, half fraudulent and half genuine. All participants are exposed to the same 20 scenarios, with the aim of discerning genuine from fraudulent communications. The scenarios are shown to each participant in a random order. A description of scenarios presented is given in Table 3.27, along with illustrative examples in Figures 3.36 and 3.37. We measure both their raw performance (whether they identified the scenario correctly as fraudulent or genuine) as well as their confidence level in their answer. An endline survey asks participants for their trust and willingness to engage in digital financial services.

### Randomisation and groups

Participants are randomly assigned at three independent stages during the lifetime of the experiment. Firstly, participants are randomly assigned into one of three educational interventions or the control group. Secondly, participants are assigned to either receive an automatically generated 5-digit UCC, or to personalise their own 5-digit code. Finally, participants are assigned to either receive an SMS which has their UCC embedded, or which does not contain their UCC.

### 3.3.3 Empirical strategy

To estimate the causal effect of treatments on participant ability to distinguish between fraudulent and genuine communications, we perform the following empirical specification:

$$Y_{is} = \alpha + \beta_1(\textit{Treatment 1})_i + \beta_2(\textit{Treatment 2})_i + \beta_3(\textit{Treatment 3})_i + \epsilon_i \quad (3.1)$$

where we define  $Y_{is}$  to be one of the outcome variables described in table 1 or table 2 for participant  $i$  (and scenario  $s$ ),  $(\textit{Treatment 1})_i$ ,  $(\textit{Treatment 2})_i$ , and  $(\textit{Treatment 3})_i$  denote treatment arms described in Table 3.1.  $\beta_1$ ,  $\beta_2$ , and  $\beta_3$  estimate the corresponding treatment effects.

We perform the following hypothesis tests after running the regression using accuracy (i.e., accurate identification of fraud scenarios) as an outcome. We additionally partition the accuracy outcome to separately analyse accuracy in respect of genuine scenarios (true positives) and accuracy in respect of fraudulent scenarios (true negatives).

- Hypothesis 1.1:  $H_0 : \beta_1 \leq 0$ . Providing MSEs with a general warning message about fraud alone (with no further educational intervention) improves their ability to distinguish between genuine and fraudulent communications (T1 vs. C).
- Hypothesis 1.2:  $H_0 : \beta_2 \leq \beta_1$ . Providing MSEs with seven key warning signs for potential fraud in a simple format (written/audio) improves their ability to distinguish between genuine and fraudulent communications, still further than can be achieved by a general warning message alone (T2 vs. T1).
- Hypothesis 1.3:  $H_0 : \beta_3 \leq \beta_2$ . Illustrating applied examples of fraudulent communications in a simple format (written/audio) improves MSEs' ability to distinguish between genuine and fraudulent communications, still further than can be achieved with simple warning signs alone (T3 vs. T2).

The other hypotheses relating to confidence, trust of DFS, and likelihood of using DFS (outlined in Table 3.3) are tested in the same manner. We also test if any of the treatments improve the ability to distinguish between genuine and fraudulent communications by running a pooled specification:

$$Y_{i,s} = \alpha + \beta_1(T)_i + \epsilon_i \quad (3.2)$$

where  $T_i$  indicates that participant  $i$  receives any of the three treatments. This specification is used to test  $H_0 : \beta \leq 0$ .

### **Heterogeneous treatment effects**

The individuals targeted by scammers exhibit significant heterogeneity in their skills and abilities. For example, the ability (e.g., sophistication, skepticism, or experience) of potential victims may matter for detection of fraud and scams (Vohs et al., 2007; Holtfreter et al., 2010). We explore how treatment effects may vary along important experiential, attitudinal, and demographic dimensions, which are elicited as part of the baseline survey which precedes the experimental task. The baseline survey collects information relating to the participants' level of experience with information communication technology (ICT) and DFS, as well as exposure to fraud and scams, levels of trust, propensity to self-control, and attitude towards risk.

For most of these outcomes, a standardized index is computed then split into types by those who are above or below average according to that index.



We additionally investigate the extent to which relevant socio-demographic factors interact with the learning interventions: specifically sector of employment, age, education, and gender.

To estimate these heterogeneous treatment effects, we perform the following specification:

$$Y_{is} = \alpha + \beta_1(T\ 1)_i + \beta_2(T\ 2)_i + \beta_3(T\ 3)_i \\ + \delta_1(T\ 1 * Moderator)_i + \delta_2(T\ 2 * Moderator)_i + \delta_3(T\ 3 * Moderator)_i + \epsilon_i \quad (3.3)$$

Where  $(T\ 1 * Moderator)_i$ ,  $(T\ 2 * Moderator)_i$ , and  $(T\ 3 * Moderator)_i$  represent the interaction of each treatment arm with a given moderating variable listed above, and  $\delta_1$ ,  $\delta_2$ , and  $\delta_3$  represent the change in the slope of the corresponding simple effects captured by  $\beta_1$ ,  $\beta_2$ , and  $\beta_3$ .

We test the null hypothesis that the change in the slope is equal to zero in each case, i.e.  $H_0 : \delta_1, \delta_2, \delta_3 = 0$ .

### UCC follow-up exercise

We evaluate the potential for the UCC to act a signal of authenticity in user communications and to elicit engagement. First, we assess how the likelihood of user engagement through SMS response is impacted by the presence of the UCC embedded in the outreach. We estimate the specification:

$$Y_i = \alpha + \eta(UCC)_i + \epsilon_i$$

where  $Y_i$  is a variable indicating if the participant responded and  $UCC_i$  indicates that the UCC was included in the communication. We test the null hypothesis that the coefficient is less than or equal to zero, i.e.  $H_0 : \eta \leq 0$ . Additionally, we will explore whether having personalised the UCC at the close of the lab session, as distinct from having one automatically assigned, strengthens the signal of authentication and further elicits user engagement. We estimate the specification:

$$Y_i = \alpha_1 + \eta_1(UCC)_i + \eta_2(UCC * Personalised)_i + \epsilon_i$$

We test the null hypothesis that the coefficient is less than or equal to zero, i.e.  $H_0 : \eta_2 \leq 0$ .

## Learning by doing

In addition to considering heterogeneity in learning effects by fraud experience, we explore learning by doing within the experiment, to assess whether performance improves and confidence grows over the course of the experimental task. To evaluate this, we test whether those scenarios that appeared later in the order were more often correctly identified by participants. In the same manner we test whether scenarios appearing later in the order were judged with a higher degree of confidence.

$$Y_{is} = \alpha + \theta(Order)_{is} + \epsilon_i \quad (3.4)$$

where  $Order_{is}$  is the order scenario  $s$  was presented to participant  $i$ , and  $Y_{is}$  represents the firstly accuracy, and secondly confidence in scenario judgements. Again, we test the null hypothesis that the coefficient is less than or equal to zero, that is  $H_0 : \theta \leq 0$  to test for learning in both cases.

## Knowledge retention quiz

We additionally investigate the degree to which participant recall of the key signs of fraud varies across treatment arms in a knowledge retention follow-up quiz administered over the phone four weeks following the completion of the main experimental task in the lab. Participants are presented with three multiple-choice questions in which they are asked to correctly identify which item represents a key sign of fraud. In this extension, we assess the relative durability of educational interventions administered in the lab, with the specification:

$$Y_i = \alpha + \beta_1(Treatment\ 1)_i + \beta_2(Treatment\ 2)_i + \beta_3(Treatment\ 3)_i + \epsilon_i \quad (3.5)$$

where  $Y_i$  denotes the quiz score of participant  $i$ ,  $(Treatment\ 1)_i$ ,  $(Treatment\ 2)_i$ , and  $(Treatment\ 3)_i$  denote treatment arms described in Table 3.1.  $\beta_1$ ,  $\beta_2$ , and  $\beta_3$  estimate the corresponding treatment effects. We test the null hypothesis that each coefficient is less than or equal to zero, i.e.  $H_0 : \beta_1, \beta_2, \beta_3 \leq 0$ .

## Descriptive statistics

Table 3.4 reports descriptive statistics for key demographic, experiential, and attitudinal characteristics across treatment cells. In order to attribute any observed difference in specified outcomes to the impact of the interventions under evaluation, it is important that randomisation was performed effectively, with the result that treatment groups are well-balanced in key covariates at the outset.

Following McKenzie (2015), Table 3.5 shows a pairwise regression of treatment status (control vs. each of our treatment groups) on the same vector of covariates included in Table 3.4, which may be correlated with our outcome variable of interest, to ascertain whether these factors differ systematically and help to predict treatment status. While we find a high degree of statistical balance in most cases, we do observe some evidence of significant imbalance, most notably in gender. Following Mutz et al. (2019), and to adjust our estimation for potentially confounding influence, and to obtain more precise treatment effect estimates, we include as a vector of controls in our estimation of treatment effects all those pre-specified prognostic variables of interest listed in Table 3.2. As an alternative approach to the selection of relevant control variables, we use partialling out lasso linear regression, which selects relevant control variables for inclusion in the estimation regression, and find our estimation is unchanged.

## 3.4 Results

### 3.4.1 Effects of fraud education interventions

#### Main effects

In this section, we outline the impact of our experimental fraud education interventions on recipients' fraud detection performance, confidence in their performance, trust in DFS providers and their likelihood to use these providers in the future, when compared against the control group.

We firstly show in Table 3.6 that no intervention succeeds in significantly impacting upon the overall level of accuracy across fraudulent and genuine scenarios. Separating performance by true positives and true negatives, we find no specific impact on true negatives, but that Treatment 3 yields a significant negative impact on true positive performance of approximately 8%. This may reflect a heightened level of skepticism engendered by exposure to treatment, with the result that genuine scenarios are more likely to be rejected as fraudulent (that is, a higher rate of 'false alarms'). Table 3.6

also reports tests of our pre-specified hypotheses that each successive treatment arm delivers incremental added value compared to the preceding arm in the sequence (i.e. Treatment 1 dominates Control, Treatment 2 dominates Treatment 1, and Treatment 3 dominates Treatment 2). In each case, we fail to reject the null hypothesis of no incremental benefit. We additionally investigate whether treatment effects are in evidence when treatment arms are pooled (all treatments together, or just the more intensive treatment arms 2 and 3) in Table 3.24, and still find no effect.<sup>25,26</sup>

It is important to note that, all estimated coefficients reported in Table 3.6 are below our estimated ex-post minimum detectable effect in each primary outcome, on the basis of observed outcomes in the control condition (see Appendix section 3.6.3). As such, we are not sufficiently powered to estimate with confidence true treatment effects that fall in the region of coefficient magnitudes reported here. However, we do not regard effects that fall so far below our minimum detectable effects as being economically meaningful for the purposes of this experiment.

We observe in Table 3.7 that notwithstanding the absence of any positive impact on overall detection ability, we do find consistent positive impacts from Treatments 2 and 3 on the level of confidence that participants report in their judgements over the presented scenarios overall, and both in respect of genuine and fraudulent scenarios of 3-5%. These effects are modest, but illustrate the troubling potential for learning interventions to engender a subjective feeling of confidence with respect to a given task, without necessarily delivering any actual improvement in underlying ability and performance. Again, we additionally test for incremental impacts from each treatment arm compared against the preceding arm in the sequence. For each outcome, we fail to reject the null hypothesis of no incremental benefit from Treatment 1 compared against the control. We do, however, for each outcome reject the null hypothesis of no incremental benefit associated with Treatment 2 when set against Treatment 1, indicating that providing participants with warning signs for potential fraud does impact upon confidence, beyond what can be achieved by a general warning message alone. Finally, we fail to reject the null hypothesis of no further incremental benefit associated with Treatment 3 when set against Treatment 2.<sup>27,28</sup>

---

<sup>25</sup>While we do not find consistent significant evidence of treatment effects on accuracy, Figure 3.1 depicts how the directional pattern of estimated effects points to a dis-improvement in the rate of true positives (i.e. a higher likelihood of false alarms), and an improvement in the rate of true negatives (i.e. a higher percentage of hits) in proportion with the intensity of the treatment administered, consistent with an overall increase in skepticism towards inbound communications.

<sup>26</sup>In Table 3.14, we re-estimate results depicted in Table 3.6, but using partialling out lasso linear regression, and find consistent treatment effect estimates.

<sup>27</sup>We additionally reported aggregated treatment effects when treatment arms are pooled (all treatments together, and just the more intensive treatment arms 2 and 3) in Table 3.25.

<sup>28</sup>In Table 3.15, we re-estimate results depicted in Table 3.7, but using partialling out lasso linear regression, and find consistent treatment effect estimates.

In Table 3.8 we evaluate the impact of our interventions on an inverse-correlation weighted matrix of reported trust in DFS.<sup>29</sup> In respect of Treatment 2, we find a significant and positive effect, with an approximately 0.2 standard deviation increase in the standardised inverse-correlation weighted index of trust, when compared against the control group. This finding is shown to be robust in a specification which additionally adjusts for the trust index measured ex-ante as part of the baseline survey. This suggests that once armed with experience of the anti-fraud learning interventions, participants are more likely to trust DFS providers. This may reflect a heightened confidence on the part of treated participants in their ability to discern fraud, successfully navigate the digital financial landscape, and as such engage with confidence with legitimate digital financial counterparties.

To complement our analysis of intervention impacts on endline trust, we additionally examine the impact of treatment on the likelihood of future use of various specific financial service entities in Table 3.9. We find evidence for a heightened likelihood of the future use of banks and mobile banking of approximately 10% and 7% respectively, with no such significant effect observed in respect of mobile money operators, online platforms, or agents. For each outcome, we fail to reject the null hypothesis of no incremental benefit from Treatment 1 compared against the control. We do, however, in respect of banks and mobile banking reject the null hypothesis of no incremental benefit associated with Treatment 2 when set against Treatment 1, beyond what can be achieved by a general warning message alone. Finally, we fail to reject the null hypothesis of no further incremental benefit associated with Treatment 3 when set against Treatment 2. It is possible that heightened likelihood of future use may reflect participants' increased confidence in filtering out fraudulent communications, enabling them to engage with confidence with relevant DFS providers.

### **Heterogeneous effects**

We next examine the extent to which the impact of our learning interventions varied across subgroups of interest, including relevant experiential, behavioural, as well as demographic characteristics. Specifically, we test whether our treatment effects are more pronounced across levels of fraud, DFS, or ICT experience, self-control, risk appetite, or trust. We additionally test whether sector of employment (agriculture vs. non-agriculture), age, education, or gender significantly interacts with our interventions. We test for these effects across our four primary outcomes of interest: overall accuracy, true positives, true negatives, and confidence.

---

<sup>29</sup>See Table 3.28 for complete definition of all variables used.

Do these experiential and demographic factors act as substitutes or complements to the anti-fraud campaign? That is, firstly, do we find differential performance in level terms across these subgroups in primary outcomes, and secondly, do particular subgroups respond more intensively to the learning interventions tested? We test these questions by exploring heterogeneous treatment effects in overall accuracy, true positives, true negatives, and confidence. These results are reported in Tables 3.16-3.23).

For the most part, we find no strong evidence for significant level or interaction effects across these subgroups of interest. However, in some instances, we do find evidence of significant coefficient estimates in interactions across Tables 3.16-3.23. Mindful of the difficulty in direct interpretation of interaction effects in regression tables, to assist with interpretation, we take this subset of factors where we seem to observe significant differences, and depict the relevant relationships graphically in Figures 3.4-3.33.

First considering level differences in outcomes across relevant subgroups, we find evidence that DFS experience may act as a substitute for the learning interventions, with participants with high DFS experience performing better in overall accuracy than those with low experience (Figures 3.4 and 3.5). This is an intuitive result, likely reflecting an advantage conferred by experience in recognising what are plausible and more suspect communications. We find tentative evidence that lower trust levels may also partially substitute for the learning interventions, with generally higher performance in accuracy among those with low trust when compared against higher trust counterparts (see Figures 3.20-3.25). It is likely that lower trust participants show a higher degree of skepticism towards inbound communications purporting to originate with a genuine service provider, with a lower threshold for flagging fraud. This interpretation is supported by the fact that the higher accuracy for low trust participants is driven by a higher rate of true negatives, and not true positives. We additionally find that low trust participants report a higher confidence in level terms in their decisions than high trust counterparts. We observe some evidence in level terms across treatment arms that men report a higher degree of confidence in their judgements than women, and more tentative evidence that men have a higher level of overall accuracy than women (Figures 3.28-3.31). Those with high levels of self-control, and with higher risk appetite appear to demonstrate a higher level of confidence in their judgements over presented scenarios across treatment arms (Figures 3.26 and 3.27, and Figures 3.32 and 3.33).

Considering next how relevant characteristics may significantly interact with our treatment interventions, we find for the most part, no evidence of significant interaction effects across these experiential, behavioural, and demographic factors of interest. However, we do find isolated instances of apparent meaningful divergence in outcomes across subgroups, where significant interaction effects in treatment can be observed.

With the interpretative aide of graphical representation, it is clear that while we do observe isolated instances where coefficients associated with a given treatment arm do differ significantly within partitioned subgroups, largely, these outcomes do not meaningfully differ from their respective controls. As such, in these cases, we can say we do not find strong or consistent evidence that certain subgroups respond more intensively to treatment than others.<sup>30</sup>

One tentative trend observed is an apparent convergence in confidence as treatments become more intense for certain subgroups that show an initial deficit. Lower risk appetite participants (Figure 3.32), females (Figure 3.30), and to a lesser extent low self-control (Figure 3.26) and low DFS experience participants (Figure 3.12) show signs of catch-up from an initial deficit in confidence when compared against their opposing indexed counterparts (i.e. high risk appetite, males, high self-control, and high DFS experience participants respectively). While the evidence in this regard is not strong, it is intuitive that those initially lacking in confidence in comparative terms may show a higher marginal return in the confidence-engendering effect of targeted learning interventions.

### **Learning by doing**

We assess additionally whether performance improves and confidence grows as participants proceed through scenarios presented as part of the experimental task. Specifically, whether those scenarios that appear later in the sequence are more accurately and confidently called. In Table 3.10, we find no evidence for such learning effects in accuracy or confidence, failing to reject the null hypothesis of no positive effect. Figures 3.2 and 3.3 graphically depict these outcomes over the sequence of scenarios presented. The absence of learning effects of this sort may reflect the degree of variety in scenarios presented, such that no systematic patterns or heuristics quickly establish themselves in the minds of participants.

### **Knowledge retention quiz**

In Table 3.12, we test whether scores obtained by respondents in a knowledge retention quiz designed to assess durability of the learning interventions varied by treatment

---

<sup>30</sup>To take a case in point, we find in Table 3.16 (and corresponding Figures 3.4 and 3.5) a significant divergence in accuracy scores across those with and without prior fraud experience who were exposed to Treatment 1. However, in neither case does the outcome diverge significantly from the respective control condition, indicating the absence of a meaningful heterogeneity in treatment effects properly understood. Similar piecemeal patterns are observed widely in our data. As such, we do not claim to find strong evidence of heterogeneous treatment effects, notwithstanding the presence of some noteworthy level effects described previously.

group. In a poisson model, we estimate the impact from treatment on the count of correct answers reported in the quiz. We find modest evidence of heightened performance among treated participants. Respondents who received any treatment perform only between 1.15 and 1.17 times better than control subjects in the count of correct answers.<sup>31</sup>

### 3.4.2 Results of UCC intervention

In Table 3.11, we report results from a final experimental manipulation, the UCC follow-up exercise. This tests the impact on the likelihood of user engagement of a unique pre-specified authentication code embedded in SMS communication. We find that the mere presence of a UCC code in communication does not significantly increase the likelihood of engagement. In addition, neither does the effectiveness of the code as a signal of authenticity enhanced when the recipient has personalised their own code, instead of having it automatically generated and assigned. In both cases we fail to reject the null hypothesis of no incremental impact on the likelihood of engagement.

---

<sup>31</sup>As part of the follow-up contact, experimental participants are asked about their hypothetical future preference regarding the format of authentication codes, considering the options of a numerical code, a word, or a sentence/phrase. 82% prefer a numerical code, 12% prefer a word, and 6% prefer a sentence/phrase. However, it is likely that the anchoring effect of recent experience with a numerical code over the course of their participation in the study influences these choices. As such, they should not be seen as clean or organic measures of preference.



## 3.5 Conclusion

In this paper, we demonstrate the difficulty in meaningfully improving detection ability between genuine and fraudulent communications. We test three learning interventions, which vary in intensity from simple warning messages about the risk of fraud and the importance of vigilance, to a much deeper illustration of key signs of fraud with applied illustrative scenarios lasting 25 minutes. None of these treatments are successful in significantly improving performance in an experimental task where participants are asked to judge the authenticity or otherwise of 20 fictionalised communications scenarios. We fail to reject the null hypothesis of no incremental benefit from each treatment arm, in respect of overall accuracy, true positives, or true negatives. In this, we join the null effect column in the contested accounting over the remedial value of educational interventions in the financial and digital fraud domains. In view of the direct and specific nature of the learning interventions embodied in Treatments 2 and 3, the contextually-adjusted and engaging nature of the interventions, and the immediacy of the experimental task that followed, this failure is nonetheless surprising and disappointing. We do not find evidence of the returns to intensity in learning interventions observed elsewhere, and nor does the fact that interventions are delivered just in time, and at a relevant teachable moment yield the beneficial impact that the existing literature would lead us to hope for.

Our results cast some doubt over the promise of light-touch targeted educational interventions in moving the dial on stubborn and persistent vulnerabilities faced by users in navigating the digital financial landscape today. More immediately, our results highlight the severity of the challenge which MSEs in Nigeria are likely to face in avoiding the pitfalls presented by pervasive fraud.

One important caveat to register against these results is that the coefficients estimated for our treatment arms with respect to accuracy-related outcomes are all below our ex-post estimated minimum detectable effects, which range from 5% in overall accuracy, to 11% in true positives. As such, we cannot exclude the possibility that true effects from our learning intervention may have fallen in this range. However, the real-world economic meaning of treatment effects in this range, when set against the time cost of administering the learning intervention, and the fact that performance is tested immediately in a lab environment, with no competing demands on participants' attention, can be said to be limited.<sup>32</sup>

---

<sup>32</sup>However, the lack of ecological validity in our lab setting, where participants do not face real financial or other exposure from erroneous judgement of communications, may limit degree to which ours can be regarded as a definitive evaluation of the potential impact of these learning interventions.

Neither do we find evidence that a technical solution to the problem of authentication shows promise. In a separate experimental manipulation, the presence of a unique communications code embedded in follow-up communication with participants does not increase the likelihood of engagement, and neither does personalisation of the UCC strengthen its effectiveness as a signal of authenticity to increase engagement. While the failure to elicit positive treatment effects from the UCC is notable, it should be regarded as a less conclusive verdict than the failure of our learning interventions, where the experimental task represented a direct test of discriminant ability, and in a lab setting where achieving meaningful treatment effects should be comparatively easy. To illustrate the contrasting environments, in a follow-up phone survey, the most common reason cited among participants for non-response (both among those with and without the UCC embedded) was that participants were too busy, or simply forgot. Given the likely low real-world salience of the UCC follow-up task in the minds of participants, our UCC result should not be viewed as a final verdict on the technology.

While we do not find evidence of positive impacts from our learning interventions on the primary targeted objective in fraud detection performance, we do find that they produce effects on the degree of confidence which participants report in their judgements over presented scenarios. This result is troubling when set against the absence of any corresponding improvement in actual performance, and speaks to the risk associated with ineffectual learning interventions, where recipients may feel stronger in the given field, having undergone some targeted training, without actually being any stronger. That is to say, it speaks to the danger of false confidence effects. We do, however, find positive treatment effects on other associated outcomes: trust in DFS, the likelihood of future use of banks and mobile banking, and improved knowledge regarding the signs of fraud.

We find some tentative evidence of divergent patterns in performance across relevant subgroups in several areas. However, this evidence is subtle and secondary, and as such is not trumpeted loudly. We do not find evidence of significant heterogeneity in treatment effects across subgroups defined by relevant experiential, behavioural, and socio-demographic characteristics, but we do find some evidence of significant level effects in performance across treatment conditions. Intuitively, we observe that higher DFS experience, and lower trust participants are more accurate in their judgements over fictionalised scenarios than their less experienced and higher trust indexed counterparts. Another suggestive result relates to apparent catch-up growth in confidence levels reported among subgroups initially showing a deficit in confidence when compared against their indexed counterparts, as treatments increase in intensity. This subtle trend is evident for lower risk-appetite, female, lower self-control, and lower

DFS-experienced participants, pointing to a potentially higher marginal return in engendered confidence from learning interventions.

Notwithstanding the failure to yield encouraging results in the improvement of fraud detection, there remains great urgency in devising effective interventions to combat digital fraud due to the large and growing costs associated with digital deception. Future research should not be discouraged from this challenge, but integrate the lessons from past efforts, taking particular cautionary note of associated risks.<sup>33</sup>

---

<sup>33</sup>Instructive lessons will be taken forward from the results of this lab trial to inform the design of a future field trial aimed at combating susceptibility to digital fraud in the real-world, using a sample of new and existing users of the Amana Market platform in Nigeria.

Table 3.2: Variables of interest for heterogeneous treatment effects

Variable of interest	Detail
ICT experience	Standardized index of experiences with information communication technologies. After indexing, individuals are split into high and low experience types.
DFS experience	Standardized index of experiences with digital financial services. After indexing, individuals are split into high and low experience types.
Fraud experience	Respondents are split into two types: those who have directly encountered fraud previously, and those who have not.
Gender	An indicator variable equal to one if the business owner is a woman, zero otherwise.
Occupation	A set of indicator variables (and a left-out group) for the following occupations: Agriculture, Non-Agriculture
Self-Control	A standardized index of self-control, impulsiveness, attentiveness. After indexing, individuals are split into those who have above or below average self-control.
Risk Preference	A standardized index of risk preferences built from two question: a simple elicitation of risk preferences and a self-reported assessment of risk preferences. After indexing this may be split into high and low risk types.
Generalized Trust and Skepticism	A standardized index of variables associated with generalized trust and scepticism, including questioning mind. After indexing, participants are split into a high and low trust types.

Notes: See Table 3.28 for complete definition of all variables used.

Table 3.3: Research hypotheses for core research questions

Research Question	No.	Hypothesis
Do anti-fraud interventions increase the ability to distinguish between fraud and legitimate communications?	1.0	Providing MSEs with the anti-fraud campaign improves their ability to distinguish between genuine and same fraudulent communications (T1, T2, and T3 vs. C).
	1.1	Providing MSEs with a general warning message about fraud alone (with no further educational intervention) improves their ability to distinguish between genuine and fraudulent communications (T1 vs. C).
	1.2	Providing MSEs with warning signs for potential fraud in a simple format improves their ability to distinguish between genuine and fraudulent communications, still further than can be achieved by a general warning message alone (T2 vs. T1).
	1.3	Illustrating applied examples of fraudulent communications in a simple format improves MSEs' ability to distinguish between genuine and fraudulent communications, still further than can be achieved with simple warning signs alone (T3 vs. T2).
Do anti-fraud interventions increase confidence in the ability to distinguish between fraud and legitimate communications?	2.0	Providing MSEs with the anti-fraud campaign improves their confidence in their ability to distinguish between fraudulent and legitimate communications.
Do anti-fraud interventions increase trust in digital financial services?	3.0	Providing MSEs with the anti-fraud campaign improves their trust in DFS.
Does a simple anti-fraud intervention increase usage of digital financial services?	4.0	Providing MSEs with the anti-fraud campaign improves their likelihood of using DFS in the future.
Is the UCC suitably deployed?	5.1	How does the presence (absence) of a pre-specified authentication code affect the degree of confidence recipients place in customer outreach?
	5.2	Is the effectiveness of a pre-specified authentication code as a signal of authenticity enhanced when the recipient has specified their own code, as against when it is automatically generated and assigned?
Is knowledge from educational interventions effectively retained over a short time horizon?	6.0	How does performance in a knowledge retention quiz relating to key signs of fraud administered at +3 weeks vary in accordance with the intensity of the original educational intervention administered?
Do participants learn by doing?	7.0	Does accuracy improve and confidence grow in respect of scenarios presented later in the sequence when compared against those at the start?

Table 3.4: Descriptive statistics across treatment cells

VARIABLES	(1) Control	(2) Treatment 1	(3) Treatment 2	(4) Treatment 3
Age (Years)	26.78 (7.11)	26.74 (6.73)	26.40 (7.19)	28.09 (7.35)
Female (%)	0.65 (0.47)	0.41 (0.49)	0.43 (0.50)	0.45 (0.50)
Third level education (%)	0.57 (0.50)	0.57 (0.50)	0.53 (0.50)	0.51 (0.50)
Married (%)	0.50 (0.50)	0.51 (0.50)	0.47 (0.50)	0.53 (0.50)
Agricultural employment (%)	0.31 (0.46)	0.38 (0.49)	0.41 (0.49)	0.37 (0.49)
Contacted by scammer (%)	0.64 (0.48)	0.66 (0.48)	0.73 (0.44)	0.74 (0.44)
Access to smartphone (%)	0.94 (0.24)	0.90 (0.30)	0.97 (0.17)	0.93 (0.25)
Business owner (%)	0.90 (0.30)	0.89 (0.32)	0.89 (0.32)	0.87 (0.34)
Has formal financial account (%)	0.85 (0.36)	0.83 (0.38)	0.88 (0.32)	0.83 (0.38)
Used online platforms (%)	0.34 (0.48)	0.33 (0.47)	0.39 (0.49)	0.40 (0.49)
Trusting (%)	0.48 (0.50)	0.58 (0.49)	0.53 (0.50)	0.53 (0.50)
Risk averse (%)	0.35 (0.48)	0.25 (0.43)	0.32 (0.47)	0.27 (0.47)
Observations	195	195	195	195

Notes: Table reports means and standard deviations in parentheses of mortgage borrower characteristics in each treatment and control group.

Table 3.5: Covariate balance by treatment

	(1)	(2)	(3)
	Treatment 1	Treatment 2	Treatment 3
Age	-0.007 (0.005)	-0.008* (0.005)	0.005 (0.005)
Female	-0.239*** (0.054)	-0.221*** (0.055)	-0.174*** (0.055)
Third level education	0.038 (0.054)	0.010 (0.056)	0.016 (0.055)
Married	0.065 (0.061)	0.018 (0.062)	0.006 (0.063)
Agricultural employment	0.026 (0.055)	0.062 (0.055)	0.011 (0.056)
Contacted by scammer	-0.001 (0.056)	0.063 (0.057)	0.085 (0.059)
Access to smartphone	-0.110 (0.093)	0.155 (0.120)	-0.022 (0.104)
Business owner	-0.011 (0.081)	0.019 (0.083)	-0.101 (0.080)
Has formal financial account	-0.052 (0.070)	0.068 (0.076)	-0.124* (0.074)
Used online platforms	-0.001 (0.055)	0.031 (0.054)	0.061 (0.055)
Trusting	0.095* (0.051)	0.060 (0.052)	0.079 (0.052)
Risk averse	-0.089 (0.055)	0.004 (0.055)	-0.052 (0.056)
Constant	0.873*** (0.177)	0.492** (0.195)	0.544*** (0.174)
Observations	390	390	390
R-squared	0.083	0.074	0.066

Notes: Table reports linear prediction of treatment status (for each treatment arm) compared against the control group, across a range of important descriptive characteristics. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.6: Overall effect

	(1)	(2)	(3)
	Overall	True positive	True negative
Treatment 1	-0.010 (0.013)	-0.021 (0.026)	0.001 (0.025)
Treatment 2	-0.001 (0.012)	-0.026 (0.026)	0.023 (0.023)
Treatment 3	-0.009 (0.013)	-0.049* (0.026)	0.032 (0.023)
Constant	0.654*** (0.016)	0.591*** (0.033)	0.718*** (0.030)
Observations	780	780	780
R-squared	0.085	0.019	0.062
p-value (T1≤C)	0.791	0.794	0.483
p-value (T2≤T1)	0.237	0.563	0.162
p-value (T3≤T2)	0.725	0.810	0.340

Notes: Table reports results from two-sided test for treatment effects on overall accuracy, true positives, and true negatives. Also reported are one-sided tests of pre-specified hypotheses for incremental positive treatment effects from each treatment arm compared against the preceding arm in the sequence. Regression includes vector of controls listed in Table 3.2. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$



Table 3.7: Confidence effect

	(1)	(2)	(3)
	Overall	True positive	True negative
Treatment 1	0.006 (0.115)	-0.010 (0.116)	0.017 (0.120)
Treatment 2	0.207* (0.112)	0.200* (0.115)	0.214* (0.114)
Treatment 3	0.256** (0.115)	0.207* (0.118)	0.283** (0.118)
Constant	5.901*** (0.143)	5.858*** (0.147)	5.935*** (0.148)
Observations	780	780	780
R-squared	0.164	0.152	0.159
p-value (T1≤C)	0.480	0.533	0.443
p-value (T2≤T1)	0.0265	0.0246	0.0325
p-value (T3≤T2)	0.316	0.476	0.248

Notes: Table reports results from two-sided test for treatment effects on overall accuracy, true positives, and true negatives. Also reported are one-sided tests of pre-specified hypotheses for incremental positive treatment effects from each treatment arm compared against the preceding arm in the sequence. Regression includes vector of controls listed in Table 3.2. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.8: Effect on trust (ICW trust index)

	(1)	(2)
Treatment 1	0.135 (0.097)	0.055 (0.076)
Treatment 2	0.211** (0.095)	0.171** (0.081)
Treatment 3	0.166 (0.102)	0.084 (0.083)
Baseline trust index		0.549*** (0.038)
Constant	0.062 (0.118)	0.019 (0.096)
Observations	780	780
R-squared	0.182	0.432
p-value (T1≤C)	0.0819	0.233
p-value (T2≤T1)	0.192	0.0615
p-value (T3≤T2)	0.686	0.858

Notes: Table reports results from two-sided test for treatment effects on a standardised inverse-correlation weighted index of trust in DFS. Also reported are one-sided tests of pre-specified hypotheses for incremental positive treatment effects from each treatment arm compared against the preceding arm in the sequence. Regression includes vector of controls listed in Table 3.2. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.9: Effect on likelihood of future use

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Banks	Banks	Mobile banking	Mobile banking	Mobile money	Mobile money	Online platforms	Online platforms	Agents	Agents
Treatment 1	-0.078 (0.187)	-0.165 (0.171)	-0.182 (0.175)	-0.258 (0.162)	-0.125 (0.186)	-0.226 (0.177)	0.053 (0.175)	-0.001 (0.161)	0.272 (0.178)	0.179 (0.174)
Treatment 2	0.387** (0.181)	0.417** (0.167)	0.248 (0.166)	0.301* (0.155)	0.071 (0.181)	-0.017 (0.175)	0.109 (0.174)	0.096 (0.165)	0.126 (0.178)	0.017 (0.173)
Treatment 3	0.138 (0.184)	0.077 (0.165)	-0.011 (0.177)	-0.056 (0.157)	0.095 (0.188)	-0.022 (0.179)	0.090 (0.176)	0.051 (0.165)	0.177 (0.182)	0.072 (0.180)
Baseline likely use		0.309*** (0.035)		0.287*** (0.033)		0.254*** (0.035)		0.296*** (0.036)		0.173*** (0.034)
Constant	5.641*** (0.225)	4.081*** (0.286)	5.449*** (0.207)	4.060*** (0.258)	4.880*** (0.219)	3.751*** (0.265)	5.428*** (0.214)	3.929*** (0.292)	4.667*** (0.226)	4.006*** (0.261)
Observations	780	780	780	780	780	780	780	780	780	780
R-squared	0.136	0.246	0.132	0.238	0.117	0.188	0.107	0.209	0.091	0.128
p-value (T1≤C)	0.662	0.833	0.851	0.944	0.749	0.900	0.381	0.502	0.0635	0.151
p-value (T2≤T1)	0.003	0.000	0.003	0.000	0.138	0.115	0.361	0.255	0.808	0.842
p-value (T3≤T2)	0.934	0.985	0.944	0.990	0.446	0.510	0.547	0.614	0.385	0.371

Notes: Table reports results from two-sided test for treatment effects on the likelihood of future use of a series of a range of entities. Also reported are one-sided tests of pre-specified hypotheses for incremental positive treatment effects from each treatment arm compared against the preceding arm in the sequence. Regression includes vector of controls listed in Table 3.2. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.10: Learning by doing effects

	(1)	(2)
	Accuracy	Confidence
Order	-0.003*** (0.001)	-0.009*** (0.002)
Constant	0.643*** (0.008)	5.714*** (0.045)
Observations	15,600	15,600
Number of participants	780	780
p-value ( $\beta \leq 0$ )	1	1

Notes: Table reports results from two-sided test for ‘learning by doing’ effects on overall accuracy, and confidence, by regressing these outcomes on the order variable capturing the sequence of scenarios presented. Also reported are one-sided tests of hypotheses for positive learning effects. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.11: Impact of authentication on engagement probability

	(1)	(2)
UCC present	0.034 (0.032)	0.057 (0.041)
UCC personalised		0.103** (0.043)
UCC present # UCC personalised		-0.049 (0.063)
Constant	0.246*** (0.022)	0.198*** (0.027)
Observations	780	780
R-squared	0.001	0.010
p-value ( $\beta \leq 0$ )	0.142	0.779

Notes: Table reports from an OLS regression model predicting whether the participant responded to the outreach with the requested personal information. Column 1 includes a simple treatment condition recording when the UCC was embedded in the SMS. Column 2 adds an interaction recording when that UCC was personalised, as distinct from having been automatically generated. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.12: Impact of treatment on subsequent quiz score

	(1)
Treatment 1	0.159*** (0.061)
Treatment 2	0.136** (0.061)
Treatment 3	0.137** (0.061)
Constant	0.514*** (0.049)
Observations	519

Notes: Table reports results from a poisson regression, estimating the number out of 3 quiz questions relating to key signs of fraud correctly answered in a follow-up call. Under a poisson specification, the exponentiated coefficient gives the multiplicative term with which to calculate the expected quiz score when the given treatment has been administered, relative to the control condition. As such, Treatment 1 participants are estimated to have  $e^{0.159} = 1.17$  as a Rate Ratio, the multiplicative increase in the expected quiz score compared to Control participants. For Treatment 2:  $e^{0.136} = 1.15$ . For Treatment 3:  $e^{0.137} = 1.15$ . \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

## Appendix

### 3.6 Empirical methodology - extended

#### 3.6.1 Standard error adjustments

Treatment assignment is at the individual level, therefore for outcomes with multiple observations per participant, we will apply cluster robust standard errors at the individual level. For any outcomes with only one observation per treatment unit, we will apply heteroskedasticity robust standard errors.

#### 3.6.2 Multiple hypothesis testing

As described in the sections above, we opt to reduce the number of tests in each outcome group as opposed to adjusting for multiple hypothesis testing. Specifically, we test a single outcome for each primary outcome group. Where multiple outcomes are of interest, we will construct a standardized index of the outcomes to serve as the primary outcome for that group as in Anderson (2008).

#### 3.6.3 Statistical power

To assess the sample size requirement for the lab experiment, we estimate the minimum detectable effects (MDE) under a range of alternative design scenarios (varying by sample size, power, and number of treatment arms) in Table 3.13.

The table provides an estimate of the smallest treatment effect that could be detected with statistical confidence were it to be achieved by one of our educational interventions. For a given scenario, treatment effects smaller than that reported would not be detectable with statistical confidence. The results in the table are calculated relative to a baseline unconditional probability of unaided detection ability of 50% (i.e. chance).<sup>34</sup>

For example, starting with the top left scenario, if only one treatment were administered and the number of participants was 250, it would not be possible to detect treatment effects smaller than a 20.58% improvement over the baseline detection accuracy (with 90% power). The power level refers to an acceptable level of probability that the experiment will detect an effect when the effect is present. In this example, if we were to repeat the experiment over and over, we would detect an impact at least as big as this 90% of the time.

Ex-ante, in our chosen design, our estimated MDE with  $N = 780$  and 3 treatment arms was 24.90% (with 80% power), or 29.3% (with 90% power). Following McKenzie and Ozier (2019), we calculate ex-post MDEs using actual realised data from the control group. Ex-post, using the observed accuracy and standard deviation in the control group ( $\mu = 0.61$ ,  $\sigma = 0.13$ ), we estimate our MDE for treatment effects in accuracy to have been much lower, at 5.24% (with 80% power), or 6.16% (with 90% power). In true positives, with observed outcomes in the control group ( $\mu = 0.55$ ,  $\sigma = 0.25$ ), our

<sup>34</sup>Unaided human deception detection capacity has been estimated to be little better than chance (Hartwig and Bond Jr, 2011).

MDE is 11.28% (with 80% power), or 13.27% (with 90% power). In true negatives, with observed outcomes in the control group ( $\mu = 0.67$ ,  $\sigma = 0.24$ ), our MDE is 9.15% (with 80% power), or 10.76% (with 90% power). Tr.: number of treatment arms.

Table 3.13: Estimated statistical power under alternative design scenarios

Outcome: detection accuracy									
	N	Power 90%				Power 80%			
Tr.		1	2	3	4	1	2	3	4
	250	31.54%	38.76%	44.90%	50.08%	37.12%	45.62%	52.86%	58.94%
	500	22.26%	27.34%	31.54%	35.28%	26.22%	32.18%	37.12%	41.52%
	750	18.18%	22.26%	25.76%	28.78%	21.40%	26.22%	30.32%	33.86%
	780	17.60%	21.58%	24.90%	27.86%	20.72%	25.40%	29.30%	32.78%
	1000	15.74%	19.28%	22.26%	24.90%	18.52%	22.68%	26.22%	29.32%

Notes: Table reports ex-ante power calculations giving minimum detectable effect sizes under a range of sample size and treatment arm scenarios. We assumed baseline accuracy of chance ( $\mu = 0.50$ ,  $\sigma = 0.50$ ). Each scenario estimates the minimum detectable effect, expressed as a percentage increase over baseline scores.



## 3.7 Departures from pre-analysis plan

The preregistration for this lab experiment was filed with the American Economic Association’s registry for randomised controlled trials in July 2022 (RCT ID: AEARCTR-0009470), before data collection began (Byrne et al., 2022). We depart from the research plan pre-specified in the pre-analysis plan (PAP) in a number of areas.

### 3.7.1 Pre-specified but not included

- PAP hypotheses 6.1 and 6.2 undertook to investigate how urgency (i.e. time pressure) in scenarios affected confidence and accuracy. Due to challenges in data reporting relating to the time limit imposed in different scenarios, it was not possible to perform these tests.
- PAP hypotheses 7.1 and 7.2 undertook to assess knowledge retention from our learning interventions, by testing for decay in performance between two time points: a quiz administered at the close of the endline survey, and a follow-up quiz administered at +3 weeks. We additionally undertook how this rate of decay varied in accordance with the intensity of the original treatment administered. Due to oversight in implementation, the first quiz was not administered as part of the endline survey, making it impossible to perform the pre-specified tests of decay. As an alternative, we test instead how performance in the follow-up quiz alone varies by treatment.
- Our PAP envisaged that our participant sample would be comprised partly of a supplemental pool of students from the partnering university in Nigeria, creating an occupational subcategory of ‘student’. However, the opportunity to supplement our sample instead with additional users from the Amana Market platform arose. This was deemed preferable, corresponding more closely to the target population of interest.

### 3.7.2 Included but not pre-specified

- We pre-specified that we would examine learning by doing effects within the experiment, testing whether if those scenarios that appeared later in the order were more often correctly identified by participants. We did not additionally specify that we would examine whether confidence grows through the sequence of presented scenarios, but this outcome is also evaluated in the paper. Confidence is, however, pre-specified as one of our primary outcomes related to the ability to distinguish between genuine and fraudulent communications. We viewed this as sufficiently interesting and important to merit inclusion.

### 3.8 Additional tables and figures

Table 3.14: Overall effect: partialling out lasso approach

	(1)	(2)	(3)
	Overall	True positive	True negative
Treatment 1	-0.010 (0.013)	-0.021 (0.026)	0.001 (0.025)
Treatment 2	-0.001 (0.012)	-0.025 (0.026)	0.023 (0.023)
Treatment 3	-0.009 (0.012)	-0.049* (0.026)	0.032 (0.023)
Observations	780	780	780

Notes: Table reports an the same effects reported in Table 3.6, but using an alternative approach to the selection of relevant control variables for the purpose of robustness: partialling out lasso linear regression. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.15: Confidence effect: partialling out lasso approach

	(1)	(2)	(3)
	Overall	True positive	True negative
Treatment 1	0.005 (0.114)	-0.010 (0.115)	0.017 (0.119)
Treatment 2	0.207* (0.111)	0.200* (0.115)	0.214* (0.113)
Treatment 3	0.255** (0.114)	0.206* (0.117)	0.283** (0.117)
Observations	780	780	780

Notes: Table reports an the same effects reported in Table 3.7, but using an alternative approach to the selection of relevant control variables for the purpose of robustness: partialling out lasso linear regression. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.16: Heterogeneous effects in accuracy: experience and behaviour

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment 1	0.005 (0.017)	-0.022 (0.017)	-0.002 (0.016)	0.008 (0.018)	0.000 (0.019)	-0.001 (0.018)
Treatment 2	-0.003 (0.016)	-0.012 (0.019)	0.002 (0.016)	0.005 (0.018)	0.005 (0.019)	-0.006 (0.017)
Treatment 3	-0.010 (0.016)	-0.022 (0.017)	-0.004 (0.016)	-0.005 (0.018)	-0.029 (0.018)	-0.005 (0.017)
No direct experience	-0.014 (0.018)					
Treatment 1# No direct experience	-0.042* (0.025)					
Treatment 2# No direct experience	0.012 (0.025)					
Treatment 3# No direct experience	0.008 (0.025)					
Low self-control		-0.029 (0.018)				
Treatment 1# Low self-control		0.025 (0.025)				
Treatment 2# Low self-control		0.021 (0.024)				
Treatment 3# Low self-control		0.027 (0.025)				
Low risk aversion			0.012 (0.017)			
Treatment 1# Low risk aversion			-0.030 (0.026)			
Treatment 2# Low risk aversion			-0.009 (0.024)			
Treatment 3# Low risk aversion			-0.013 (0.026)			
Low trust				0.037** (0.018)		
Treatment 1# Low trust				-0.033 (0.025)		
Treatment 2# Low trust				-0.013 (0.024)		
Treatment 3# Low trust				-0.008 (0.025)		
Low DFS experience					-0.038** (0.018)	
Treatment 1# Low DFS experience					-0.020 (0.025)	
Treatment 2# Low DFS experience					-0.012 (0.024)	
Treatment 3# Low DFS experience					0.045* (0.025)	
Low ICT experience						-0.007 (0.017)
Treatment 1# Low ICT experience						-0.016 (0.025)
Treatment 2# Low ICT experience						0.013 (0.025)
Treatment 3# Low ICT experience						-0.008 (0.025)
Constant	0.652*** (0.018)	0.664*** (0.018)	0.650*** (0.018)	0.648*** (0.018)	0.655*** (0.019)	0.652*** (0.018)
Observations	780	780	780	780	780	780
R-squared	0.091	0.087	0.086	0.087	0.095	0.087

Notes: Table explores heterogeneous treatment effects in accuracy across experiential and behavioural characteristics, using interaction terms. Regression includes vector of controls listed in Table 3.2. Standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.17: Heterogeneous effects in accuracy: demographic

	(1)	(2)	(3)	(4)
Treatment 1	-0.006 (0.015)	-0.011 (0.016)	0.001 (0.020)	-0.018 (0.020)
Treatment 2	-0.002 (0.014)	-0.008 (0.015)	0.009 (0.019)	-0.011 (0.019)
Treatment 3	-0.002 (0.015)	0.001 (0.017)	-0.009 (0.019)	-0.010 (0.020)
Agriculture	0.001 (0.020)			
Treatment 1#Agriculture	-0.012 (0.027)			
Treatment 2#Agriculture	0.001 (0.027)			
Treatment 3# Agriculture	-0.019 (0.027)			
Above median age		0.006 (0.018)		
Treatment 1# Above median age		0.002 (0.025)		
Treatment 2# Above median age		0.017 (0.025)		
Treatment 3# Above median age		-0.016 (0.025)		
Lower education			-0.020 (0.018)	
Treatment 1# Lower education			-0.017 (0.025)	
Treatment 2# Lower education			-0.019 (0.025)	
Treatment 3# Lower education			0.002 (0.025)	
Female				-0.040** (0.019)
Treatment 1#Female				0.014 (0.026)
Treatment 2# Female				0.019 (0.026)
Treatment 3# Female				-0.000 (0.025)
Constant	0.651*** (0.017)	0.654*** (0.017)	0.666*** (0.020)	0.659*** (0.020)
Observations	780	780	780	780
R-squared	0.086	0.087	0.097	0.086

Notes: Table explores heterogeneous treatment effects in accuracy across demographic characteristics, using interaction terms. Regression includes vector of controls listed in Table 3.2. Standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.18: Heterogeneous effects in true positives: experience and behaviour

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment 1	-0.002 (0.032)	-0.032 (0.032)	-0.014 (0.031)	-0.015 (0.040)	0.007 (0.037)	-0.016 (0.039)
Treatment 2	-0.058* (0.031)	-0.027 (0.037)	-0.044 (0.033)	-0.018 (0.038)	-0.036 (0.038)	-0.048 (0.036)
Treatment 3	-0.087*** (0.032)	-0.046 (0.033)	-0.053* (0.032)	-0.047 (0.039)	-0.081** (0.037)	-0.044 (0.038)
No direct experience	-0.068* (0.036)					
Treatment 1# No direct experience	-0.058 (0.055)					
Treatment 2# No direct experience	0.110** (0.055)					
Treatment 3# No direct experience	0.129** (0.054)					
Low self-control		-0.051 (0.035)				
Treatment 1# Low self-control		0.026 (0.053)				
Treatment 2# Low self-control		0.003 (0.050)				
Treatment 3# Low self-control		-0.006 (0.051)				
Low risk aversion			-0.016 (0.037)			
Treatment 1# Low risk aversion			-0.033 (0.057)			
Treatment 2# Low risk aversion			0.055 (0.051)			
Treatment 3# Low risk aversion			0.012 (0.056)			
Low trust				0.017 (0.035)		
Treatment 1# Low trust				-0.012 (0.052)		
Treatment 2# Low trust				-0.015 (0.050)		
Treatment 3# Low trust				-0.005 (0.052)		
Low DFS experience					-0.030 (0.035)	
Treatment 1# Low DFS experience					-0.053 (0.052)	
Treatment 2# Low DFS experience					0.022 (0.050)	
Treatment 3# Low DFS experience					0.071 (0.052)	
Low ICT experience						-0.003 (0.036)
Treatment 1# Low ICT experience						-0.008 (0.051)
Treatment 2# Low ICT experience						0.054 (0.051)
Treatment 3# Low ICT experience						-0.013 (0.052)
Constant	0.605*** (0.035)	0.594*** (0.034)	0.595*** (0.035)	0.587*** (0.037)	0.595*** (0.037)	0.594*** (0.038)
Observations	780	780	780	780	780	780
R-squared	0.039	0.020	0.023	0.020	0.027	0.022

Notes: Table explores heterogeneous treatment effects in true positives across experiential and behavioural characteristics, using interaction terms. Regression includes vector of controls listed in Table 3.2. Standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.19: Heterogeneous effects in true positives: demographic

	(1)	(2)	(3)	(4)
Treatment 1	0.007 (0.032)	-0.036 (0.034)	-0.014 (0.037)	-0.050 (0.038)
Treatment 2	-0.001 (0.033)	-0.001 (0.033)	-0.014 (0.036)	-0.061 (0.038)
Treatment 3	-0.004 (0.032)	-0.013 (0.039)	-0.077** (0.036)	-0.081** (0.039)
Agriculture	0.073** (0.037)			
Treatment 1#Agriculture	-0.091* (0.054)			
Treatment 2#Agriculture	-0.080 (0.052)			
Treatment 3# Agriculture	-0.135** (0.053)			
Above median age		0.029 (0.035)		
Treatment 1# Above median age		0.032 (0.052)		
Treatment 2# Above median age		-0.055 (0.051)		
Treatment 3# Above median age		-0.069 (0.053)		
Lower education			0.011 (0.035)	
Treatment 1# Lower education			-0.015 (0.051)	
Treatment 2# Lower education			-0.022 (0.050)	
Treatment 3# Lower education			0.053 (0.051)	
Female				-0.047 (0.037)
Treatment 1#Female				0.046 (0.053)
Treatment 2# Female				0.062 (0.051)
Treatment 3# Female				0.052 (0.052)
Constant	0.565*** (0.034)	0.583*** (0.035)	0.588*** (0.037)	0.617*** (0.040)
Observations	780	780	780	780
R-squared	0.027	0.026	0.023	0.021

Notes: Table explores heterogeneous treatment effects in true positives across demographic characteristics, using interaction terms. Regression includes vector of controls listed in Table 3.2. Standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.20: Heterogeneous effects in true negatives: experience and behaviour

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment 1	0.011 (0.030)	-0.012 (0.032)	0.011 (0.029)	0.031 (0.039)	-0.006 (0.035)	0.015 (0.035)
Treatment 2	0.052* (0.027)	0.002 (0.033)	0.048* (0.028)	0.027 (0.034)	0.045 (0.032)	0.036 (0.031)
Treatment 3	0.067** (0.028)	0.002 (0.032)	0.045 (0.028)	0.036 (0.036)	0.023 (0.033)	0.035 (0.031)
No direct experience	0.041 (0.037)					
Treatment 1# No direct experience	-0.026 (0.052)					
Treatment 2# No direct experience	-0.087* (0.051)					
Treatment 3# No direct experience	-0.114** (0.050)					
Low self-control		-0.007 (0.035)				
Treatment 1# Low self-control		0.023 (0.049)				
Treatment 2# Low self-control		0.039 (0.045)				
Treatment 3# Low self-control		0.059 (0.046)				
Low risk aversion			0.040 (0.037)			
Treatment 1# Low risk aversion			-0.027 (0.054)			
Treatment 2# Low risk aversion			-0.074 (0.047)			
Treatment 3# Low risk aversion			-0.038 (0.050)			
Low trust				0.058* (0.035)		
Treatment 1# Low trust				-0.054 (0.049)		
Treatment 2# Low trust				-0.011 (0.045)		
Treatment 3# Low trust				-0.011 (0.047)		
Low DFS experience					-0.045 (0.035)	
Treatment 1# Low DFS experience					0.013 (0.049)	
Treatment 2# Low DFS experience					-0.047 (0.045)	
Treatment 3# Low DFS experience					0.019 (0.046)	
Low ICT experience						-0.010 (0.035)
Treatment 1# Low ICT experience						-0.024 (0.048)
Treatment 2# Low ICT experience						-0.027 (0.045)
Treatment 3# Low ICT experience						-0.004 (0.047)
Constant	0.699*** (0.032)	0.734*** (0.033)	0.706*** (0.032)	0.709*** (0.036)	0.716*** (0.034)	0.711*** (0.034)
Observations	780	780	780	780	780	780
R-squared	0.070	0.064	0.065	0.064	0.065	0.062

Notes: Table explores heterogeneous treatment effects in accuracy across experiential and behavioural characteristics, using interaction terms. Regression includes vector of controls listed in Table 3.2. Standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$



Table 3.21: Heterogeneous effects in true negatives: demographic

	(1)	(2)	(3)	(4)
Treatment 1	-0.019 (0.032)	0.014 (0.032)	0.015 (0.032)	0.014 (0.037)
Treatment 2	-0.004 (0.029)	-0.015 (0.030)	0.033 (0.031)	0.039 (0.035)
Treatment 3	-0.000 (0.029)	0.014 (0.031)	0.060** (0.029)	0.060 (0.037)
Agriculture	-0.072** (0.036)			
Treatment 1#Agriculture	0.066 (0.049)			
Treatment 2#Agriculture	0.081* (0.046)			
Treatment 3# Agriculture	0.097** (0.047)			
Above median age		-0.016 (0.035)		
Treatment 1# Above median age		-0.029 (0.049)		
Treatment 2# Above median age		0.089** (0.045)		
Treatment 3# Above median age		0.037 (0.046)		
Lower education			-0.051 (0.035)	
Treatment 1# Lower education			-0.020 (0.047)	
Treatment 2# Lower education			-0.016 (0.044)	
Treatment 3# Lower education			-0.049 (0.046)	
Female				-0.034 (0.038)
Treatment 1#Female				-0.019 (0.051)
Treatment 2# Female				-0.024 (0.048)
Treatment 3# Female				-0.052 (0.048)
Constant	0.738*** (0.032)	0.725*** (0.032)	0.745*** (0.034)	0.701*** (0.037)
Observations	780	780	780	780
R-squared	0.067	0.071	0.083	0.063

Notes: Table explores heterogeneous treatment effects in true negatives across demographic characteristics, using interaction terms. Regression includes vector of controls listed in Table 3.2. Standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.22: Heterogeneous effects in confidence: experience and behaviour

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment 1	-0.019 (0.135)	0.060 (0.133)	-0.061 (0.134)	0.211 (0.189)	-0.099 (0.152)	0.167 (0.161)
Treatment 2	0.186 (0.129)	0.210 (0.133)	0.158 (0.127)	0.232 (0.186)	0.120 (0.154)	0.372** (0.161)
Treatment 3	0.148 (0.137)	0.059 (0.150)	0.162 (0.134)	0.499*** (0.183)	0.029 (0.154)	0.440*** (0.163)
No direct experience	-0.219 (0.179)					
Treatment 1# No direct experience	0.069 (0.254)					
Treatment 2# No direct experience	0.035 (0.257)					
Treatment 3# No direct experience	0.368 (0.256)					
Low self-control		-0.343** (0.167)				
Treatment 1# Low self-control		-0.161 (0.243)				
Treatment 2# Low self-control		0.002 (0.212)				
Treatment 3# Low self-control		0.403* (0.225)				
Low risk aversion			-0.600*** (0.189)			
Treatment 1# Low risk aversion			0.209 (0.256)			
Treatment 2# Low risk aversion			0.131 (0.246)			
Treatment 3# Low risk aversion			0.296 (0.255)			
Low trust				0.638*** (0.168)		
Treatment 1# Low trust				-0.384* (0.231)		
Treatment 2# Low trust				-0.070 (0.221)		
Treatment 3# Low trust				-0.477** (0.227)		
Low DFS experience					-0.416** (0.171)	
Treatment 1# Low DFS experience					0.204 (0.233)	
Treatment 2# Low DFS experience					0.163 (0.219)	
Treatment 3# Low DFS experience					0.478** (0.231)	
Low ICT experience						0.075 (0.171)
Treatment 1# Low ICT experience						-0.292 (0.224)
Treatment 2# Low ICT experience						-0.326 (0.220)
Treatment 3# Low ICT experience						-0.359 (0.232)
Constant	5.939*** (0.155)	5.936*** (0.146)	5.961*** (0.153)	5.788*** (0.175)	6.004*** (0.159)	5.758*** (0.170)
Observations	780	780	780	780	780	780
R-squared	0.167	0.172	0.166	0.171	0.169	0.168

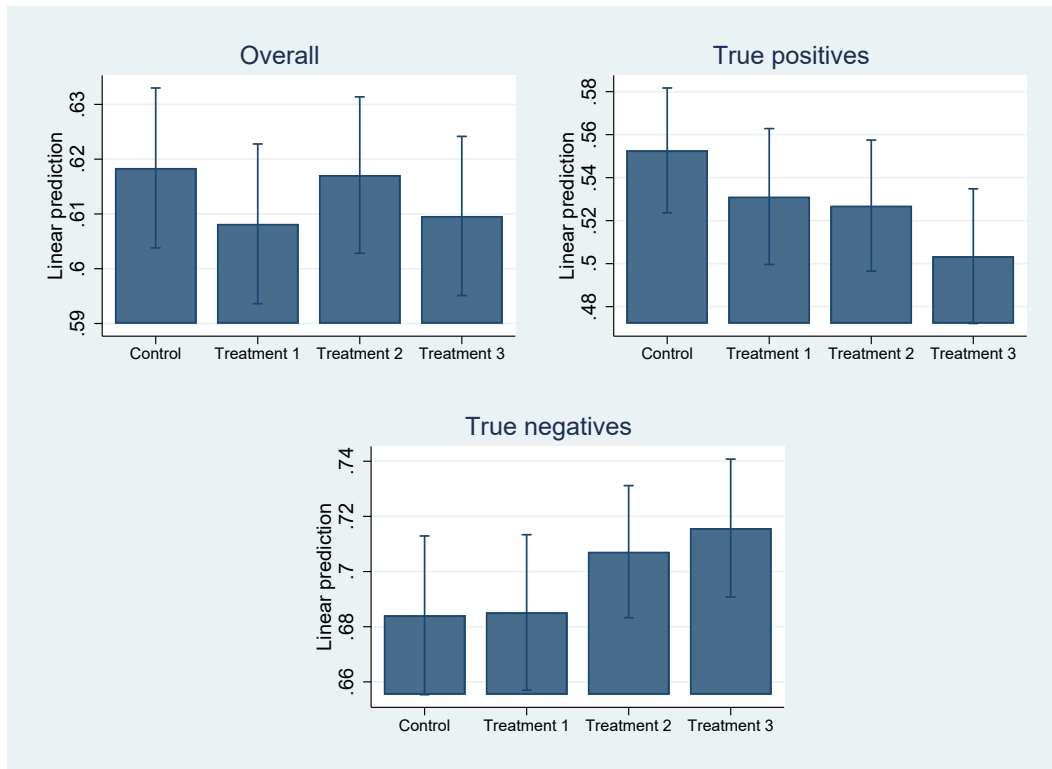
Notes: Table explores heterogeneous treatment effects in confidence across experiential and behavioural characteristics, using interaction terms. Regression includes vector of controls listed in Table 3.2. Standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.23: Heterogeneous effects in confidence: demographic

	(1)	(2)	(3)	(4)
Treatment 1	-0.018 (0.146)	0.028 (0.150)	-0.043 (0.153)	-0.211 (0.167)
Treatment 2	0.210 (0.145)	0.109 (0.153)	0.166 (0.148)	0.090 (0.156)
Treatment 3	0.396*** (0.139)	0.403*** (0.156)	0.176 (0.144)	-0.035 (0.169)
Agriculture	0.090 (0.177)			
Treatment 1#Agriculture	0.043 (0.239)			
Treatment 2#Agriculture	-0.035 (0.221)			
Treatment 3# Agriculture	-0.390 (0.239)			
Above median age		0.011 (0.171)		
Treatment 1# Above median age		-0.041 (0.235)		
Treatment 2# Above median age		0.237 (0.219)		
Treatment 3# Above median age		-0.260 (0.229)		
Lower education			-0.233 (0.168)	
Treatment 1# Lower education			0.092 (0.225)	
Treatment 2# Lower education			0.071 (0.216)	
Treatment 3# Lower education			0.149 (0.225)	
Female				-0.497*** (0.173)
Treatment 1#Female				0.380 (0.232)
Treatment 2# Female				0.137 (0.227)
Treatment 3# Female				0.528** (0.231)
Constant	5.872*** (0.160)	5.882*** (0.150)	6.037*** (0.159)	6.084*** (0.171)
Observations	780	780	780	780
R-squared	0.169	0.169	0.168	0.171


Notes: Table explores heterogeneous treatment effects in confidence across demographic characteristics, using interaction terms. Regression includes vector of controls listed in Table 3.2. Standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Figure 3.1: Direction of treatment effects in overall accuracy




Notes: Figure reports marginal effects corresponding to Table 3.6, depicting the direction of treatment effects in accuracy outcomes. While we do not see significant treatment effects, directional patterns are evident.

Figure 3.2: Accuracy in judgements through the sequence of scenarios




Notes: Figure plots the mean level of accuracy in scenario judgements by the order in which they appear in the experimental task. The placement of any individual scenario in the sequence is randomised. Figure graphically evaluates the hypothesis that performance shows a ‘learning by doing’ effect.

Figure 3.3: Level of confidence reported through the sequence of scenarios




Notes: Figure plots the mean level of confidence reported in scenario judgements by the order in which they appear in the experimental task. The placement of any individual scenario in the sequence is randomised.

Figure 3.4: Predictive margins in accuracy (fraud experience interaction)




Notes: Figure reports predicted outcomes in accuracy, from an interaction of treatment status with fraud experience found in Table 3.16. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation.

Figure 3.5: Difference in predicted values (no direct experience:1, direct experience:0)




Notes: Figure plots the relationship shown in Figure 3.4, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.6: Predictive margins in true positives (fraud experience interaction)




Notes: Figure reports predicted outcomes in true positives, from an interaction of treatment status with fraud experience found in Table 3.18. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation.

Figure 3.7: Difference in predicted values (no direct experience:1, direct experience:0)




Notes: Figure plots the relationship shown in Figure 3.6, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.8: Predictive margins in true negatives (fraud experience interaction)



Notes: Figure reports predicted outcomes in true negatives, from an interaction of treatment status with fraud experience found in Table 3.20. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation.

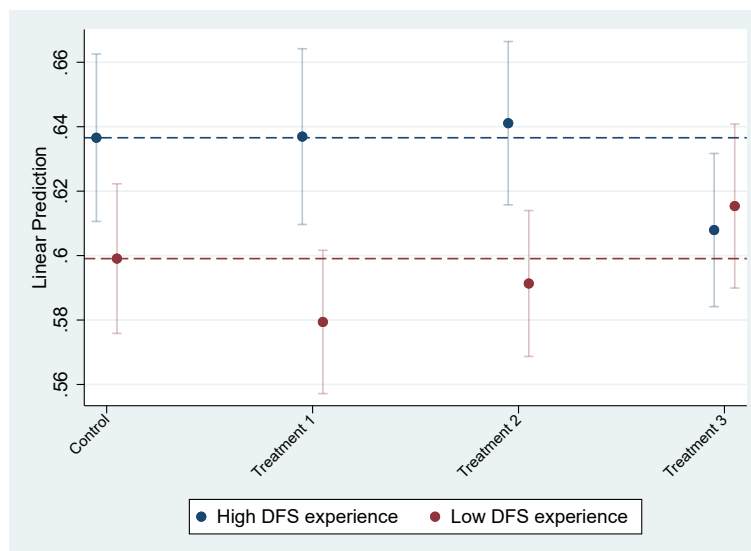
Figure 3.9: Difference in predicted values (no direct experience:1, direct experience:0)



Notes: Figure plots the relationship shown in Figure 3.8, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.




Figure 3.10: Predictive margins in accuracy (DFS experience interaction)




Notes: Figure reports predicted outcomes in accuracy, from an interaction of treatment status with DFS experience found in Table 3.20. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation, albeit with level difference evident.

Figure 3.11: Difference in predicted values (low DFS experience:1, high DFS experience:0)




Notes: Figure plots the relationship shown in Figure 3.10, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.12: Predictive margins in confidence (DFS experience interaction)




Notes: Figure reports predicted outcomes in confidence, from an interaction of treatment status with DFS experience found in Table 3.22. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation.

Figure 3.13: Difference in predicted values (low DFS experience:1, high DFS experience:0)




Notes: Figure plots the relationship shown in Figure 3.12, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.14: Predictive margins in true positives (employment sector interaction)




Notes: Figure reports predicted outcomes in true positives, from an interaction of treatment status with employment sector found in Table 3.19. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation.

Figure 3.15: Difference in predicted values (agriculture:1, non-agriculture:0)



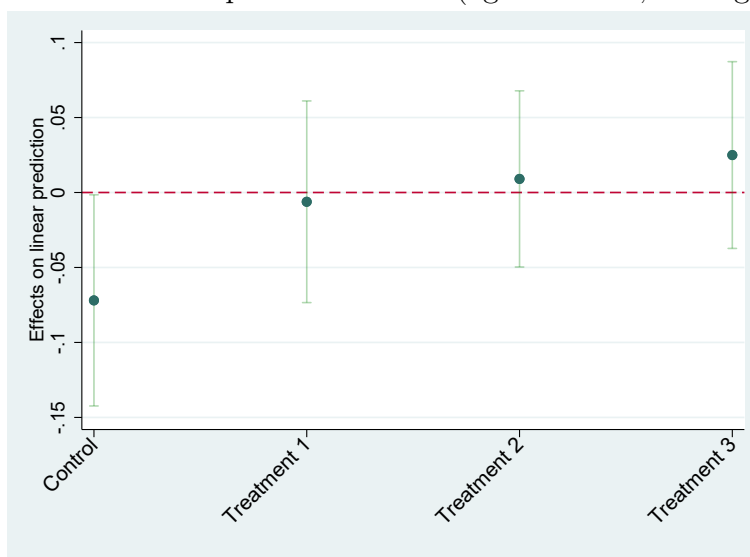
Notes: Figure plots the relationship shown in Figure 3.14, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.16: Predictive margins in true negatives (employment sector interaction)



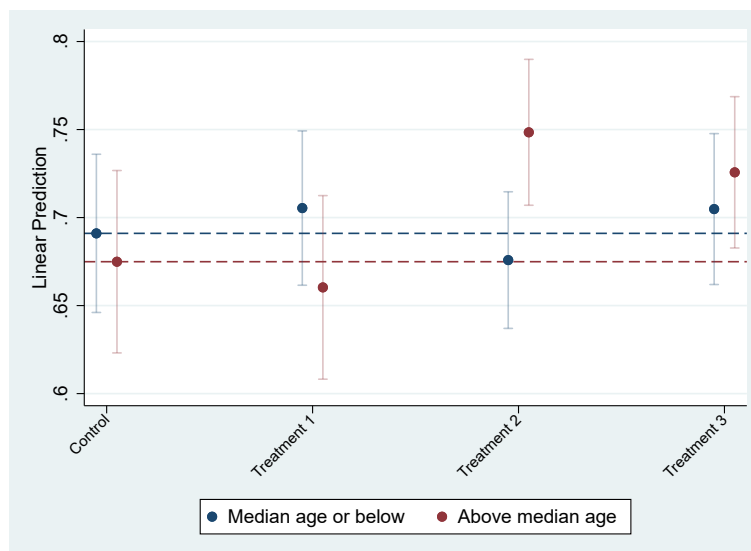
Notes: Figure reports predicted outcomes in true negatives, from an interaction of treatment status with employment sector found in Table 3.21. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation.

Figure 3.17: Difference in predicted values (agriculture:1, non-agriculture:0)




Notes: Figure plots the relationship shown in Figure 3.16, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.18: Predictive margins in true negatives (age interaction)



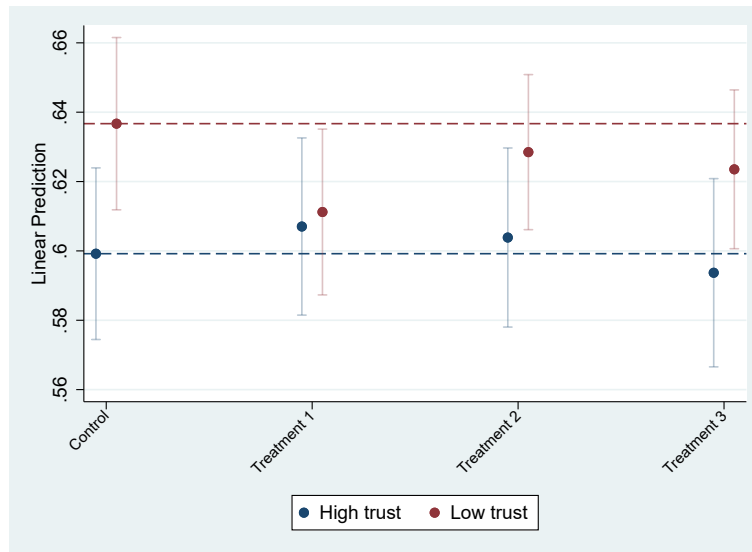
Notes: Figure reports predicted outcomes in true negatives, from an interaction of treatment status with age found in Table 3.21. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation.

Figure 3.19: Difference in predicted values (above median age:1, median age or below:0)




Notes: Figure plots the relationship shown in Figure 3.18, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.20: Predictive margins in accuracy (trust interaction)




Notes: Figure reports predicted outcomes in accuracy, from an interaction of treatment status with trust found in Table 3.16. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation, albeit with tentative level difference evident.

Figure 3.21: Difference in predicted values (low trust:1, high trust:0)




Notes: Figure plots the relationship shown in Figure 3.20, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.22: Predictive margins in true negatives (trust interaction)




Notes: Figure reports predicted outcomes in true negatives, from an interaction of treatment status with trust found in Table 3.20. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation.

Figure 3.23: Difference in predicted values (low trust:1, high trust:0)




Notes: Figure plots the relationship shown in Figure 3.22, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.24: Predictive margins in confidence (trust interaction)



Notes: Figure reports predicted outcomes in confidence, from an interaction of treatment status with trust found in Table 3.22. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation, albeit with level difference evident.


Figure 3.25: Difference in predicted values (low trust:1, high trust:0)



Notes: Figure plots the relationship shown in Figure 3.24, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.




Figure 3.26: Predictive margins in confidence (self-control interaction)




Notes: Figure reports predicted outcomes in confidence, from an interaction of treatment status with trust found in Table 3.22. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation, albeit with level difference evident.

Figure 3.27: Difference in predicted values (low self-control:1, high self-control:0)




Notes: Figure plots the relationship shown in Figure 3.26, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.28: Predictive margins in accuracy (gender interaction)




Notes: Figure reports predicted outcomes in accuracy, from an interaction of treatment status with gender found in Table 3.17. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation, albeit with tentative level difference evident.

Figure 3.29: Difference in predicted values (female:1, male:0)




Notes: Figure plots the relationship shown in Figure 3.28, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.30: Predictive margins in confidence (gender interaction)




Notes: Figure reports predicted outcomes in confidence, from an interaction of treatment status with gender found in Table 3.23. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation, albeit with tentative level difference evident.

Figure 3.31: Difference in predicted values (female:1, male:0)




Notes: Figure plots the relationship shown in Figure 3.30, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Figure 3.32: Predictive margins in confidence (risk appetite interaction)



Notes: Figure reports predicted outcomes in confidence, from an interaction of treatment status with risk appetite found in Table 3.22. Superficial appearance of significant heterogeneity in treatment effect found in the regression table is shown to be immaterial on graphical representation, albeit with level difference evident.

Figure 3.33: Difference in predicted values (low risk appetite:1, high risk appetite:0)



Notes: Figure plots the relationship shown in Figure 3.32, but focuses on the difference across the moderating variable at each treatment cell, which has its own error term. This is to establish whether that difference is statistically different to zero.

Table 3.24: Overall effect - pooled treatments

	(1)	(2)	(3)	(4)	(5)	(6)
	Overall	Overall	True positive	True positive	True negative	True negative
Treated	-0.007 (0.010)		-0.032 (0.021)		0.019 (0.020)	
Treated (2 or 3)		-0.005 (0.011)		-0.036 (0.022)		0.026 (0.021)
Constant	0.654*** (0.016)	0.642*** (0.018)	0.591*** (0.033)	0.582*** (0.035)	0.718*** (0.030)	0.702*** (0.033)
Observations	780	585	780	585	780	585
R-squared	0.084	0.077	0.018	0.018	0.059	0.073
p-value ( $\beta \leq 0$ )	0.746	0.671	0.939	0.945	0.174	0.103


Notes: Table reports pooled treatment effects in overall accuracy, true positives, and true negatives, as an aggregated counterpart to Table 3.6. Columns 1, 3, and 5 report all treatment arms pooled against the control group, while Columns 2, 4, and 6 pool only Treatments 2 and 3 against the control group. Table reports results from two-sided test for pooled treatment effects on overall accuracy, true positives, and true negatives. Also reported are one-sided tests of pre-specified hypotheses for incremental positive treatment effects from each treatment arm compared against the preceding arm in the sequence. Regression includes vector of controls listed in Table 3.2. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Table 3.25: Confidence effect - pooled treatments

	(1)	(2)	(3)	(4)	(5)	(6)
	Overall	Overall	Genuine	Genuine	Fraudulent	Fraudulent
Treated	0.157 (0.096)		0.134 (0.098)		0.173* (0.100)	
Treated (2 or 3)		0.222** (0.102)		0.193* (0.105)		0.241** (0.105)
Constant	5.896*** (0.144)	5.848*** (0.154)	5.853*** (0.147)	5.785*** (0.158)	5.930*** (0.148)	5.893*** (0.161)
Observations	780	585	780	585	780	585
R-squared	0.158	0.175	0.146	0.159	0.153	0.173
p-value ( $\beta \leq 0$ )	0.0512	0.0152	0.0864	0.0332	0.0416	0.0109


Notes: Table reports pooled treatment effects in overall confidence, confidence in reported in genuine calls, and confidence reported in fraudulent calls, as an aggregated counterpart to Table 3.7. Columns 1, 3, and 5 report all treatment arms pooled against the control group, while Columns 2, 4, and 6 pool only Treatments 2 and 3 against the control group. Table reports results from two-sided test for pooled treatment effects on overall accuracy, true positives, and true negatives. Also reported are one-sided tests of pre-specified hypotheses for incremental positive treatment effects from each treatment arm compared against the preceding arm in the sequence. Regression includes vector of controls listed in Table 3.2. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Figure 3.34: Accuracy and confidence: scenario averages



Notes: Figure describes a scatter plot of mean accuracy achieved in each scenario against mean confidence reported by participants in their judgement on that scenario, and shows a positive correlation.

Figure 3.35: Accuracy and confidence: scenario averages, separated by status



Notes: Figure describes a scatter plot of mean accuracy achieved in each scenario against mean confidence reported by participants in their judgement on that scenario, separating by genuine and fraudulent scenarios. The positive overall correlation observed in Figure 3.34 is shown to be driven by true negatives.

Table 3.26: Key signs of fraud

Key sign	Description
Fabricated sense of urgency	Pressuring you to “act now” or else a deal will go away, your account will be closed, or you will experience other negative consequences.
Random outreach	You are contacted out of the blue, e.g., the message comes from an unfamiliar email address, behind what looks like a genuine sender name, or phone call etc. and it is hard to understand why you are being contacted.
Unfamiliar but genuine looking email	The message comes from an unfamiliar email address, behind what looks like a genuine sender name.
Poorly written message	The message is poorly written with misspellings and incorrect grammar, or a familiar company name is misspelt.
Personal information theft	Asking for personal information and access to your money-such as your ATM cards, bank accounts, credit cards, or investment account, or for you to confirm personal information they claim to have.
Suspicious call	Calling or emailing you, claiming to be from the government and asking you to pay money.
Suspicious offer	The offer seems too good to be true.

Notes: Table lists the seven key signs of fraud included in Treatments 2 and 3.

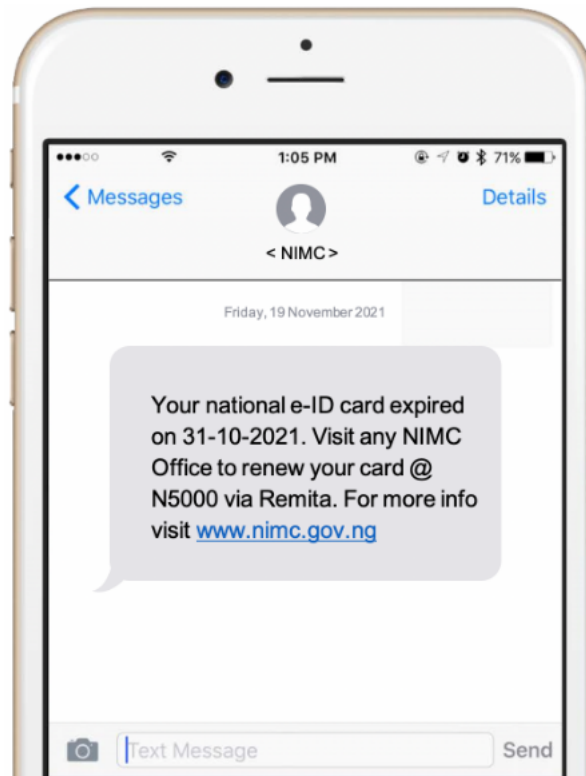


Table 3.27: Overview of scenarios

Scenario	Genuine/Fraudulent	Description
1	F	Bank email account update
2	F	Pop-up window link to claim prize
3	G	Bank text account update
4	F	Mobile company call for sensitive info
5	F	Call re investment opportunity, seeks transfer and personal info
6	F	WhatsApp lawyer inheritance, seeks processing fee and personal info
7	G	SMS chance to win
8	G	SMS gov ID expired
9	F	Email delivery update fee
10	F	Call+SMS re. accidental cash transfer
11	F	Call re. prize giveaway, seeks processing fee
12	G	Email annual account statement
13	G	Email survey
14	F	Email investment opportunity
15	G	Email order collection
16	G	Call re. gov loan scheme
17	F	WhatsApp offer with fee
18	G	Email reward offer
19	G	Email cash offer
20	G	Email customer survey

Notes: Table reports an overview of the 20 scenarios which participants are asked to evaluate as part of the experimental task.

Figure 3.36: Example of a genuine scenario in the experimental task

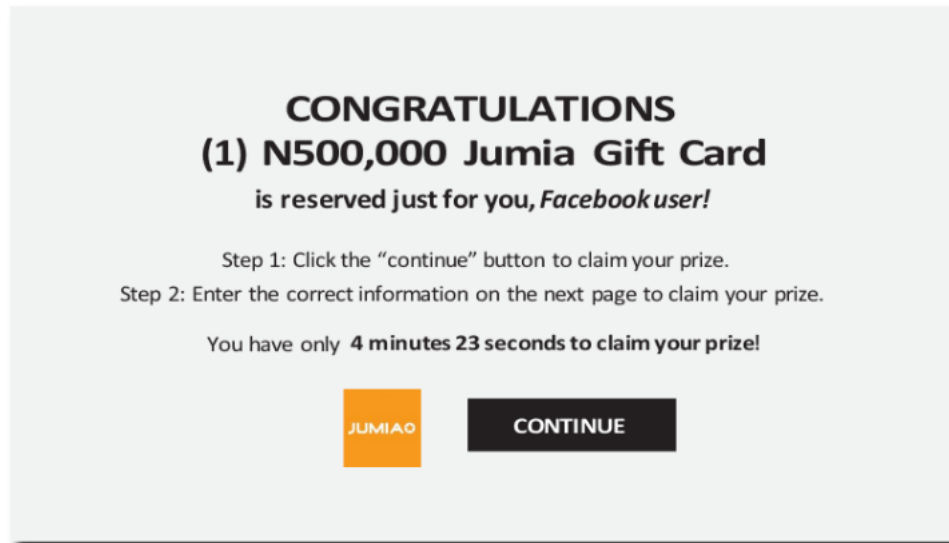


Do you think this is a genuine or fraudulent scenario?

Kana/kina ganin wannan labari ba na damfara bane ko kuwa na damfara ne?

- Genuine  
Ba damfara bace
- Fraudulent  
Damfara ce

Figure 3.37: Example of a fraudulent scenario in the experimental task



Do you think this is a genuine or fraudulent scenario?

Kana/kina ganin wannan labari ba na damfara bane ko kuwa na damfara ne?

Genuine  
Ba damfara bace

Fraudulent  
Damfara ce

Figure 3.38: Elicitation of confidence in judgements following each presented scenario

How confident are you in your answer?  
 Minene matakin tabbacin amsar da ka/kika bada?

<b>1</b> <b>Very</b> <b>Uncertain</b> <b>Ba nida</b> <b>tabbaci</b> <b>sosai</b>	<b>2</b> <b>Uncertain</b> <b>Ba nida</b> <b>tabbaci</b>	<b>3</b> <b>Somewhat</b> <b>Uncertain</b> <b>Rashin</b> <b>tabbaci na</b> <b>kadan ne</b>	<b>4</b> <b>Neutral</b> <b>Tsakatsakiya</b>	<b>5</b> <b>Somewhat</b> <b>Confident</b> <b>Ina da</b> <b>tabbaci</b> <b>kadan</b>	<b>6</b> <b>Confident</b> <b>Ina da</b> <b>tabbaci</b>	<b>7</b> <b>Very</b> <b>Confident</b> <b>Ina da</b> <b>tabbaci</b> <b>sosai</b>
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Next

Table 3.28: Definition of variables used

Variable	Definition
No direct experience	<p data-bbox="703 309 1430 383">Binary variable recording that the participant reports ‘No’ to the question:</p> <ul data-bbox="751 427 1430 555" style="list-style-type: none"> <li data-bbox="751 427 1430 555">• Have you ever experienced someone contacting you pretending to be someone else to steal money or sensitive information?</li> </ul>
Low self-control	<p data-bbox="703 633 1430 808">Binary variable recording that the participant falls below the median value of a standardised index constructed from the participant’s level of agreement with the following questions:</p> <ul data-bbox="751 853 1430 1357" style="list-style-type: none"> <li data-bbox="751 853 1430 927">• I spend too much in the moment and let the future take care of itself</li> <li data-bbox="751 949 1430 1023">• Financial services are complicated and confusing to me</li> <li data-bbox="751 1046 1430 1120">• Convenience plays an important role in the decisions I make</li> <li data-bbox="751 1142 1430 1216">• I often act without thinking through all the alternatives</li> <li data-bbox="751 1238 1430 1267">• I am optimistic about my future</li> <li data-bbox="751 1290 1430 1364">• If I work hard today, I will be more successful in the future</li> </ul>

---

Variable	Definition
Low risk appetite	<p data-bbox="592 250 1310 421">Binary variable recording that the participant falls below the median value of a standardised index constructed from the participant's responses to the following questions:</p> <ul data-bbox="639 465 1310 878" style="list-style-type: none"><li data-bbox="639 465 1310 779">• Suppose you're offered a business investment that returns 5,000 Naira on average. Half the time the investment returns 10,000 Naira. However, half of the time the investment returns nothing. What is the maximum you personally would be willing to pay to make this investment?</li><li data-bbox="639 801 1310 878">• Indicate your level of agreement for the following statement: I am a person who takes risks</li></ul>
Low trust	<p data-bbox="592 958 1310 1128">Binary variable recording that the participant falls below the median value of a standardised index constructed from the participant's responses to the following questions:</p> <ul data-bbox="639 1173 1310 1348" style="list-style-type: none"><li data-bbox="639 1173 1310 1209">• In general, most people can be trusted</li><li data-bbox="639 1227 1310 1303">• I often reject statements unless I have proof that they are true</li><li data-bbox="639 1321 1310 1348">• I frequently question things that I see or hear</li></ul>

---

---

Variable	Definition
Low DFS experience	<p data-bbox="703 248 1430 421">Binary variable recording that the participant falls below the median value of a standardised index constructed from the participant's responses to the following questions:</p> <ul data-bbox="751 465 1430 1928" style="list-style-type: none"> <li data-bbox="751 465 1430 539">• Do you have an account at a formal financial institution?</li> <li data-bbox="751 562 1430 636">• Have you used a formal bank account in the last 90 days?</li> <li data-bbox="751 658 1430 687">• When did you first get a formal bank account?</li> <li data-bbox="751 710 1430 784">• Can you access your formal bank account or bank application on your phone?</li> <li data-bbox="751 806 1430 880">• Have you accessed your formal bank account on your phone in the last 90 days?</li> <li data-bbox="751 902 1430 931">• Do you use a phone for conducting business?</li> <li data-bbox="751 954 1430 1028">• Have suppliers contacted you on your personal phone or business phone in the last 90 days?</li> <li data-bbox="751 1050 1430 1211">• Do you have a mobile money account? E.g. Paga Mobile, MTN Momo, First Banks Firstmonie, Kudi Mobile, UBA Moni Agent or Polaris Sure Padi</li> <li data-bbox="751 1234 1430 1308">• Have you ever transferred money to another individual or business using your phone?</li> <li data-bbox="751 1330 1430 1449">• Have you used mobile money or any other digital payments provider to send or receive payments in the last 90 days?</li> <li data-bbox="751 1471 1430 1590">• When did you first use mobile money or any other digital payment services to send or receive payments?</li> <li data-bbox="751 1612 1430 1731">• Do you usually access mobile money or other digital payment services using an app, short-code menu or SMS?</li> <li data-bbox="751 1753 1430 1827">• Have you ever bought or sold goods using an online platform (e.g. Jumia)?</li> <li data-bbox="751 1850 1430 1928">• When was the first time you bought or sold goods using an online platform?</li> </ul>

---

---

Variable	Definition
Low ICT experience	<p>Binary variable recording that the participant falls below the median value of a standardised index constructed from the participant's responses to the following questions:</p> <ul style="list-style-type: none"><li>• Do you have access to a smartphone?</li><li>• Are you the primary user for your smartphone?</li><li>• On average over the past 30 days, how often have you used a smartphone to do any of the following:<ul style="list-style-type: none"><li>– To make phone calls?</li><li>– To send SMS messages?</li><li>– Use messenger apps (Facebook messenger, WhatsApp, etc.)</li><li>– Browse social media?</li><li>– To conduct purchasing transactions</li><li>– To conduct banking transactions</li></ul></li></ul>
Agriculture	<p>Binary variable recording whether the participant is employed in the agricultural sector (1), or a non-agricultural sector (0).</p>
Above median age	<p>Participant age is greater than 25 years.</p>
Lower education	<p>Binary variable recording whether the participant's level of educational attainment is secondary or below (1), or tertiary (0).</p>
Female	<p>Binary variable recording whether the participant identifies as female (1), or not (0).</p>

---



---

Variable	Definition
ICW trust index	<p data-bbox="703 248 1425 472">Inverse correlation weighted index of trust constructed from the participant’s responses to the following questions (computed at baseline and endline): How likely are you to use these types of DFSs (Digital Financial Services) in the future?</p> <ul data-bbox="751 517 1390 734" style="list-style-type: none"> <li>• Banks</li> <li>• Mobile banking</li> <li>• Mobile money operators</li> <li>• Online platforms for buying or selling goods</li> <li>• Agents</li> </ul> <p data-bbox="703 779 1425 904">Indicate your level of agreement for the following statement: In general, I trust that my financial information is kept safe by -</p> <ul data-bbox="751 949 1390 1167" style="list-style-type: none"> <li>• Banks</li> <li>• Mobile banking</li> <li>• Mobile money operators</li> <li>• Online platforms for buying or selling goods</li> <li>• Agents</li> </ul> <p data-bbox="703 1211 1425 1337">Indicate your level of agreement for the following statement: In general, I trust that my money is kept safe from fraud or theft by:</p> <ul data-bbox="751 1382 1390 1599" style="list-style-type: none"> <li>• Banks</li> <li>• Mobile banking</li> <li>• Mobile money operators</li> <li>• Online platforms for buying or selling goods</li> <li>• Agents</li> </ul>
Age	Reported age of participant.
Third level education	Binary inverse of ‘Lower education’ variable.
Married	Binary variable recording whether the participant is married (1) or not married (i.e. divorced, separated, single, or widowed) (0).
Contacted by scammer	Binary inverse of ‘No direct experience’ variable.

---

---

Variable	Definition
Access to smartphone	Binary variable recording whether the participant has access to a smartphone (1), or not (0).
Business owner	Binary variable recording whether the participant is a business owner (1), or not (0).
Has formal financial account	Binary variable recording whether the participant answered has an account at a formal financial institution (1), or not (0).
Used online platforms	Binary variable recording whether the participant has ever bought or sold goods using an online platform (1), or not (0).
Trusting	Binary inverse of 'Low trust' variable.
Risk averse	See 'Low risk appetite'.

---

Notes: Table reports a description of each variable used in the analysis.

# Chapter 4

## A Lender of Last Resort: Access to Emergency Financial Support and Household Financial Risk Taking

### 4.1 Introduction

A longstanding and central question in household finance is why so many households do not participate in the stock market, and take advantage of the equity premium and opportunity for wealth accumulation that financial markets offer. While it is clear that many households' circumstances are such that holding any substantial exposure to the stock market is not the optimal course, a key tenet of portfolio theory holds that all expected utility maximising households should hold at least some fraction of their wealth in riskier assets to benefit from the equity premium (Haliassos and Bertaut, 1995).<sup>1</sup> Despite this, 4 in 5 European households<sup>2</sup> opt-out entirely from stock market participation. Participation rates in the United States are consistently higher, at over 50% (Parker and Fry, 2020), but non-participation still presents a significant puzzle.

Explanations for cross-country variation in stock market participation typically cite institutional and cultural factors such investor protection, trust, and gender norms (Ke, 2018). Addressing variation in stock market participation between the US and

---

<sup>1</sup>In the basic two-period, two asset expected-utility model, and in the presence of an equity premium, households should always be willing to invest at least a small amount in the asset offering higher expected return. At the optimum, households should be unable to do better by shifting a marginal sum from one asset to the other, a condition which is violated at zero stock holding (Haliassos and Bertaut, 1995).

<sup>2</sup>Author's calculation from *Household Finance and Consumption Survey*.

European countries, Guiso et al. (2003) highlights several important factors, including a stronger perceived ‘carrot’ and ‘stick’ in the United States through higher average yearly returns on the one hand, and comparatively inadequate social security provisioning for the future on the other; differential tax incentives; the predominance of high-impact employer-sponsored seminars on stock market participation and retirement planning in the US; and higher average actual and perceived transaction costs in Europe.

Research in the under-participation of households in stock markets has been motivated by a diverse set of factors from political economy: including the opportunity cost in the failure to take advantage of the equity premium (Haliassos, 2003), and the consequent implication for long-term patterns of wealth-inequality (Piketty, 2014), the under-diversification of household portfolios and the implied concentration risk in labour income and real estate assets (Campbell, 2006), the elasticity of intertemporal substitution (Vissing-Jørgensen, 2002), and even the political economy of financial regulation, by engendering divergent attitudes towards corporate and investment income tax policy as well as risk-sharing and redistribution among those with and without financial assets (Cole et al., 2014).<sup>3,4</sup>

Piketty (2014) famously establishes the central role that can be ascribed to higher investment returns in driving patterns of wealth inequality, summarised by the inequality “ $r > g$ ”, that the rate of return on capital is greater than the growth rate of the economy. This wealth wedge mechanism is partly levered through differential stock market participation rates among households, resulting in differential access to the equity premium which is known in turn to be instrumental in driving patterns of wealth inequality (Guiso et al., 2003). In this light, sharply divergent patterns of stock market participation can be regarded as an important lacuna in the democratic access to capital growth.

Several direct estimates of the welfare loss accruing to households from non-participation have been put forward. Calvet et al. (2007), using government records covering all Swedish households over the period from 1999 to 2002, estimate that the return loss

---

<sup>3</sup>As Guiso et al. (2003) underlines, while wider stock market participation promises to expand access to financial instruments bearing higher expected returns, facilitate household portfolio diversification, and ultimately contribute to a reduction in wealth inequality, it also raises concerns regarding the financial sophistication of new entrants, and their vulnerability to volatile swings in asset values. This risk highlights the importance of effective consumer protection regimes, information disclosure, and suitability and appropriateness assessment in mediating the sale of investment products and any expansion of participation in the population.

<sup>4</sup>Somewhat against the tide, Calvet et al. (2007) argue that standard analysis may overstate the cost of non-participation by failing to take account of the likely underdiversification of the portfolios of non-participant households were they to enter the market. In this light, non-participation may be regarded as a smaller ‘investment mistake’ than it usually is.

from non-participation of 4.3% (return units on the household portfolio) under the assumption that such households would invest efficiently and earn the equity premium.<sup>5</sup> In a separate study using US data from the *Panel Study of Income Dynamics*, Cocco et al. (2005) estimate the cost from non-participation to be 2.1% of annual consumption. Larsen and Munk (2021) estimate that the expected wealth at retirement of a ‘stock avoider’ is more than 20% lower and her expected consumption at age 80 is 31% lower when compared against the optimal life-cycle stock-bond allocation.

Explanations put forward to help account non-participation have examined, *inter alia*, the role of demographic and socio-economic factors (Van Rooij et al., 2011; Almenberg and Dreber, 2015; Guiso et al., 2003), personality traits (Conlin et al., 2015), social network and community effects (Brown et al., 2008; Hung, 2021), institutions and trust (Asgharian et al., 2019; Guiso et al., 2008; Georgarakos and Pasini, 2011), social norms (Ke, 2018), and informational and entry costs (Bertaut, 1998; Peress, 2005). Collectively, factors such as these can go a long way empirically to accounting for asymmetrical patterns of stock market participation among households.

This paper examines a previously unexplored dimension which builds upon the social and network strand of literature: the role of private lenders of last resort, specifically, the ability of some households to call on financial support from friends or family in an emergency.<sup>6</sup> Building from a theory of moral hazard (Kunreuther and Michel-Kerjan, 2014) and insurance-induced consumption whereby agents act more incautiously or consume more of a service under the umbrella of insurance coverage, it is intuitive that those households who are privately insured against adverse financial shocks will be more willing to assume financial risk. With emergency support, and in a downside investment scenario, the investor “*personally must pay only a fraction of the costs of the damage*” (Stiglitz, 1983).

A distinct literature has addressed the various impacts from financial and economic support mechanisms on household financial behaviour and outcomes. These include the impact of inheritances on reducing labour supply, driving wealth inequality, and aiding the transition to entrepreneurship (Brown et al., 2010; Basiglio et al., 2019; Holtz-Eakin et al., 1993), the impact of informal support networks in meeting the economic needs of low income families (Lee et al., 2020; Harknett, 2006; Henly, 2002), and the impact of public safety nets on recipients’ economic outcomes such as bankruptcy, credit access

---

<sup>5</sup>Under another counterfactual scenario where non-participants invest relatively inefficiently (i.e. adjusting for potential underdiversification of the hypothetical portfolio of the non-participating household), the return cost falls to 2.3%.

<sup>6</sup>This paper uses data from the Eurosystem *Household Finance and Consumption Survey*. With thanks to Tara McIndoe-Calder and Laura Boyd for providing access.

and performance, poverty, and self-sufficiency (Deshpande et al., 2017; Kanz, 2016; Hoynes et al., 2016).

This paper adds to this stock of existing literature by offering one of the first examinations of how the *option* of financial support, as distinct from the actual delivery of financial support from friends and family influences the household financial behaviour.<sup>7</sup> Furthermore, the paper contributes to our understanding of advantage, portfolio choice, and access to financial markets which play an important role in driving patterns of widening income and wealth inequality. Most importantly, ours is the first paper to address how the option of financial support relates to the likelihood of stock market participation and to risk appetite.

While the option of emergency financial support arguably offers a highly practical facility which would influence household financial behaviour, this research should also be seen as helping to illuminate an important aspect of the psychology of money, namely, how the peace of mind brought about by emergency financial buffers may relate to financial risk-taking.<sup>8</sup>

We use comprehensive data from the *Household Finance and Consumption Survey* (HFCS) covering 21 countries, in tandem with propensity score matching methods in an effort to establish observably comparable cohorts of households, differing only in their access to emergency financial support. With this approach, we evaluate the impact of the option of financial assistance from friends or family on the likelihood of holding financial investments in the form of publicly traded shares, mutual funds, money market mutual funds or hedge funds, and on self-reported financial risk appetite.<sup>9</sup>

Our results show that those households that enjoy the option of emergency financial support are approximately 6% more likely to participate in the stock market, and approximately 2% more likely to report higher financial risk appetite. We find a high degree of statistical consistency in these effects across subgroups of interest, including credit constrained and unconstrained households, and according to the gender and age of the head of household. We observe tentative evidence that the mechanism is not

---

<sup>7</sup>Crucially, our findings are not country-specific, but encompass instead a broad sample of 21 euro-area countries.

<sup>8</sup>The private bailout facility is not an insurance contract which formally reduces downside risk for the benefitting household, but as we discuss in Section 4.3, the mechanism which we hypothesise to be at work in this setting operates in exactly this manner psychologically. Recent years have seen an explosion of interest in the manifold ways in which psychology plays a role in determining the shape of our financial lives, and this paper should be viewed in that context.

<sup>9</sup>By way of orientation: 57% of respondents in our data enjoy the option of emergency financial support, 20% hold stocks or mutual funds, and 6% report having a willingness to take higher than average financial risks. The median value of investment holdings among participating households is €10,000. 48% hold less than €10,000, 60% hold less than €20,000, and 75% hold less than €50,000.

as strong among older households, consistent with a life cycle pattern with younger households facing greater immediate financial pressures while also holding a higher latent risk appetite, such that the shadow price on financial insurance is greater. Our main results are found to be robust to a wide range of alternative specifications and matching algorithms.

The results point towards a mechanism of advantage compounding advantage.<sup>10</sup> In doing so, they highlight the importance of effective and sustainable policies of financial inclusion. Stronger targeted financial literacy and education initiatives should be considered in order to equip less advantaged households with the tools to consider, evaluate, and navigate the opportunities, risks, as well as the opportunity costs associated with different paths in the household financial landscape. The paper's findings highlight the potential for existing market tools such as downside investment protection in replicating the insurance role provided by private safety nets to extend participation among European households. As well as constituting a potential market opportunity, this possibility also carries an implication for public policy. Specifically, for the intelligent, agile and robust enforcement of consumer protection regimes in the domain of retail investment and any associated add-on products.<sup>11,12</sup>

A series of plausible unobserved household and background network characteristics do have the potential to confound our estimated impact from the emergency financial support option.<sup>13</sup> While it is possible to narrow this channel with credible proxy measures such as those relating to education, field of employment, and inheritance history, it is not possible to definitively exclude the possibility that certain unobserved background network effects create some confounding influence over our estimates. This must be

---

<sup>10</sup>An effect also referred to as 'the Matthew effect of accumulated advantage' (Rigney, 2010).

<sup>11</sup>Our results do not necessarily imply that governments should step in to insure stock market participation in an effort to push non-participating households into the stock market. However, in recognition of the implication of widespread non-participation for widening wealth inequality, under-diversification of household portfolios, as well as political economy, government should understand the issue, and foster the conditions for effective financial inclusion through education and literacy initiatives as well as consumer protection regimes which enable households to make fully informed decisions within the financial landscape for themselves.

<sup>12</sup>Financial literacy initiatives and consumer protection regimes represent, of course, nontrivial costs on public resources. Education campaigns to foster inclusion should be carefully weighed in terms of the opportunity cost in time, money, and attention, as well as carefully evaluated in their actual efficacy. In most jurisdictions, robust consumer protection will form part of the discharge of existing statutory public responsibilities. The paper highlights, however, the importance of vigilance and agility in this regard.

<sup>13</sup>We take account of a broad set of relevant observable factors affecting comparability across households with and without the option of emergency financial support (including income, education, circumstances of employment, net wealth, age, marital status, gender, number of children, tenure status, and a full set of country controls). We also subject our main results to a series of robustness exercises which variously adjust the specification used to compute propensity scores and the matching algorithm used to link treated and untreated observations.

transparently recognised, with the implication that our main estimated treatment effect cannot be regarded as a directly identified causal estimate of the impact of the emergency financial support option, but rather as a composite which likely includes some potential influence attributable to network-based exposure to financial markets.

The paper proceeds as follows: Section 4.2 provides a review of relevant literature and outlines the contribution of the paper, Section 4.3 describes the hypothesised mechanism at work, Section 4.4 describes our propensity score matching methodology, and Section 4.5 describes in detail the data and model choices underpinning the analysis. Section 4.6 presents empirical results, Section 4.7 provides further discussion of key results and policy implications, and Section 4.8 concludes.

## 4.2 Literature

This paper can be seen at the juncture of two distinct strands of the household financial literature. Firstly, the long-running research focus exploring the determinants and obstacles to stock market participation, and secondly, the literature which addresses the impact of financial and economic support mechanisms on household financial behaviour and outcomes.

### 4.2.1 Stock market participation

An array of demographic factors have been shown to play an important role in predicting asymmetrical patterns of stock market participation. Van Rooij et al. (2011) explore how financial literacy affects the probability of stock market participation, and find that a lack of understanding of economics and finance act as a significant deterrent to stock ownership. This relationship is further demonstrated by Arrondel et al. (2015), who find that, controlling for other relevant factors, financial literacy is a significant predictor of participation, but not of the portfolio share allocated to stocks conditional on participation (i.e. literacy helps to overcome the entry costs that obstruct individuals from holding stocks, but does not influence portfolio allocation once the decision to participate in the stock market has been made).

Using household survey data in Sweden, Almenberg and Dreber (2015) document a significant gender gap in stock market participation. However, the authors find that this gap is significantly reduced once adjustment is made for financial literacy. By contrast, a documented gender gap in risk appetite, with women found to be less risk-seeking than men, remains after adjusting for financial literacy.

Guiso et al. (2003) offer a comprehensive review of differential direct stock market participation patterns across France, Germany, Italy, the Netherlands, Sweden, the



UK, and the US. In a regression setting which adjusts for other relevant factors, the authors confirm that the likelihood of direct participation is strongly increasing in income, wealth, and education levels (a pattern additionally documented by Poterba and Samwick (2003)).

The authors observe that the effect of wealth is stronger than the effect of income. They report that the probability of participation in Italy is over 30 percentage points higher in the fourth wealth quartile than in the lowest quartile. In France it is over 40 points higher. In the Netherlands, Sweden, Germany and the UK it is about 50 points higher, and 63 points higher in the US. In France, Germany, Italy and the Netherlands the probability of direct stockholding increases by at most 13 percentage points, moving from the first to the fourth income quartile, while in Sweden, the UK and the US the increase is 22 points. The authors find that in all countries included in their analysis, that college education significantly increases the likelihood of stockholding, ranging from 4 points increased likelihood in Germany, to 8 points in the US. Notably, in most countries under analysis, age-related coefficients are not found to be statistically different from zero, notwithstanding a univariate descriptive finding of concave age-participation profiles.

Kaustia and Torstila (2011) even postulate a role for ‘value-expressive’ political abstention from the stock market. Using data from Finland, the authors show how political preferences help to predict patterns of participation. Controlling for income, wealth, education, and other relevant factors, a change of one point to the left on a 1-10 right left scale is associated with a 5-6% reduced likelihood of equity market participation.

Pushing beyond the more obvious socio-economic and demographic explanations, considerable research has focused on the role of social network and community effects. Brown et al. (2008) provides the first paper which documents the causal importance of the community effect in predicting stock market participation, with individuals whose community has a higher fraction of investors being more likely themselves to participate, conditional on a rich set of controls. The authors proffer evidence for a ‘word of mouth’ effect by showing the result to be stronger in more sociable communities.<sup>14</sup> The paper further demonstrates a ‘local firm effect’ whereby, independent of a direct employment-ownership channel, the probability of stock ownership is positively impacted by the presence of publicly traded firms in the community.

Guiso et al. (2008) provide seminal research documenting the role of trust in shaping patterns of stock market participation, using micro data from the Netherlands and

---

<sup>14</sup>This social network and word of mouth effect is further demonstrated more recently by Hung (2021), who employs novel data from Facebook and a social connectedness index to estimate the impact of peer participation on individual participation decisions.

Italy. They find that individuals who report higher levels of trust are significantly more likely to buy stocks: trusting others increases the likelihood of participation by 50% of the average sample probability. The authors additionally document how the share of stockholders across countries correlates positively with the average levels of prevailing trust, even controlling for legal enforcement and legal origin (i.e. common law systems).

Georgarakos and Pasini (2011) confirm the relationship between trust and participation, but interestingly find that sociability (as measured by participation in voluntary activities, educational training courses, sport and political clubs, etc.) can partly offset the discouragement effect on stock-holding induced by low prevailing regional trust. The authors contend that social interactions enable information sharing, which serves to reduce information costs and lower the stock market participation hurdle. In earlier research, Hong et al. (2004) similarly observe that stock-market participation can be influenced by social interaction, with those who interact with neighbours or attend church being substantially more likely to invest in the market than non-social households.

A variety of cognitive and behavioural explanations have also been put forward. These include ‘narrow-framing’ in how people evaluate risk, which refers to the tendency for an agent to evaluate a new gamble in isolation, rather than in the holistic context of the existing suite of risks the agent already faces, with the result that such gambles are too often rejected (Barberis et al., 2006). Similarly, ambiguity aversion (Dimmock et al., 2016), and the inertial influence of costly information (Haliassos and Bertaut, 1995) are commonly advanced to help to account for the non-participation puzzle.

In this paper, I focus on a new dimension, which builds upon this social network strand of the literature: the importance of financial support from friends and family.

### **4.2.2 Impact of support mechanisms**

While the paper can be located in broad terms within the literature on stock market participation, the specific mechanism under analysis situates the paper more directly within a second strand of the literature: the literature addressing the impact of financial and economic support mechanisms on household financial behaviour and outcomes. This literature has traditionally focused on three broad areas: (1) the impact of inheritances in shaping the transition to homeownership, and patterns of wealth accumulation and inequality, (2) the impact of informal support networks in meeting the economic needs of low income families, and (3) the impact of public safety nets on recipients’ economic outcomes. Very little if any research, however, has addressed how the option of emergency financial support from friends and family influences household financial behaviour in the general population.

## Inheritances

Extensive research attention has focused on the multifaceted impacts inheritance receipt can bring to bear on household financial decision making. Using Dutch household panel survey data, Basiglio et al. (2019) show how forward-looking inheritance expectations lead to a reduction in savings and expected labour supply: with a 1 percentage point increase in probability of receiving an inheritance giving rise to a 4 percentage points higher probability of dissaving. In addition, those who anticipate receiving large inheritances report an expected probability of working beyond the age of 62 that is 20 per cent lower than their counterparts.

Joulfaian and Wilhelm (1994) also explore the extent of a labour disincentive arising from inheritance, and find only very modest impacts on the number of hours worked. The authors report that the elasticity of consumption with respect to inheritance is greater than that for hours worked, and that the result is highly robust, but the magnitude of the effect is small.

Brown et al. (2010) comprehensively examine the impact of inheritance receipt on the probability of retirement. Using data from the 1992-2002 *Health and Retirement Study*, the authors find receiving inheritance is associated with a 12 per cent increase in the likelihood of retirement in the subsequent two-year period, and by 7 per cent relative to the baseline over an eight-year period. They find that the impact is stronger for larger inheritances, and when the inheritance is unexpected.<sup>15</sup> Amedah et al. (2017) similarly document how, conditional on a set of potential confounding variables, inheritance receipt increases the retirement hazard by approximately 20 per cent relative to non-recipients.

Holtz-Eakin et al. (1993) examine the impact of inheritance receipt on the probability of transitioning to entrepreneurship, using matched US federal estate and personal income tax returns. They find that a \$100,000 inheritance increases the probability of becoming an entrepreneur by 3.3 percentage points. They also find that this effect is more pronounced for those with low initial wealth when compared against those with high initial wealth, consistent with an interpretation that the incremental impact of inheritance should be lower for those who are less liquidity-constrained ex-ante. The entrepreneurship channel is further documented in Blanchflower and Oswald (1998), who find small but significant inheritance effects on the probability of subsequent self-employment, especially so for the young who, unlike their older counterparts, may have fewer alternative avenues to acquiring capital.

<sup>15</sup>The authors show that the effect of a given dollar amount of inheritance on the probability of early retirement is over twice as large when unexpected: the effect of raising the inheritance value by \$100,000 is to increase the probability of early retirement by 10.3 percentage points if the inheritance is unexpected and by 4.3 percentage points if it is expected.

## Informal support networks

Numerous papers have explored the diverse angles through which ongoing familial support mechanisms or ‘private safety nets’ influence economic outcomes.<sup>16</sup> For instance, Lung et al. (2020) show medical students who reported that they “very likely” could turn to parents or family for help with an unforeseen expense also showed greater subjective happiness compared to those who perceived they were less than very likely to receive financial support, while adjusting for other relevant factors. The finding illustrates the possibility of psychological impacts from economic security.

Lee et al. (2020) document the impact of parental financial assistance on the transition to homeownership, reporting that receiving a transfer greater than \$5,000 for any purpose increases the probability of transitioning to homeownership by 15.1%. Köppe (2018) shows how in-kind support from parents through co-residence benefits young people who are able to purchase their first home earlier than independent mortgagors saving for a deposit while renting.

Using data from in-depth interviews with a diverse group of people who experienced job loss in the period 2007-2011, Gould-Werth (2018) illustrates how rainy day support from ‘private resource banks’ including reserves of personal resources and social connections amassed during more favourable times assist in smoothing the return to work following job loss, with the uneven distribution of such resources contributing to a magnification of racial inequality arising from the employment interruption.

Harknett (2006) outlines how access to private safety nets for low-income single mothers (in this case, the option to draw upon family and friends for material or emotional support) correlates with human capital deficits, depressive symptoms, and low self-efficacy, indicating how social network disadvantages can reinforce individual-level disadvantage. Combining survey data with administrative records, the author finds that controlling for potentially confounding individual and household characteristics, mothers with the greatest social support also worked on average almost a full quarter (3 months) more than mothers with the least social support over a 3-year evaluation period.<sup>17</sup>

---

<sup>16</sup>The term ‘private safety net’ has been frequently employed to refer to a network of informal and formal relationships that families can draw upon to meet their needs, frequently childcare-based (see for instance Edin and Lein (1998) and Katras et al. (2004)).

<sup>17</sup>Henly (2002) similarly illustrates the role of informal systems of support (social networks of relatives, friends, and acquaintances) in facilitating the labour market involvement of lower-skilled, low-wage mothers.

## Public safety nets

A final strand of the literature on the impact of support mechanisms in household finance addresses the diverse impacts of from public/social safety nets in expanding both opportunities and risks for recipients.<sup>18</sup> Related to our study, Gormley et al. (2010) document a positive correlation between the generosity of unemployment insurance programmes stock market participation at the state level in the US, illustrating a role for public safety nets in household portfolio decisions. The authors postulate that that households which might struggle to weather large adverse economic shocks may rationally hold their wealth outside the stock market where it would be required in its entirety just to preserve subsistence consumption in such an adverse shock. Greater state-level unemployment insurance, the authors argue, should insulate more from subsistence-level income shocks, and induce higher participation.

Deshpande et al. (2017) estimates the impact of government cash assistance on household bankruptcy, using a natural experiment presented by a 1996 welfare reform law in the US. The authors note that the likely effect is theoretically unclear, as payments may stabilise income and reduce the likelihood of bankruptcy arising from unexpected income shocks, while on the other hand, stable payments may facilitate greater access to credit leading to exposure to bankruptcy where recipients struggle to service the debt properly. The authors find that the loss of cash assistance causes a sharp reduction in household bankruptcy rates, consistent with an access to credit interpretation.

Kanz (2016) exploits a natural experiment arising from India's 'Agricultural Debt Waiver and Debt Relief Scheme' to examine the impact of debt relief on financial outcomes, and presents a series of cautionary results. Firstly, debt relief did not have the effect of reintegrating recipient households into formal lending relationships but to an increased reliance on informal credit. Secondly, relief did not increase investment or agricultural productivity among beneficiaries. And finally, relief significantly influences recipients' expectations regarding the reputational consequences of future default on a bank loan, indicating a potential link between relief and induced moral hazard.<sup>19</sup>

In a study which exploits state-level variation in the generosity of the exemption threshold for the forfeiture of assets in personal bankruptcy filings, Gropp et al. (1996) identify the potential for distortionary impacts. The authors note that generous state-level

<sup>18</sup>A wider and less directly relevant literature addresses the impacts of public safety nets on measures of non-financial human development, such as Hoynes et al. (2016) who show how access to food stamps in childhood leads to a significant reduction in the incidence of metabolic syndrome and, for women, an increase in economic self-sufficiency in the long-run. Similarly, the impact of social safety net programmes in protecting against poverty and hardship have also been widely demonstrated (see for instance Singh et al. (2014), Ravallion et al. (1995).

<sup>19</sup>On the other hand, the authors find that recipients are significantly more concerned that debt relief will result in more binding borrowing constraints in future.

bankruptcy exemptions are typically regarded as benefiting less well-off borrowers, however their results indicate that they have the effect of shifting credit away from low-asset borrowers and towards high-asset borrowers. Borrowers in the lower half of the asset distribution are found to have less debt and face higher interest rates on car loans in states with high bankruptcy exemptions than borrowers in low-exemption states, while the probability of being turned down for credit or discouraged from borrowing is 5.5 percentage points higher in states with unlimited exemptions than in states in the bottom quartile of the exemption distribution.

### 4.2.3 Contribution

To this stock of existing literature, the present paper represents an important contribution for four primary reasons. It is one of few papers to address the question of how the option of financial support, as distinct from the actual delivery of financial support influences the household financial behaviour.<sup>20</sup> Secondly, it is the first paper to address how the option of financial support relates to the likelihood of stock market participation and to risk appetite. Thirdly, in doing so, the paper contributes to our understanding about advantage, portfolio choice, and access to finance which play an important role in driving patterns of widening income and wealth inequality. Further, this research question helps to illuminate an important aspect of the psychology of money, namely, how the peace of mind brought about by emergency financial buffers may relate to financial risk-taking. Finally, our findings relate not just to one country in isolation, but are incorporate a Europe-wide sample through the 21 participating countries in our chosen dataset.

## 4.3 Mechanism

To formalise our conceptual framework, it is useful to draw upon the theory of moral hazard and insurance-induced consumption. Fundamental theoretical mechanisms underpinning risk-taking and insurance have been understood within economics for many decades. Stiglitz (1983) summarises the central dynamic succinctly:

*“..This, then, is the fundamental conflict: the more and better insurance that is provided against some contingency, the less incentive individuals have to avoid the insured event, because the less they bear the full consequences of their actions”.*

This captures the central mechanism underpinning the hypothesis of this paper: the presence of emergency financial support acts as a type of indirect insurance which

---

<sup>20</sup>As described above, numerous studies have analysed the impact of actual support such as inheritance receipt, and public cash transfers on a range of economic outcomes, very few have studied how the *option* of support affects outcomes.

attenuates the downside penalty associated with a risky financial prospect such as stock market participation. In so doing, it induces the beneficiary along the probability distribution towards participation. We can represent this dynamic with a very simple system of equations - a simplified participation equation, and the standard formula for expected loss. At the household-level, we can say without controversy that the probability of stock market participation is increasing in the net expected gain from participation (Equation 1).

Focusing specifically on the expected loss term, and utilising the standard formula in finance for the calculation of expected losses (Equation 2), we recall that expected losses can be expressed as the multiple of the exposure at default (i.e. the market exposure in the adverse outcome), the probability of default (i.e. the inherent riskiness in the portfolio), and the loss given default (i.e. the financial loss accruing to the individual in the adverse outcome). In our conceptual framework and proposed mechanism, the bailout insurance facility acts to reduce the beneficiary's loss given default (*"he personally must pay only a fraction of the costs of the damage"*). Following the system of equations, when we reduce  $LGD_i$ , we reduce  $E(Loss)_i$ . When we reduce  $E(Loss)_i$ , we increase the net expected gain from participation, and as such, the probability of participation.

$$Probability(Participation)_i = f(E(Gain)_i - E(Loss)_i) \quad (4.1)$$

$$E(Loss)_i = EAD(i) \times PD_i \times LGD_i^{21} \quad (4.2)$$

The dynamics described here relate closely to the concept of moral hazard, where agents behave less cautiously as a result of having purchased insurance (see for instance Kunreuther and Michel-Kerjan (2014), and insurance-induced consumption, whereby insurance coverage encourages beneficiaries to consume more healthcare<sup>22</sup> than they would if they were uninsured (see for instance Frick and Chernew (2009), Joseph (1972)). In this paper, we are describing instead a relationship of insurance-induced participation.<sup>23</sup>

---

<sup>21</sup>EAD: Exposure at default; PD: Probability of default; LGD: Loss given default.

<sup>22</sup>The concept is most frequently employed in health economics, but is generalisable to other settings.

<sup>23</sup>Gormley et al. (2010) similarly highlight an insurance channel in predicting stock market participation for those households that might struggle to preserve subsistence consumption in an adverse shock, and for whom it may be rational to save all current wealth on a precautionary basis and not participate. The authors document a positive correlation between the generosity of state-level unemployment insurance and participation rates, as well as between private health insurance (in the form of employee-sponsored health insurance, life insurance, and long-term care insurance) and participation rates.

Behavioural change of this sort is typically regarded in economic theory as representing a welfare-diminishing deadweight loss in the efficient operation of markets (or ‘excess-consumption’, see for instance Feldstein (1973)). By contrast, Frick and Chernew (2009) posit that the presence of flaws in decision making and misinformation, which may result in the under-utilisation of services in the uninsured state, help to make the case for why insurance-induced consumption may improve efficiency.<sup>24</sup> The authors refer to this as ‘beneficial moral hazard’. Basic frictions in the uninsured state may include a simple lack of information about the consequences associated with different treatment conditions, cognitive limitations associated with processing complex information, or hyperbolic discounting of future welfare. To the extent that barriers such as these contribute towards inefficient underutilisation of services, moral hazard effects can be regarded as potentially beneficial. Indeed, Cocco et al. (2005) demonstrate using a realistically calibrated life-cycle model how the welfare loss associated with stock market non-participation can be significant.

The Frick and Chernew (2009) framework can be easily adapted to our current setting, where we can substitute the inefficient underutilisation of healthcare services with the original stock market under-participation puzzle. Similarly, we can substitute the utilisation frictions in the healthcare setting with the full range of economic and non-economic frictions discussed in Section 4.1 and 4.2 which help to explain dis-inclination towards stock market participation.

In adapting the instructive conceptual framework of insurance-induced consumption to illustrate the essential underlying dynamics hypothesised to be at play in our context, it is not necessary to make any definitive judgement as to the welfare-implications, i.e. whether the additional consumption represents an efficiency improvement or dis-improvement.

## 4.4 Methodology

In the absence of an exogenous source of variation or means of randomly assigning the option of financial support across households to allow for a direct causal evaluation of the impact of coverage on household financial behaviour, we utilise a second-best estimation strategy in the form of propensity score matching (PSM). Recognising that financial support coverage is not randomly assigned within the population, and that households benefiting from the option of emergency financial support are likely to differ systematically in terms of wealth, income, education, and other important factors,

---

<sup>24</sup>The authors’ research context is health insurance, but the framework is generalizable to other settings.



the estimation strategy here can be viewed as a bias-reduction strategy which endeavours insofar as possible to establish a comparable control group against which we can estimate the role of emergency financial coverage.

In an effort to simulate the ideal identification conditions of randomised assignment into treatment and control, PSM matches treated and untreated observations on the basis of the proximity of their estimated propensity scores. This allows an aggregated treatment effect to be computed by comparing the outcomes of treated observations paired with counterfactuals that are just as likely to be treated based on their observable characteristics. The method overcomes the ‘curse of dimensionality’ whereby we would struggle to find exact matches on each observable characteristic that affects treatment probability by collapsing these into a single dimension: the propensity score. This method therefore facilitates more precise and reliable matched comparisons which, when aggregated, provide an estimate of the treatment effect of interest.

The identifying assumptions that underpin the validity of the PSM estimate are the conditional independence and common support assumptions. The conditional independence assumption requires that conditional on the vector of observables used to calculate the propensity scores, allocation into treated or control groups is as good as random. With propensity score matching, the researcher can make no claim as to the role of unobservable factors in influencing treatment probability. This is an important difference with the randomised controlled trial, where randomisation guarantees balance on unobservable as well as observable factors. The common support condition stipulates that conditional on the vector of observables used in the matching exercise, the probability of each observation falling within the treatment group lies strictly within the unit interval, ensuring adequate overlap in the features of the two groups to allow for the formation of quality matches (Baum, 2013).<sup>25</sup>

To build confidence in our results, we run a series of robustness exercises which variously adjust our specification for computing propensity scores, and the matching algorithm used to link treated and untreated observations. We additionally seek to narrow the potential channels of confounding variation through detailed exploration and incorporation of credible proxy measures (see Section 4.6.2).

While PSM provides a valuable tool to substantially shrink the selection problem insofar as can be achieved with the data by collapsing observable differences between treated and untreated observations, it is important to recognise and freely acknowledge that this is a limited and imperfect statistical strategy. We cannot exclude the

---

<sup>25</sup>While all PSM estimators follow this basic logic of aggregated pairwise comparisons, there are many ways in which the comparison observations may be found and weighted. We make use of three such matching algorithms: nearest-neighbour matching, radius matching, and kernel matching.

possibility that unobserved and uncontrolled background characteristics such as family wealth and networks may simultaneously influence likelihood of having recourse to emergency financial support and the likelihood of holding shares or a household's attitude to financial risk. These effects may continue to be present in our estimated treatment effects from emergency financial coverage. As such, this paper is not presented as a direct causal identification of the impact of emergency financial support, but rather an investigation of how an important manifestation of background financial advantage relates to patterns of financial risk after taking account of a broad set of relevant factors affecting comparability.

## 4.5 Data and model

To pursue this research question, we make use of the third wave of the HFCS. The HFCS is a cross-national survey covering 21 countries in Europe, collected in 2017 from participating 84,829 households. The data contains information regarding socio-demographic variables, assets, liabilities, income and consumption for a sample of households that is representative both at the national and the euro area level. A set of population weights is provided in order to ensure the representativeness of the sample (see weighting discussion in Section 4.9.1).<sup>26,27</sup>

Our key treatment variable of interest is having access to emergency financial support, measured as those respondents who respond yes to the question: *In an emergency, could (you/your household) get financial assistance of say €5,000 from friends or relatives who do not live with you?* Our key outcome variable of interest is whether a household does or does not hold stocks or mutual funds, measured as those respondents who answer yes to the question: *Do you/does anyone in your household own stock shares in any publicly traded companies?* Or to the question: *Do you/Does anyone in your*

---

<sup>26</sup>As well as offering the advantage of broad scope and interest, it is important to recognise that cross-country survey data poses certain disadvantages. For instance, cross-country cultural differences affecting honesty in responding to surveys, etiquette in utilising financial support, institutional differences in bankruptcy procedures, and ease of accessing the stock market are likely to vary across countries. Such effects can be largely captured by the inclusion of fixed effect controls for level differences, but it is not possible to take account of slope (marginal) effects across countries in factors such as these within our empirical framework.

<sup>27</sup>HFCS data are provided as imputed in relation to household assets, liabilities, and income. For each missing value, five imputed values are provided. The propensity score matching procedure utilised here does not facilitate integrated multiple imputation estimation, and as such, estimation is performed on individual implicates. Treatment effects are identical across all five implicates. For expositional purposes and to avoid graphical and tabular duplication, results are presented based on the first implicate.

*household have any investments in mutual funds, money market mutual funds or hedge funds?*<sup>28,29</sup>

Our secondary outcome variable of interest relates to subjectively reported financial risk appetite, and is based upon the survey question: *Which of the following statements comes closest to describing the amount of financial risk that you (and your husband/wife/partner) are willing to take when you save or make investments? 1 - Take substantial financial risks expecting to earn substantial returns; 2 - Take above average financial risks expecting to earn above average returns; 3 - Take average financial risks expecting to earn average returns, or 4 - Not willing to take any financial risk.* We construct a binary outcome variable which serves as our dependent variable in the financial risk appetite model, and is coded as 1 where the respondent reports above average, or substantial financial risk tolerance, and 0 otherwise.

In estimating propensity scores at the household level for the likelihood of enjoying the option of emergency financial support, our main specification adjusts for the full set of available likely confounders within the dataset, including income, education, circumstances of employment, net wealth, age, marital status, gender, number of children, tenure status, and a full set of country controls. Variable definitions are provided in Table 4.13.

In matching observations on the basis of the proximity of their propensity score, our preferred specification applies nearest-neighbour matching while imposing a strict maximum tolerable caliper distance of 0.01 standard deviations.

We describe in greater detail modelling choices relating to the computation of propensity scores and the choice over alternative matching algorithms in Appendix 4.9.2, and the robustness of our results to alternative specifications in Section 4.6.2.

### 4.5.1 Descriptive composition

It is important to take stock of some simple descriptive facts about households in our dataset in terms of key dependent and independent variables, in order to orient our analysis. On average, 57% of households in the survey data report having access to emergency financial support. Figure 4.1 provides a breakdown across countries

---

<sup>28</sup>Importantly, our treatment status is subjectively reported, and as such may involve some element of subjective bias. For instance, respondents could be overly optimistic (or pessimistic) regarding the likely availability of emergency financial support from friends or relatives. It is additionally conceivable that overconfidence in this regard may be correlated with financial risk appetite. We cannot control for such reporting bias with the available data, though the possibility of its presence is acknowledged.

<sup>29</sup>In 94% of cases, the survey respondent corresponds to the reference person in the household, where the reference person is considered to be the person who is most knowledgeable about the financial situation of the household and provides the financial information for the whole household.

included in the survey. Belgium, Germany, Ireland, Luxembourg, the Netherlands, and Portugal all report support rates above 60%, while Estonia, Croatia, Lithuania, and Latvia have the lowest coverage rates, all below 40%. 20% of households captured in the survey report holding stocks or mutual funds, but there is substantial variation evident across countries. Figure 4.2 illustrates that the highest rate of participation is Finland, with over half of households holding stocks or mutual funds, while Greece, Lithuania, Poland, and Slovakia report the lowest shares of participation, all falling below 5%.

Turning to financial risk appetite, we find that 6% of household respondents subjectively report having a willingness to take higher than average financial risks. Figure 4.3 shows how this metric varies by country, with Hungary, Italy, Austria, and Lithuania reporting the highest share of ‘risk-loving’ households, all reporting rates above 10%, while Slovenian households report the lowest tolerance, with just 1.4% of household respondents reporting a willingness to take higher than average financial risks.

In Figure 4.4 we plot the distribution of the value of investment holdings among stock market participants. Among those with positive holdings, 48% hold less than €10,000, 60% hold less than €20,000, and 75% hold less than €50,000.

We can observe in simple descriptive terms that stock market participation and risk appetite within our sample is broadly increasing in income and wealth, and that at every level, those households benefiting from the option of emergency financial support are more likely to participate in the stock market than their counterparts. Stock market participation is increasing in age up to a point, before declining for those of retirement age, and again the participation rates is substantially greater for households benefiting from the option of emergency financial support at all ages. By contrast, we can see that younger respondents are most likely to report higher financial risk appetite, and that this declines linearly with age. Again, households benefiting from the option of emergency financial support are more likely to report a higher risk appetite at all ages (see Figure 4.6).

In Figure 4.5, we report how the probability of benefiting from the option of emergency financial support varies by wealth, income, and age. We can see that the likelihood is increasing linearly in wealth and income, and decreasing in age.

In Table 4.7 (Column 1), we report the descriptive correlates of access to emergency financial support from a simple regression framework. We find that household income, net wealth, self-employment, and higher education are all positively associated with access to emergency support. Relative to employees, the unemployed and retired/other

respondents are significantly less likely to have access to emergency support. Respondents with more children are more likely to benefit from access to financial support. Households holding financial assets or reporting higher financial risk appetite are more likely to benefit from access to emergency support, and any household which has a prior history of receiving a substantial gift or inheritance, or having received financial support previously, is more likely to be in a position to call on emergency support in the future.

We find that older respondents are slightly less likely to report access to the safety net. Curiously, we find that those who own all or part of their home (relative to those renting or having free use of their property) are slightly less likely to report access to the safety net.

Column 2 reports the descriptive correlates of stock market participation, showing that income, net wealth, higher education, and higher financial risk appetite are all positively associated with participation. Those reporting access to emergency financial support, previously having received a substantial gift or inheritance, or having previously received financial assistance are more likely to report participation. The self-employed (relative to employees), older, female, and married respondents are less likely to report participation. Those who own all or part of their home (relative to those renting or having free use of their property) are slightly less likely to report participation.

In Column 3, we report the descriptive correlates of higher financial risk appetite. We find that those reporting higher financial risk appetite are more likely to have higher net wealth, to be self-employed, have higher education, report access to emergency financial support, hold financial assets, and have received prior financial assistance. This cohort is also less likely to report higher incomes, be unemployed or retired, be advanced in age, be female, have a greater number of children, to own all or part of their home, or to have received a substantial gift or inheritance in the past.

## 4.6 Empirical results

### 4.6.1 Assessing model performance

Several diagnostic criteria have been put forward to assess the success or otherwise of the matching procedure.

Firstly, in order to establish the validity of the matching procedure, it is essential that our model achieves balance across observable covariates between the treated and control

groups. Most importantly, we can refer to the standardised differences, which is found by dividing the difference in means of the covariate between the treated and untreated group by the standard deviation of the treated group. For a successfully balanced procedure, the standardised mean difference between the treated and untreated in the matched sample should be as close as possible to 0.<sup>30</sup> In our case, we observe standardised differences far below commonly accepted thresholds in the literature, at less than 1% in the majority of cases, and never exceeding 2%, indicating a high degree of covariate balance. Tables 4.8, 4.9, 4.10 and 4.11 report covariate balance statistics before and after the matching procedure, and show a significant collapsing of the difference in each covariate between treated and control groups following the matching procedure. We can graphically observe the reduction in covariate imbalance from the raw data to the matched sample in the visual representation found in Figures 4.11 and 4.14.

We can also refer to some additional assessment criteria set out by Rubin (2001). Specifically, these are Rubin's Bias (the absolute standardized difference of the means of the linear index of the propensity score in the treated and (matched) non-treated group) and Rubin's Ratio (the ratio of treated to (matched) non-treated variances of the propensity score index). It is recommended that Rubin's Bias be less than 25 and that Rubin's Ratio fall between 0.5 and 2 for the samples to be considered sufficiently balanced. For the participation model, we observe a Rubin's Bias value of 2.6 and Ratio value of 1.02, while in the financial risk appetite model, a Bias value of 4 and an Ratio value of 1.05 (see Table 4.12).

Sianesi (2004) suggests, as a means to establish the success of the matching procedure, to re-estimate the propensity score on the matched sample. That is, only on participants and matched non-participants, and then compare the pseudo- $R^2$  before and after matching. The pseudo- $R^2$  indicates how well the regressors explain the participation probability. After matching there should be no systematic differences in the distribution of covariates between both groups and therefore, the pseudo- $R^2$  should be fairly low. Furthermore, the joint significance test of all regressors should not be rejected before, and should be rejected after matching. In Table 4.12 we report the success of the preferred specification on these joint-significance measures: we find a pseudo- $R^2$  after matching of 0 in both models, and that the probability associated with our chi-squared

---

<sup>30</sup>No firm consensus has been established as to the acceptable threshold below which we can consider a covariate balanced, but commonly proposed rules of thumb for adjudicating adequacy in covariate balance range from 0.25 (25%) (see for instance Harder et al. (2010); Cochran (1968); Rubin (2001)) to 0.1 (10%) (Stuart et al., 2013). Caliendo and Kopeinig (2008) note that most empirical studies deem a bias reduction to below 3% or 5% as sufficient. However, Imai et al. (2008) make the case for why there can be no threshold below which the level of imbalance is always acceptable, arguing instead that the difference between the groups should be minimized without limit.

statistic is 0 before matching in both models, and 1 (0.983) in our participation (financial risk attitude) model. This provides strong supportive evidence for the success of our matching procedure.

Another important diagnostic indicator is the variance ratio of the treated and untreated groups, which provides a means of assessing the satisfaction of the balancing condition. The variance ratio refers to the mean ratio of the variance of a variable in treated subjects to the variance of the same variable in untreated subjects. In a successfully balanced sample, we additionally expect to observe that the variance ratio approaches 1, indicating no difference in the second central moment across treated and untreated groups in the matched sample.<sup>31</sup> In Figures 4.11 and 4.14 we see that the variance ratio is almost exactly 1 in the majority of cases, and always comfortably below 2, providing reassurance as to the comparability of treated and untreated units in the matched sample. This diagnostic result provides further evidence for the satisfaction of the balancing condition and the success of the matching procedure.

A further diagnostic criterion to satisfy for the validity of the matching procedure is to establish the common support. This can be done by visual inspection of the densities of propensity scores of treated and non-treated groups. It is recommended that if there are sizable differences between the maxima and minima of the density distributions, we should remove cases that lie outside the support of the other distribution, bearing in mind that results are strictly valid only for the region of common support. Figures 4.12, 4.13, 4.15, and 4.16 illustrate how the common support condition is comfortably satisfied, and as well as sufficient overlap, shows a high degree of balance in the distribution of the propensity scores across treated and comparison groups.

Based on the cumulative weight of this battery of diagnostic criteria, we are confident of the success of the matching procedure in the preferred specification in terms of the balancing and common support conditions.

However, as noted previously, while the balancing condition for observable factors is satisfied, and the common support condition holds, it is not possible to exclude the possibility that unobserved background factors which simultaneously influence the likelihood of emergency coverage and attitude towards financial risk continue to differ between treated and matched control observations, and may feature in our estimated treatment effects. As such, the estimated treatment effects should not be regarded as direct causal estimates of the impact of emergency financial support. Rather, estimated effects should be interpreted as estimates of how an important manifestation of background financial advantage relates to patterns of financial risk after taking account of a broad set of relevant factors affecting comparability.

---

<sup>31</sup>A variance ratio of 1 indicates perfect balance, while a variance ratio below 2 is generally acceptable (Zhang et al., 2019).

## 4.6.2 Estimated effects

In the participation model, the preferred specification yields a positive and significant average treatment effect on the treated observations. Those households which do benefit from the option of emergency financial coverage are 6.1% more likely to participate in the stock market. In the financial risk appetite model, we find a lesser effect, albeit still positive and significant. Households with the option of emergency financial support are 2.2% more likely to report a higher than average financial risk appetite (see Tables 4.1 and 4.2).

In these tables, we report results from additional robustness checks which assess the sensitivity of the primary estimates from the preferred specification to tweaks to our specification which variously include additional controls. The observed treatment effects are found to be stable across these alternative specifications (see also Figures 4.7 and 4.9 for graphical exposition).

We additionally investigate how the treatment effect from our main specification varies across subgroups of interest, namely, by the degree of credit constraint faced by the household, the age of the head of household, and the gender of the head of household (see also Figures 4.8 and 4.10). Within the participation model, we find a high degree of stability, with sparse evidence of statistically significant divergence across subgroups. Notably, we do not observe that our estimated treatment effects are stronger for households who report being credit constrained, or according to the gender of the head of household. We observe tentative evidence that the mechanism is not as strong among older households than it is among younger households.<sup>32</sup>

In Tables 4.3 and 4.4 we report results from further robustness exercises which test whether results are robust to the choice of matching algorithm. In the preferred specification, the matching is conducted on the basis of 1 nearest neighbour, with a maximum tolerated caliper distance imposed at 0.01 SD. We consider 8 alternative matching algorithms, where the number of neighbours and caliper distance is adjusted. We additionally test a kernel matching algorithm with tight and loose caliper distances imposed. The results in both the participation model and financial risk attitude model show a high degree of consistency across these alternative matching algorithms.

### Discussion on unobservable characteristics

As noted previously, there is a series of plausible unobserved characteristics that have the potential to confound our estimated impact from the emergency financial support

---

<sup>32</sup>See Table 4.6 for an overview of the sample composition in each of these subgroups.



option. In this section, we aim to analyse and mitigate this problem systematically and transparently. This is with a view to narrowing down the potential channels of confounding influence, to enable a more clear-eyed assessment and interpretation of our main empirical effects.

In Table 4.5, we decompose five potential channels of confounding influence in two separate categories: firstly, confounding variation arising from household characteristics, and secondly, confounding variation arising from what can be broadly termed ‘background network effects’. For each factor, we report evidence for an empirical relationship documented previously within the literature, we consider the extent to which we can address the problem by adjusting the model using proxy measures, and we report the net result of that adjustment. In addition to their correlation with our outcome variables of interest, it is by virtue of their potential simultaneous correlation with our primary treatment variable of interest (having the option of emergency financial support) that these factors qualify as potential confounders of concern.<sup>33</sup>

The first factor of interest as a potential confounder is the field of education of the respondent (e.g. does the individual have a business, financial, or economic educational background?). Prior evidence has shown that this can have a significant effect on the likelihood of stock market participation, particularly where this is in a business or financial field. This variable is unfortunately not present in our dataset. However, we can observe and adjust for the level of education, and area of employment, which is typically correlated (though clearly imperfectly) with the field of education (see for instance Robst (2007)). Level of education is already included in the main specification, and Tables 4.1 and 4.2 show that the main result is unchanged with the addition of field of employment. It is worth noting, however, that while being potentially correlated with household financial behaviours, there is little theoretical reason to expect that field of education should be correlated with the option of emergency financial support from friends or family, reducing concern regarding the potential confounding influence from this channel.

The area of employment variable, however, is measured at less than ideal granularity, at the 21 category classification of NACE Rev. 2, meaning it is not possible to clearly identify the precise field of employment, but being reliant on broader classifications such as ‘Professional, scientific and technical activities’, and ‘Financial and insurance

---

<sup>33</sup>Of less concern as a source of confounding variation in our average treatment effect, but nonetheless important to acknowledge is a set of additional private insurance mechanisms available to the household which could plausibly moderate the distribution of conditional average treatment effects. These include the option to flexibly increase the number of hours worked, or to turn first to the household’s own available buffer savings. While it is not possible to capture the labour hours flexibility channel using our data, our results are robust to the addition of a control in the matching algorithm capturing the total value of household savings.

activities'. To take account of the additional likelihood of stock market participation arising from business or financial sector employment, we include a broad sector of employment dummy recording whether the reference person is employed in a financial/business/professional scientific role, or not. Our main results are robust to this inclusion (see Tables 4.1 and 4.2).

Financial literacy is another potentially relevant unobserved factor in our setting, with evidence pointing to the increased likelihood of stock market participation among individuals with greater levels of financial literacy (Van Rooij et al., 2011). While financial literacy is not directly measured within our dataset, a reasonable proxy in the form of the level of educational attainment is available, a variable known to be highly correlated with financial literacy (see for instance Lusardi and Mitchell (2014)). As noted above, this proxy measure is already included within our main specification.

Moving to a consideration of background network effects, we identify two primary channels through which confounding influence may be brought to bear in our estimation strategy: peer and parental participation effects. Brown et al. (2008) document a role for within-community peer effects in predicting individual stock market participation, with the level of stock market participation within one's community positively influencing the likelihood of an individual's decision to participate in the stock market. This factor is not observed within our dataset, where the best available proxy metric is level of educational attainment. Evidence has shown that educational attainment is both positively correlated with stock market participation (see for instance Athreya et al. (2017)), and with more extensive social networks (see for instance Andersson (2018)). As noted previously, educational attainment is already included as a control within our main specification. The adequacy of this proxy measure is, however, clearly limited, and this limitation is recognised. It follows that peer participation effects must continue to be regarded as a potential confounding influence captured within our main estimated treatment effect.

A similar situation pertains in relation to the influence of parental stock market participation. Li (2014), and Zhao and Min (2021) both show that individuals whose parents participated in the stock market are significantly more likely to participate themselves. Parental financial habits are not recorded within our dataset, but several practical proxies do present themselves. We observe whether a household has benefited from prior receipt of inheritance, and we additionally observe whether this bequest took the form of stocks or securities, as distinct from other asset categories such as housing. We also observe the total contemporary wealth of the respondent household, known to be serially correlated across generations (Pfeffer and Killewald, 2018), and correlated with stock ownership (Andersen and Nielsen, 2011). Household wealth is already captured

within our main specification, and our results are additionally robust to the addition of both inheritance indicators (see Tables 4.1 and 4.2). These indicators represent the best available proxies for the parental channel of influence over contemporary financial behaviours, albeit imperfectly. As a result, it is not possible to exclude the possibility that confounding variation in the form of parental influence from prior participation is partially captured within our main estimated treatment effect.

Returning to the broad headings of potential confounding variation, it is possible to conclude that background network effects represent a more salient channel of confounding influence over our estimates than household-specific characteristics, where we have some success in identifying credible proxy measures and finding robust model estimates. Background network effects are found to be more difficult to adjust for in our setting using available data. As such, it should be transparently recognised that our main estimated treatment effect must be regarded not as a directly identified causal estimate of the impact of the emergency financial support option, but rather as a composite which includes some potential influence attributable to network-based exposure to financial markets.

## 4.7 Discussion

Our theory proposes that the option of emergency financial support serves as a type of informal insurance mechanism which can attenuate the downside loss inherent in risky financial propositions. The results presented here are consistent with this theory. We find that those households that enjoy the option are approximately 6% more likely to participate in the stock market, and approximately 2% more likely to report higher financial risk appetite.

A valid question arises as to the real financial meaning of the option of financial support of €5,000 when set against the sums that participating households may stake in the stock market, and by extension, as to the credibility of our proposed mechanism in insulating the household from adverse outcomes. Here, it is important to bear in mind that the median value of investment holdings among participating households is just €10,000. 48% hold less than €10,000, 60% hold less than €20,000, and 75% hold less than €50,000. Against this backdrop, the mechanism of emergency financial support should be viewed as a valuable option with a credible claim to weigh upon household financial decision making.

We investigate how the intensity of our mechanism varies across subgroups of interest, such as credit constrained and unconstrained households, and according to the gender

and age of the head of household, finding a high degree of consistency. We observe tentative evidence that the mechanism is not as strong among older households. This may reflect younger households facing greater immediate financial pressures than their older counterparts, while also holding a higher latent underlying risk appetite, such that a relatively more intensive channel of insurance-induced consumption is observed (i.e. that the shadow price on financial insurance is higher for younger households, with a small relaxation of an existing constraint (or an extension of insurance coverage over financial losses) bringing forward a greater shift along the demand curve for stock holdings).<sup>34</sup>

For policy implications arising from our findings, we can point towards two areas of consideration. Firstly, it is notable that our results capture a mechanism of advantage compounding advantage: illustrating a way in which access to capital growth in financial markets is easier for already more advantaged households. In doing so, we highlight the importance of designing and promoting effective policies of financial inclusion. Strengthening targeted financial literacy and education initiatives should be considered as a means of arming less advantaged households with the tools they need to consider, evaluate, and navigate the opportunities, risks, as well as the opportunity costs associated with different paths in the household financial landscape.

Such policies are not without cost, and range in intensity from light-touch public information campaigns to intensive educational interventions. Calls for the integration of dedicated financial education to school curricula must bear in mind the scarcity of schooling hours as a resource, the opportunity cost of classroom time dedicated to that purpose, and the actual efficacy of any learning interventions. However, absent efforts to foster greater financial inclusion and participation, the social costs associated with widespread non-participation outlined in Section 4.1 are likely to accumulate.

Secondly, the results highlight the possibility that existing market mechanisms, such as downside protection or portfolio insurance, could be applied to plug the gap implied by the absence of a private safety net. Portfolio insurance provides downside protection against losses while still preserving access to upside potential, and gained popularity in the 1980s in a variety of commonly offered retail products (Dichtl and Drobetz, 2011; Clarke and Arnott, 1987). Such tools alter the probability distribution of returns, by introducing a kinked return profile which eliminates outcomes below a specified floor, in exchange for a reduced range of upside potential. The sacrifice of upside potential can be viewed as the premium to be paid for the insurance provided, in proportion to the range of adverse outcomes for which protection is desired. The utility of portfolio insurance has already been mapped within the context of portfolio

---

<sup>34</sup>i.e. the slope of the demand curve for younger households is greater than that for older households.

theory, where its adoption and popularity is consistent with loss aversion among retail investors (Kahneman and Riepe, 1998; Dichtl and Drobetz, 2011).

In our research context, such tools can be viewed as offering the potential to replicate by way of a formal market mechanism the insurance role provided by private safety nets observed in this study, and in so doing, to expand stock market participation. If well-designed, such products have the potential to provide a beneficial function to households. While this represents a commercial opportunity, to be entered into by private parties, it carries also an implication for the public policymaker. Specifically, it highlights the importance of vigilant consumer protection regimes, as well as intelligent and agile enforcement in the domain of retail investment and any associated add-on products to protect against misleading or spurious products.<sup>35</sup> These products should be sold in accordance with high standards of transparency and clarity, as well as consumer suitability and appropriateness assessment.

It is also important to emphasise what our results do not imply for policy. Crucially, it does not follow that governments should step in to insure stock market participation as a means of adjusting for the relative disadvantage of households who do not enjoy the option of emergency private financial support, to expand access to the equity premium. Nor does it follow that policy should be deployed in an effort to push non-participating households into the stock market.<sup>36</sup> However, in recognition of the implication of widespread non-participation for widening wealth inequality, under-diversification of household portfolios, as well as political economy, public policymakers should take an interest in understanding the issue, and in fostering the conditions for effective financial inclusion through education and literacy initiatives as well as consumer protection regimes which enable households to make fully informed decisions within the financial landscape for themselves.

## 4.8 Conclusion

A wide range of factors have already been documented as helping to explain the stock market participation puzzle. These include the fact of monetary and non-monetary barriers to entry, various demographic factors including age, gender, education, income, wealth, and trust, financial literacy, and social connectedness. In this paper, we contribute a new dimension to the participation puzzle: the catalytic role of financial support from friends and family, namely, the ability to draw upon financial support to

---

<sup>35</sup>In most jurisdictions, this will form part of the normal discharge of existing public financial conduct and consumer protection responsibilities.

<sup>36</sup>As a rule, it remains important that financial market exposure is undertaken in a way and to a degree that is best suited to an individual household's means and circumstances.

the tune of €5,000 in an emergency. The impact of financial support mechanisms upon financial and economic behaviour and outcomes has been documented in more general contexts, including in the field of inheritance, ongoing financial support from parents, and wider safety nets. However, little prior research has focused on how the option of emergency financial support relates to patterns of household financial risk exposure.


With a European household survey spanning 21 countries and over 84,000 households, and making use of propensity score matching methods, we find that those households that benefit from the option of emergency financial support are 6% more likely to participate in the stock market, and 2% more likely to report higher than average financial risk tolerance. Under a propensity score matching strategy it is not possible to exclude the possibility of unobserved background factors continuing to differ systematically between treated and matched control observations, and as such, our estimated effects cannot be given direct causal interpretation. However, the results do contribute to our understanding about advantage, portfolio choice, and access to finance, and help to illuminate an important aspect of the psychology of money, namely, how an important manifestation of background financial advantage relates to patterns of household financial risk exposure.

Our results point towards the importance of designing and promoting effective and sustainable policies of financial inclusion through financial literacy and education initiatives which enable households to evaluate the opportunities, risks, and opportunity costs associated with alternative paths in their financial lives.

Our results additionally highlight the potential utility of portfolio insurance policies in replicating by way of a formal market mechanism the insurance role provided by private safety nets observed in this study. Properly designed, such downside protection tools can present a beneficial function for households. In view of the possibility that such market mechanisms could be harnessed to induce participation among otherwise circumspect households, policymakers should ensure that consumer protection in the domain of retail investment and associated add-on products is agile and robustly enforced.


Potential policy implications outlined here should be viewed on a spectrum ranging from the uncontroversial to the more contentious. Vigilant and agile enforcement of consumer protection regimes as a means of making financial markets safe for inclusive participation should be seen on the lighter end of the distribution, representing in most jurisdictions a reminder of the importance of status quo institutional protections. More costly financial education initiatives can be viewed on the other end of the spectrum, and should be carefully weighed in terms of costs, benefits, and targeting. Judgements on the appropriateness of policy interventions along this spectrum is likely in part to reflect preferences in political economy.

Figure 4.1: Households benefiting from the option of emergency financial support, percentage by country



Notes: Data for Euro area countries are drawn from the HFCS. Data for the United States is drawn from the 2019 wave of the *Survey of Consumer Finances* (SCF). In the SCF, respondents are asked whether in an emergency they could call upon financial support to the tune of \$3,000, whereas in the HFCS the level of support is €5,000.


Figure 4.2: Ownership of shares or mutual funds by country



Notes: Data for Euro area countries are drawn from the HFCS. Data for the United States is drawn from the 2019 wave of the SCF.

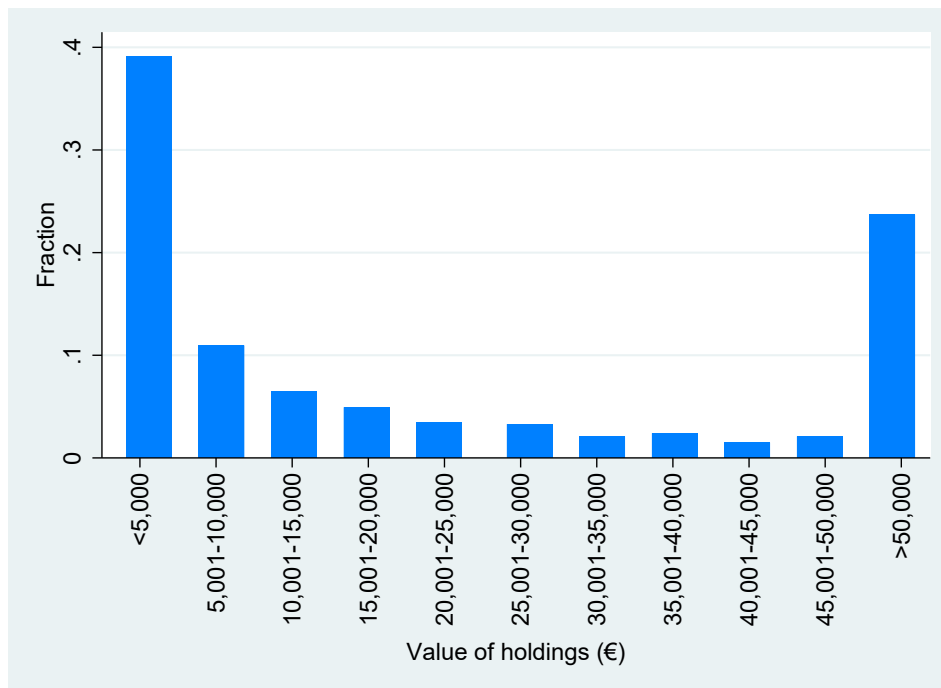


Figure 4.3: High financial risk appetite by country




Notes: Data for Euro area countries are drawn from the HFCS. Data for the United States is drawn from the 2019 wave of the SCF. Financial risk appetite data are not collected on a consistent basis in the HFCS and SCF. To align with the HFCS 1-5 coding scale as best as is possible, we classify as high risk appetite those US survey respondents who indicate 9 or 10 within the SCF 1-10 financial risk appetite scale.

Figure 4.4: Value of investment holdings among stock market participants




Notes: Figure reports the distribution of the value of investment holdings among stock market participants in the HFCS.

Figure 4.5: Probability of emergency financial support option (by wealth, income, age)



Notes: Figure reports how the probability of access to emergency financial support option varies across three dimensions: wealth, income, and age.

Figure 4.6: Participation and risk appetite across households with and without the option of emergency financial support (wealth, income, age)



Notes: Figure reports how the rate of stock market participation (upper panel), and high financial risk appetite (lower panel) varies across three dimensions: wealth, income, and age. Plots are additionally separated by access to a lender of last resort.

Table 4.1: Main effects and robustness to alternative specifications (participation model)

	ATT	SE
Baseline model	.0605021 ***	.0091315
Excluding survey weights	.0609986***	.0092151
Additional control: financial risk attitude	.0480834***	.0085523
Additional control: total savings	.0617334***	.0113785
Additional control: prior inheritance receipt	.0607175***	.0108482
Additional controls: prior inheritance receipt, inheritance receipt taking the form of shares/securities.	.0604533***	.0108657
Additional controls: Employment dummy for financial/business/professional scientific role.	.05959***	.0091618
Baseline model, by household credit constraint status:		
Credit constrained	.0656774***	.0115656
Credit unconstrained	.0613853***	.009639
Baseline model, by age:		
Age 18-35	.0717885***	.0131485
Age 36-50	.0691021***	.0106583
Age 51-64	.0681511***	.0119504
Age 65+	.0421835***	.0077826
Baseline model, by gender of head of household:		
Head of household: Female	.0502139***	.0079505
Head of household: Male	.0660253**	.0102084

Notes: Table reports main estimated effects of the impact of access to emergency financial support on the probability of stock market participation. Table additionally describes the robustness of the estimated effect from the baseline model to a range of alternative specifications, along with estimated subgroup effects.

Table 4.2: Main effects and robustness to alternative specifications (financial risk appetite model)

	ATT	SE
Baseline model	.0210302***	.0067586
Excluding survey weights	.0210203***	.0067643
Additional control: participation	.0180546**	.0064138
Additional control: total savings	.0184635***	.0048897
Additional control: prior inheritance receipt	.0237113***	.006443
Additional controls: prior inheritance receipt, inheritance receipt taking the form of shares/securities.	.0236162***	.0064534
Additional controls: Employment dummy for financial/business/professional scientific role.	.02075***	.0064686
Baseline model, by credit constraint status		
Credit constrained	.0238802***	.0085925
Credit unconstrained	.0213929***	.0068257
Baseline model, by age		
Age 18-35	.0269758**	.0128552
Age 36-50	.030388***	.008188
Age 51-64	.0132398*	.0071064
Age 65+	.0032829	.0049707
Baseline model, by gender of head of household:		
Head of household: Female	.0161336***	.0055556
Head of household: Male	.0232548***	.007729

Notes: Table reports main estimated effects of the impact of access to emergency financial support on the probability of higher financial risk appetite. Table additionally describes the robustness of the estimated effect from the baseline model to a range of alternative specifications, along with estimated subgroup effects.

## Appendix

Table 4.3: Robustness to alternative matching algorithms (participation model)

	<b>ATT</b>	<b>SE</b>
Baseline model (1 neighbour, caliper 0.01)	.0605021 ***	.0091315
5 neighbours, caliper 0.01	.0605261 ***	.0099198
1 neighbour, caliper 0.05	.0614217 ***	.0093596
5 neighbours, caliper 0.05	.0606008 ***	.0099403
1 neighbour, caliper 0.001	.0601194 ***	.0094617
5 neighbours, caliper 0.001	.0611223 ***	.0099926
Caliper 0.01, epanechnikov kernel	.0605021 ***	.0095365
Caliper 0.05, epanechnikov kernel	.0614217 ***	.0093596
Caliper 0.1, epanechnikov kernel	.0639545 ***	.0092292


Notes: Table demonstrates the robustness of the main estimated effect from the baseline participation model to alternative propensity score matching algorithms.

Table 4.4: Robustness to alternative matching algorithms (financial risk appetite model)

	<b>ATT</b>	<b>SE</b>
Baseline model (1 neighbour, caliper 0.01)	.0210302 ***	.0067586
5 neighbours, caliper 0.01	.0228668 ***	.007178
1 neighbour, caliper 0.05	.0211839 ***	.006649
5 neighbours, caliper 0.05	.0228366 ***	.0071912
1 neighbour, caliper 0.001	.021596 ***	.006855
5 neighbours, caliper 0.001	.0231527 ***	.0071694
Caliper 0.01, epanechnikov kernel	.0210302 ***	.0067586
Caliper 0.05, epanechnikov kernel	.0211839 ***	.006649
Caliper 0.1, epanechnikov kernel	.021678 ***	.0065403

Notes: Table demonstrates the robustness of the main estimated effect from the baseline financial risk appetite model to alternative propensity score matching algorithms.


Figure 4.7: Estimated effects on stock market participation: robust to alternative specifications



Notes: Figure displays coefficient estimates from the first six specifications contained in Table 4.1.




Figure 4.8: Estimated effects on stock market participation: subgroup analysis




Notes: Figure displays coefficient estimates from the latter eight specifications contained in Table 4.1.

Figure 4.9: Estimated effects on financial risk appetite: robust to alternative specifications



Notes: Figure displays coefficient estimates from the first six specifications contained in Table 4.2.

Figure 4.10: Estimated effects on financial risk appetite: subgroup analysis



Notes: Figure displays coefficient estimates from the latter eight specifications contained in Table 4.2.

Table 4.6: Total sample breakdown - heterogeneous treatment effect categories

Credit constraint status	Sample (%)	Age	Sample (%)	Gender	Sample (%)
Constrained	0.06	18-35	0.20	Female	0.39
Unconstrained	0.94	36-50	0.32	Male	0.61
		51-64	0.24		
		65+	0.24		

Notes: Table decomposes five potential channels of confounding influence arising from household characteristics, background network effects. For each factor, the table reports evidence for the relevant endogeneity documented previously within the literature, and addresses the extent to which the problem can be mitigated by adjusting the model using proxy measures.

Table 4.5: Potential confounders with proxy solutions

<b>Factor</b>	<b>Evidence for a relationship</b>	<b>Best available proxy</b>	<b>Result</b>
<i>Household/respondent characteristics</i>			
Field of education – financial market orientation	Evidence shows that business school education can increase the likelihood of stock market participation (Dong et al., 2022).	Level of education, area of employment (correlated with field of education).	Result robust to these factors.
Field of employment: financial market orientation	Girshina et al. (2019) show how the degree of exposure to financial sector employment increases the probability of stock market participation.	Cannot observe granular fields of employment, but can control for broad sector of employment (dummy for financial/business/professional scientific).	Result robust to this addition.
Financial literacy	Van Rooij et al. (2011) show evidence for the role of financial literacy in predicting stock market participation.	Financial literacy is unobservable. Best proxy is level of education. Education is highly correlated with financial literacy.	Result robust to this.
<i>Background network effects</i>			
Peer participation	Brown et al. (2008) document a role for peer effects within communities in predicting individual stock market participation.	Unobservable. Best proxy is level of education.	Result robust to best available proxy, but peer participation remains a potential confounder.
Parental participation	Li (2014), Zhao and Min (2021) document a role for parental participation in predicting participation of offspring, pointing to the importance of information sharing and parental risk-taking affecting children’s risk taking.	Unobservable. Best proxies are (a) prior receipt of inheritance where the source is parents, (b) prior inheritance of stocks/securities, and (c) contemporary household wealth.	Result robust to these best available proxies (included individually and simultaneously), but parental participation remains a potential confounder.

Notes: Table decomposes five potential channels of confounding influence arising from household characteristics, background network effects. For each factor, the table reports evidence for the relevant endogeneity documented previously within the literature, and addresses the extent to which the problem can be mitigated by adjusting the model using proxy measures.

Table 4.7: Descriptive correlates of key dependent and independent variables.

VARIABLES	(1) Safety net	(2) Holds stocks/ mutual funds	(3) Higher financial risk
Gross income (log)	0.0197*** (0.00237)	0.0895*** (0.00201)	-0.0102*** (0.00110)
Net wealth (log)	0.0493*** (0.00147)	0.0629*** (0.00135)	0.00508*** (0.000750)
Employment status: Employee	0.0128** (0.00580)	-0.0311*** (0.00462)	0.00758*** (0.00268)
Employment status: Self-employed	0.0442*** (0.00751)	-0.0594*** (0.00530)	0.0295*** (0.00379)
Employment status: Unemployed	-0.106*** (0.0113)	-0.0181* (0.0105)	-0.00588 (0.00479)
Age	-0.00343*** (0.000179)	-0.00143*** (0.000148)	-0.00135*** (8.90e-05)
Female	0.000149 (0.00401)	-0.0236*** (0.00311)	-0.0202*** (0.00199)
Married	-0.00419 (0.00411)	-0.0117*** (0.00318)	-0.000645 (0.00201)
Higher education	0.0919*** (0.00411)	0.0508*** (0.00303)	0.0232*** (0.00197)
No. children	0.00653** (0.00268)	0.000976 (0.00201)	-0.00272** (0.00115)
Tenure status	-0.0585*** (0.00574)	-0.0372*** (0.00452)	-0.00864*** (0.00270)
Holds shares	0.00107 (0.00608)		0.0393*** (0.00247)
Holds bonds	0.0840*** (0.0109)		0.0331*** (0.00398)
Holds mutual funds	0.0235*** (0.00577)		0.0213*** (0.00238)
Higher financial risk	0.0897*** (0.00847)	0.0933*** (0.00566)	
Gift/inheritance received	0.0463*** (0.00411)	0.0684*** (0.00291)	-0.00619*** (0.00198)
Financial assistance received	0.131*** (0.00987)	0.0467*** (0.00805)	0.0113*** (0.00423)
Safety net		0.0191*** (0.00308)	0.0195*** (0.00200)
Observations	65,578	65,578	65,578

Notes: Table reports descriptive correlates of access to emergency financial support, stock market participation, and higher financial risk appetite.

Table 4.8: Covariate balance pre and post matching procedure: standardised differences (participation model)

Means	Raw			Matched (ATE)		
	Treated	Untreated	Std. Dif.	Treated	Untreated	Std. Dif.
Household weight	1750.239	1795.966	-0.01	1760.255	1773.375	0.00
Gross HH income	33964.19	23253.57	0.24	30012.19	29134.34	0.02
Net assets	450528.8	260618.8	0.10	386289.3	365335.3	0.01
Self-employed	0.1204086	0.072952	0.16	0.100653	0.1001772	0.00
Unemployed	0.0254735	0.053838	-0.15	0.038196	0.0367231	0.01
Employee	0.5177272	0.439366	0.16	0.481379	0.4773643	0.01
Age	52.38557	55.37771	-0.19	53.82614	54.04632	-0.01
Female	0.3766333	0.406099	-0.06	0.389392	0.3904838	0.00
Married	0.589402	0.530843	0.12	0.565914	0.5656711	0.00
Higher education	0.4212386	0.242734	0.39	0.344761	0.3430896	0.00
No. children	0.387508	0.296909	0.12	0.344274	0.3352196	0.01
Tenure status	0.7754629	0.696505	0.18	0.742311	0.7448453	-0.01
AT	0.0384337	0.035481	0.02	0.036966	0.0366544	0.00
BE	0.0313471	0.021636	0.06	0.026647	0.0257035	0.01
CY	0.0148117	0.017012	-0.02	0.015565	0.0156405	0.00
DE	0.0662694	0.051232	0.06	0.060526	0.0625077	-0.01
EE	0.0185784	0.049438	-0.17	0.031258	0.0306723	0.00
FI	0.1215365	0.113282	0.03	0.117478	0.1187328	0.00
FR	0.1525857	0.169726	-0.05	0.157721	0.1580335	0.00
GR	0.0321345	0.041955	-0.05	0.037181	0.0366971	0.00
HR	0.0090019	0.02612	-0.13	0.017642	0.0161352	0.01
HU	0.0718238	0.065217	0.03	0.068418	0.0681206	0.00
IE	0.0664397	0.044421	0.10	0.05636	0.0556686	0.00
IT	0.0870185	0.093355	-0.02	0.089222	0.0906897	-0.01
LT	0.0104916	0.029624	-0.14	0.020355	0.0183818	0.01
LU	0.0248138	0.012612	0.09	0.019502	0.0184629	0.01
LV	0.0079166	0.024271	-0.13	0.015111	0.0144712	0.01
MT	0.0125771	0.011575	0.01	0.012064	0.0122858	0.00
NL	0.0306235	0.02556	0.03	0.028118	0.0278583	0.00
PL	0.0695893	0.062078	0.03	0.067476	0.0700079	-0.01
PT	0.0939774	0.04061	0.21	0.071682	0.0726634	0.00
SI	0.0197276	0.030464	-0.07	0.02408	0.0240806	0.00

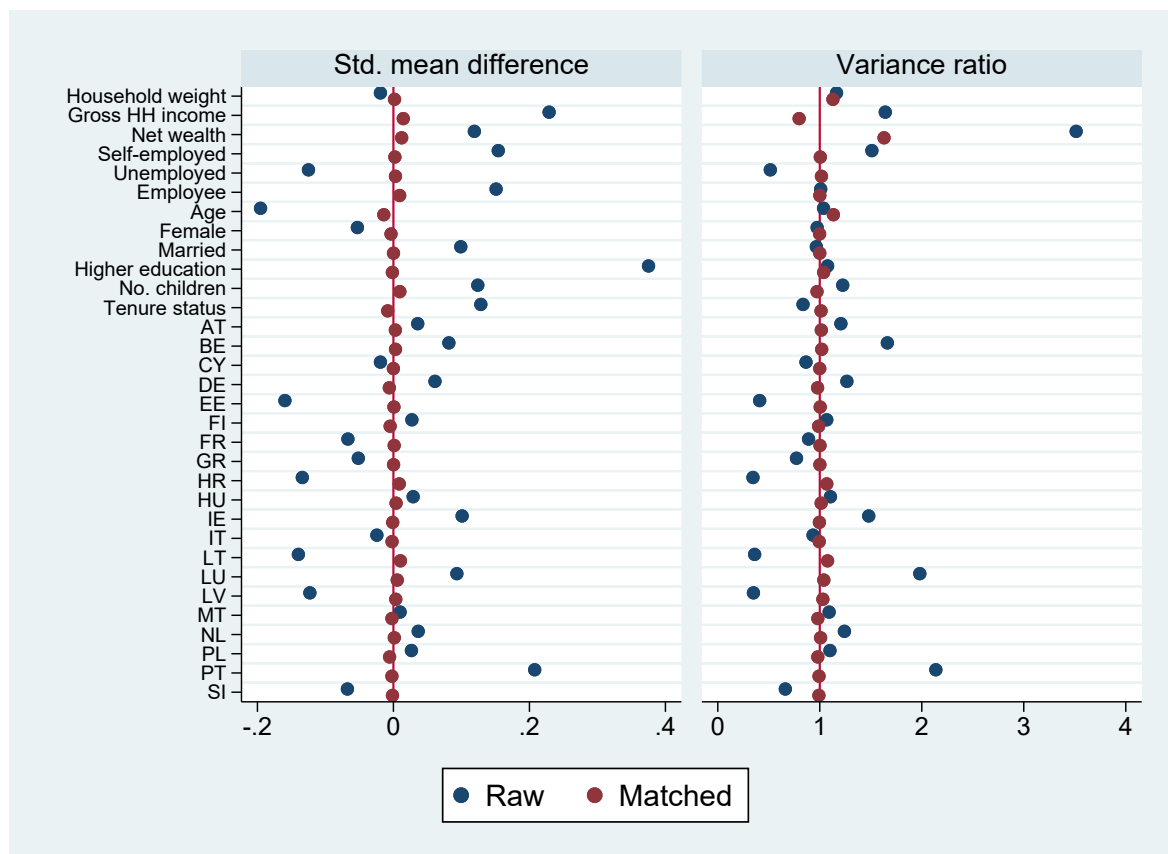
Notes: Table describes the degree of covariate balance in the participation model across treated and untreated observations in the raw data, and then following the propensity score matching procedure, focusing on the standardised differences in mean values.

Table 4.9: Covariate balance pre and post matching procedure: variance ratios (participation model)

Means	Raw			Matched (ATE)		
	Treated	Untreated	Ratio	Treated	Untreated	Ratio
Household weight	1.30E+07	1.14E+07	1.14	1.28E+07	1.17E+07	1.09
Gross HH income	2.59E+09	1.55E+09	1.67	2.09E+09	2.53E+09	0.82
Net assets	4.86E+12	1.80E+12	2.70	3.95E+12	3.18E+12	1.24
Self-employed	0.1059126	0.067632	1.57	0.090524	0.0901442	1.00
Unemployed	0.0248251	0.050941	0.49	0.036737	0.0353755	1.04
Employee	0.2496911	0.24633	1.01	0.249659	0.2494946	1.00
Age	264.593	254.7099	1.04	277.6942	245.0467	1.13
Female	0.2347857	0.241189	0.97	0.237771	0.2380129	1.00
Married	0.2420124	0.249056	0.97	0.245661	0.2456942	1.00
Education	0.2438018	0.18382	1.33	0.225906	0.2253855	1.00
No. children	0.6348053	0.524959	1.21	0.572923	0.5819903	0.98
Tenure status	0.1741239	0.211392	0.82	0.191289	0.1900561	1.01
AT	0.0369573	0.034223	1.08	0.0356	0.0353118	1.01
BE	0.0303651	0.021169	1.43	0.025938	0.0250435	1.04
CY	0.0145926	0.016723	0.87	0.015323	0.0153963	1.00
DE	0.0618791	0.048608	1.27	0.056864	0.0586021	0.97
EE	0.0182337	0.046995	0.39	0.030281	0.0297323	1.02
FI	0.1067676	0.100452	1.06	0.103679	0.1046382	0.99
FR	0.129306	0.140923	0.92	0.132848	0.1330627	1.00
GR	0.0311025	0.040196	0.77	0.035799	0.0353514	1.01
HR	0.0089211	0.025439	0.35	0.017331	0.0158753	1.09
HU	0.0666666	0.060965	1.09	0.063738	0.063482	1.00
IE	0.0620268	0.042449	1.46	0.053185	0.0525711	1.01
IT	0.079448	0.084642	0.94	0.081264	0.0824674	0.99
LT	0.0103817	0.028747	0.36	0.019941	0.0180444	1.11
LU	0.0241986	0.012453	1.94	0.019122	0.0181225	1.06
LV	0.0078541	0.023682	0.33	0.014883	0.0142622	1.04
MT	0.0124192	0.011441	1.09	0.011919	0.0121352	0.98
NL	0.0296864	0.024907	1.19	0.027328	0.027083	1.01
PL	0.064748	0.058226	1.11	0.062925	0.0651086	0.97
PT	0.0851475	0.038962	2.19	0.066545	0.0673853	0.99
SI	0.0193388	0.029537	0.65	0.0235	0.0235014	1.00

Notes: Table describes the degree of covariate balance in the participation model across treated and untreated observations in the raw data, and then following the propensity score matching procedure, focusing on the variance ratio.


Figure 4.11: Covariate balance: graphical representation (participation model)



Notes: Figure illustrates the graphical complement of Tables 4.8 and 4.9. It shows the standardised mean difference and the variance ratio across covariates before and after the matching procedure is completed. In both cases the differences collapse towards 0 and 1 respectively, pointing towards the success of the propensity score matching procedure.




Figure 4.12: Density balancing plot (participation model)



Notes: Figure displays the satisfaction of an important visual diagnostic criterion: common support. Figure shows the densities of propensity scores of treated and non-treated groups in the participation model from the raw data, and in the matched sample. As well as demonstrating sufficient overlap, the figure shows a high degree of balance in the distribution of the propensity scores across treated and comparison groups.

Figure 4.13: Cumulative distribution balancing plot (participation model)



Notes: Figure displays the satisfaction of an important visual diagnostic criterion: common support. Figure shows the densities of propensity scores of treated and non-treated groups in the participation model from the raw data, and in the matched sample. As well as demonstrating sufficient overlap, the figure shows a high degree of balance in the distribution of the propensity scores across treated and comparison groups.

Table 4.10: Covariate balance pre and post matching procedure: standardised differences (financial risk appetite model)

Means	Raw			Matched (ATE)		
	Treated	Untreated	Std. Dif.	Treated	Untreated	Std. Dif.
Household weight	1753.598	1800.495	-0.01	1764.904	1777.092	0.00
Gross HH income	34012.86	23374.56	0.23	30114.41	29246.51	0.02
Net assets	451483.3	261869.2	0.10	387907.1	366013	0.01
Self-employed	0.1203503	0.073269	0.16	0.10075	0.1002114	0.00
Unemployed	0.0252693	0.053132	-0.14	0.037591	0.0362347	0.01
Employee	0.5178705	0.438791	0.16	0.481416	0.4775776	0.01
Age	52.39945	55.42687	-0.19	53.8519	54.06579	-0.01
Female	0.3762983	0.405114	-0.06	0.388792	0.3899086	0.00
Married	0.5895454	0.531799	0.12	0.566413	0.5661944	0.00
Education	0.4214189	0.242932	0.39	0.345265	0.3436904	0.00
No. children	0.3875624	0.296376	0.12	0.344298	0.335166	0.01
Tenure status	0.7754674	0.697025	0.18	0.742404	0.7452592	-0.01
AT	0.0386749	0.036009	0.01	0.037334	0.0370122	0.00
BE	0.0313725	0.021275	0.06	0.026578	0.0256127	0.01
CY	0.0149046	0.017265	-0.02	0.015716	0.0157978	0.00
DE	0.0666424	0.051994	0.06	0.061079	0.0630066	-0.01
EE	0.0186736	0.050117	-0.17	0.031515	0.0309283	0.00
FI	0.1215924	0.113431	0.03	0.117584	0.1188591	0.00
FR	0.1530291	0.170146	-0.05	0.158175	0.1584273	0.00
GR	0.0323361	0.042579	-0.05	0.037531	0.0370545	0.00
HR	0.0090584	0.026509	-0.13	0.017917	0.0162849	0.01
HU	0.0709039	0.063883	0.03	0.067326	0.0670326	0.00
IE	0.0665353	0.044656	0.10	0.056574	0.0558277	0.00
IT	0.0875645	0.094744	-0.02	0.090142	0.0915811	0.00
LT	0.0091012	0.024774	-0.12	0.017015	0.0155177	0.01
LU	0.0249695	0.012799	0.09	0.019693	0.0186568	0.01
LV	0.0079663	0.024603	-0.13	0.0154	0.0145617	0.01
MT	0.0126561	0.011747	0.01	0.012183	0.0124001	0.00
NL	0.0301732	0.024689	0.03	0.027576	0.0273262	0.00
PL	0.0692122	0.061921	0.03	0.06722	0.0697285	-0.01
PT	0.094353	0.0411	0.21	0.072206	0.0732706	0.00
SI	0.0198514	0.030918	-0.07	0.024307	0.0243227	0.00

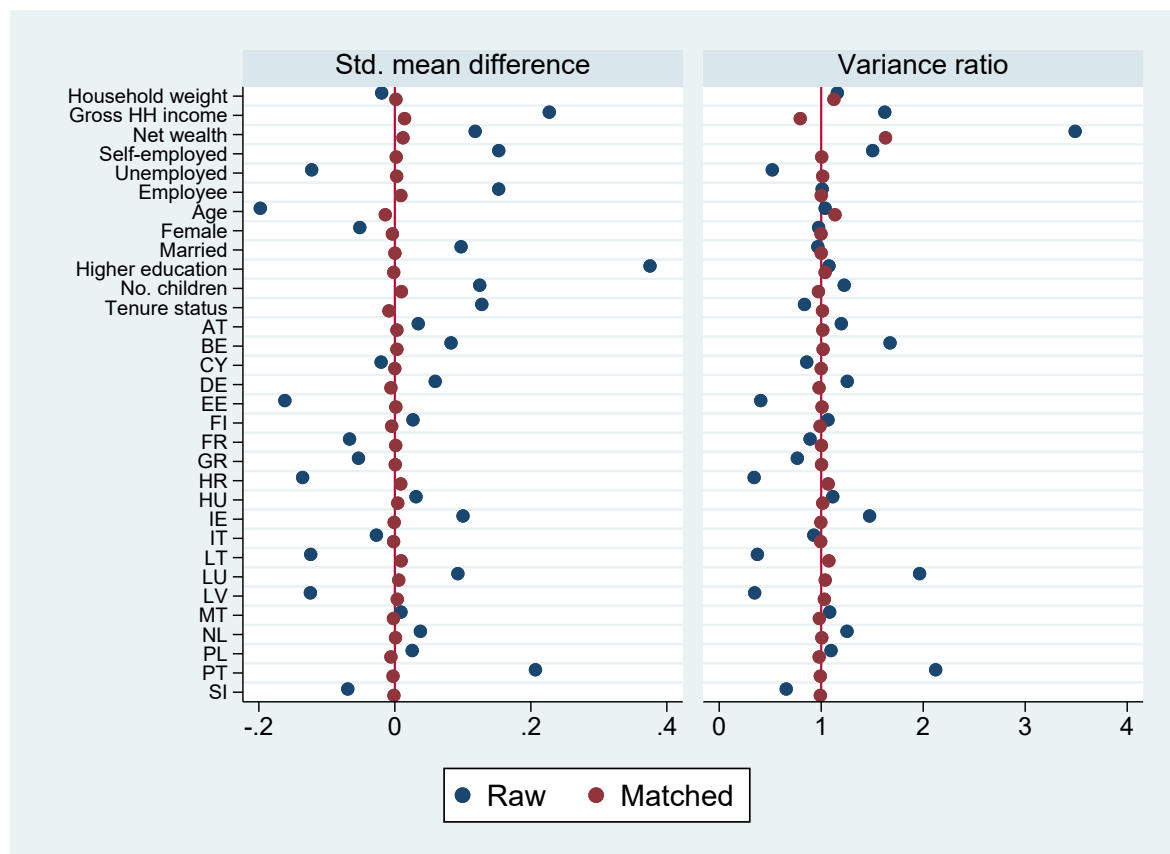
Notes: Table describes the degree of covariate balance in the financial risk appetite model across treated and untreated observations in the raw data, and then following the propensity score matching procedure, focusing on the standardised differences in mean values.

Table 4.11: Covariate balance pre and post matching procedure: variance ratios (financial risk appetite model)

Means	Raw			Matched (ATE)		
	Treated	Untreated	Ratio	Treated	Untreated	Ratio
Household weight	1.31E+07	1.15E+07	1.13	1.29E+07	1.18E+07	1.09
Gross HH income	2.58E+09	1.56E+09	1.65	2.09E+09	2.54E+09	0.82
Net assets	4.89E+12	1.81E+12	2.71	4.01E+12	3.14E+12	1.28
Self-employed	0.1058684	0.067903	1.56	0.090601	0.0901716	1.00
Unemployed	0.0246313	0.05031	0.49	0.036178	0.0349228	1.04
Employee	0.249686	0.24626	1.01	0.24966	0.2495043	1.00
Age	264.4487	254.2744	1.04	277.6989	244.7304	1.13
Female	0.2347029	0.241004	0.97	0.237638	0.2378866	1.00
Married	0.2419868	0.248996	0.97	0.245595	0.2456253	1.00
Education	0.2438302	0.183921	1.33	0.226062	0.2255737	1.00
No. children	0.63517	0.524953	1.21	0.573481	0.5827836	0.98
Tenure status	0.1741215	0.211187	0.82	0.191245	0.1898533	1.01
AT	0.0371799	0.034713	1.07	0.035941	0.0356433	1.01
BE	0.0303889	0.020823	1.46	0.025872	0.0249574	1.04
CY	0.0146828	0.016967	0.87	0.015469	0.0155487	0.99
DE	0.0622025	0.049292	1.26	0.05735	0.0590384	0.97
EE	0.0183253	0.047606	0.38	0.030522	0.0299726	1.02
FI	0.10681	0.100567	1.06	0.10376	0.1047346	0.99
FR	0.129614	0.141201	0.92	0.133158	0.1333319	1.00
GR	0.0312912	0.040767	0.77	0.036123	0.0356825	1.01
HR	0.0089765	0.025807	0.35	0.017596	0.0160202	1.10
HU	0.065878	0.059804	1.10	0.062795	0.062541	1.00
IE	0.0621097	0.042663	1.46	0.053374	0.0527125	1.01
IT	0.0798987	0.08577	0.93	0.082018	0.0831964	0.99
LT	0.0090186	0.024161	0.37	0.016726	0.0152773	1.09
LU	0.0243465	0.012636	1.93	0.019305	0.0183092	1.05
LV	0.007903	0.023999	0.33	0.015163	0.0143501	1.06
MT	0.0124962	0.011609	1.08	0.012035	0.0122467	0.98
NL	0.0292634	0.02408	1.22	0.026816	0.0265802	1.01
PL	0.0644232	0.058088	1.11	0.062703	0.0648683	0.97
PT	0.0854523	0.039412	2.17	0.066993	0.0679039	0.99
SI	0.0194577	0.029963	0.65	0.023717	0.0237318	1.00


Notes: Table describes the degree of covariate balance in the financial risk appetite model across treated and untreated observations in the raw data, and then following the propensity score matching procedure, focusing on the variance ratio.

Figure 4.14: Covariate balance: graphical representation (financial risk appetite model)




Notes: Figure illustrates the graphical complement of Table 4.10 and 4.11. It shows the standardised mean difference and the variance ratio across covariates before and after the matching procedure is completed. In both cases the differences collapse towards 0 and 1 respectively, pointing towards the success of the propensity score matching procedure.

Figure 4.15: Density balancing plot (financial risk appetite model)



Notes: Figure displays the satisfaction of an important visual diagnostic criterion: common support. Figure shows the densities of propensity scores of treated and non-treated groups in the financial risk appetite model from the raw data, and in the matched sample. As well as demonstrating sufficient overlap, the figure shows a high degree of balance in the distribution of the propensity scores across treated and comparison groups.

Figure 4.16: Cumulative distribution balancing plot (financial risk appetite model)



Notes: Figure displays the satisfaction of an important visual diagnostic criterion: common support. Figure shows the densities of propensity scores of treated and non-treated groups in the financial risk appetite model from the raw data, and in the matched sample. As well as demonstrating sufficient overlap, the figure shows a high degree of balance in the distribution of the propensity scores across treated and comparison groups.

Table 4.12: Additional diagnostic assessment of matching procedure

	<b>Participation model</b>	<b>Financial risk attitude model</b>
<b>Pseudo R-squared</b>		
Before matching	0.097	0.096
After matching	0.000	0.000
<b>P &gt; chi-squared</b>		
Before matching	0.000	0.000
After matching	1.000	0.983
<b>Rubin's bias (B)</b>		
Before matching	75.9	75.6
After matching	2.6	4.0
<b>Rubin's ratio (R)</b>		
Before matching	0.86	0.86
After matching	1.02	1.05
<b>Median bias</b>		
Before matching	12.2	12.2
After matching	0.3	0.5
<b>Mean bias</b>		
Before matching	9.6	9.6
After matching	0.2	0.5

Notes: Table reports additional diagnostic criteria set out by Rubin (2001). Table shows Rubin's Bias (the absolute standardized difference of the means of the linear index of the propensity score in the treated and (matched) non-treated group) and Rubin's Ratio (the ratio of treated to (matched) non-treated variances of the propensity score index).



Table 4.13: Definition of variables used

Variable	Definition
Household weight	A variable to compensate for unequal probability of the household being selected into the sample and differential unit non-response.
Gross income	Total gross annual household income aggregate, equalised using the modified OECD equivalence scale (number of economically active members in household).
Net financial position	Total assets (excl. public and occupational pension plans) less total outstanding liabilities.
Employment status	Categorical variable recording employment status as employee, self-employed, unemployed, or retired/other. The reference category used is retired/other.
Age	Age of the reference person.
Female	Gender of the reference person (0=male; 1=female).
Married	Binary outcome variable indicating whether the marital status of the reference person is married/in a consensual union on a legal basis, or not.
Higher education	Dummy variable recording that the highest level of education of the reference person in the household is tertiary or above.
No. children	Number of dependent children. Dependent children defined as all persons aged 0-15 and person aged 16-24 not at work and living with a parent.
Tenure status	Dummy variable recording the reference person's tenure status for the main place of residence: =1 where the reference person owns all or part of the main residence, and =0 where the reference person is renting/subletting, or has free use of the property.
Country dummies	20 country dummies included, omitted country dummy is Slovakia.
Total savings	Total value of savings accounts.

---

Variable	Definition
Higher financial risk	Dummy variable recording the financial risk appetite reported by the reference person: =0 where the respondent reports they are not willing to take any financial risk, or a willingness to take average financial risks expecting to earn average returns, and =1 for respondents who report a willingness to take above average financial risks expecting to earn above average returns, or a willingness to take substantial financial risks expecting to earn substantial returns.
Invested	Binary outcome variable indicating whether the household owns stock shares in any publicly traded companies/has any investments in mutual funds, money market mutual funds or hedge funds, or not.
Safety net	Dummy variable recording that in an emergency, the household could get financial assistance of €5,000 from friends or relatives who do not live with them.
Has shares	Dummy variable recording that the household owns publicly traded shares.
Has bonds	Dummy variable recording that the household owns bonds (corporate, government, bills, or notes).
Has mutual funds	Dummy variable recording that the household owns investments in mutual funds, money market mutual funds, or hedge funds.
Gift/inheritance received	Dummy variable recording that any member of the household has received an inheritance or substantial gift, including money or any other assets from someone who is not a part of the current household.
Prior financial assistance	Dummy variable indicating that in the past 12 months the household has received financial help (money or help with bills or other expenses) from persons outside the household such as relatives, friends, or other.

---

---

Variable	Definition
Credit constrained	Dummy variable indicating that in the last three years, the respondent (or another member of the household) considered applying for a loan or credit but decided not to, thinking that the application would be rejected.
Employment dummy	Dummy variable indicating that the head of household's field of employment is 'professional, scientific and technical activities', or 'financial and insurance activities', and zero otherwise.
Inheritance dummy	Dummy variable indicating that the household received inheritance taking the form of securities or shares, and zero otherwise.

---

## 4.9 Methodology and data - extended

### 4.9.1 Weighting

The HFCS data set has a variable on cross-sectional weights to compensate for unequal probability of the household being selected into the sample and differential unit non-response. Weights are adjusted for information on external sources using the calibration to margins method. The method consists of adjusting weights to permit extrapolating survey results and obtaining totals consistent with information known over the whole population from external sources (such as distribution by age and sex). Weights play a critical role in the interpretation of survey data. However, there is some debate on when weights should be used in regressions.

In the context of propensity score matching, Zanutto (2006) recommends that the propensity score model itself should not be survey-weighted, as the propensity score model is used only to match treated and control units with similar background characteristics together in the sample and not to make inferences about the population-level propensity score model. However, it may be necessary to include survey weights as a regressor in the final outcome model depending on the study goal (DuGoff et al., 2014).

For good measure, we follow the common strategy of running specifications with and without survey weights, and find consistent results.<sup>37</sup>

---

<sup>37</sup>ATE refers to the average treatment effect, and ATT refers to the average treatment effect on the treated. The ATE is estimated for each subject, regardless of whether they fall within the treated or untreated cohort. Averages of these effects over all the subjects in the data estimate the ATE. The ATT, by contrast, uses only the subjects observed in the treatment group. Here, I am interested in estimating the population ATT. Frequently the objective in the analysis of complex survey data

## 4.9.2 Model

### Variable choice

Several important considerations must be borne in mind in the selection of variables for inclusion in the propensity score estimation model. Our choice is governed in the first instance by the conditional independence assumption, which requires that our outcome variable of interest is independent of treatment status once we condition upon the propensity score (Caliendo and Kopeinig, 2008), as well as by the objective of satisfying the balancing property in the matched sample. As such, it is important to include variables that simultaneously influence the outcome variable and treatment variable under analysis, informed by theory and previous research in our domain. In order to satisfy the balancing property, it may also be necessary to include higher order terms of selected variables. It should be borne in mind here that the estimation of the propensity score does not itself require a behavioural interpretation (World Bank, 2018). It is recommended that variables which are themselves likely to be affected by treatment status should not be included as controls, for the avoidance of post-treatment bias (Ho et al., 2007). As such, we begin with a parsimonious model to estimate the propensity score, inspect for balancing properties and common support, and iteratively amend the model to achieve the best possible degree of balance while satisfying the considerations outlined above concerning the choice of proper covariates.

The final estimation model we arrive at includes all of the factors listed in Table 4.10, including the full set of available likely confounders within the dataset, including income, education, circumstances of employment, net wealth, age, marital status, gender, number of children, tenure status, and a full set of country controls. As part of robustness exercises, we also consider specifications which include additional control variables in the form of prior inheritance receipt, financial employment, financial risk appetite (in the participation model), and stock market participation (in the financial risk appetite model).

### Matching algorithms

A critical consideration relates to the basis upon which matching observations within the comparison group are selected and paired with observations in the treated group - i.e. the choice over matching algorithms. A range of potential algorithms are available.

Nearest-neighbour matching is the simplest method, and involves pairing a treatment observation with a comparison observation with the closest propensity score. Here, the researcher may also choose to oversample, by obtaining a comparison observation that is the weighted average of several nearest-neighbours. This brings the advantage of reduced variance through the utilisation of a greater amount of information, but

---

is to make population-level inferences (i.e. the population ATE or population ATT). As such, it is necessary to account for weights in the final outcome model if the association of interest may vary between the population and the sample. Alternatively, where the goal is to estimate effects for the survey sample itself (the sample ATT or sample ATE), the survey weights are not needed. Here, the objective is to estimate effects which pertain to the general population. We therefore need to account for survey weights to adjust for the fact that the HFCS oversamples wealthy households.

the potential disadvantage of comparison observations whose average match quality is lower in virtue of the wider sampling of neighbours.

Radius (or caliper) matching, by contrast, imposes a tolerable distance between propensity scores in the formation of matches. This allows us to overcome the problem of low quality matches where the nearest-neighbour may lie at a far remove from a treatment observation, and so to improve the average quality of matches. The tolerable distance is defined by the caliper chosen by the researcher, with narrower calipers implying stricter matching rules.

Offering a further refinement, kernel matching involves the comparison of treated observations with a weighted composite of comparison observations within a bandwidth of the propensity score, where observations with closer propensity scores to each treated observation receive a higher weighting. This method takes the benefits and attenuates the costs occurring with radius matching by harnessing the information of many observations but weighting in favour of more proximate observations to preserve match quality (Garrido et al., 2014).

In our preferred specification, we make use of nearest-neighbour matching while imposing a strict maximum tolerable caliper distance of 0.01 standard deviations. However, for completeness and to demonstrate the robustness of findings to the choice of matching algorithm, we additionally report results for a range of alternative matching algorithms which variously adjust the number of neighbouring observations included in the matching exercise, and the tolerated caliper range for matched propensity scores (see Tables 4.3 and 4.4).

# Chapter 5

## Conclusion

This thesis makes important contributions to the body of understanding in relation to consumer and household financial decision making, why we fail to take full advantage of the opportunities presented by the financial system, and slip into costly pitfalls that punctuate our financial lives. Navigating the financial landscape is not a walk in the park. As well as the inherent complexity of financial instruments, the overwhelming array of choices and vast accompanying volumes of consumer disclosures with which we are confronted, the intrinsic uncertainty of our economic, financial, and personal lives, and the unhelpful presence of malicious scam artists, we are also subject to an array of behavioural and cognitive biases, social and environmental influences, and face innumerable additional competing demands on our limited attention every day. These realities can frustrate our efforts to make the right choices for our circumstances and realise the full opportunities of the contemporary financial system offers. They leave a landscape littered with apparent puzzles and paradoxes, predictable errors and missed opportunities, pockets consumer detriment and unrealised gains.

This thesis sets out to tackle three of these. Three challenges in behavioural consumer finance where the promise of the contemporary financial system is not being realised: the failure of large numbers of mortgage holders to take advantage the open goal which is the opportunity for material reductions in household debt burdens through refinancing; the frustration of the development of strong and trusted digital financial ties for small businesses in less developed countries by pervasive digital fraud, and the widespread opting-out of European households from stock market participation who forego the opportunity for capital growth, reinforcing patterns of wealth inequality. The thesis seeks not only to contribute insight and understanding in these areas, but also to put forward practical policy solutions that actually help to remediate some of the harms incurred. In this we have some success.

Chapter 2 shows how simple and unobtrusive changes to the way in which consumers are notified about refinancing opportunities, intelligently designed and experimentally evaluated, can interrupt inertia and inattention, and unlock one of the costliest and most pervasive financial mistakes. This is a problem which has been stubbornly entrenched, and impervious to persistent regulatory attention and information campaigns. What's more, helping to unblock refinancing inertia offers a useful policy lever to improve interest rate pass-through to the real economy, stimulate consumption, improve competitiveness, and alleviate household debt service burdens. This chapter reports results from the first mortgage refinancing field experiment undertaken outside of the United States, the first such exercise targeted at a general population of mortgage holders, and the first to report statistically and economically meaningful impacts, with the strongest treatment delivering a 76% increase in refinancing activity. This chapter should be read in the context of a growing body of evidence that demonstrates the value of behaviourally informed approaches in delivering effective consumer protection in essential product markets.

Chapter 3 moves to tackle the pervasive problem of non-institutional fraud which so perniciously clogs up the effective operation of digital financial services in low and middle income countries. As a counterpoint to the success reported in Chapter 2, this chapter provides a somewhat cautionary tale for policymakers embarking on corrective interventions in consumer financial or digital markets. Here, none of the experimentally tested interventions succeeded in meaningfully improving participants' ability to detect and discern genuine from fraudulent communications presented, notwithstanding positive impacts on certain adjacent outcomes. More troubling are the resulting false confidence effects engendered, with participants feeling more confident in their judgements, notwithstanding the absence of any corresponding improvement in actual underlying ability. The chapter offers a valuable contribution in highlighting the subtle dynamics of light-touch learning interventions, and the interplay between trust, confidence, and ability in respect of deception detection in a digital financial context. It further highlights the severity of the challenge faced by small business owners in successfully navigating the menace of fraud, and casts doubt on the utility of light-touch, quick-fix learning interventions as meaningful antidotes. No one said it would be easy.

In Chapter 4, attention turns to another longstanding puzzle in consumer and household finance: the opting out of large swathes of the population from the opportunity for capital growth offered by stock-market participation. The importance of this question stems not only from the pure financial opportunity cost implied for non-participant households, but also from the consequent reinforcement of long term patterns of wealth inequality, under-diversification of household portfolios, and political economy effects arising from the partition of financial and economic interests within populations. The

chapter introduces a previously unexplored dimension to the manifold existing theories and explanations advanced to account for non-participation: the unseen background influence of private emergency financial support mechanisms. Consistent with a theory of insurance-induced consumption, this chapter shows that households benefiting from such a safety net or ‘lender of last resort’ are more likely to participate and to report higher financial risk appetites. The chapter highlights a mechanism of advantage compounding advantage, and points towards the potential utility of downside protection policies in replicating by way of a market mechanism this social and family network effect, and facilitating wider access and participation.

The fortunes of the separate studies in light-touch intervention documented in Chapter 2 (where we see a large return on a modest intervention) and Chapter 3 (where our interventions do not gain useful traction) can loosely trace the outline of the fundamental question, do nudges work? This is a question which has been the subject of significant attention, and has prompted considerable soul-searching among practitioners in recent years. Initial exuberance regarding the transformative potential for such interventions, driven by early successes (Thaler and Sunstein, 2009) and integration at the heart of public policymaking (Halpern, 2015) has somewhat given way to a crisis of confidence in the discipline. Competing meta-analyses examining the overall effectiveness of aggregated nudge interventions have produced estimates which range from substantial to null average effect sizes (Mertens et al., 2022; Maier et al., 2022). What is clear is that the record is complicated. Many interventions will have no effect, and publication bias may inflate the perceived utility of the toolkit. Chapters 2 and 3 microcosmically reflect the hit and miss reality of the art: we can achieve enormously disproportionate results from small targeted interventions, but we are involved in a process of scientific trial and error, and not everything we nudge turns to gold. That the record is nuanced in this way should not be a cause for discouragement. On the contrary, it should reinvigorate our interest in thinking imaginatively about solutions, our commitment to the experimental method for pre-testing interventions, and our determination to understand the subtle and complex dynamics of consumer decision making. Chapter 4 offers an instructive example in respect of this latter challenge.

Above all, what this thesis endeavours to demonstrate is that if we actually harness the enormous stock of knowledge that already exists about human behaviour, think creatively about solutions, take seriously the cognitive, behavioural, and environmental context in which financial decisions are made, place humans at the centre of design, and commit to rigorous pre-testing, we have the necessary tools at our disposal to have a material impact. We can and should adjust our policy instruments to the realities of the people they are there to serve, make markets function more effectively in the



consumer interest, and help to unlock more and more of the unrealised potential that consumer financial markets offer to transform lives.

# Bibliography

- Adams, P. and Hunt, S. (2013). Encouraging consumers to claim redress: evidence from a field trial. *FCA Occasional Paper*, (2).
- Adams, P., Hunt, S., Palmer, C., and Zaliauskas, R. (2021). Testing the effectiveness of consumer financial disclosure: Experimental evidence from savings accounts. *Journal of Financial Economics*, 141(1):122–147.
- Adams, P., Hunt, S., Smart, L., and Zaliauskas, R. (2015). Stimulating interest: Reminding savers to act when rates decrease. *FCA Occasional Paper*, (7).
- Agarwal, S., Driscoll, J. C., and Laibson, D. I. (2013). Optimal mortgage refinancing: a closed-form solution. *Journal of Money, Credit and Banking*, 45(4):591–622.
- Agarwal, S., Rosen, R. J., and Yao, V. (2016). Why do borrowers make mortgage refinancing mistakes? *Management Science*, 62(12):3494–3509.
- Aghion, P. and Bolton, P. (1987). Contracts as a barrier to entry. *The American Economic Review*, pages 388–401.
- Akerlof, G. A. (1970). The market for lemons: quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.
- Almenberg, J. and Dreber, A. (2015). Gender, stock market participation and financial literacy. *Economics Letters*, 137:140–142.
- Altavilla, C., Canova, F., and Ciccarelli, M. (2020). Mending the broken link: Heterogeneous bank lending rates and monetary policy pass-through. *Journal of Monetary Economics*, 110:81–98.
- Amedah, S. A., Fougère, M., et al. (2017). Impact of inheritance and housing wealth on retirement decision. *Open Journal of Social Sciences*, 5(07):1.
- Amromin, G., Bhutta, N., and Keys, B. J. (2020). Refinancing, monetary policy, and the credit cycle. *Annual Review of Financial Economics*, 12(1):67–93.
- Andersen, S., Campbell, J. Y., Nielsen, K. M., and Ramadorai, T. (2020). Sources of inaction in household finance: Evidence from the danish mortgage market. *American Economic Review*, 110(10):3184–3230.
- Andersen, S. and Nielsen, K. M. (2011). Participation constraints in the stock market: Evidence from unexpected inheritance due to sudden death. *The Review of Financial Studies*, 24(5):1667–1697.
- Anderson, G., Bunn, P., Pugh, A., and Uluc, A. (2014). The potential impact of higher interest rates on the household sector: evidence from the 2014 nmg consulting survey. *Bank of England Quarterly Bulletin*, page Q4.

- Anderson, K. (2003). Off the hook: Reducing participation in telemarketing fraud. *Washington, DC: AARP Foundation*.
- Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association*, 103(484):1481–1495.
- Andersson, M. A. (2018). Higher education, bigger networks? differences by family socioeconomic background and network measures. *Socius*, 4:2378023118797217.
- Andries, N. and Billon, S. (2016). Retail bank interest rate pass-through in the euro area: An empirical survey. *Economic Systems*, 40(1):170–194.
- Argyle, B., Nadauld, T., and Palmer, C. (2022). Real Effects of Search Frictions in Consumer Credit Markets. *The Review of Financial Studies*. hhac088.
- Arrondel, L., Debbich, M., and Savignac, F. (2015). Stockholding in france: the role of financial literacy and information. *Applied Economics Letters*, 22(16):1315–1319.
- Ascari, G. and Haber, T. (2021). Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy. *The Economic Journal*, 132(641):37–57.
- Asgharian, H., Liu, L., and Lundtofte, F. (2019). Institutional quality, trust and stock-market participation: Learning to forget. *Trust and Stock-Market Participation: Learning to Forget*.
- Athreya, K. B., Ionescu, F., Neelakantan, U., et al. (2017). College or the stock market, or college and the stock market? Technical report, Board of Governors of the Federal Reserve System (US).
- Australian Competition and Consumer Commission (2018). Residential mortgage price inquiry.
- Ba, B. (2018). Going the extra mile: The cost of complaint filing, accountability, and law enforcement outcomes in chicago. Technical report, Working paper.
- Bajo, E. and Barbi, M. (2018). Financial illiteracy and mortgage refinancing decisions. *Journal of Banking & Finance*, 94:279–296.
- Bandura, A. (1977). Self-efficacy: toward a unifying theory of behavioral change. *Psychological review*, 84(2):191.
- Bandura, A., Freeman, W. H., and Lightsey, R. (1999). Self-efficacy: The exercise of control.
- Banerjee, A., Breza, E., Duflo, E., and Kinnan, C. (2019). Can microfinance unlock a poverty trap for some entrepreneurs? Technical report, National Bureau of Economic Research.
- Banerjee, A. V., Duflo, E., Glennerster, R., and Kothari, D. (2010). Improving immunisation coverage in rural india: clustered randomised controlled evaluation of immunisation campaigns with and without incentives. *British Medical Journal*, 340.
- Bar-Gill, O. (2012). *Seduction by contract: Law, economics, and psychology in consumer markets*. Oxford University Press.
- Barberis, N., Huang, M., and Thaler, R. H. (2006). Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American Economic Review*, 96(4):1069–1090.
- Basiglio, S., Rossi, M. C., van Soest, A., et al. (2019). Subjective inheritance expectations and economic outcomes. Technical report.
- Behavioural Insights Team (2014). East: Four simple ways to apply behavioural insights.

- Ben-Shahar, O. and Schneider, C. E. (2014). *More than you wanted to know*. Princeton University Press.
- Beraja, M., Fuster, A., Hurst, E., and Vavra, J. (2019). Regional heterogeneity and the refinancing channel of monetary policy. *The Quarterly Journal of Economics*, 134(1):109–183.
- Berger, D., Milbradt, K., Tourre, F., and Vavra, J. (2021). Mortgage prepayment and path-dependent effects of monetary policy. *American Economic Review*, 111(9):2829–78.
- Berger, D., Milbradt, K., Tourre, F., and Vavra, J. (2022). Refinancing frictions, mortgage pricing and redistribution. Unpublished Working Paper.
- Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L. F., and Palmer, C. (2019). Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice. NBER Working Paper No. 26164.
- Bertaut, C. C. (1998). Stockholding behavior of us households: Evidence from the 1983–1989 survey of consumer finances. *Review of Economics and Statistics*, 80(2):263–275.
- Binder, C. (2017). Fed speak on main street: Central bank communication and household expectations. *Journal of Macroeconomics*, 52:238–251.
- Biros, D. P., George, J. F., and Zmud, R. W. (2002). Inducing sensitivity to deception in order to improve decision making performance: A field study. *MIS quarterly*, pages 119–144.
- Blanchflower, D. G. and Oswald, A. J. (1998). What makes an entrepreneur? *Journal of Labor Economics*, 16(1):26–60.
- Blinder, A. S., Ehrmann, M., De Haan, J., and Jansen, D.-J. (2022). Central bank communication with the general public: Promise or false hope?
- Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., and Jansen, D.-J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature*, 46(4):910–45.
- Brown, J. R., Coile, C. C., and Weisbenner, S. J. (2010). The effect of inheritance receipt on retirement. *The Review of Economics and Statistics*, 92(2):425–434.
- Brown, J. R., Ivković, Z., Smith, P. A., and Weisbenner, S. (2008). Neighbors matter: Causal community effects and stock market participation. *The Journal of Finance*, 63(3):1509–1531.
- Burgoon, J. K. (2015). Deception detection accuracy. *The International Encyclopedia of Interpersonal Communication*, pages 1–6.
- Burgoon, J. K., Buller, D. B., Ebesu, A. S., and Rockwell, P. (1994). Interpersonal deception: V. accuracy in deception detection. *Communications Monographs*, 61(4):303–325.
- Burke, J., Kieffer, C. N., Mottola, G. R., and Pérez-Arce, F. (2020). Can educational interventions reduce susceptibility to financial fraud? Available at SSRN 3747165.
- Byrne, S., Devine, K., and McCarthy, Y. (2020). Room to improve: A review of switching activity in the irish mortgage market. *Economic Letters*, 12/EL, 20.
- Byrne, S., Putman, D., King, M., and Jang, C. (2022). Strategies for reducing non-institutional fraud and building trust in a digital market platform: A behavioral lab experiment in nigeria. *AEA RCT Registry*.

- Caliendo, M. and Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22(1):31–72.
- Calvet, L. E., Campbell, J. Y., and Sodini, P. (2007). Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy*, 115(5):707–747.
- Calza, A., Monacelli, T., and Stracca, L. (2013). Housing finance and monetary policy. *Journal of the European Economic Association*, 11(s1):101–122.
- Campbell, J. Y. (2006). Household finance. *The Journal of Finance*, 61(4):1553–1604.
- Caribou Digital (2019). Returns to investment in higher education.
- Central Bank of Ireland (2017a). Addendum to the consumer protection code 2012. Technical report.
- Central Bank of Ireland (2017b). Mortgage switching research. Technical report.
- Central Statistics Office (2020). Household Finance and Consumption Survey: Debt and Credit.
- Clarke, R. G. and Arnott, R. D. (1987). The cost of portfolio insurance: Tradeoffs and choices. *Financial Analysts Journal*, 43(6):35–47.
- Cloyne, J., Ferreira, C., and Surico, P. (2020). Monetary policy when households have debt: new evidence on the transmission mechanism. *The Review of Economic Studies*, 87(1):102–129.
- Cocco, J. F., Gomes, F. J., and Maenhout, P. J. (2005). Consumption and portfolio choice over the life cycle. *The Review of Financial Studies*, 18(2):491–533.
- Cochran, W. G. (1968). The effectiveness of adjustment by subclassification in removing bias in observational studies. *Biometrics*, pages 295–313.
- Cole, S., Paulson, A., and Shastry, G. K. (2014). Smart money? the effect of education on financial outcomes. *The Review of Financial Studies*, 27(7):2022–2051.
- Conlin, A., Kyröläinen, P., Kaakinen, M., Järvelin, M.-R., Perttunen, J., and Svento, R. (2015). Personality traits and stock market participation. *Journal of Empirical Finance*, 33:34–50.
- Conroy, J. (2017). Digital authentication: New opportunities to enhance the customer journey. *Aite Group, Inc.* <https://aite-novarica.com/report/digital-authentication-new-opportunities-enhance-customer-journey>.
- Darby, M. R. and Karni, E. (1973). Free competition and the optimal amount of fraud. *The Journal of Law and Economics*, 16(1):67–88.
- De Bondt, G. (2002). Retail bank interest rate pass-through: new evidence at the euro area level. *Available at SSRN 357380*.
- DeFusco, A. A. and Mondragon, J. (2020). No job, no money, no ref: Frictions to refinancing in a recession. *The Journal of Finance*, 75(5):2327–2376.
- DeLiema, M., Mottola, G. R., and Deevy, M. (2017). Findings from a pilot study to measure financial fraud in the united states. *Available at SSRN 2914560*.
- Deshpande, M., Gross, T., and Wang, J. (2017). The effect of government cash assistance on household financial outcomes. Technical report, NBER DRC Working Paper No. NB17-03. Cambridge, MA: National Bureau of Economic Research.
- Di Maggio, M., Kermani, A., Keys, B. J., Piskorski, T., Ramcharan, R., Seru, A., and Yao, V. (2017). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review*, 107(11):3550–88.

- Di Maggio, M., Kermani, A., and Palmer, C. J. (2020). How quantitative easing works: Evidence on the refinancing channel. *The Review of Economic Studies*, 87(3):1498–1528.
- Dichtl, H. and Drobetz, W. (2011). Portfolio insurance and prospect theory investors: Popularity and optimal design of capital protected financial products. *Journal of Banking & Finance*, 35(7):1683–1697.
- Dimmock, S. G., Kouwenberg, R., Mitchell, O. S., and Peijnenburg, K. (2016). Ambiguity aversion and household portfolio choice puzzles: Empirical evidence. *Journal of Financial Economics*, 119(3):559–577.
- Dong, T., Eugster, F., and Nilsson, H. (2022). Business school education, motivation, and young adults’ stock market participation. *Journal of Accounting and Public Policy*, page 106958.
- DuGoff, E. H., Schuler, M., and Stuart, E. A. (2014). Generalizing observational study results: applying propensity score methods to complex surveys. *Health Services Research*, 49(1):284–303.
- Edin, K. and Lein, L. (1998). The private safety net: The role of charitable organizations in the lives of the poor. *Housing Policy Debate*, 9(3):541–573.
- Ehrmann, M., Holton, S., Kedan, D., and Phelan, G. (2021). Monetary policy communication: perspectives from former policy makers at the ecb.
- Eichenbaum, M., Rebelo, S., and Wong, A. (2022). State-dependent effects of monetary policy: The refinancing channel. *American Economic Review*, 112(3):721–61.
- Ericson, K. M. (2017). On the interaction of memory and procrastination: Implications for reminders, deadlines, and empirical estimation. *Journal of the European Economic Association*, 15(3):692–719.
- European Commission (2015). Green paper on retail financial services: Better products, more choice, and greater opportunities for consumers and businesses. Technical report.
- European Commission (2019). Study on switching of financial services and products – final report. Technical report, Directorate-General for Financial Stability, Financial Services and Capital Markets Union.
- Evans, W. N., Kolka, S., Sullivan, J. X., and Turner, P. S. (2023). Fighting poverty one family at a time: Experimental evidence from an intervention with holistic, individualized, wrap-around services. NBER Working Paper No. 30992.
- Farrell, J. and Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. *Handbook of Industrial Organization*, 3:1967–2072.
- Feldstein, M. S. (1973). The welfare loss of excess health insurance. *Journal of Political Economy*, 81(2, Part 1):251–280.
- Fernandes, D., Lynch Jr, J. G., and Netemeyer, R. G. (2014). Financial literacy, financial education, and downstream financial behaviors. *Management Science*, 60(8):1861–1883.
- Financial Conduct Authority (2016). Increasing consumer engagement in the annuities market: can prompts raise shopping around? Technical report.
- Financial Conduct Authority (2019). Mortgages market study: Final report. Technical report.
- Finkelstein, A. and Notowidigdo, M. J. (2019). Take-Up and Targeting: Experimental Evidence from SNAP\*. *The Quarterly Journal of Economics*, 134(3):1505–1556.

- Fisher, J., Gavazza, A., Liu, L., Ramadorai, T., and Tripathy, J. (2022). Refinancing cross-subsidies in the mortgage market. SSRN Working Paper No. 3962433.
- Florêncio, D. and Herley, C. (2013). Sex, lies and cyber-crime surveys. In *Economics of information security and privacy III*, pages 35–53. Springer.
- Francois, P. and Zabochnik, J. (2005). Trust, social capital, and economic development. *Journal of the European Economic Association*, 3(1):51–94.
- Frick, K. D. and Chernew, M. E. (2009). Beneficial moral hazard and the theory of the second best. *INQUIRY: The Journal of Health Care Organization, Provision, and Financing*, 46(2):229–240.
- Fu, J. and Mishra, M. (2022). Combatting fraudulent and predatory fintech apps with machine learning.
- Garrido, M. M., Kelley, A. S., Paris, J., Roza, K., Meier, D. E., Morrison, R. S., and Aldridge, M. D. (2014). Methods for constructing and assessing propensity scores. *Health Services Research*, 49(5):1701–1720.
- Garz, S., Giné, X., Karlan, D., Mazer, R., Sanford, C., and Zinman, J. (2021). Consumer protection for financial inclusion in low-and middle-income countries: Bridging regulator and academic perspectives. *Annual Review of Financial Economics*, 13:219–246.
- Genakos, C., Roumanias, C., and Valletti, T. (2015). Loss aversion on the phone.
- Georgarakos, D. and Pasini, G. (2011). Trust, sociability, and stock market participation. *Review of Finance*, 15(4):693–725.
- George, J. F., Marett, K., and Tilley, P. (2004). Deception detection under varying electronic media and warning conditions. In *Proceedings of the 37th Annual Hawaii International Conference on System Sciences, 2004.*, pages 9–pp. IEEE.
- Gerardi, K., Goette, L., and Meier, S. (2013). Numerical ability predicts mortgage default. *Proceedings of the National Academy of Sciences*, 110(28):11267–11271.
- Gerardi, K., Willen, P. S., and Zhang, D. H. (2023). Mortgage prepayment, race, and monetary policy. *Journal of Financial Economics*, 147(3):498–524.
- Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.
- Gilchrist, S., López-Salido, D., and Zakrajšek, E. (2015). Monetary policy and real borrowing costs at the zero lower bound. *American Economic Journal: Macroeconomics*, 7(1):77–109.
- Giordana, G. and Ziegelmeyer, M. (2020). Stress testing household balance sheets in luxembourg. *The Quarterly Review of Economics and Finance*, 76:115–138.
- Girshina, A., Mathä, T. Y., and Ziegelmeyer, M. (2019). Peer effects in stock market participation: Evidence from immigration. *Available at SSRN 3515394*.
- Gormley, T., Liu, H., and Zhou, G. (2010). Limited participation and consumption-saving puzzles: A simple explanation and the role of insurance. *Journal of Financial Economics*, 96(2):331–344.
- Gormsen, N. J. and Huber, K. (2023). Equity factors and firms’ perceived cost of capital. SSRN Working Paper No. 4321285.
- Gould-Werth, A. (2018). After the rainy day: how private resources shape personal trajectories following job loss and amplify racial inequality. *Sociology of Race and Ethnicity*, 4(3):371–385.

- Grazioli, S. and Wang, A. (2001). Looking without seeing: understanding unsophisticated consumers' success and failure to detect internet deception.
- Gropp, R., Scholz, J. K., and White, M. J. (1996). Personal bankruptcy and credit supply and demand.
- Grubb, M. D. (2014). Consumer Inattention and Bill-Shock Regulation. *The Review of Economic Studies*, 82(1):219–257.
- Grubb, M. D. and Osborne, M. (2015). Cellular service demand: Biased beliefs, learning, and bill shock. *American Economic Review*, 105(1):234–71.
- Guiso, L., Haliassos, M., and Jappelli, T. (2003). Household stockholding in europe: where do we stand and where do we go? *Economic Policy*, 18(36):123–170.
- Guiso, L., Sapienza, P., and Zingales, L. (2008). Trusting the stock market. *The Journal of Finance*, 63(6):2557–2600.
- Guryan, J., Ludwig, J., Bhatt, M. P., , et al. (2023). Not too late: Improving academic outcomes among adolescents. *American Economic Review*, 113(3):738–765.
- Haldane, A. and McMahon, M. (2018). Central bank communications and the general public. In *AEA papers and proceedings*, volume 108, pages 578–83.
- Haliassos, M. (2003). Stockholding: recent lessons from theory and computations. In *Stockholding in Europe*, pages 30–49. Springer.
- Haliassos, M. and Bertaut, C. C. (1995). Why do so few hold stocks? *The Economic Journal*, 105(432):1110–1129.
- Halpern, D. (2015). *Inside the nudge unit: How small changes can make a big difference*. Random House.
- Harder, V. S., Stuart, E. A., and Anthony, J. C. (2010). Propensity score techniques and the assessment of measured covariate balance to test causal associations in psychological research. *Psychological Methods*, 15(3):234.
- Harknett, K. (2006). The relationship between private safety nets and economic outcomes among single mothers. *Journal of Marriage and Family*, 68(1):172–191.
- Harold, J., Cullinan, J., Lyons, S., et al. (2019). Stick or switch? consumer switching in 14 retail markets across europe. Technical report.
- Hartwig, M. and Bond Jr, C. F. (2011). Why do lie-catchers fail? a lens model meta-analysis of human lie judgments. *Psychological bulletin*, 137(4):643.
- Henly, J. R. (2002). Informal support networks and the maintenance of low-wage jobs. *Laboring below the line: The new ethnography of poverty, low-wage work, and survival in the global economy*, pages 179–203.
- Herley, C. (2009). So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 workshop on new security paradigms*, pages 133–144.
- Herley, C. (2012). Why do nigerian scammers say they are from nigeria? In *WEIS*. Berlin.
- Hernández-Murillo, R., Shell, H., et al. (2014). The rising complexity of the fomc statement. *Economic Synopses*, (23).



- Hilal, W., Gadsden, S. A., and Yawney, J. (2022). Financial fraud:: A review of anomaly detection techniques and recent advances.
- Hirani, J. L.-J. (2021). Inattention or reluctance? parental responses to vaccination reminder letters. *Journal of Health Economics*, 76:102439.
- Ho, D. E., Imai, K., King, G., and Stuart, E. A. (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis*, 15(3):199–236.
- Holtfreter, K., Reisig, M. D., Leeper Piquero, N., and Piquero, A. R. (2010). Low self-control and fraud: Offending, victimization, and their overlap. *Criminal Justice and Behavior*, 37(2):188–203.
- Holtz-Eakin, D., Joulfaian, D., and Rosen, H. S. (1993). Entrepreneurial decisions and liquidity constraints.
- Hong, H., Kubik, J. D., and Stein, J. C. (2004). Social interaction and stock-market participation. *The Journal of Finance*, 59(1):137–163.
- Hoynes, H., Schanzenbach, D. W., and Almond, D. (2016). Long-run impacts of childhood access to the safety net. *American Economic Review*, 106(4):903–34.
- Hung, C.-C. (2021). Does social connectedness affect stock market participation? Available at SSRN 3724407.
- Imai, K., King, G., and Stuart, E. A. (2008). Misunderstandings between experimentalists and observationalists about causal inference. *Journal of The Royal Statistical Society: Series A (Statistics in Society)*, 171(2):481–502.
- Innovations for Poverty Action (2021a). Consumer protection in digital finance surveys. <https://www.poverty-action.org/consumer-protection-digital-finance-surveys>.
- Innovations for Poverty Action (2021b). Social media usage by digital finance consumers: Analysis of consumer complaints in Kenya, Nigeria and Uganda. July 2019 - July 2020. <https://www.poverty-action.org/sites/default/files/publications/Social-Media-Usage-by-Digital-Finance-Consumers-April-2021.pdf>.
- Johnson, E. J., Meier, S., and Toubia, O. (2019). What’s the catch? suspicion of bank motives and sluggish refinancing. *The Review of Financial Studies*, 32(2):467–495.
- Joseph, H. (1972). The measurement of moral hazard. *Journal of Risk and Insurance*, pages 257–262.
- Joulfaian, D. and Wilhelm, M. O. (1994). Inheritance and labor supply. *Journal of Human Resources*, pages 1205–1234.
- Kabanda, S., Tanner, M., and Kent, C. (2018). Exploring sme cybersecurity practices in developing countries. *Journal of Organizational Computing and Electronic Commerce*, 28(3):269–282.
- Kahneman, D. and Riepe, M. W. (1998). Aspects of investor psychology. *Journal of Portfolio Management*, 24(4).
- Kaiser, T., Lusardi, A., Menkhoff, L., and Urban, C. (2021). Financial education affects financial knowledge and downstream behaviors. *Journal of Financial Economics*.
- Kaiser, T. and Menkhoff, L. (2017). Does financial education impact financial literacy and financial behavior, and if so, when? *The World Bank Economic Review*, 31(3):611–630.

- Kanz, M. (2016). What does debt relief do for development? evidence from india's bailout for rural households. *American Economic Journal: Applied Economics*, 8(4):66–99.
- Karlan, D., McConnell, M., Mullainathan, S., and Zinman, J. (2016). Getting to the top of mind: How reminders increase saving. *Management Science*, 62(12):3393–3411.
- Katras, M. J., Zuiker, V. S., and Bauer, J. W. (2004). Private safety net: Childcare resources from the perspective of rural low-income families. *Family Relations*, 53(2):201–209.
- Kaustia, M. and Torstila, S. (2011). Stock market aversion? political preferences and stock market participation. *Journal of Financial Economics*, 100(1):98–112.
- Ke, D. (2018). Cross-country differences in household stock market participation: The role of gender norms. In *AEA Papers and Proceedings*, volume 108, pages 159–62.
- Kell, P. (2016). Asic and behavioural economics: Regulating for real people. Address given at The Impacts of Behavioural Economics on Financial Markets and Regulations Symposium.
- Keys, B. J., Pope, D. G., and Pope, J. C. (2016). Failure to refinance. *Journal of Financial Economics*, 122(3):482–499.
- Kim, E. T. (2022). The digital divide and refinancing inequality. Working Paper, University of California, Los Angeles.
- King, M., Singh, A. P., et al. (2018). Conned by a cashback? disclosure, nudges and consumer rationality in mortgage choice. Technical report, Trinity College Dublin, Department of Economics.
- Köppe, S. (2018). Passing it on: inheritance, coresidence and the influence of parental support on homeownership and housing pathways. *Housing Studies*, 33(2):224–246.
- Kraft, M. A. (2020). Interpreting effect sizes of education interventions. *Educational Researcher*, 49(4):241–253.
- Kunreuther, H. and Michel-Kerjan, E. (2014). Economics of natural catastrophe risk insurance. In *Handbook of the Economics of Risk and Uncertainty*, volume 1, pages 651–699. Elsevier.
- Lamla, M. and Vinogradov, D. (2019). Central bank announcement: Big news for little people? *Journal of Monetary Economics*, 108:21–38.
- Larsen, L. S. and Munk, C. (2021). The design and welfare implications of mandatory pension plans. Available at SSRN 3415132.
- Lee, H., Myers, D., Painter, G., Thunell, J., and Zissimopoulos, J. (2020). The role of parental financial assistance in the transition to homeownership by young adults. *Journal of Housing Economics*, 47:101597.
- Lee, J. and Hogarth, J. M. (2000). Consumer information search for home mortgages: who, what, how much, and what else? *Financial Services Review*, 9(3):277–293.
- Li, G. (2014). Information sharing and stock market participation: Evidence from extended families. *Review of Economics and Statistics*, 96(1):151–160.
- Lung, J., Tsen, A., Evans, E., and Dennis, J. A. (2020). Financial safety net and subjective happiness in medical students. *College Student Journal*, 54(3):317–326.
- Lunn, P., Bohacek, M., Somerville, J., Ni Choisdealbha, A., McGowan, F., et al. (2016). Price lab: An investigation of consumers' capabilities with complex products. *Economic and Social Research Institute (ESRI) Research Series*.

- Lunn, P. and Lyons, S. (2018). Consumer switching intentions for telecoms services: evidence from ireland. *Heliyon*, 4(5).
- Lusardi, A. (2015). “Just in time education”? Just in time is too late. *Wall Street Journal*.
- Lusardi, A. and Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, 52(1):5–44.
- Madrian, B. C. and Shea, D. F. (2001). The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly Journal of Economics*, 116(4):1149–1187.
- Maier, M., Bartoš, F., Stanley, T., Shanks, D. R., Harris, A. J., and Wagenmakers, E.-J. (2022). No evidence for nudging after adjusting for publication bias. *Proceedings of the National Academy of Sciences*, 119(31).
- Mayer, J. A., Lewis, E. C., Slymen, D. J., Dullum, J., Kurata, H., Holbrook, A., Elder, J. P., and Williams, S. J. (2000). Patient reminder letters to promote annual mammograms: a randomized controlled trial. *Preventive Medicine*, 31(4):315–322.
- McKay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, 106(10):3133–58.
- McKenzie, D. (2015). Tools of the trade: a joint test of orthogonality when testing for balance. *World Bank Development Impact Blog*.
- McKenzie, D. and Ozier, O. (2019). Why ex-post power using estimated effect sizes is bad, but an ex-post mde is not. *World Bank Development Impact Blog*.
- Mertens, S., Herberz, M., Hahnel, U. J., and Brosch, T. (2022). The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *Proceedings of the National Academy of Sciences*, 119(1).
- Miles, S. and Pyne, D. (2017). The economics of scams. *Review of Law & Economics*, 13(1).
- Milkman, K. L., Gandhi, L., Patel, M. S., et al. (2022). A 680,000-person megastudy of nudges to encourage vaccination in pharmacies. *Proceedings of the National Academy of Sciences*, 119(6):e2115126119.
- Miller, M., Reichelstein, J., Salas, C., and Zia, B. (2015). Can you help someone become financially capable? a meta-analysis of the literature. *The World Bank Research Observer*, 30(2):220–246.
- Mutz, D. C., Pemantle, R., and Pham, P. (2019). The perils of balance testing in experimental design: Messy analyses of clean data. *The American Statistician*, 73(1):32–42.
- Nash, R., Bouchard, M., and Malm, A. (2013). Investing in people: The role of social networks in the diffusion of a large-scale fraud. *Social networks*, 35(4):686–698.
- Nguyen, B., Steel, P., and Ferrari, J. R. (2013). Procrastination’s impact in the workplace and the workplace’s impact on procrastination. *International Journal of Selection and Assessment*, 21(4):388–399.
- Parker, K. and Fry, R. (2020). More than half of u.s. households have some investment in the stock market.
- Peress, J. (2005). Information vs. entry costs: What explains us stock market evolution? *Journal of Financial and Quantitative Analysis*, 40(3):563–594.

- Pfeffer, F. T. and Killewald, A. (2018). Generations of advantage. multigenerational correlations in family wealth. *Social Forces*, 96(4):1411–1442.
- Piketty, T. (2014). Capital in the twenty-first century. In *Capital in the twenty-first century*. Harvard University Press.
- Poterba, J. M. and Samwick, A. A. (2003). Taxation and household portfolio composition: Us evidence from the 1980s and 1990s. *Journal of Public Economics*, 87(1):5–38.
- Raval, D. (2020). Whose voice do we hear in the marketplace? evidence from consumer complaining behavior. *Marketing Science*, 39(1):168–187.
- Ravallion, M., Van de Walle, D., and Gautam, M. (1995). Testing a social safety net. *Journal of Public Economics*, 57(2):175–199.
- Rigney, D. (2010). *The Matthew effect: How advantage begets further advantage*. Columbia University Press.
- Robst, J. (2007). Education and job match: The relatedness of college major and work. *Economics of Education Review*, 26(4):397–407.
- Rubin, D. B. (2001). Using propensity scores to help design observational studies: application to the tobacco litigation. *Health Services and Outcomes Research Methodology*, 2(3):169–188.
- Schacter, D. L. (1999). The seven sins of memory: insights from psychology and cognitive neuroscience. *American Psychologist*, 54(3):182.
- Scheibe, S., Notthoff, N., Menkin, J., Ross, L., Shadel, D., Deevy, M., and Carstensen, L. L. (2014). Forewarning reduces fraud susceptibility in vulnerable consumers. *Basic and applied social psychology*, 36(3):272–279.
- Shcherbakov, O. (2016). Measuring consumer switching costs in the television industry. *The RAND Journal of Economics*, 47(2):366–393.
- Sianesi, B. (2004). An evaluation of the swedish system of active labor market programs in the 1990s. *Review of Economics and Statistics*, 86(1):133–155.
- Singh, A., Park, A., and Dercon, S. (2014). School meals as a safety net: an evaluation of the midday meal scheme in india. *Economic Development and Cultural Change*, 62(2):275–306.
- Stankov, L. and Crawford, J. D. (1996). Confidence judgments in studies of individual differences. *Personality and Individual Differences*, 21(6):971–986.
- Steel, P. and König, C. J. (2006). Integrating theories of motivation. *Academy of Management Review*, 31(4):889–913.
- Stiglitz, J. E. (1983). Risk, incentives and insurance: The pure theory of moral hazard. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 8(1):4–33.
- Stiglitz, J. E. and Weiss, A. (1981). Credit rationing in markets with imperfect information. *The American Economic Review*, 71(3):393–410.
- Stone-Gross, B., Abman, R., Kemmerer, R. A., Kruegel, C., Steigerwald, D. G., and Vigna, G. (2013). The underground economy of fake antivirus software. In *Economics of information security and privacy III*, pages 55–78. Springer.
- Stuart, E. A., Lee, B. K., and Leacy, F. P. (2013). Prognostic score-based balance measures can be a useful diagnostic for propensity score methods in comparative effectiveness research. *Journal of Clinical Epidemiology*, 66(8 Suppl):S84–S90.

- Tade, O. (2021). Covid-‘419’: Social context of cybercrime in the age of covid-19 in nigeria. *African Security*, 14(4):460–483.
- Tenreyro, S. and Thwaites, G. (2016). Pushing on a string: Us monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, 8(4):43–74.
- Thaler, R. H. and Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.
- Titus, R. M., Heinzemann, F., and Boyle, J. M. (1995). Victimization of persons by fraud. *Crime & Delinquency*, 41(1):54–72.
- Van Rooij, M., Lusardi, A., and Alessie, R. (2011). Financial literacy and stock market participation. *Journal of Financial Economics*, 101(2):449–472.
- Vancouver, J. B. and Kendall, L. N. (2006). When self-efficacy negatively relates to motivation and performance in a learning context. *Journal of Applied Psychology*, 91(5):1146.
- Vancouver, J. B., Thompson, C. M., and Williams, A. A. (2001). The changing signs in the relationships among self-efficacy, personal goals, and performance. *Journal of Applied Psychology*, 86(4):605.
- Velásquez, I., Caro, A., and Rodríguez, A. (2018). Authentication schemes and methods: A systematic literature review. *Information and Software Technology*, 94:30–37.
- Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy*, 110(4):825–853.
- Vissing-Jørgensen, A. (2019). Central banking with many voices: the communications arms race. *Conference Proceedings, 23rd Annual Conference of the Central Bank of Chile*.
- Vohs, K. D., Baumeister, R. F., and Chin, J. (2007). Feeling duped: Emotional, motivational, and cognitive aspects of being exploited by others. *Review of General Psychology*, 11(2):127–141.
- Wang, J. and Burke, K. (2022). The effects of disclosure and enforcement on payday lending in texas. *Journal of Financial Economics*, 145(2, Part B):489–507.
- Wong, A. et al. (2019). Refinancing and the transmission of monetary policy to consumption. Working Paper, Princeton University.
- Woodman, T., Akehurst, S., Hardy, L., and Beattie, S. (2010). Self-confidence and performance: A little self-doubt helps. *Psychology of Sport and Exercise*, 11(6):467–470.
- Woodman, T. and Hardy, L. (2003). The relative impact of cognitive anxiety and self-confidence upon sport performance: A meta-analysis. *Journal of Sports Sciences*, 21(6):443–457.
- World Bank (2018). Matching Techniques. available at <https://thedocs.worldbank.org/en/doc/899821526315082994-0090022018/original/MatchingCairoEnglish.pdf>
- World Bank Group (2019). Nigeria digital economy diagnostic report.
- Xiao, B. and Benbasat, I. (2015). Designing warning messages for detecting biased online product recommendations: An empirical investigation. *Information Systems Research*, 26(4):793–811.
- Yang, Y. (2014). Understanding household switching behavior in the retail electricity market. *Energy Policy*, 69:406–414.
- Yin, W. and Matthews, K. (2016). The determinants and profitability of switching costs in chinese banking. *Applied Economics*, 48(43):4156–4166.

- Zanutto, E. L. (2006). A comparison of propensity score and linear regression analysis of complex survey data. *Journal of Data Science*, 4(1):67–91.
- Zhang, D. (2022). Closing costs, refinancing, and inefficiencies in the mortgage market. Working paper, Harvard University.
- Zhang, Z., Kim, H. J., Lonjon, G., Zhu, Y., et al. (2019). Balance diagnostics after propensity score matching. *Annals of Translational Medicine*, 7(1).
- Zhao, Z. and Min, C. (2021). Investing like my parents: Do parents affect children’s risk taking behavior? *Available at SSRN 3949539*.