



UvA-DARE (Digital Academic Repository)

Understanding political behavior: Essays in experimental political economy

Gago Guerreiro de Brito Robalo, P.M.

Publication date

2014

Document Version

Final published version

[Link to publication](#)

Citation for published version (APA):

Gago Guerreiro de Brito Robalo, P. M. (2014). *Understanding political behavior: Essays in experimental political economy*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.



Understanding Political Behavior

Essays in Experimental Political Economy

Pedro Robalo

Explaining individual political behavior is one of the big challenges in the social sciences. The work contained in this thesis uses the tools of experimental economics, game theory and decision theory to shed light on political choices. Relaxing the neoclassical assumptions of self-interested preferences and full rationality, this work investigates whether group identity and altruism matter for political participation, what the role of reciprocity and normative appeals in the response to political mobilization is, and whether the costs of information influence the way it is incorporated in (political) decision-making. The methodology of experimental economics is crucial to obtain an answer to these questions. The evidence presented in this thesis shows that group identity has a mild effect on one's decision to participate and that more altruistic people participate more in politics; it underscores the importance of normative appeals for the effectiveness of political mobilization; and it demonstrates that, contrary to the standard assumption, the cost of information influences the way it is incorporated in decisions.

Pedro Robalo (1984) studied Economics at Universidade Nova de Lisboa, Portugal, where he obtained the degrees of B.Sc. (2007) and M.Sc. (2008). In 2008 he started a Ph.D. at the Tinbergen Institute/University of Amsterdam, where he earned an M.Phil. degree in 2010. In the same year he joined the Center for Research on Experimental Economics and Political Decision Making of the University of Amsterdam, where he conducted research in the areas of Experimental Economics and Political Economy. He has taught courses in Microeconomics and Game Theory. Currently, he is a research fellow of the Max Planck Institute for Research on Collective Goods in Bonn, Germany.

Understanding Political Behavior

Pedro Robalo



Universiteit van Amsterdam



UNDERSTANDING POLITICAL BEHAVIOR

ISBN 978 90 316 0412 8

Cover design: Crasborn Graphic Designers bno, Valkenburg a.d. Geul

Cover picture: Bram Wisman/MAI.

This book is no. **598** of the Tinbergen Institute Research Series, established through cooperation between Thela Thesis and the Tinbergen Institute. A list of books which already appeared in the series can be found in the back.

**UNDERSTANDING POLITICAL
BEHAVIOR**
Essays in Experimental Political Economy

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. D.C. van den Boom
ten overstaan van een door het college voor promoties
ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op woensdag 3 December 2014, te 10:00 uur

door

Pedro Miguel Gago Guerreiro de Brito Robalo
geboren te Faro, Portugal

PROMOTIECOMMISSIE

PROMOTOR:

Prof. dr. A.J.H.C. Schram

OVERIGE LEDEN:

Prof. dr. C. Engel

Prof. dr. R.B. Morton

Prof. dr. J.H. Sonnemans

Dr. J. van de Ven

Prof. dr. F.A.A.M. van Winden

Faculteit Economie en Bedrijfskunde

To the memory of my father

Acknowledgments

Writing a Ph.D. thesis turned out to be a much less solitary endeavor than I first thought. This thesis would not have been possible without the help and inspiration of many people.

First, I would like to thank my supervisor, Arthur Schram, for his invaluable guidance and encouragement. Arthur is a natural-born motivator, cheering you up when things don't go as you hoped for and bringing you down to earth when necessary. He showed me the importance of working on your own ideas and communicating your research effectively. He was available at all times, yet never pushed me through - the most valuable lesson of the past five years was delivered in the process: in academia you owe it to yourself to keep pace. Arthur was also very understanding and supportive in all non-academic matters. Chapter 2 of this thesis is joint work with him.

I would also like to thank my other co-authors, Joep Sonnemans (also Chapter 2) and Rei Sayag (Chapter 4). Working with Joep has been a pleasure and I have benefited greatly from his encyclopaedic knowledge of the behavioral sciences. Rei is not only a co-author but also a good friend. We started working together on the first week of our Ph.D. and we haven't stopped since. His energy and tenacity have been a continuous source of inspiration. I hope we will continue to work together in the future and, perhaps more importantly, to share a good laugh over beers.

I consider myself lucky for having been able to conduct the research in these pages at the Center for Research on Experimental Economics and Political Decision Making (CREED) of the University of Amsterdam (UvA). CREED is an exceptional research group thanks to a passionate, friendly and dynamic group of people. I had the privilege of often discussing my research with Frans, who gives invaluable advice to starting Ph.D.s (he once told us that trying to contribute to long-standing research topic is like jumping on a moving train - it is easier to make an impact with something new). I learnt a great deal about teaching with Theo, who also delivers comments that are always useful and to the point. Sharing an office with Marcelo and Roel was both fun and instructive, as they were equally eager to discuss experiments and

football. Above all, I am indebted to the everyday interactions with all CREED and UvA colleagues, who were warm and challenging in equal parts: Aaron, Adrian, Ailko, Aljaž, Anita, Audrey, Ben, Boris, David, Francisco, Gönül, Jeroen, Jindi, Joel, Jona, Jos, Julian, Matthijs, Matze, Max, Noemi, Thomas, Sander, Simin and Yang. The cooperative spirit of CREED persists whenever you encounter former CREED-ers. I won't forget when Jens Großer sat down an entire afternoon in New York with me to discuss a preliminary version of what is now Chapter 2. I would also like to thank Karin Breen and Robert Helmink for making things run smoothly at UvA.

I am also indebted to many people at the Tinbergen Institute (TI). TI is a unique institution, combining graduate training of the highest level with a laid-back and cooperative atmosphere. This is possible because of the utmost dedication of people like Arianne de Jong and Judith van Kronenburg, among others. I was personally involved in the Educational Board and the TI Magazine, which provided me with enriching experiences. Indelible are also the memories from the Roetersstraat times, when we graduate students would practically live at the institute. At TI or outside of it, I'm happy to have shared good moments with Agbeko, Eva, Gosia, Guilherme, Inez, Lerby, Marcin, Mark, Natalya, Olivier, Ona, Patrick, Philippe and Stan.

The late stages of this thesis were completed at the Max Planck Institute for Research on Collective Goods, to where I moved in the Spring of 2014. I would like to thank Christoph Engel for the opportunity. I would also like to leave a word of appreciation to my colleagues at the Max Planck Institute who made me feel at home right away. In particular, I would like to thank Franzi Tausch, the other 'new arrival' at Max Planck, for teaching me the ropes and for tips on typesetting a thesis.

Amsterdam wouldn't have been the same without the friends that made life fun and exciting. They would often ask how I was progressing and address me the daring questions only friends can ask. Among the many people that came and went, I particularly miss some of those who happened to leave town before me. Alice was a true *gezelligheid*-maker and someone who could always cheer you up. Chana shook the habitual with her sharp wit and challenging ways. Liv was a ray of sunshine in our lives and is one of the most fun persons I've met. The list wouldn't be complete without a word of appreciation to my Portuguese friends, the people I grew up with and to whom I always feel close despite the geographic distance. Among them, I was happy to have had Moritz and Pedro as a Portuguese presence in the Netherlands.

A big thank you goes to my 'proxy' Amsterdam family, Heiner and Ben. Living with economists is an exhilarating experience, contrary to what some might have you

believe. Discussions on the fixed and variable costs of lifetime coffee consumption or how to behave strategically as a tenant were part of our everyday lives. Heiner has been a friend since the day we started our Ph.D., and I cannot thank him enough for all the good times and for also being there in the not so good ones. Ben's reliable presence and unique sense of humor made everyday life different for the better. He's an innate economist and I benefited from our discussions. When Ben left, Mehmet took over and was also an excellent flatmate.

At last, I would like to thank the people who are dearest to me. Jozefina's support was very important, as was her ability to bring joy and fun to everything we do. My grandparents Ru-ru and Nibinho were always interested in how I was doing. They are the example of how hard work leads not only to success but also accomplishment. I owe a big thank you to my sister Ana, who is someone I can always depend on. My mother has been a constant guidance in my life, and I wouldn't have finished this Ph.D. if it was not for her generosity and unwavering support. As I finished this thesis, I had to think of the family members that were so important to me and who left during its writing. My grandmother Natália's memory shall help me never to compromise on style and manners. My uncle Pedro showed me how to be a maverick and yet be dedicated to your close ones.

More than anyone, I am thankful to my father, whose example will go unmatched for as long as I live. He was someone with a strong character and solid principles, but also a person of rare sensibility. I thank him for bringing me up in a house with books and for always nurturing my interests. He was a cultured man and a mighty debater. We would often discuss economic issues - and disagree fiercely. I thought that one day, when I got my Ph.D., perhaps I could convince him of the virtue of my arguments. That day is approaching and I realize that I would not waste a second on that; as in many other things, I share more and more of his wise views. I miss him dearly.

Pedro
Bonn, August 2014

I acknowledge the financial support of a Doctoral Grant from Fundação para a Ciência e Tecnologia (the Portuguese National Science Foundation); all experiments have been funded or co-funded by the Research Priority Area Behavioral Economics of the University of Amsterdam, to which I am also grateful.

Contents

1	Introduction	1
1.1	Methodology: Game Theory, Decision Theory and Laboratory Experiments	2
1.2	Overview	4
	Appendix 1.A Participation Games	8
	Appendix 1.B Quantal Response Equilibrium	9
2	Other-regarding Preferences, Group Identity and Political Participation	13
2.1	Introduction	13
2.2	Related Literature	16
2.3	Conceptual Framework and Hypotheses	18
2.4	Experimental Design	21
2.5	Experimental Results	26
2.5.1	Subject Classification	26
2.5.2	Group Identity Induction	29
2.5.3	Other-regarding Preferences and Participation Behavior	32
2.5.4	Group Identity and Participation Behavior	33
2.5.5	Multivariate Analysis	36
2.6	Conclusion	39
	Appendix 2.A Equilibria of the Participation Game	41
	2.A.1 In-group Bias and Aggregate Participation	43
	2.A.2 Altruism, In-group Bias and Individual Participation	44
	Appendix 2.B Value Orientation Test - Ring Test	48
	Appendix 2.C Auxiliary Tables	48
	Appendix 2.D Other-regarding Preferences and Group Identity	51
	Appendix 2.E Experiment Instructions	52
3	Why Does Political Mobilization Work? The Role of Norms and	

Reciprocity	57
3.1 Introduction	57
3.1.1 Candidate Explanations	59
3.2 The Model	63
3.2.1 Implemented Parameterization and Equilibrium Predictions	65
3.2.2 Comparative Statics	67
3.3 Experimental Design and Hypotheses	68
3.3.1 Treatments	71
3.3.2 Hypotheses	72
3.4 Results	75
3.4.1 The Impact of Mobilization	76
3.4.2 Participation Appeals and Group Norms	80
3.4.3 Test of Equilibrium Predictions	82
3.4.4 Regression Analysis	86
3.5 Conclusion	88
Appendix 3.A The Model - Details	90
3.A.1 Participation Probabilities for Beta Players	90
3.A.2 Quantal Response Equilibrium	92
Appendix 3.B Experiment Instructions	99
3.B.1 Comprehension Questions	102
Appendix 3.C Trust Game Results	102
Appendix 3.D Subject Ratings	104
Appendix 3.E Independence of Observations Across Blocks	106
Appendix 3.F Transcript of Alphas' Messages	110

4 Paying is Believing: The Effect of Costly Information on Bayesian Updating	113
4.1 Introduction	114
4.1.1 Sunk Information Costs	115
4.2 Experimental Design	119
4.2.1 Choice Framework	120
4.2.2 Treatments	121
4.3 Experimental Results	124
4.3.1 Treatment effects	125
4.3.2 Interpreting the effect of cost	129
4.3.3 Selection effects	132
4.4 Pricing Information	136
4.5 Conclusion	139

Appendix 4.A Details on the Choice Framework and Type Classification .	141
4.A.1 Choice Framework	141
4.A.2 Type Classification	142
Appendix 4.B Experiment Instructions	143
Appendix 4.C Additional results	146
Appendix 4.D Biases in Bayesian updating	151
Bibliography	153
Summary	165
Samenvatting (Summary in Dutch)	169
Sumário (Summary in Portuguese)	173

List of Figures

2.1	Sequence in the experiment	23
2.2	Experimental treatments	26
2.3	Typical choice in the ring test	28
2.4	Other-regarding preference types	30
2.5	Group identity induction across treatments	30
2.6	Evolution of participation and other-regarding preferences	33
2.7	Individual participation rates and other-regarding preferences.	34
2.8	Participation across treatments	35
2.9	In-sample model predictions	39
2.10	Other-regarding preferences and group identity	51
3.1	QRE and NE predictions for Alphas	67
3.2	Sequence in one of the blocks of the experiment.	71
3.3	Activation decisions of Alphas through sub-blocks.	77
3.4	Aggregate participation across treatments	78
3.5	Participation and individual costs	79
3.6	Participation by reciprocator type	80
3.7	Participation and normative appeals	82
3.8	QRE equilibrium participation probabilities for different values of μ	95
3.9	Amount sent in the trust game.	103
3.10	Amount returned in the trust game.	103
3.11	Ratings distribution.	105
4.1	The scoring function	120
4.2	Outline of the experimental design.	122
4.3	Average decisions by blocks of 5 periods.	126
4.4	Distributions of individual decision making.	128
4.5	Average decision by subject type.	134
4.6	Implied cost levels	138
4.7	Decisions and the effect of biases on decision making.	151

List of Tables

2.1	Other-regarding preference types	29
2.2	Participation and other-regarding preferences	34
2.3	Panel regression model	38
2.4	Other-regarding preferences and individual participation	46
2.5	In-group bias and individual participation	47
2.6	Ring test.	48
2.7	Reported motivations in the questionnaire.	49
2.8	Participation, personal traits and other-regarding preference types . .	49
2.9	Regression of average participation on a linear trend	50
2.10	Conditional distribution of types across treatments, in percentage. . .	50
3.1	Parameter values.	65
3.2	Point predictions	66
3.3	Activation decisions	76
3.4	Model predictions and observed results	84
3.5	Panel regression results	87
3.6	Payoff matrix of Alphas for $\mu^{Beta}=0.4$	94
3.7	Panel regression results	105
3.8	Panel regression results	106
3.9	Determinants of participation across sub-blocks	108
3.10	Determinants of participation across blocks	109
4.1	Decision block: priors and lottery distributions (urns).	121
4.2	Subject types.	123
4.3	Summary statistics	125
4.4	Decision averages by treatment and information condition.	127
4.5	Representativeness bias and treatment effects.	132
4.6	Type distribution by treatment	133
4.7	Regression results by ball draw.	135
4.8	Parameterizations A and B.	141

4.9	Values for parameterizations A and B.	142
4.10	Information acquisition rates of types across treatments.	143
4.11	p-values for Kolmogorov-Smirnov tests of decision distributions	146
4.12	Representativeness bias and treatment effects	147
4.13	Decision averages by treatment, information and type.	148
4.14	Regression results by information condition (ball draw)	149
4.15	Regression results by information condition (no draw)	150

Chapter 1

Introduction

*Each person wants to participate and at the same time to be left alone. And because it is not possible to have it both ways, there is always a conflict.*¹

Thomas Bernhard, *Aus Gesprächen mit Thomas Bernhard* (1991)

Understanding political participation is one of the big challenges in the social sciences. The above quote alludes to the tension that individuals face whenever they are called upon to intervene in political affairs: on the one hand, a desire to participate - to vote, protest, campaign, or speak out; on the other hand, a propensity to be “left alone” and let others take responsibility and bear the costs of the political process. In this respect, political participation constitutes a collective action problem (Olson 1965): outcomes depend on the sum of costly individual efforts, which leads to an incentive to free-ride on others’ actions. What is remarkable about political participation instances is how often they successfully overcome the collective action problem. The fact that most democratic elections, small and large, exhibit substantial turnout runs counter to the inescapable nature of collective action. Understanding what motivates people to take part in politics can enlighten the conditions under which collective action broadly understood is poised for failure or success.

This thesis explores how non-standard preferences (Chapters 2 and 3) and decision making biases (Chapter 4) might influence the decision to participate and the optimality of political choices, respectively. The standard models in political economy and public choice have mirrored the central tenets of the neoclassical view, namely self-interested preference orderings and full rationality (Rowley et al. 1993). This paradigm has yielded substantial insights into the political process. However, some phenomena seem to elude the paradigm, chiefly among them the inability to

¹This translation is borrowed from Alexander Düttmann, to whom I’m thankful for the original German reference (Hofmann 1991).

fully explain individual political participation. The reasons underlying this failure have been the subject of lengthy arguments and counter-arguments (e.g. Green and Shapiro 1994 and Friedman 1996). This thesis aims at extending the neo-classical model by providing directions in which its microfoundations can be enriched. The ultimate goal is to increase its explanatory content and predictive power.

The work presented in this thesis follows the spirit of much of behavioral and experimental economics: it seeks to improve the psychological realism of political economy through the extension of models and their empirical testing (Camerer et al. 2004). The improved realism is achieved by extending the game theoretic workhorse of political participation (Palfrey and Rosenthal 1983, 1985) to incorporate social preferences (Chapters 2 and 3) and by questioning an assumption shared by all models of voting under uncertainty (Chapter 4).

Relaxing the self-interest assumption is important because other-regarding motives may play an important role in political choices. In fact, social scientists agree that some motives matter more than others depending on context. In the political realm, where collective choices are the product of countless individual inputs, other-regarding motives are likely to play a prime role. This is explained in part by the nature of collective action, as one's actions have little instrumental power and considerations other than self-interest naturally come into play (Caplan 2007, Feddersen et al. 2009). In fact, as actions map imperfectly into outcomes, individuals will be more willing to entertain expressive or other-regarding motivations. This should not overshadow the fact that collective decisions bring people together and can make them more public-spirited, as they consider how political outcomes might affect others. The extension presented in Chapter 2 investigates whether other-regarding and group-discriminating preferences influence the decision to participate. Chapter 3 asks whether reciprocity and compliance with normative appeals drive a response to mobilization efforts. Chapter 4 questions the realism of a common assumption in most models in political economy, namely that the cost of information should not affect the way it is used in situations of choice under risk, like elections and other instances of political choice.

1.1 Methodology: Game Theory, Decision Theory and Laboratory Experiments

Despite looking for new directions to improve the microfoundations of political economy models, the work in this thesis retains the economic approach's "conventional techniques and goals: formal theoretical and empirical analysis using tractable mod-

els, with a focus on prediction and estimation” (Rabin 2013a, 2013b). All chapters use a combination of theoretical models and laboratory experiments. However, primacy is given to what can be learned from the experimental data - the models are developed to the extent that they can make predictions regarding what to expect in the laboratory. In this sense, these are experiments with a theoretical underpinning rather than experimental tests of theoretical models. Chapters 2 and 3 present a game theoretic framework, while Chapter 4 employs a decision theoretic framework. All chapters use the data of laboratory experiments as the source of empirical evidence.

A game theoretic framework has obvious advantages and disadvantages for the study of political participation. The disadvantages have been the object of lengthy discussions (see Green and Shapiro 1994). Here I would like to single out a few of the numerous advantages. First and foremost, the game theoretic formulation stays true to the structure of the actual phenomenon. One can add infinite layers of complexity to the participation problem, but it ultimately boils down to the question of which group manages to gather more support (e.g. which candidate obtains more votes in a plurality rule election). Misrepresenting the situation is likely to create more problems than solve the existing ones. This is not to say that the model in its primitive form is satisfactory (it does not seem to be). Rather, we must find ways to extend it meaningfully. A second and corollary advantage is that the precise nature of the game theoretic formulation helps us achieve both clarity and parsimony, which is key to isolate causal links in empirical tests. A third advantage is that it takes into account the strategic interaction that takes place among individuals. This point is important because equilibrium behavior might often run counter to intuitive or educated guesses.

Despite substantial resistance by many social scientists, the experimental method is by now recognized as one of the great sources of knowledge in the social sciences (Falk and Heckman 2009). Its use is increasing in both economics and political science. Laboratory experiments are important because they are the research method where controlled variation is most easily achieved. By means of random assignment of subjects to treatment conditions, one can rule out unobservable characteristics as the explanation for the correlation among variables of interest. In other words, laboratory experiments allow us to establish causality in a convincing way. As an illustration, Chapter 2 reports an experiment that manipulates the group identity attachment of different groups. This is done by creating groups that differ along a measured personality characteristic. Groups are defined by this characteristic, and nothing else. In the field, any intervention leading to the same partition of

groups would likely entail varying other confounding factors, like the group-specific interaction between its members that is precluded in the laboratory.

Aside from the standard justifications, there is a feature of the laboratory experimental method that I find particularly compelling and which is present throughout this thesis: artificiality. The artificiality of the laboratory environment allows for manipulations that would be hard to achieve in the field, or would be at least ethically reproachable to pursue. The last two chapters provide an illustration. Chapter 3 decomposes mobilization efforts into a material effort component and a normative appeal component, implementing environments in which both or only one is present. This allows us to assess which plays a more prominent role in the success of political mobilization. In the field it is impossible to separate these two components. Chapter 4 implements a treatment in which a costly piece of information is imposed on subjects, i.e. money is deducted from a subject's earnings for something he or she did not agree to acquire. In the context of a laboratory experiment this is a deontologically reasonable practice, but the same cannot be said of an analogous intervention outside of the laboratory. For many research questions, the artificiality of the laboratory can be an advantage, not the drawback it is considered to be.

1.2 Overview

In what follows I will briefly describe the research questions, methodology and main results of each chapter.² Chapters 2 and 3 deal with political participation, while Chapter 4 investigates a phenomenon with potential relevance for political choices.

Political participation in this thesis is understood in a broad sense: any situation in which an individual is given the opportunity to make a costly effort to help the group of which he or she is a member achieve an advantage over a competing group. Examples include large and small scale elections, two-sided rallies and campaigning for a candidate. The institutional features and the incentives of such situations is captured by the participation game of Palfrey and Rosenthal (1983), which is described in detail in Appendix 1.A. In brief terms, each player has to decide whether to participate or to abstain. Participation is costly for the individual. The group who wins the game is the one in which more players chose to participate. Players in the winning group get a reward that is higher than the one that accrues to players in the losing group. Rewards are irrespective of individual participation, i.e. only depend on victory or defeat.

²For the sake of brevity I will avoid providing references as much as possible in what remains of this chapter. All arguments and the supporting references can be found in the referenced chapters.

Chapter 2 deals with the effect of group identity and altruism on the decision to participate and is based on joint work with Arthur Schram and Joep Sonnemans. When groups with diverging interests settle their disputes via democratic politics, which is the case in *any* election, the allegiances each individual has to the group should matter for her decision to participate. More broadly, how much more she cares about an individual of her group than an individual of the other group should determine her willingness to endure the costs of participation. In fact, group identity seems to be a driving force of participation: in the United States, African-Americans participate at higher rates than their socioeconomic status would predict (Leighley and Vedlitz 1999). This constitutes a puzzle because socioeconomic status is typically the best predictor of individual political participation. One of the candidate explanations for this puzzle is a heightened sense of group identification. However, field data make it extremely hard to identify the causal effect of group identification per se on participation, as it evolves concomitantly with mobilization and socialization processes. A laboratory experiment grants us the control necessary for this investigation.

We allow for group-directed other-regarding preferences and implement them in the participation game. We manage to induce different levels of group identification in our treatments. Competition in the participation game takes place either between groups which are composed of dissimilar subjects (in which case group identity is high), or between groups of similar subjects (in which case group identity is low). Based on the theoretical results and the extant evidence, we hypothesize that both individual and aggregate participation should be increasing in the level of individual group identification and treatment-level group identity, respectively. Concurrently, our experiment also allows us to test whether individuals with altruistic preferences tend to participate more often. At the aggregate level, no differences in participation are observed between high group identity and low group identity environments. At the individual level, there is a modest effect of group identification on participation. A more robust effect is found for non-group-specific altruistic concerns. This is in line with altruism theories of voter turnout (e.g. Evren 2012). A by-product contribution of this study is methodological, as we propose a novel procedure that manages to induce different levels of group identity in the laboratory without resorting to natural groups.

Alongside a concern for others, having been asked by others to participate also seems to provide a reason to do so. In fact, changing mobilization patterns were thought to be the key to the diminishing participation observed in the United States in the last three decades of the 20th century. In their influential work, Rosenstone

and Hansen (1993) claim that a disintegration of mobilization activities is responsible for the decline in participation of a population who was getting richer and more educated (and should as a consequence get more involved in politics). The gist of their argument is that citizens who had been reached by campaigns or activist groups participate at higher rates. The changing nature of campaigning - from labor-intensive methods like canvassing to capital-intensive methods like mass-media advertising - could then explain the decrease in turnout. This study has a major methodological flaw as it ignores the endogeneity of strategic contact and participation, i.e. those more likely to be contacted are also more likely to participate. This flaw is underlined and addressed by a large field experimental literature (spanned by the seminal work of Gerber and Green 1999), which shows more convincingly that the old tactics of mobilization (e.g. door-to-door canvassing) indeed work much better than the contemporary automated ones (e.g. mass mailings). However, the question remains as to what drives people to respond to mobilization efforts. Chapter 3 investigates the psychological mechanisms underlying this phenomenon. The starting point is the observation that all mobilization efforts involve a material effort and a normative appeal. Mobilization could then work either via reciprocity concerns, i.e. as a token of appreciation for the material effort; and/or via compliance with normative appeals, i.e. to avoid the disutility associated with violating the participation norm.

The participation game is extended to allow for mobilization and participation. This requires that one of the subjects in each group is appointed with the task of mobilizing others in his or her group. This subject is assigned a budget that she can either keep or use to mobilize others. Mobilizing others increases the chances of winning the participation game at the expense of a forgone budget. For a given activation patten, the remaining subjects play a standard participation game. The proposed extension is devised having in mind the laboratory implementation, but can also serve as a first step towards a model that includes both mobilization and participation. Existing models of group mobilization eschew the collective action dimension of the participation decision, because they focus on the behavior of leaders and do not confront the fact that individuals may react strategically to different levels of mobilization. In particular, mobilizing an extra citizen will affect the participation probabilities of all others, which is something that the framework I propose takes into account.

The experimental treatments consist of varying the mobilization method (human-driven or automated) and the normative appeal conveyed by the mobilizing subject to others in his or her group (present or absent). The main results show that the

normative appeal is successful in increasing participation, in particular when it is coupled with the mobilization effort. Mobilization alone is not enough to increase participation, which disconfirms the reciprocity conjecture. I also carry out an assessment of the model's point and comparative statics predictions, to conclude that the behavior of leaders is not in line with the point predictions, while most comparative statics results seem to hold.

Information is a crucial determinant of correct decisions, be they individual or collective. Chapter 4, which is based on joint work with Rei Sayag, proposes a first approach to the question of whether the costs of information affect the way it is incorporated in decision making. The findings are potentially relevant for two well-known and related results in political economy: rational ignorance and the Condorcet jury theorem. Rational ignorance involves individuals not acquiring costly information prior to an election because their vote is unlikely to be pivotal, rendering an informed vote and an uninformed as virtually identical. On the other hand, the basic formulation of the Condorcet jury theorem argues that increasing the size of an electorate in which each member has access to free information leads to better collective decisions via more efficient information aggregation.

The literature dealing with these topics, as most of the literature on individual decision making, assumes that information is incorporated in individuals' judgments via Bayes' Rule. It is further assumed that Bayes' rule is uniformly applied irrespective of the cost of information. In our study we ask whether this presumption is legitimate. We construct an individual decision making task under risk and vary the way in which information is made available to subjects: for free, optionally at a cost or imposed at a cost. The laboratory allows us to circumvent the problematic selection issues present in the field, where the subjects who acquire information are the ones most likely to benefit from it. We find that the assumption that Bayesian updating does not depend on information's cost should be questioned: subjects overweight both the signals they choose to acquire and the signals that they were forced to acquire. In sum, costly information is weighted more heavily than free information, and leads to more extreme shifts in posterior beliefs. Whether this leads to more optimal decision making depends on how far the posterior under free information lies from the normative optimum. Future work should investigate whether this phenomenon is also observed when the strategic complexities of a voting situation are introduced.

Appendix

1.A Participation Games

Palfrey and Rosenthal (1983) put forward the following game theoretic model of political participation. There is a finite set of players and each player belongs to one of two groups. Define the set of players as $I = [1, \dots, M_i, M_i + 1, \dots, M_i + M_j]$, where M_i and M_j are the number of players in groups G_i and G_j , respectively. The action space of a player has two elements: participation and abstention. All players decide simultaneously (and individually) whether or not to participate. Participation is costly (c), abstention is not. The group where more players participate, wins. Players on the winning side obtain a payoff (B^W) that is higher than that accruing to players on the losing side (B^L). In case of a tie the winner is decided by a fair coin toss. The structure and payoffs of the game are common knowledge to all players.

From the perspective of a player i , define m_i as the number of other members in i 's group who participate, and m_j as the number of players in the other group who participate. The expected utility of participation and abstention are, respectively:

$$E[U_i^{Part.}] = \Pr[m_i + 1 > m_j]B^W + \Pr[m_i + 1 = m_j]\frac{(B^W + B^L)}{2} + \Pr[m_i + 1 < m_j]B^L - c \quad (1.1)$$

and

$$E[U_i^{Abst.}] = \Pr[m_i > m_j]B^W + \Pr[m_i = m_j]\frac{(B^W + B^L)}{2} + \Pr[m_i < m_j]B^L \quad (1.2)$$

In words, an individual who chooses to participate increases the chances of creating or breaking a tie but has to pay a cost. In equilibrium, player i must be indifferent between participating and abstaining, and therefore we equate equations 1.1 and 1.2, which simplifies to:

$$\Pr[m_i = m_j] + \Pr[m_i = m_j - 1] = \frac{2c}{(B^W - B^L)} \quad (1.3)$$

This condition defines a Nash Equilibrium (NE) in the participation game. It tells us that a player will participate if the probability that she breaks ($\Pr[m_i = m_j]$) or creates ($\Pr[m_i = m_j - 1]$) a tie, multiplied by the expected benefit, $(B^W - B^L)/2$, equals the cost of participation, c . NE existence depends on how costs relate to

benefits. For $c > (B^W - B^L)/2$ the only equilibrium is pure and has no player participating. For $c \leq (B^W - B^L)/2$ there exist NE, either in pure strategies, mixed strategies, or both.

The two most relevant classes of NE are pure strategy equilibria and ‘totally quasi-symmetric equilibria’, i.e. equilibria in which all players use mixed strategies and players in the same group employ the same strategy. Namely, for $(M_i, M_j) = (1, 1)$ there exists a unique NE where both players choose to participate. The same equilibrium exists for all cases where $M_i = M_j \gg 1$. When $M_i \geq 1$ and $M_j = 0$ (and vice-versa) the game reduces to a public goods game (Isaac and Walker 1988). For such games there exist M_i pure strategy NE in which one player in i participates and the others abstain. For $M_i \geq 2$ and $M_j = 0$ (and vice-versa) there also exist mixed strategy NE. In the more general case of $(M_i, M_j) \gg 1$, $M_i \neq M_j$, several mixed strategy NE exist (see Palfrey and Rosenthal 1983 for an extensive description).

The computation of equilibria requires a specification of the probability terms in equation 1.3. This will be done for the model’s implementation in Chapters 2 and 3. In the laboratory, the participation game was framed as a ‘disc buying game’, following part of the literature (e.g. Schram and Sonnemans 1996a). Buying a disc corresponds to participating. This formulation makes the game easy to grasp, while avoiding potentially charged words like ‘participation’ and ‘abstention’, which would involve losing control as to what attitudes and beliefs subjects hold with respect to them.

1.B Quantal Response Equilibrium

As for many other games, the experimental data obtained from participation games does not always coincide with NE predictions (e.g. Schram and Sonnemans 1996a). An equilibrium concept that usually performs better in explaining experimental data is Quantal Response Equilibrium (QRE, McKelvey and Palfrey 1995). QRE has shown to predict the data obtained from tests of political participation models particularly well (Goeree and Holt 2005).

QRE is an extension of NE that accommodates bounded rationality, and is therefore more attuned to the assessment of laboratory data. QRE extends the payoffs of the game by an additive stochastic component, which can be seen as a way of incorporating statistical noise into players’ choices. Subjects facing somewhat complex decisions in an unfamiliar environment (the laboratory) are prone to making mistakes. An equilibrium concept that models potential mistakes explicitly can thus

help produce more accurate predictions. In a QRE, best responses are played with higher probability than worse responses, but not with certainty, as in NE. In other words, best response functions, which are deterministic in NE, become probabilistic in QRE.

The main advantage of QRE within our methodological framework is better predictive power. A secondary advantage has to do with equilibrium selection. As mentioned above, participation games typically have several NE. In a game where several classes of NE exist, QRE helps to select the one which tends to have high empirical verification. A third advantage is that QRE retains most of the important features of NE, e.g. the probability of choosing a certain action is increasing in the payoff difference to the alternative(s), beliefs are consistent in equilibrium, etc.

In what follows I provide a sketch of the procedure. As we saw, in a NE player i participates whenever $E[U_i^{Part.}] \geq E[U_i^{Abst.}]$, given the actions of all other players $j \neq i$. We add a stochastic element to each expected utility term, $\mu\varepsilon_i^{Part.}$ and $\mu\varepsilon_i^{Abst.}$, respectively, which can be regarded as a random utility shock (e.g. due to unobserved preferences) or bounded rationality (e.g. mistakes on the path to equilibrium). In a QRE, player i will participate whenever $E[U_i^{Part.}] + \mu\varepsilon_i^{Part.} \geq E[U_i^{Abst.}] + \mu\varepsilon_i^{Abst.} \Leftrightarrow (E[U_i^{Part.}] - E[U_i^{Abst.}]) / \mu \geq \varepsilon_i^{Abst.} - \varepsilon_i^{Part.}$. For a given μ and an assumed distribution of $\varepsilon_i^{Abst.} - \varepsilon_i^{Part.}$, the probability of participation is $p = F[(E[U_i^{Part.}] - E[U_i^{Abst.}]) / \mu]$, where $F(\cdot)$ is the cumulative distribution function of $\varepsilon_i^{Abst.} - \varepsilon_i^{Part.}$.

The parameter μ governs the extent of bounded rationality (noise) in players' decisions, and the ε_i represent i.i.d. realizations of a random variable. The parameter μ and the distribution of the ε_i are the primitive elements of QRE. The parameter governing the amount of bounded rationality, μ , determines how the participation decision responds to expected payoffs. The equilibrium predictions will change for different values of μ . This value is usually estimated from experimental data obtained in identical or similar conditions. A value of $\mu = 0$ corresponds to a situation of no stochastic noise in decisions, i.e. full rationality. A NE results in this case. As $\mu \rightarrow \infty$, the amount of noise increases and players become indifferent between alternatives, playing each with equal probability. Following Goeree and Holt (2005) and Levine and Palfrey (2007), among others, I assume that $\varepsilon_i^{Abst.}$ and $\varepsilon_i^{Part.}$ follow independent extreme value distributions with parameter μ , which results in the difference $\varepsilon_i^{Abst.} - \varepsilon_i^{Part.}$ following a logistic distribution. McKelvey and Palfrey (1995) refer to this as a logit equilibrium, or logit QRE.

In a logit QRE a player's probability of participation is

$$\begin{aligned}
p_i &= \frac{\exp\left(\frac{E[U_i^{Part.}]}{\mu}\right)}{\exp\left(\frac{E[U_i^{Part.}]}{\mu}\right) + \exp\left(\frac{E[U_i^{Abst.}]}{\mu}\right)} \\
&\Rightarrow p_i = \frac{1}{1 + \exp\left(\frac{E[U_i^{Abst.}] - E[U_i^{Part.}]}{\mu}\right)} \\
&\Rightarrow p_i = \frac{1}{1 + \exp\left[\frac{c - \frac{(BW - BL)}{2}(\Pr[m_i = m_j] + \Pr[m_i = m_j - 1])}{\mu}\right]} \quad (1.4)
\end{aligned}$$

This formulation closely follows the standard approach in modeling discrete probabilistic choice (see McFadden 2001 for an overview).

Chapter 2

Other-regarding Preferences, Group Identity and Political Participation¹

This chapter studies the relationship between other-regarding preferences, group identity and political participation. In doing so, we propose a novel group identity induction procedure that succeeds in creating environments in the laboratory where in-group bias is either high or low. At the individual level, we find that both altruistic subjects and group identifiers participate above average. The most competitive subjects participate much less often than other types, while the most altruistic subjects manage to sustain high participation levels. At the aggregate level, participation is higher in environments where group identity is high than when it is low. This suggests that the higher participation observed in field settings for close-knit (political) groups might be due to a heightened sense of group identification as opposed to being solely caused by group mobilization processes.

2.1 Introduction

In June 2012, the Spanish village of Guijo de Galisteo held an unusual referendum. Located in a region where unemployment rates topped 30%, the mayor decided to let the population choose how to use the 15,000 Euro of the municipality budget traditionally allocated to summer festivities: either fund a series of bullfights or create new jobs. Typically, a minority of Spanish citizens appreciates bullfighting.²

¹This chapter is based on Robalo et al. (2013).

²A 2008 poll by the newspaper El País finds that only 37% of the Spanish claim to appreciate bullfights (<http://goo.gl/1vcDdE>). Approval rates are relatively homogenous across Spanish regions according to a Gallup poll conducted in Spain (<http://goo.gl/Z8fcCY>).

Yet, with a recorded turnout of 24% (in an electorate of 1764 voters), bullfight partisans outvoted the supporters of job creation by a margin of approximately 14 percentage points. We thus conjecture that supporters of the bullfight proposal turned out at a higher rate than those in favor of spending the money on job creation. What could cause such a difference in turnout?

Consider what may have motivated citizens to go out and vote in this referendum. Think first about the citizens who support the proposal to spend the money on job creation. A motivation to make the effort to vote may be that one is unemployed, but it is also possible that employed voters wish to aid their fellow unemployed citizens. In other words, pro-job voters may have decided to cast a vote either out of self-interest or as a consequence of other-regarding preferences. Hence, individual preferences may affect the decision to vote or abstain. On the other hand, it seems less straightforward to attribute casting a vote in favor of bullfights to other-regarding preferences. One may however speculate that the bullfight partisans participated in higher numbers due to a relatively higher sense of group identity, irrespective of their motivation. If a sense of group identity provides a motivation to go out and vote, a minority may manage to outnumber the majority at the polls.

As illustrated by this example, field data are likely to provide us with interesting case studies. Yet, though we may ‘conjecture’ that bullfight supporters voted at higher rates because of a sense of group identity, our conclusions are likely to be muddled by unobservability and endogeneity problems. Moreover, the arguments put forward with respect to preference types and group identity are interrelated. A careful study of these issues calls for more control than is available in the field. This chapter investigates the issues raised in this example using an experimental framework.

We study participation in the context of an interaction between groups competing for benefits. Participation can be anything from voting in a referendum to showing up at a rally. We address the question of how such participation is affected by the interaction between the two elements introduced in the example: a sense of group identity that members may have and the extent to which members have preferences that take into account the well-being of others. This is an important question because the outcome of group conflicts can have severe consequences for the members of the groups concerned, irrespective of an individual member’s decision to participate in it. If certain individuals or groups participate more than others, this might bias policies in a direction that is not representative of the majority’s preferences. If, for example, altruists participate more, policies may develop that are perceived to be more ‘altruistic’. Or, if some groups manage to create a stronger

feeling of group identity than others, this could put them in an advantageous position that is unrelated to the conflict at hand. Either of these effects could harm the efficient use of an economy's resources because they yield an allocation that is biased towards the preferences of the political participators (see Lijphart 1997 for a similar argument with respect to election turnout).

We conjecture that the individual participation decision in such environments takes into account the ties that bind the group. Moreover, an individual may more generally take the consequences for others into account when deciding on her actions. In other words, an individual may have other-regarding preferences, but these may be specifically directed towards her group's members and their strength may depend on the extent to which she identifies with the group. How much an individual cares for the individuals in her group and how much she cares about people in the other group are both likely to influence the sacrifices she is willing to make. This work addresses this conjecture by studying the effect of other-regarding preferences and group identity (and their interaction) on participation in group action.

An important goal of our experimental design is therefore to create environments with distinct levels of group identity in order to study its influence on individual and aggregate participation. In addition, we want to know whether participation depends on other-regarding motivations, both in general and in interaction with group identity. For this purpose the design includes a measurement of such motivations, using a so-called value orientation test. We measure political participation by studying individual choices in the participation game introduced in Chapter 1. Hence, our experiment induces distinct levels of group identity, measures other-regarding preferences and allows us to link (combinations of) these variables to political participation.

In order to derive hypotheses on individual and aggregate behavior, we combine insights from a theoretical analysis of the participation game with the available empirical evidence. First, we hypothesize that other-regarding subjects will participate more often than selfish ones. Second, we expect environments with a high bias towards the in-group to foster fiercer competition, and therefore generate higher aggregate participation. Third, we hypothesize that subjects who exhibit higher levels of identification with their group will participate more often.

Our results may be summarized as follows. First, they provide support for the hypothesis that individual participation is higher for other-regarding subjects. In particular, we observe that the most uncooperative subjects stand out from the rest by abstaining much more often. The estimated model predicts a 50 percentage points difference in participation between the most selfish and the most altruistic

subjects. Second, we were successful in inducing distinct levels of in-group bias across treatments, something that bears methodological relevance given the notorious difficulty of creating group identity in the laboratory (Eckel and Grossman 2005). This allows us to conclude that aggregate participation is higher in environments where group identification is high, albeit modestly. Third, regarding the individual impact of group identity, individuals who tend to feel more attached to their group in the first place participate more often. Our experimental inducement of further group identity crowds out this relationship, however.

The organization of this chapter is as follows. The next section discusses the literature that relates political participation to other-regarding preferences or group identity. Section 2.3 presents the conceptual analysis of the participation game and derives our hypotheses. Section 2.4 describes our experimental design. In section 2.5 we present and analyze our data. A final section concludes.

2.2 Related Literature

Both other-regarding preferences and group identification have been the subject of recent attention within the rational choice approach to political participation. This approach has traditionally struggled with the so-called ‘paradox of participation’³, but has by now uncovered various factors that help explain why rational individuals may participate in a group effort (see Palfrey 2009 for a recent overview).

The addition of other-regarding preferences to the calculus of participation has led to models that escape the prediction of low mobilization. For individuals with such preferences participation becomes instrumentally rational if the benefits derived from one’s group winning (which now include the benefits to co-members) are not overcome by the low probability of being pivotal. Models in this vein have been proposed by Jankowski (2002), Edlin et al. (2007), Feddersen et al. (2009), and Evren (2012). There is some field (Knack 1992, Jankowski 2007) and experimental (Fowler 2006, Fowler and Kam 2007, Dawes et al. 2011) evidence supporting a positive relationship between social preferences and participation. Our results add to this stream of evidence by relating a direct measure of an individual’s level of other-regarding preferences to the frequency of participation in intergroup competition. Moreover, we contribute with novel evidence on the interaction between an individual’s other-regarding concerns and the extent to which she identifies with her co-members. To some extent, this analysis supplements the work conducted by

³This paradox confronts empirically observed high rates of participation (e.g., in large-scale elections) with the theoretical observation that participation is seemingly irrational.

Fowler (2006), who uses a combination of field and experimental data to show that social identity (proxied by party identification) acts as a catalyst on the positive impact of altruistic motivations on political participation. Though the first to interact other-regarding preferences and group identity, Fowler's methodology has some shortcomings that are mainly related to the lack of control in the field.⁴ Our laboratory control allows us to measure other-regarding preferences and induce group identity in ways that rule out priming and response bias effects that are likely to occur in a situation where measurements are based on survey questions in a political context.

As for group identity, we first note that groups act as agents of intentional and unintentional mobilization (e.g., Pollock 1982) and therefore play a crucial role in the study of political participation. In this chapter we are especially interested in the impact of different levels of group identification on the spontaneous mobilization of groups. There is a large literature dealing with the relationship between material resources and spontaneous voter mobilization (e.g., the seminal work by Verba and Nie 1972). The empirical literature has established a number of solid findings, such as a positive correlation between income and participation. However, some puzzles remain. For example, the positive correlation between income, education and participation is much weaker for African-American voters, who participate beyond what their socioeconomic status would predict. Leighley and Vedlitz (1999) provide a number of candidate explanations for this phenomenon: psychological resources (e.g. political interest and participation efficacy beliefs), social connectedness, and group identity. All of these explanations have theoretical appeal and have received some empirical support. However, the mechanisms at work are difficult to identify in the field. For example, it is hard to disentangle the effect of group identity from the impact of social connectedness on participation. Do the members of a group voluntarily participate because of their strong sense of group-belonging, or because their environment encourages participation? There is a large literature studying the effects of social context and networks on turnout, whose main conclusions are that participative social environments induce higher individual participation (Kenny 1992), and that exposure to similar views within one's social network leads to higher individual participation (Mutz 2002). Notwithstanding, the magnitude of the effect

⁴Fowler uses survey questions regarding election participation, party identification and political knowledge. Subsequently, subjects play a dictator game, either against someone with the same political preference, with a different political preference or an unknown preference. These dictator choices are poorly incentivized, however. The observed distribution of giving is at odds with recent meta-studies (Engel 2011), but in line with non-incentivized studies. His results show that more generous people do not participate more, but those who are strong party identifiers and also more generous do participate more.

of a citizen’s social network on his participation decision is relatively small (Kenny 1992), and social interaction and social pressure mediate the impact of social group membership on individual political behavior.

This overview shows how difficult it is to isolate the effects of group membership on the participation decision. Furthermore, social context, social networks and participation behavior might be endogenously determined, making it difficult to elicit the direction of causality. In contrast, an easy and clean test of the group identity effect can be obtained in the controlled laboratory environment. By comparing the behavior of groups that differ only with respect to their sense of group-belonging we can isolate the effect of group identity on participation.

The so-called group identity paradigm studies the influence of ‘group-belonging’ sentiments on how individuals make decisions in instances of intergroup behavior. In short, it studies the “group self” (Ellemers, 2012) or the influence of the “group in the individual” (Hogg and Abrams 1988). Tajfel (1982) gives two criteria that must be satisfied for group identification to arise; the first is cognitive, and requires that members are aware of group membership; the second is evaluative, in that “awareness is related to some value connotations”. Group-based behavior then arises from categorization processes that partition the social world into an ‘in-group’ and an ‘out-group’; relative attachment to the in-group over the out-group (the ‘in-group bias’) is then assumed to drive intergroup relations. The body of knowledge on group identity that has developed over the past few decades is quite extensive and has produced a number of robust findings (see Brewer 2007). Recent experimental work has shown that group identity and its salience impacts strategic behavior (Charness et al. 2007) and that individuals tend to be more altruistic towards in-group members (Chen and Li 2009), for example. Our experimental design builds on the accumulated knowledge because it allows us to create an environment where group identity can be induced, measured, and controlled. Moreover, it allows us to investigate the role of group identity in an environment of intergroup competition. To the best of our knowledge, we are the first to do so in a controlled laboratory environment.

2.3 Conceptual Framework and Hypotheses

We study participation behavior using the participation game detailed in the appendix of Chapter 1. This section provides an outline of this framework’s application to our research questions and the main results that follow from the implementation (Appendix 2.A presents a detailed analysis).

We assume that players have a utility function that allows for other-regarding (altruistic) and group-discriminating components:⁵

$$U_i = u_i + \alpha_i \left(\beta_i \sum_{h \in \{G_i \setminus i\}} U_h + \gamma_i \sum_{k \in G_j} U_k \right) \quad (2.1)$$

where u_i is the individual material payoff, α_i is the weight put on other players' welfare, $\beta_i \geq 0$ is the weight put on the welfare of players in the same group, $\gamma_i \geq 0$ is the weight put on the welfare of players in the other group, G_i is player i 's group (the in-group), and G_j is the group against which i 's group competes (the out-group). These preferences express an interdependent utility function which is increasing in other individuals' utilities, but which allows the utility of individuals in the in-group to be given higher weights. We normalize β_i and γ_i such that $\beta_i + \gamma_i = 1$, which is possible to obtain from any β'_i and γ'_i : $\beta_i = \beta'_i / (\beta'_i + \gamma'_i)$ and $\gamma_i = \gamma'_i / (\beta'_i + \gamma'_i)$.

In equation 1.3, the benefits of winning and losing the participation game were B^W and B^L . At this point we make them more precise. Following the notation of Appendix 1.A, let the superscripts "W" and "L" denote being on the winning or the losing side, respectively. Accordingly, we define U_i^W as the utility benefit accruing to a player who has utility function 2.1 and is on the winning side. We define U_i^L in an analogous way. The theoretical analysis is thus carried out by incorporating B^W and B^L into U_i^W and U_i^L , and by equating $B^W = u_i^W$ and $B^L = u_i^L$. In words, B^W is the individual material payoff accruing to a player; this would be all a player would take into account if she had selfish preferences, as in Appendix 1.A. The participation game is extended to allow for the preference structure posited in 2.1, which results in an equilibrium condition similar to equation 1.3:

$$\Pr[m_i = m_j] + \Pr[m_i = m_j - 1] = \frac{2c}{(U_i^W - U_i^L)} \quad (2.2)$$

As we can see from 2.2, for constant c , the equilibria will be a function of the cost-benefit ratio, which in turn depends on α , β and γ (in addition to B^W and B^L , which are also held constant). For example, for a situation where the in-group is preferred to the out-group ($\beta > \gamma$) the utility difference is increasing in α .

As mentioned in Appendix 1.A, the analysis of participation games usually starts with quasi-symmetric Nash equilibria, i.e. equilibria in which all members of a group employ the same strategy. Given that our preference structure is richer than that of most previous studies, it is necessary to derive equilibria in which probabilities may

⁵Henceforth, we use the terms 'other-regarding preferences' and 'altruism' interchangeably, as altruistic preferences are the ones we conceptualize in our theory and measure in our experiment.

differ across players. Allowing for heterogeneity is another source of Nash equilibria multiplicity. We therefore derive QRE and present the results in Appendix 2.A, which help us inform the equilibrium behavior of players and put forward a number of hypotheses.

First, consider the relationship between an individual’s altruism level (as measured by α) and her participation decision. Intuitively, we expect individuals with a higher sense of altruism to be more willing to sacrifice themselves for their group, provided they prefer the in-group to the out-group, which is a weak assumption. This is another way of saying that there is more at stake for an individual who values the welfare of others in her group, and therefore stronger altruism will lead to more frequent participation.⁶ The theoretical analysis of the game indeed provides evidence that the (quantal response) equilibrium level of participation is increasing in other-regarding concerns (α) in a broad parameter range, including parameters that are compatible with our data. The existing empirical evidence provides further support for the conjecture that other-regarding concerns foster individual participation. Relating self-stated motivations to participation game behavior, Schram and Sonnemans (1996b) found that subjects with individualistic goals were less likely to participate, whereas subjects with cooperative goals were more likely to participate.⁷

This yields our first hypothesis:

Hypothesis 1: *Individual participation is increasing in the level of other-regarding concerns, i.e. more altruistic subjects participate at higher rates.*

Next, we consider the effects of group identity on participation. We are mainly concerned whether participation is higher in scenarios where in-group bias is more pronounced. The QRE that we obtain show that aggregate participation is increasing in in-group bias levels. This is supported by the empirical regularities mentioned in the previous section, in particular the fact that in-group favoritism leads to more cooperation within the group. We therefore expect higher aggregate participation when group identification is induced. In line with this conjecture, Schram and Sonnemans (1996b) studied the effect of group identity on participation behavior by

⁶Extremely altruistic individuals, i.e. those who value the welfare of an anonymous person more than their own, may prefer the out-group to win. In such cases, non-participation is the individual best-response. However, such levels of altruism are hardly observed in the laboratory, let alone in the field.

⁷In fact, this is precisely what experimental subjects will tell you. The post-experiment questionnaire asked subjects what they thought moved a participant who participated often. More than 70% responded that this was either cooperation towards the in-group or cooperation towards both groups. Moreover, a participant who participated rarely was attributed a selfish motivation by 77.5% of the subjects. For details, see Table 2.7 in the Appendix.

implementing different matching protocols in a participation game.⁸ They induce group identity using the minimal group paradigm and find that the effect of group identity is significant, though not pronounced. Moreover, various studies using the participation game framework (Bornstein et al. 1989, Bornstein 1992, Schram and Sonnemans 1996a,b, Goren and Bornstein 2000) have experimentally explored the role of communication within the in-group. We conjecture that the exchange of non-binding promises (cheap talk) between group members reinforces the sense of group identity. These studies have shown that such communication significantly increases participation levels.⁹ This allows us to formulate our second hypothesis:

Hypothesis 2a: *Higher in-group bias leads to higher levels of aggregate participation.*

We further consider situations where individuals within groups are heterogeneous in terms of their in-group bias. This allows us to address how group identity operates at the individual level. Do subjects who feel more attached to their group participate more often than subjects with lower levels of group identity? The theoretical results show that subjects who identify more with their group will tend to participate with a higher probability in a relevant parameter range. Intuitively, the reason may be that individuals who identify more with their group are more willing to incur sacrifices for it, and therefore participate at higher rates. Our third hypothesis follows:

Hypothesis 2b: *Individual participation is increasing in group identification, i.e. subjects exhibiting higher in-group bias participate at higher rates.*

2.4 Experimental Design

Our experiment is composed of three main parts, each to be explained in detail below.¹⁰ In the first part, we measure subjects' other-regarding preference type using a value orientation test commonly referred to as the 'ring test'. This test measures how a subject trades off her own welfare for the welfare of another individual. In the second part we manipulate the group attachment of subjects in order to obtain

⁸The authors implement three treatment conditions which were conceived to yield increasing levels of group identity: i) group composition varied from period to period, and both subject identity and choices were anonymous; ii) group composition remained constant, and both identity and choices were anonymous; iii) group composition remained constant, identity was revealed but choices remained anonymous. Participation in ii) was higher than in i) but also higher than in iii).

⁹Goren and Bornstein (2000) show that without communication players associate high participation levels to cooperation towards the in-group and do not associate low levels of participation to inter-group cooperation.

¹⁰See Appendix 2.E for a transcript of the instructions.

environments where the in-group bias is either high or low. Additional allocation decisions and survey questions were used to measure the degree of in-group bias. In the third part, subjects interact in the participation game. The value orientation test and the participation game are identical across treatments, whereas the group identity induction procedure varies in order to implement different treatments.

The experimental sessions were run at the CREED laboratory of the University of Amsterdam (UvA). Participants were recruited from the CREED laboratory subject pool using an online registration system. The subject pool consists of approximately 2000 students, mainly UvA undergraduates from various disciplines. A total of 160 subjects (44% of which were female) participated in 8 sessions (3 each for two group identity treatments, 2 for a control treatment), which took place in June and October 2011. On average, participants earned 28.5 Euros, which included a 7 Euro show-up fee. The experiment was programmed and conducted in z-Tree (Fischbacher 2007). Payoffs in the experiment are expressed in tokens, exchanged to Euros at a rate of 0.005 Euros per token. For the first and third part (ring test and participation game) we administer practice questions that check subjects' understanding of the tasks. The typical experimental session lasted around two hours.

We will first describe the two parts that are common to all treatments (see Figure 2.1 for a diagram showing the sequence of parts in the experiment). The test used to measure other-regarding preferences uses decomposed games (Liebrand 1984), a tool applied by social psychologists to assess an individual's social motives. More precisely, the ring test measures the rate at which an individual trades off her own welfare for the welfare of another individual. For a discussion of this test see Liebrand (1984) or Offerman et al. (1996). The version used in the experiment was proposed by van Dijk et al. (2002). Each participant is anonymously paired with two other participants; her choices affect one of them, and the choices of the other one affects her in an identical way. The two participants with whom a subject is paired remain constant throughout the first part of the experiment, but subjects are informed that identities will not be disclosed at any point. The test consists of 32 pairwise dictator choices, each presenting the participant with two alternative own-other allocations of monetary payoffs (see Appendix 2.B for a list of all choices). For each pairwise choice the participant has to pick her most preferred allocation. Each of the 32 pairwise choices is shown on the screen, both in text and bar graphics. Participants are informed that they will only learn the earnings or losses from this part of the experiment at the end of the session.

For the participation game, individuals are allocated to groups of five participants, with two groups constituting an 'electorate' of ten participants. The pa-

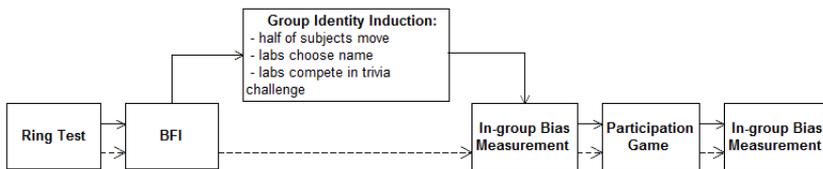


Figure 2.1: Sequence in the experiment

Notes: Solid lines indicate the sequence in the group identity treatments (see below), while dashed lines indicate the sequence in the control treatment. BFI: Big Five Inventory.

parameter values used throughout the second part of the experiment are $B^W = 120$, $B^L = 30$, and $c = 30$. Groups remain constant and repeatedly play the game for 40 rounds. At the end of each round, participants are informed of how many others participated in each group, their own earnings (in tokens) in that round, and their cumulative earnings (in tokens) in that part of the experiment.

Our treatments consist of variations in the second part of the experiment, where groups are formed and group identity is induced. A crucial choice concerns the variable (characteristic) used to differentiate between groups. The minimal group paradigm has shown that, in some situations, a mere awareness of belonging to a group, together with group competition for a prize, generates behavior consistent with group discrimination (Diehl 1990). In a laboratory setting the minimal group paradigm has not always been successful in producing such results, as pointed out by Charness et al. (2007). For one, the salience of groups in the laboratory, where interaction takes place via computers, is low. In order to induce strong feelings of group identity we therefore need to both differentiate groups along a dimension that matters to subjects and make this difference salient.

The number of potential differentiating variables is quite vast. Political groups may differ along many dimensions, including (but not limited to) ideology, income, education, religion, occupation and race. The relevance of specific variables depends on the political situation one is interested in. For example, opposing groups in a general election may differ along different dimensions than groups on either side of a gun rights rally. To avoid obvious links to specific group conflicts, while using a variable that bears relevance for political choices, we distinguish between groups based on a personality trait: openness to experience (openness, for short).¹¹

¹¹We measure personality traits under the Big Five taxonomy. This framework has shown that factor analysis performed on a multitude of personality data generally recovers five personality characteristics (openness, conscientiousness, extraversion, agreeableness and neuroticism). We

The relationship between personality traits and political ideology has been widely studied. The literature has reached a broad consensus in that liberals tend to score higher than conservatives on self-report measures of openness as measured by the Big Five (Carney et al. 2008 and the references therein).¹² These authors further show that the distinction between liberals and conservatives in terms of self-reported openness translates into “objective behavioral indicators” associated with openness, namely nonverbal behavior in a conversation (facial expressions, nonverbal signals, interaction style) and the contents of personal bedrooms and work offices (furniture and decoration style, personal belongings). For example, liberals tend to smile more during a conversation, while the bedrooms of conservatives tend to look more organized. Jost (2006) uses American state-level personality estimates to show that openness scores were the strongest regional personality predictor of the state vote share cast for Democrats and Republicans in the Clinton-Dole, Gore-Bush and Kerry-Bush races. Jost et al. (2003) show that the relationship between openness and ideology extends to non-American samples.

In sum, openness is one of the best proxies for ideological dispositions and has been shown to matter for political choices. Groups with contrasting openness score composition are thus composed of individuals who would make different political choices, and therefore draw a parallel to what would distinguish groups in most political conflicts. The obvious advantage of using a personality trait instead of self-reported ideology is to avoid the confounds implied by the meaning of ideology at a certain point in time or within a particular party system.

For our inference to be valid, openness should neither be correlated with other-regardingness nor with choices in the participation game.¹³ Table 2.8 in Appendix

measure personality traits using the Big Five Inventory (John et al. 2008), a highly validated questionnaire consisting of 44 short sentences based on trait adjectives known to be prototypical markers of the Big Five. This test provides a 1-to-5 score of each personality trait. A clear advantage of using a Big Five personality trait in our context is that, as stressed by Gerber et al. (2011), relative to other psychological constructs “the Big Five are measured with minimal references to political content, and are therefore less likely to be confounded by the political outcomes they may predict.”

¹²There is also some evidence that conservatives tend to score higher on conscientiousness.

¹³Some studies have investigated the relationship between openness and political participation. The overall evidence does not support a robust relationship between the two. Mondak et al. (2010) have shown that openness correlates positively with many forms of political participation if one does not account for political knowledge. In other words, individuals with higher openness scores tend to seek more information, and therefore participate more in politics; however, openness is not causal, as controlling for political knowledge produces a non-significant relationship between participation and openness. Gerber et al. (2011) corroborate this, as they find no relationship between openness and recorded voter turnout. Openness is significantly correlated with reported, but not with measured, turnout (and even then modestly and with marginal significance), campaign participation and local political participation (contacting a local officer and attending or speaking at a local political meeting). The latter forms of participation involve incentives that differ from

2.C provides statistical evidence that complies with both requisites. Namely, only Agreeableness seems to be significantly correlated with participation behavior, and no personality trait seems to be significantly correlated with other-regarding preferences as measured by the ring test.

For all treatment configurations, the group identity induction procedure in the second part of the experiment starts by requesting participants to answer the Big Five Inventory (BFI). Subsequently, they are explained what openness is and how different openness scores translate into personal characteristics and behavior. They are also told their own score. This allows us to implement three treatment configurations: high group identity (High), low group identity (Low), and no group identity (Control). All procedures are known to participants.

In the High and Low treatments, the half of participants whose openness scores are highest are asked to move to a second laboratory, while the half scoring lowest remain in the laboratory where the experiment started.¹⁴ After all participants have settled at their new computer stations, they are asked to jointly decide on a name to identify their laboratory. Participants are presented with three pre-determined options. They can discuss their choice with the other participants in the same laboratory via a chat interface. Each participant submits a choice, and the most chosen option becomes the name that identifies their laboratory for the remainder of the experiment. Next, the two laboratories compete in a trivia challenge. Each participant is presented with five timed trivia questions; a correct answer is worth one point and an incorrect one is worth zero points. The individual scores are aggregated by laboratory, and the laboratory with the highest score wins 2000 tokens to be equally distributed among its members.

The High and Low treatments differ only with respect to the groups that interact in the participation game (see Figure 2.2). To start, all members of a group are in the same laboratory. In the High treatment, a group of five participants interacts with a group of five participants drawn from the other laboratory. In the Low treatment, the interacting groups are drawn from the same laboratory. The Control treatment differs from High and Low in that the group identity induction procedure consists of the BFI alone (with an underlying group assignment protocol that mimics High).¹⁵ Hence, subjects in the Control treatment do not know that they are allocated to

the ones found in the type of situation under study - competition between groups.

¹⁴The second laboratory room is part of the same laboratory, and therefore is located close-by. Most subjects who stayed in the laboratory room where the experiment started also moved to a different computer station, such that all subjects in each laboratory are seated next to each other (separated by partitions).

¹⁵Since openness does not influence participation it does not matter whether the Control session mimics High or Low.

groups based on their openness score.

Finally, we use dictator allocation decisions to measure subjects' in-group bias. In particular, we asked each subject to divide 200 tokens between a random participant of his or her group (except himself or herself) and a random participant of the other group. This allocation decision was administered twice, just before and just after the participation game. In addition, the final questionnaire included an item for which subjects had to rate, on a 1-10 scale, how attached they felt to their own group and to the other group.

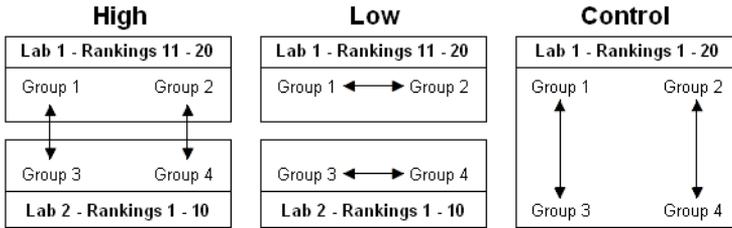


Figure 2.2: Experimental treatments

Notes: Arrows indicate competition in the participation game.

2.5 Experimental Results

Sub-sections 2.5.1 and 2.5.2 present preliminary steps to the analysis of our results, which is carried out in sub-sections 2.5.3 -2.5.5. In 2.5.1 we put forward a classification of subjects according to their other-regarding preference type. In section 2.5.2 we investigate the validity of our group identity manipulation procedure. In 2.5.3 and 2.5.4 we present results on bilateral relationships between other-regarding preferences, group identity and participation. These analyses provide partial support for our hypotheses. Stronger support is reported in section 2.5.5, where we present a multivariate analysis explaining the participation decision.

2.5.1 Subject Classification

Hypothesis 1 concerns differences in participation behavior across individuals with distinct other-regarding concerns. To enable this comparison, we divide subjects into categories representing different other-regarding preference types (henceforth, ‘types’). We start with a brief characterization of our measure. The application of the ring test to measure such preferences presupposes the existence of a motivational

vector for each subject. This vector represents the individual’s trade-off between own welfare and the welfare of another subject in a two-dimensional vector space, where one dimension indexes the own payoff and the other indexes the payoff accruing to the other individual (see Figure 2.3). Each own-other allocation can be represented by a vector whose origin coincides with the origin of the coordinate system. For each of the 32 pairwise choices in the ring test, a participant optimizes by choosing the allocation that is closest to her motivational vector. Averaging over an individual’s 32 choices yields an approximation of her motivational vector.

The ring test measures other-regarding preferences with respect to distributive outcomes. It does not take into account reciprocity concerns, but it can accommodate inequity-averse preferences as in e.g. Fehr and Schmidt (1999). For example, a subject who experiences no disutility from a disadvantageous position and places equal weight on the own payoff and the disadvantageous position of others (in Fehr and Schmidt’s terminology, $\alpha = 0$ and $\beta = 1$), has a motivational vector with slope of 1. However, the ring test’s power is limited with respect to the identification of inequity-aversion parameters, as its focus (and strength) lies in measuring how a subject trades off her own payoff for that of an anonymous subject.

A subject’s motivational vector can be fully described by its length and direction. The length can be interpreted as the degree of choice consistency. We will restrict our sample to the 152 subjects (95% of the total) with a reasonable degree of consistency.¹⁶ The slope of the motivational vector – which can also be expressed as the angle formed by the vector and the horizontal axis – describes how the individual trades off own welfare for the welfare of others.¹⁷ For example, one can think of an individual whose vector has an angle of 26.6° – corresponding to a slope of 0.5 – as someone willing to give away 50 Euro cents to another individual for each Euro she keeps for herself. The slope of the vector provides a measure of α in equation 2.1: the marginal rate of substitution of i ’s utility of money for j ’s utility. The average angle of the motivational vector in our sample is 6.77° , with a standard deviation of 19.98° ; the two widest angles are -68.93° and 73.25° .¹⁸

The ring test typically makes use of a standard set of categories to classify individuals (Liebrand 1984), assigning them one of six labels (‘aggressive’, ‘competitive’,

¹⁶In our application of the ring test, each vector (allocation) has a length equal to 1000. If a subject always chooses the option closest to her motivational vector, its length is also 1000. We exclude from the sample subjects whose vector has length smaller than 600. For comparison, a random sequence of choices yields a motivational vector with length equal to 500. The same consistency criterion was used by van Dijk et al. (2002).

¹⁷In our analysis we use the average vector’s angle (and not the slope) to represent a subject’s other-regarding preference type.

¹⁸This corresponds to an average slope of 0.12 with standard deviation of 0.36. The steepest slopes are -2.60 and 3.32.

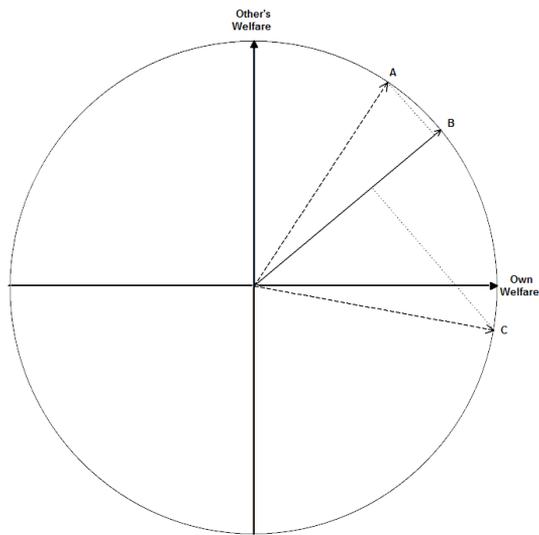


Figure 2.3: Typical choice in the ring test

Notes: Suppose B is the motivational vector of a given (rational) subject. When asked to choose between A and C, she will choose A as it is closest to her motivational vector (it has the highest projection on it).

‘individualistic’, ‘cooperative’, or ‘altruistic’). Each category corresponds to an area of the circle. One problem with this classification is that it makes for a poor distribution of data across categories, since individuals tend to massively concentrate on the categories ‘individualistic’ and ‘cooperative’. In our sample, 93.13% of subjects would fall within these two categories. We therefore put forward a new classification, which balances a good categorization of the data with an empirically relevant set of categories. This is presented in Table 2.1 and Figure 2.4.

	Angle(°)	Slope	% Subjects
Competitive	<0	<0	19.74
Individualistic	0	0	23.03
Weakly Altruistic	(0,8.53]	(0,0.15]	20.39
Mildly Altruistic	(8.53,21.8]	(0.15,0.4]	18.42
Strongly Altruistic	>21.8	>0.4	18.42

Table 2.1: Other-regarding preference types

Notes: Rows define the types we distinguish, based on angle (column 2) or slope (column 3) of the estimated motivational vector. The final column shows the distribution of subjects across types. The distribution of subjects across types per treatment is given in Table 2.10 in Appendix 2.C.

The first category (‘competitive’) comprises those individuals who are willing to sacrifice part of their gains to decrease the other individual’s earnings. The second is composed of individualistic types: subjects whose only motive is to maximize personal gains, regardless of the trade-off imposed on the other. In contrast, altruists are willing to give up some of their personal gains in order to increase the gains of an anonymous other individual. If the rate at which this sacrifice is made is below 0.15 we call them ‘weak’ altruists; if the rate is between 0.15 and 0.4 we call them ‘mild’ altruists; and if it is above 0.4 we call them ‘strong’ altruists. Of course, this classification is no less ad hoc than the standard one. To be sure, *a priori* there exist no meaningful cut-off values to distinguish between types.

2.5.2 Group Identity Induction

In order to assess the extent to which group identity feelings were successfully induced, and to know how they vary across treatments, we consider measured in-group bias. As explained in the previous section, our measurement was done in two ways: first, through two dictator allocations of monetary endowments between a random member of the in-group and the out-group (one before and one after the participation game); and second, through self-reported attachment to in-group and out-group (on a scale from 0 to 10). Figure 2.5 presents results from these measures.

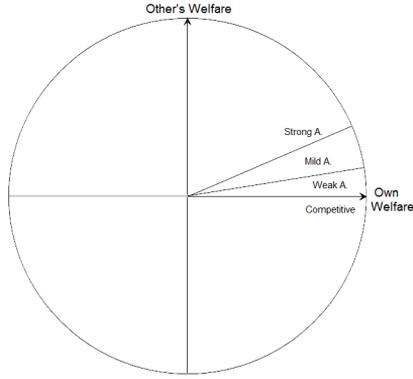


Figure 2.4: Other-regarding preference types

Notes: The type 'Individualistic' coincides with the horizontal axis.

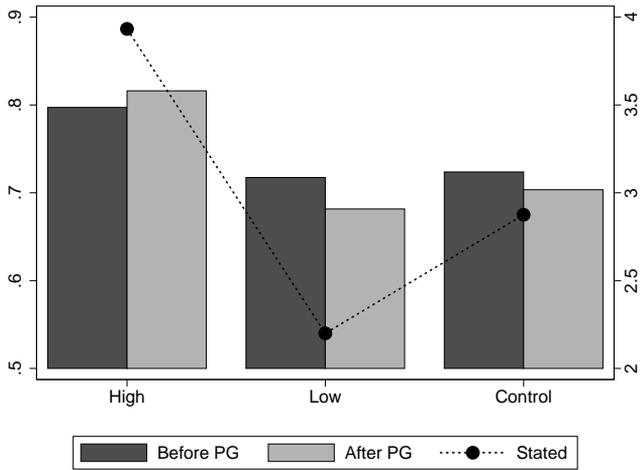


Figure 2.5: Group identity induction across treatments

Notes: Bars show the fraction of the endowment allocated to the member of the in-group (left axis). Dark gray (light gray) gives the measurement before (after) the participation game. Stated in-group bias is represented by the connected dots (right axis).

The percentage allocated to the in-group member is highest in the High treatment. This value is (marginally) significantly different than the percentages allocated in the Low and Control treatments, both before (two-sided Mann-Whitney U-test, $p = 0.06$ and $p = 0.08$; MW hereafter) and after (MW, $p=0.10$ and $p=0.03$) the participation game. The average of the two allocation decisions made by each subject yields significant differences between High and the other two treatments (MW $p=0.01$ and $p=0.03$ for comparisons with Low and Control, respectively). Allocation decisions in the Low treatment are not statistically different from those of Control (MW, $p > 0.75$ for separate and average comparisons). In the High treatment, a subject allocates approximately 80% of the total amount to the member of his or her own group before the participation game; in the Low and the Control treatments this figure is lower (approximately 72%). This pattern widens after groups have interacted in the participation game: the average allocation in the High treatment increases to 82%, while in the Low and Control treatments it decreases to 68% and 70%, respectively. Allocations before and after the participation game are not statistically different, neither overall nor for any specific treatment (MW, all $p > 0.59$). Finally, note that the results for the Low and Control treatments provide some support for the ‘minimal group paradigm’ (Tajfel 1982); subjects give more to the own group member than to someone from the other group, even when no group identity is induced.

The results of the allocation decisions are corroborated by the second indicator of group identity. In the questionnaire, subjects were asked to report their attachment to the in-group and the out-group on a 1-to-10 scale. Computing the difference between these two values yields a measure of self-reported in-group bias on a -10-to-10 scale (see Figure 2.5). Average stated in-group bias is 3.9 in High, 2.2 in Low, and 2.9 in Control, with standard deviations of 2.87, 3.42, and 4.50, respectively. The difference between High and Low is statistically significant (MW, $p=0.01$), while those between High and Control, and Low and Control, are not (MW, $p=0.37$ and $p=0.27$, respectively).

The purpose of group identity induction was to create distinct levels of group identity between the High and the Low treatments. The results presented in Figure 2.5 and the corresponding statistical tests show that the procedure was successful. This analysis is disaggregated for the different other-regarding preference types in Appendix 2.D.¹⁹

¹⁹With respect to other-regarding preferences and group identity, two questions can naturally be raised: what types are most likely to show a high degree of in-group favoritism, and what types are more likely to be influenced in their group attachment by interaction in the participation game. In sum, Appendix 2.D yields three main findings about the interaction between other-regarding

2.5.3 Other-regarding Preferences and Participation Behavior

We now turn to our main research questions, which is how political participation is affected by other-regarding preferences and group identity. We start by making the link between motivational vectors and choices in the participation game. Figure 2.6 presents average participation rates per other-regarding preference type throughout the participation game. We observe that competitive individuals clearly participate less often than any other type. The difference between the individual average participation of competitors and any other category is statistically significant (MW $p < 0.01$ for all comparisons).²⁰ As for the other categories, it is interesting to note that strong altruists participate less often than the other groups in the first twenty rounds, a tendency that is reversed in the second half of the participation game. Though this pattern is consistent with an initial attempt to lower the aggregate participation in the electorate in order to save on voting costs, the effect is statistically insignificant (there are no statistical differences in any pairwise comparison that does not include competitors).

Consistent with previous evidence, there is a tendency for participation levels to decrease as the game unfolds (e.g. Schram and Sonnemans 1996a). Regressing average participation on a linear trend per type yields a negative and significant relationship for all types except strong altruists, who exhibit a positive, albeit non-significant, increase in participation over time (see Table 2.9 in Appendix 2.C). We conclude that strong altruists are the only type that succeeds in stabilizing in-group cooperation over time.

In order to see in more detail how participation depends on other-regardingness, Figure 2.7 shows a scatter plot of individual participation rates within each type and a fitted least squares trend line. This shows that the relationship between the individual participation rate (i.e., the fraction of the 40 periods that a subject chose to participate) and other-regardingness (measured by the angle of the motivational vector) is positive for most categories (by definition, there is no such relationship for the Individualistic category). A regression of the participation rate on the degree of altruism produces a positive coefficient for each category, even though statisti-

preferences and in-group bias. First, on aggregate, in-group bias does not differ across types, except that Competitors are less prone to it than Individualists. Second, the differences in in-group bias we observe across treatments can to a large extent be attributed to altruistic types responding differently to distinct levels of induced group identity. Third, except for Competitors, in-group bias is not affected by the interaction with others in the participation game.

²⁰Non-parametric tests of type behavior use individual average participation as the observation unit. Non-parametric tests of aggregate behavior use the average participation of a pair of competing groups as the observation unit.

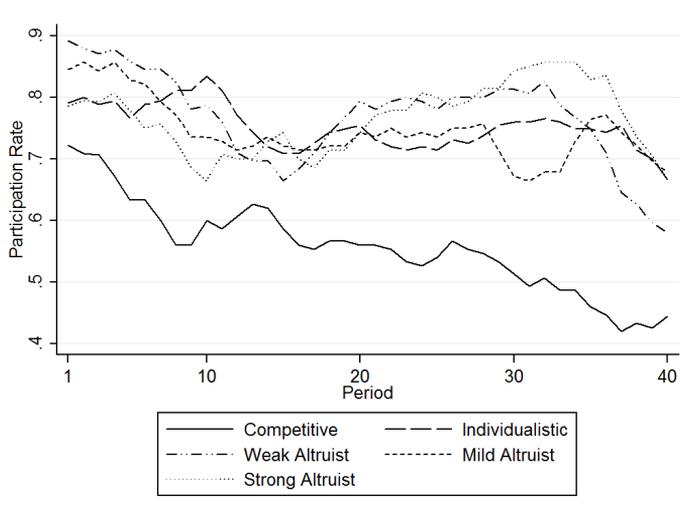


Figure 2.6: Evolution of participation and other-regarding preferences

Notes: Lines connect the average participation rates of participants of a given type, across all group identity treatments. Five-period moving average (two lags, two leads, where available) time series are plotted.

cal significance is only achieved when considering the full sample (see Table 2.2). As conjectured, individual average participation is increasing in a subject's other-regarding preferences. This result is corroborated by panel data regressions to be presented in subsection 2.5.5.

All in all, our analyses show that there is a positive relationship between other-regarding preferences and participation behavior. The effect is statistically strong at the aggregate level and appears to be present for each of the categories we distinguished. A pronounced difference is observed for the category of competitors, who significantly abstain more than other types. The evidence lends support to Hypothesis 1. Sub-section 2.5.5 provides a further robustness check on this relationship and elaborates on the magnitude of the observed effect.

2.5.4 Group Identity and Participation Behavior

Our second and third hypotheses concern the relationship between group identity and participation behavior. As formulated in Hypothesis 2a, we expect groups with stronger group attachment to exhibit higher levels of aggregate participation. At the individual level, Hypothesis 2b predicts that subjects reporting a higher sense

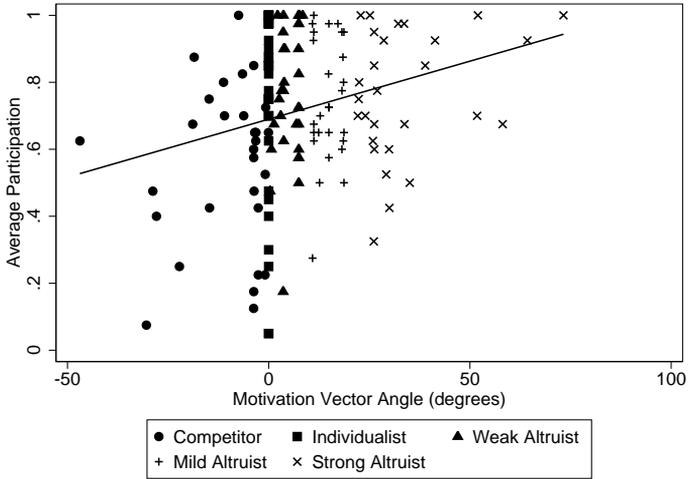


Figure 2.7: Individual participation rates and other-regarding preferences.

	All	Competitor	Weak A.	Mild A.	Strong A.
Motivational Vector ($^{\circ}$)	0.346*** (3.52)	0.220 (0.59)	1.840 (1.37)	0.412 (0.32)	0.383 (1.99)
Constant	0.690*** (31.49)	0.580*** (10.20)	0.672*** (8.28)	0.683*** (3.30)	0.635*** (7.62)

Table 2.2: Participation and other-regarding preferences

Notes: The first number in each cell reports the coefficients of an OLS regression, multiplied by 10^2 , where the dependent variable is the fraction of times participated; $N=152$ for “All”. Robust standard errors, t-statistics in parentheses. *** indicates significance at the 1% level.

of group identity participate more often.

Figure 2.8 shows aggregate participation levels for each of the three treatments across the forty periods of the participation game. Aggregate participation rates vary between 46% and 90%, with a higher overall participation rate in High (74.3%) than in Low (69.3%) and Control (69.2%). Participation variance is lowest in High, followed by Low and Control (standard deviations equal to 4.3%, 7.1% and 10.9%, respectively). The participation patterns we observe are similar to previous evidence in some respects: a decrease in participation as interaction is repeated, and an abrupt decline in the last couple of periods (see Schram and Sonnemans 1996a,b). In the first 10 periods participation is remarkably similar in the three treatments. After this point, we observe a departure of participation in High from the levels observed in Low, while Control exhibits a more erratic pattern. Despite the 5 percentage point difference between High and the other treatments, neither mean nor median participation in High is statistically different from Low and Control. Similarly, we observe no treatment differences in the last 20 or 10 rounds. Hence, based solely on a non-parametric analysis, we cannot reject a null hypothesis of no differences in favor of Hypothesis 2a.

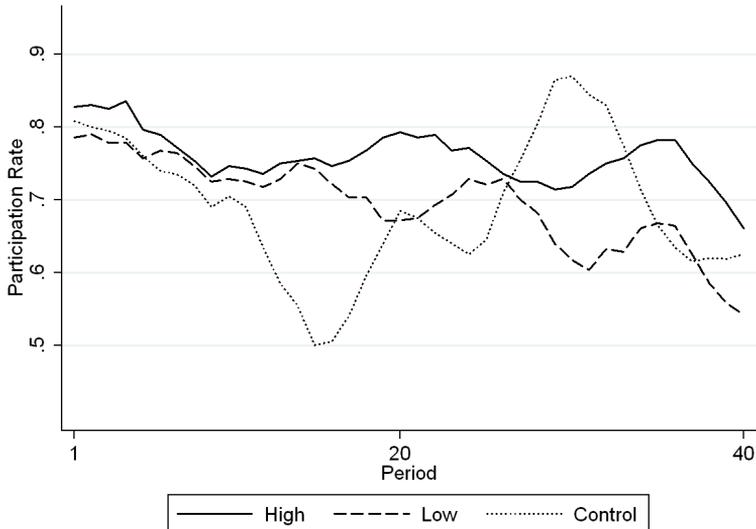


Figure 2.8: Participation across treatments

Notes: For each treatment the corresponding line shows the five-period moving-average participation rate (two lags and two leads, where available).

At the individual level, we first compare average participation across all treatments. Splitting the sample in terciles according to average allocation decision in the two dictator allocation tasks, we find that the first tercile (those with the lowest in-group bias) participates less often (69.6%) than the second (73%) and third (72.8%). The difference is not statistically significant, however. Considering each treatment separately, the only marginally significant effect is that in Control, the third tercile participates more than the first (MW, $p=0.08$). At this level of aggregation, we find no support for Hypothesis 2b. Below, we will see that more support for the alternative Hypothesis 2b is obtained when using a multivariate analysis of individual behavior.

The evidence we have provided so far has basically been founded on partial analysis with little room for interaction between the relevant variables. A more detailed analysis is carried out in the next subsection, where we present results from a multivariate regression model, which allows us to provide a more definite test of these hypotheses.

2.5.5 Multivariate Analysis

For the multivariate analysis of the participation decision we use a regression model that takes the panel structure of our data into account and corrects for individual heterogeneity in the participation game. We employ a logit specification in order to better accommodate those individuals with relatively extreme degrees of altruism and group identification.

Two points related to our empirical strategy are in order. First, since we observe decreasing participation rates across repetitions, with a steep decrease in the last few periods, we estimate models that include trend and squared trend terms. Second, we employ the average of the two allocation decisions as a measure for a subject's in-group bias (recall that there is a great degree of consistency between the first and the second allocation decision). Table 2.3 provides the results of a logit regression explaining the decision to participate. The trend coefficient is negative and significant, which reflects the decreasing participation rate in the later rounds of the game. We accept this result as given, but will elaborate on three other results that provide support for each of our hypotheses.

First, we note that the level of altruism as measured by the angle of the motivational vector significantly and positively affects the likelihood to participate. The 'average subject' (the one assigned the sample average value of each independent variable) is predicted to participate approximately 80% of the times by the model. A marginal increase in altruism leads to an increase in the probability of participating

equal to 0.36%-points, an effect that is statistically significant (Wald, $p < 0.01$). This means, for example, that an individual moving from the category Weakly Altruistic to Mildly Altruistic (a difference of approximately 10° in the ring test) increases the probability of participation by approximately 3.6 percentage points. The difference between the two widest vectors is 142.18° , which implies a predicted difference of 49.8 percentage points in participation probabilities. The significant effect of altruism on participation provides further support for Hypothesis 1. This effect is also observed when considering the High and Low treatments separately.²¹

Second, in support of Hypothesis 2a we observe a positive aggregate effect of group identity on participation, compared to the control treatment, that is significant at the 5%-level. This follows from the coefficient estimate for the dummy variable High. Hence, being in an electorate with high group identity raises everyone's probability of participating in group action, independent of the own level of identification with the group.

Third, the coefficients for the in-group bias are supportive of Hypothesis 2b (an individual effect) for the Control treatment. The effect is significant at the 10%-level.²² The effect is not significant in High or in Low, however. A similar regression for decisions in High gives a coefficient for in-group bias that is equal to -0.363 ($z = -0.64$) with a marginal effect of -0.046 ($z = -0.64$). For Low, the coefficient is 0.030 ($z = 0.06$) with marginal effect 0.005 ($z = 0.06$). We conclude that there is only a relationship between individual in-group bias and participation in groups where we have not induced group identity in any way. Indeed, in Control, the coefficient for in-group bias is 1.062 ($z = 1.83^*$) with a marginal effect of 0.189 ($z = 1.81^*$). We will return to this point in the concluding discussion.

Finally, we consider the participation levels implied by the levels of altruism and in-group bias in our sample, i.e. we derive the models' in-sample predictions. The results of this exercise are reported in Figure 2.9 and show that our data describe a clear relationship between other-regarding concerns and participation. The models' predictions based on in-group bias also show a positive, albeit less clear, relationship with participation.

²¹The effect is also positive for the control treatment, but not statistically different than zero at conventional levels.

²²If we use the alternative measure of group identity (based on the responses in the questionnaire), this effect is significant at the 5%-level.

	Coefficient	Marginal Effects
Motivational Vector (°)	0.022*** (2.74)	0.004*** (2.69)
In-group bias	0.980* (1.72)	0.158* (1.71)
High	1.096** (2.07)	0.162** (2.23)
Low	0.397 (0.89)	0.062 (0.92)
Trend	-0.032*** (2.66)	-0.005*** (2.61)
Trend ²	0.000 (1.19)	0.000 (1.19)
In-group bias*High	-1.336* (1.72)	-0.215* (1.72)
In-group bias*Low	-0.961 (1.31)	-0.155 (1.31)
Constant	1.083*** (2.99)	— —

Table 2.3: Panel regression model

Notes: Cells present the logit estimation (with random effects at the individual level) coefficients (column 2) and marginal effects (column 3); N=152. High and Low are dummy variables representing these treatments. In-group bias is measured as the average of the two dictator allocation decisions, re-scaled to the interval $[-1, 1]$. Absolute z-scores in parentheses. */**/** indicates significance at the 10%/ 5%/ 1% level. Marginal effects for the treatment and interacted variables are calculated by restricting the estimation sample accordingly.

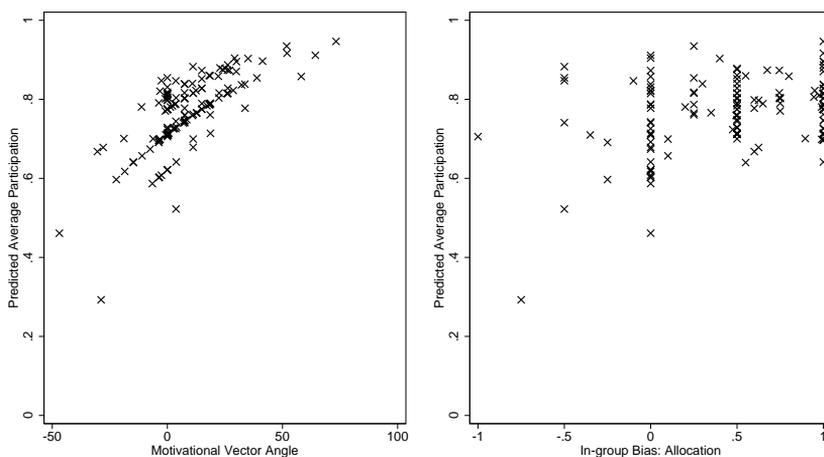


Figure 2.9: In-sample model predictions

Notes: Each data point represents a model prediction of individual participation for an individual for whom all variables take average sample values except for the variable in the axis (subjects displayed in ascending order of the relevant variable).

2.6 Conclusion

In his appraisal of the rational choice literature on election participation, Feddersen (2004) argues that “while a canonical model does not yet exist, the literature appears to be converging toward a ‘group-based’ model of turnout, in which group members participate in elections either because they are directly coordinated and rewarded by leaders as in ‘mobilization’ models or because they believe themselves to be ethically obliged to act in a manner that is consistent with the group’s interest as in ‘ethical agent’ models.” This chapter addresses the latter aspect of group membership and constitutes an attempt at providing evidence from a controlled environment that adds to the stream of literature that tries to evaluate political participation in light of other-regarding concerns and group-directed duties. In particular, we have used an experimental framework to address the influence of other-regarding motivations and group identification on political participation decisions.

The empirical literature on the impact of group membership on participation has shown that higher social connectedness leads to higher participation levels. However, the mechanism at work is difficult to identify in the field, in the sense that it is impossible to disentangle the impact of group-level mobilization processes from the impact of group identity concerns. Our laboratory control has allowed us to isolate

the latter mechanism.

We study political participation as a competition between groups, where victory depends on the sum of the individual efforts by the individuals in a group. Our experimental design allows us to measure the other-regarding concerns of subjects and implement environments with different levels of group identity.

Our main conclusions are that individual participation is increasing in other-regarding concerns and group attachment, as conjectured. We also found support for an impact of group identity on aggregate participation levels (but only in a multivariate analysis that corrects for the influence of confounding factors). This latter result implies that the higher participation levels observed in field studies for environments where group identification is high (e.g., contexts with pronounced ethnic divisions and high political participation) might be due to this heightened sense of group identification. Whether group mobilization adds something to this effect, is a question that is addressed to some extent in the next chapter.

Finally, there is only a correlation between individual-level sense of in-group bias and participation in our Control treatment, i.e., when we did not induce any group identity. In this case, people with a large bias towards the own group tend to participate more in political action. When we induce a high sense of group identity vis-à-vis the other group, individual differences still exist (at an overall higher level) but no longer matter for the participation decision. Similarly, when our procedures induce group identity towards both the own and the other group, differences still exist (at a lower level) but do not matter for participation. In other words, individual differences within a group matter only when people experience moderate differences between the groups.

All in all, we conclude that more other-regarding individuals participate more. Moreover, a common sense of identification with the group yields higher aggregate levels of political participation. As described in the previous paragraph, the effect of group identity is more complex at the individual level and depends on experienced differences between groups. Each of these results may serve as input in a canonical model as envisaged by Feddersen (2004).

Appendix

2.A Equilibria of the Participation Game

In this appendix we formally derive equilibria of the participation game, which allows us to obtain comparative static results that inform our hypotheses. We denote the in-group and the out-group of player i by G_i and G_j , respectively. We posit individual preferences that accommodate general altruism towards others as well as discrimination between in-group and out-group members:

$$U_i = U_i(u_i(x_i); U_h, h \in \{G_i \setminus i\}; U_k, k \in \{G_j\}; \alpha_i, \beta_i, \gamma_i) \quad (2.A.1)$$

In equation 2.A.1, U_i denotes i 's utility, x_i gives her monetary earnings and $u_i(x_i)$ describes her utility of wealth. α_i is a parameter describing the weight i attributes to others' utility, relative to her own, β_i is the weight she attributes to the utility of other members in her own group, and γ_i is the weight she attributes to the utility of members in the other group.

To derive comparative statics for the participation game, we make the following three assumptions:

1. If members of the own group receive higher utility from an outcome than members of the other group do, then more in-group bias leads to higher utility:

$$U_{h,h \in \{G_i \setminus i\}} > U_{k,k \in \{G_j\}} \Rightarrow (\partial U_i) / (\partial \beta_i) > 0 \quad (2.A.2)$$

2. If members of the own group receive lower utility from an outcome than members of the other group, then more in-group bias leads to lower utility:

$$U_{h,h \in \{G_i \setminus i\}} < U_{k,k \in \{G_j\}} \Rightarrow (\partial U_i) / (\partial \beta_i) < 0 \quad (2.A.3)$$

3. The utility derived from winning the participation game (and, as a consequence, $h \in \{G_i \setminus i\}$ also winning and $k \in \{G_j\}$ losing the participation game) is larger than the utility derived from losing the participation game (and, as a consequence, $h \in \{G_i \setminus i\}$ also losing and $k \in \{G_j\}$ winning the participation game):

$$\begin{aligned} U^W &\equiv U_i(x_i = B^W, x_{h,h \in \{G_i \setminus i\}} = B^W, x_{k,k \in \{G_j\}} = B^L) > \\ U^L &\equiv U_i(x_i = B^L, x_{h,h \in \{G_i \setminus i\}} = B^L, x_{k,k \in \{G_j\}} = B^W) \end{aligned} \quad (2.A.4)$$

where U_i^W (U_i^L) is the utility in case of victory (defeat). Note that 3. implies

an intuitive restriction on the parameters α_i and β_i , i.e., they are such that any individual prefers the own team winning the participation game to the other team winning.

Equations 2.A.2-2.A.4 yield:

$$(\partial U^W)/\partial\beta > 0, (\partial U^L)/\partial\beta < 0 \Rightarrow \partial(U^W - U^L)/\partial\beta > 0 \quad (2.A.5)$$

In words, 2.A.5 states that an increase in an individual's in-group bias will lead to a higher marginal benefit of her group winning the participation game.

Next, we need to determine how this increased marginal benefit affects the choice to participate. *Ceteris paribus*, this will yield a higher participation probability, simply because the benefits increase while the costs remain unchanged. This is not necessarily true in an equilibrium analysis, however, because other voters may respond to variations in an individual's incentives. We therefore proceed with equilibrium analysis. We assume complete information throughout: in addition to the rules of the game, monetary payoffs and group size, we assume that players know the utility functions of all other players. This simplification does not hinder the derivation of broad comparative statics and keeps the analysis tractable. The alternative would be to adopt incomplete information (i.e. players don't know other players' preference parameters), which would require further ad hoc assumptions on beliefs.

We use a utility function of the type defined in 2.A.1. The individual preferences put forward in equation 2.1, which we reproduce here, accommodate general altruism towards others as well as discrimination between in-group and out-group members (see Section 2.3 of the main text for an explanation of the notation):

$$U_i = u_i + \alpha_i(\beta_i \sum_{h \in \{G_i \setminus i\}} U_h + \gamma_i \sum_{k \in \{G_j\}} U_k) \quad (2.A.6)$$

The utility payoff depends on whether a player's group wins or loses the game: define $\Gamma = \{W, L\}$ as these two events. Given our preference structure, the payoffs of the game are interdependent across players. Assuming that i is in the winning group, equation 2.A.6 should be adjusted by substituting u_i for B^W and plugging in B^W and B^L where appropriate. The utility if i loses can be obtained in a similar fashion. For given preferences and for each case (winning or losing), this yields a system of $M + N$ equations (the individual utilities) in $M + N$ variables (the utility payoffs):

$$\begin{aligned}
[I - \Omega] u(\Gamma) &= b(\Gamma) \\
\Rightarrow u(\Gamma) &= [I - \Omega]^{-1} b(\Gamma)
\end{aligned} \tag{2.A.7}$$

where $I_{(M+N) \times (M+N)}$ is the identity matrix, $u(W) = (U_1^W, \dots, U_M^W, U_N^L, \dots, U_{M+N}^L)'$, $b(W) = (B^W, \dots, B^W, B^L, \dots, B^L)'$ and

$$\Omega = \begin{bmatrix} 0 & \alpha_1 \beta_1 & \cdots & \alpha_1 \beta_1 & \alpha_1 \gamma_1 & \alpha_1 \gamma_1 \\ \alpha_2 \beta_2 & & \ddots & & & \alpha_2 \gamma_2 \\ \vdots & & & & & \vdots \\ \alpha_M \beta_M & & & & & \alpha_M \gamma_M \\ \alpha_{M+1} \gamma_{M+1} & & & & & \alpha_{M+1} \beta_{M+1} \\ \vdots & & & & \ddots & \vdots \\ \alpha_{M+N} \gamma_{M+N} & \cdots & \alpha_{M+N} \gamma_{M+N} & \alpha_{M+N} \beta_{M+N} & \cdots & 0 \end{bmatrix}$$

The case $\Gamma = L$ is defined accordingly. The solution to 2.A.7 allows us to calculate the utility of a winning and losing player for any combination of B^W , B^L , other-regarding and in-group bias parameters (Ω), and to use these (together with c) to determine equilibria.

As discussed in Section 2.3, we compute QRE using this preference structure. However, QRE requires us to specify the noise parameter, μ (see equation 1.4). This parameter is typically estimated from experimental data. Goeree and Holt (2005) show that a value of $\mu = 0.8$ accommodates the data of Schram and Sonnemans (1996a), in which participation fluctuates in the 30 – 50% range. Since we observe higher participation levels, our data would possibly imply a slightly lower value of μ . For our purposes, the precise value of this parameter is not particularly relevant as only point predictions, and not comparative statics, will depend on it. For this reason, we use $\mu = 0.8$ for the numerical QRE results that follow.

2.A.1 In-group Bias and Aggregate Participation

A (mixed) strategy in the participation game is simply a probability of participation, which is denoted by p . To start, we consider totally quasi-symmetric equilibria (see Chapter 1) where all voters in G_i vote with the same probability, p_{G_i} , and all voters in G_j vote with the same probability, p_{G_j} . For participation games where both groups have equal size the probability terms are then defined as (for a player in G_i):

$$\Pr [m_i = m_j] = \sum_{k=0}^{M_i-1} \binom{M_i-1}{k} \binom{M_j}{k} (p_{G_i})^k (1-p_{G_i})^{M_i-1-k} (p_{G_j})^k (1-p_{G_j})^{M_j-k} \quad (2.A.8)$$

$$\Pr [m_i = m_j - 1] = \sum_{k=0}^{M-1} \binom{M-1}{k} \binom{M_j}{k+1} (p_{G_i})^k (1-p_{G_i})^{M-1-k} (p_{G_j})^{k+1} (1-p_{G_j})^{M_j-k-1} \quad (2.A.9)$$

The assumption that in equilibrium every player in the same group participates with the same probability can be intuitively justified by the assumption that players are homogenous in their other-regarding preferences. For our analysis, we further assume here that players in both groups have the same parameters (and therefore, $p_{G_i} = p_{G_j} = p$), which means that we will investigate how equilibria change when we vary the in-group bias parameters for all players. Our strategy is to numerically determine the equilibrium p for distinct parameters and to derive comparative static predictions from comparing these equilibria.

We first determine the effect of in-group bias for this homogenous case. With respect to the preferences put forward in 2.A.6 (with $\alpha_i = \alpha$, $\beta_i = \beta$, $\gamma_i = \gamma$, $\forall i$) we implement five parameterizations that use $\alpha = 0.12$ (the average slope of the motivational vector in our data) but have different in-group bias ratios: $\beta/\gamma = \{0, 1/3, 1, 3, \infty\}$. For each parameter configuration, we solve 2.A.7 and plug the result and the above probabilities in equation 1.4, substituting U_i^W and U_i^L for B^W and B^L . We then solve for p , which yields the QRE $p = \{0.26, 0.28, 0.32, 0.48, 0.58\}$, where each value corresponds to the listed β/γ . These are also the predicted levels of aggregate participation. We conclude that for the homogenous case, and the moderate level of other-regarding concerns found in our data ($\alpha = 0.12$), equilibrium (aggregate) participation is increasing in in-group bias.

2.A.2 Altruism, In-group Bias and Individual Participation

Next, we drop the assumed homogeneity and allow for different mixed strategies for each player. This enables an investigation of the comparative statics at the individual level. The probability terms become, for each player i :

$$\Pr [m_i = m_j] = \sum_{j=1}^{126} \prod_{h \neq i} (p_h)^{A_{jk}} (1-p_h)^{1-A_{jk}}, k = \begin{cases} h, & \text{if } h < i \\ h-1, & \text{else} \end{cases} \quad (2.A.10)$$

The results are reported in Table 2.4. For each parameterization we report whether this relationship is negative (‘-’), positive (‘+’) or constant (‘=’). We observe that the individual probability of participation is generally increasing in other-regarding concerns for in-group bias levels of $4/3$ and above. If the in-group bias is smaller than 1 (i.e., i prefers the other group), more other-regarding players will participate less. For the average level of other-regarding concerns in our data ($\alpha = 0.12$), individual participation is increasing in other-regarding concerns for all values of in-group favoritism. This relationship is reversed when a high level of other-regarding concerns is combined with very high values of in-group favoritism, though one may doubt the empirical relevance of this combination, as it is not observed in the data. In general, the results presented in Table 2.4 provide support to Hypothesis 1: individual participation is increasing in other-regarding preferences for parameter values that are empirically relevant.

α^*	$(\beta/\bar{\gamma})$	0	2/3	4/3	2	8/3	∞
-0.75		+	-	+	+	+	+
-0.5		+	-	+	+	+	+
-0.25		-	-	+	+	+	+
0		-	-	+	+	+	+
0.25		-	-	+	+	+	-
0.5		-	-	+	+	-	-
0.75		-	-	+	+	-	-

Table 2.4: Other-regarding preferences and individual participation

Notes: The parameter α takes a baseline value for each parameterization, α^* , which is incremented in steps of 0.1 for the players in each group in order to generate heterogeneity. The relationship between individual participation probabilities and α can be negative (‘-’) or positive (‘+’).

Table 2.5 employs the same procedure to induce heterogeneity in the individual in-group bias parameters, i.e. each player in one of the symmetric groups has a different β_i/γ_i . One player, call it player 1, has a baseline value of $\beta_1/\gamma_1 = \beta^*/\gamma^*$. The β (γ) increases (decreases) with an increment of 0.1 for the subsequent players, such that player 4 has $\beta_4/\gamma_4 = (\beta^* + 0.3)/(\gamma^* - 0.3)$. We compute equilibria for six different values of β^*/γ^* , the ones presented in Table 2.4 except 0 (such that in-group bias does not take negative values). The presented results show that individual participation is increasing in individual in-group bias for low values of other-regarding concerns. For values of α above 0.5 (which seem empirically irrelevant), an in-group bias above $8/3$ (i.e. $\beta > 0.72$) leads to a negative relationship. The in-group bias measurement that we implemented in the experiment does not allow for a precise correspondence between the subjects’ choices and the parameters

of our model. However, we find it plausible that someone allocating 3/4 of the endowment to the in-group member cares three times more about the in-group, and therefore has an in-group bias ratio of $\beta/\gamma = 3$. Subjects allocated an average of 148.6 out of 200 tokens to the in-group member (pooling all treatments and both decisions), which leads us to believe that such a ratio is plausible. Assuming an other-regarding parameter equal to the data average, our results support Hypothesis 2.b in the sense that individual participation is increasing in in-group bias in a parameter range that is compatible with the observed data.

$\bar{\alpha}$	(β^*/γ^*)	2/3	4/3	2	8/3	∞
-0.75		-	-	-	-	-
-0.5		-	-	-	-	-
-0.25		-	-	-	-	-
0		=	=	=	=	=
0.25		+	+	+	+	+
0.5		+	+	+	-	-
0.75		+	-	-	-	-

Table 2.5: In-group bias and individual participation

Notes: The parameters β/γ takes a baseline value for each parameterization, β^*/γ^* , which is incremented for subsequent players in each group in order to generate heterogeneity. The relationship between individual participation probabilities and β/γ can be negative ('-'), positive ('+') or constant ('=').

2.B Value Orientation Test - Ring Test

The ring test of van Dijk et al. (2002) is reproduced in Table 2.6. For each Decision, a subject has to choose between Alternative A and Alternative B.

Decision	Alternative A		Alternative B		Decision	Alternative A		Alternative B	
	Self	Other	Self	Other		Self	Other	Self	Other
1	0	500	304	397	17	0	-500	-304	-397
2	304	397	354	354	18	-304	-397	-354	-354
3	354	354	397	304	19	-354	-354	-397	-304
4	397	304	433	250	20	-397	-304	-433	-250
5	433	250	462	191	21	-433	-250	-462	-191
6	462	191	483	129	22	-462	-191	-483	-129
7	483	129	496	65	23	-483	-129	-496	-65
8	496	65	500	0	24	-496	-65	-500	0
9	500	0	496	-65	25	-500	0	-496	65
10	496	-65	483	-129	26	-496	65	-483	129
11	483	-129	462	-191	27	-483	129	-462	191
12	462	-191	433	-250	28	-462	191	-433	250
13	433	-250	397	-304	29	-433	250	-397	304
14	397	-304	354	-354	30	-397	304	-354	354
15	354	-354	304	-397	31	-354	354	-304	397
16	304	-397	0	-500	32	-304	397	0	500

Table 2.6: Ring test.

2.C Auxiliary Tables

Table 2.7 presents the input of subjects on the motivations of different strategies in the participation game. This information was collected in the post-experimental questionnaire. Note that the available options correspond to motivations that have a rough correspondence to our preference specification: individualistic ($\alpha = 0$), competitive ($\alpha < 0$), in-group cooperator ($\alpha > 0, \beta > \gamma$), overall cooperator ($\alpha > 0, \beta = \gamma$).

Table 2.8 presents OLS regression results on the relationship between personality traits and participation behavior, and between personality traits and other-regardingness.

Table 2.9 presents OLS regression results on the relationship between average individual participation and a linear trend, for each type.

Table 2.10 presents the distribution of subjects across types and treatments.

	Main goal of a participant who...	
	participated...	did not participate...
	most of the times	
Make as much money as possible for himself or herself.	27.50%	77.50%
Increase the difference between his or her earnings and the earnings of other participants.	1.88%	20.00%
Help his or her group make as much money as possible.	63.75%	1.25%
Help both his or her group and the other group make as much money as possible.	6.88%	1.25%

Table 2.7: Reported motivations in the questionnaire.

	Average Participation	Other-regardingness (motivational vector (°))
Agreeableness	-.0608* (-1.71)	0.080 (0.03)
Conscientiousness	.007 (0.24)	2.167 (0.93)
Extraversion	.029 (0.93)	-2.539 (-1.09)
Openness	-.056 (-1.64)	0.449 (0.19)
Neuroticism	.007 (0.22)	-2.321 (-0.90)
Other-regardingness (motivational vector (°))	.003** (3.02)	
Constant	.975** (4.38)	16.196 (0.98)
R^2	0.11	0.02

Table 2.8: Participation, personal traits and other-regarding preference types

Notes: OLS regression. N=152. t-statistics in parentheses. */** indicates significance at the 10%/1% level.

	Competitors	Individualists	Weak Altruists	Mild Altruists	Strong Altruists
Constant	0.682** (34.42)	0.794** (40.12)	0.845** (30.58)	0.804** (37.05)	0.734** (23.83)
Linear Trend	-0.601** (-7.21)	-0.210* (-2.49)	-0.373** (-3.18)	-0.296** (-3.21)	0.150 (1.15)
R^2	0.58	0.14	0.21	0.21	0.03

Table 2.9: Regression of average participation on a linear trend

Notes: OLS regression. N=152. t-statistics in parentheses. */** indicates significance at the 10%/1% level. Linear trend coefficients multiplied by 10^2 .

Treatment\Type	Competitive	Individualistic	Weak Alt	Mild Alt	Strong Alt
High	6%	9%	7%	7%	8%
Low	9%	6%	8%	8%	6%
Control	5%	9%	5%	3%	5%

Table 2.10: Conditional distribution of types across treatments, in percentage.

2.D Other-regarding Preferences and Group Identity

In this appendix we examine the relationship between altruism and group identity. In other words, we are interested in whether distinct motivational types respond differently to group identity manipulations. For this purpose, Figure 2.10 shows the average percentages of the endowment allocated to the in-group member – both before and after the participation game – per type and treatment.

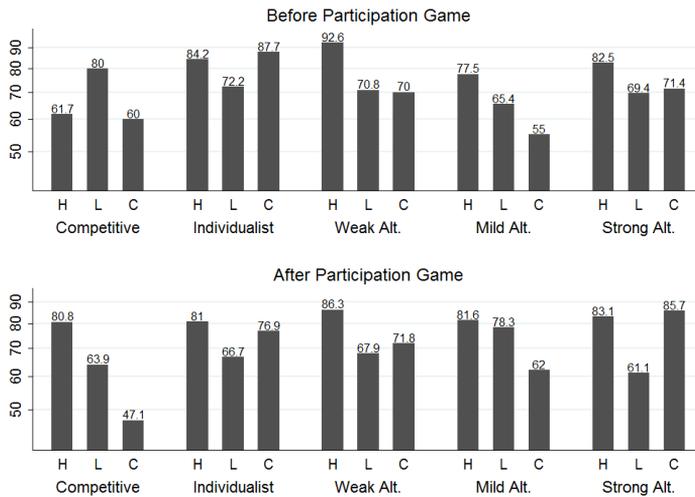


Figure 2.10: Other-regarding preferences and group identity

Notes: Bars depict, for each type, the average allocation to the member of the in-group, as measured before (upper panel) and after (lower panel) the participation game. “H”, “L” and “C” stand for High, Low and Control.

Consider first the average in-group bias across treatments. Individualists are the category showing the highest in-group bias, with an average allocation of 79.1% of the endowments to the in-group. The group showing the lowest in-group bias are Competitors, for whom the average allocation to the in-group member is 67.4%. Despite the apparent diversity in allocation behavior, the only significant difference across categories when using the average of the two decisions is between Competitors and Individualists (MW, $p=0.05$). A Pearson’s chi-square test corroborates this point: there is no significant systematic difference over categories for the average of the two decisions, neither across all treatments, nor for any particular treatment (all

$p > 0.35$). In the allocation decision before the participation game the bias is stronger for Individualists than for Competitors and Mild Altruists (MW, $p = 0.02$ for both comparisons; pooling treatments). For the allocation after the participation game, there are no statistically different decisions across motivational types. We conclude that subjects with distinct other-regarding preferences do not exhibit strong and systematic differences in in-group bias if we pool treatments.

Some types react differently to group identity manipulations, however. Considering the average of the two decisions, we find (weak) evidence of differential behavior of Weak Altruists between High and Control (MW, $p=0.07$), Strong Altruists between High and Low (MW, $p=0.04$), and Mild Altruists between both High and Low, and High and Control (MW, $p=0.05$ and $p=0.08$, respectively). Other comparisons do not reach statistical significance below 0.10. Bearing in mind that we observed a difference between the average allocation in High and in the other two treatments (section 2.5.2), this evidence suggests that differences in group attachment across treatments are mostly driven by the three altruistic types. Altruistic types not only share more with an anonymous other, they also allocate a relatively higher amount to the member of their in-group when group identity is high.

Next, we consider whether our subjects' in-group bias is affected by the interaction in the participation game. Eyeballing Figure 2.10 suggests similar patterns across the two decisions, with a possible exception for Competitors. The difference between the two decisions is not statistically significant for this group, nor for any other, however.²³ The changes between the two measurements are symmetric. We observe some instances where subjects seem to be punishing their group (21.05% of the subjects decrease their allocation to the in-group after the participation game), a majority of subjects exhibiting stable in-group bias (54.61%) and some rewarding the in-group by giving more after the participation game (24.34%).

2.E Experiment Instructions

In this appendix we reproduce the experiment's instructions. Italicized text corresponds to text that was part of High and Low but not Control.

Welcome to this experiment in decision-making. Depending on your decisions and the decisions of other subjects you may earn money. You will be paid privately at the end of

²³At this level of disaggregation the number of observations is small. The conclusions do not change if we aggregate data. Pooling across treatments, there is only evidence of different behavior for Mild Altruists across the two decisions (MW, $p=0.08$). Pooling across types, we observe no statistically significant differences for any treatment, when comparing behavior before and after the participation game.

the session. This is an anonymous experiment: your identity will not be revealed to other participants. The choices you make in early parts of the experiment may be used in later parts. Since this experiment involves gains and losses, it is possible (though very unlikely) that you make a negative amount in the experiment. In that case, your earnings will be deducted from the show-up fee. It is not possible that your losses exceed the show-up fee. This experiment is composed of three main tasks: Task 1, Task 2, and Task 3. You will receive instructions for a new task after the previous one has been completed. Note that a new task will only begin when every participant has finished the previous one.

Ring Test: In Task 1 you will be asked to make 32 decisions with monetary consequences. In each of the 32 situations you will have to choose between two options: Option A and Option B. For each option, two numbers will be displayed. The first is the number of tokens that you yourself will receive (positive amounts) or pay (negative amounts). The second is the number of tokens that the “Other” will receive or pay as a consequence of your decision. The “Other” is an anonymous person in this room, with whom you are randomly matched for the entire duration of Task 1. You will also be randomly matched with a second, different anonymous participant whose choices will affect you in the same way that your choices affect the “Other”. Note: this means that the person who receives or loses money due to your decisions is a different person than the one whose decisions make you earn or lose money.

Your total payoff is the result of both your decisions and the decisions made by the participant whose choices affect you. No participant will know with whom he or she has been paired. Participants will only be informed about the total amount they earned or lost at the end of the experiment.

BFI: In Task 2 you will be asked to rate a number of characteristics that may or may not apply to you. There are 44 statements in total, distributed over 4 screens. Please pick a number from 1 to 5 next to each statement to indicate the extent to which you agree or disagree with that statement. Most people take no more than 10 minutes to complete this task.

The Openness Score: The statements you rated in Task 2 constitute a self-report inventory of personality traits (characteristics). We employed one of the most used and reliable personality trait tests. One of the traits that was measured is ‘Openness’, whose score can range from 1 to 5.

What is Openness? Openness is a personality trait that involves active imagination, aesthetic sensitivity, attentiveness to inner feelings, preference for variety, and intellectual curiosity. It captures receptivity to novel experiences and ideas. It is not the cultural habits and knowledge acquired through education or breeding, nor is it related to intelligence or any other cognitive ability.

People whose Openness Score lies more to the left-hand side of the scale:

- tend to be more conventional and traditional in their opinions and behavior.
- prefer familiar routines to new experiences.
- generally focus on a narrower range of interests.
- are practical and down-to-earth.
- are able to more easily separate ideas from feelings.

People whose Openness Score lies more to the right-hand side of the scale

- are curious, open to unknown things and variety.
- are frequently described as imaginative, artistic, unconventional and tolerant.
- are more willing to accept the validity of astrology and esoteric phenomena.
- have more easily access to thoughts and feelings simultaneously, thus experiencing things more intensely.

We have constructed a ranking of the Openness Scores of the twenty participants of this experiment. This ranking ranges from 1 to 20, with 1 being the participant with the Openness Score more to the right, and 20 the participant with the Openness Score more to the left. We would like to ask the ten participants with rankings 1 to 10 to move to another lab. Please wait for the organizers' instructions to do so. The other ten participants can remain seated. Given your ranking, we would kindly ask you to prepare to move to the other lab/remain seated.

After Moving and Sitting Down at New Computer Stations

For the next 3 minutes, we would like the participants in each lab to pick a name to identify their lab. We provide you with three pre-defined possibilities. You can discuss this with the other participants in the same lab as you by using the chat box below. Each participant submits his preferred choice, and the most picked choice will be the name that will identify your lab for the remainder of the experiment.

Trivia Challenge

The participants in [name of the participant's lab] and [name of the other lab] will now compete in a trivia challenge. Each participant will be asked five trivia questions. You have 30 seconds to answer each question. You cannot answer after time is up. Each correct answer corresponds to one point, an incorrect answer corresponds to zero points. In the end, the points of all participants in [name of the participant's lab] will be summed, and compared to the total number of points achieved by the participants in [name of the other lab]. The lab with more points gets a total reward of 2000 tokens (10 Euros), to be equally distributed among all participants of the winning lab, i.e. each participant gets 200 tokens (1 Euro). In case the two labs achieve the same number of points, the winner is decided randomly (with equal probability).

Participation Game: In Task 3 you will be asked to make decisions in 40 rounds, with one decision per round. You will be part of a group of 5 participants: you and 4 others. *The participants that are part of your group are all drawn from [name of the participant's lab].* Group composition will remain constant for the whole of Task 3. Your group will interact with another group of 5 participants, *all of them drawn from [name of the other lab].*

In every round, each member of a group will have to decide on whether to buy a “disc” or not. A “disc” costs 30 tokens. Members of the group with more “discs” receive a higher reward: 120 tokens. Members of the group with fewer “discs” receive a lower reward: 30 tokens.

If the number of discs in the two groups is the same, the group who gets the higher reward in that round is picked with equal probability. In other words, in case of a tie each group has a 50% chance of getting the high reward. Note that if one of the groups gets the high reward the other necessarily gets the low reward.

As an example, assume that 3 people in your group buy discs, but only 2 people in the other group buy discs. In this situation, your group gets the high reward in this round. A member of your group who bought a disc gets a payoff of 90 tokens in this round. A member of your group who did not buy a disc gets a payoff of 120 tokens in this round. A member of the other group who bought a disc gets a payoff of 0 tokens in this round. A member of the other group who did not buy a disc gets a payoff of 30 tokens in this round.

Allocation Decisions: We would like to ask you to divide 200 tokens (1 Euro) between a random participant who is part of your group (excluding yourself) and a random participant who is part of the other group. *Recall that your group is composed of you and 4 other participants from [name of the participant's lab]. The other group is composed of 5 participants from [name of the other lab].*

These amounts will be paid at the end of the experiment. We will randomly select both a member of your group and a member of the other group who will receive your chosen allocation. You will be affected by the choices of two other random participants in the same way.

Chapter 3

Why Does Political Mobilization Work? The Role of Norms and Reciprocity

In Chapter 2 we saw that group identity can explain increased participation to a moderate extent. This chapter addresses a phenomenon closely linked to the activities of political groups: mobilization. I start from the observation that political mobilization works but not all mobilization efforts are equally effective. Whereas door-to-door canvassing can increase voter turnout by as much as 9 percentage points, other methods like direct mailings have a negligible impact. In this chapter I propose a theoretical model and use a laboratory experiment to investigate two non-mutually exclusive channels through which mobilization efforts can work: reciprocity and social norms. The results show that mobilization does not seem to be driven by reciprocity from citizens to the mobilizing agent alone; rather, mobilization efforts work best when they are coupled with a social norm appeal.

3.1 Introduction

Mobilization matters in politics. The last fifteen years have provided a wealth of field experimental evidence which shows that mobilization efforts work (see Green et al. 2013 for a meta-analysis and appraisal of the literature).¹ However, not all mobilization efforts are equally effective. Whereas door-to-door canvassing can in-

¹Rosenstone and Hansen (1993) define mobilization as the “process by which candidates, parties, activists and groups induce other people to participate.” Most commonly studied have been the cases of get-out-the-vote drives and partisan appeals by campaigns, but examples also include “distributing voter registration forms and absentee ballots, driving people to the polls on election day, or providing child care to free parents to attend meetings and demonstrations.”

crease voter turnout by as much as 8.7 percentage points, other methods like direct mailings have proven largely ineffective (Gerber and Green 2000). The evidence suggests that more personal methods like door-to-door canvassing and volunteer phone banks tend to mobilize citizens more effectively than anonymous and massively administered tactics like direct mailings and robot phone calls (Green and Gerber 2004). Though it seems that the ‘human factor’ plays an important role, the psychological process underlying the mobilization mechanism is still far from well understood. In this chapter I empirically test two non-mutually exclusive channels through which mobilization efforts can work: reciprocity and social norm transmission. Given the inherent difficulty in measuring reciprocity preferences in the field and the impracticality of fully observing the content of participation appeals, I conduct a laboratory experiment which extends the canonical framework in the study of laboratory elections (Levine and Palfrey 2007).

The fact that ‘old-fashioned’ mobilization tactics like door-to-door canvassing seem to be more effective than mass media appeals should not mask the fact that mobilization is today more relevant than ever. King et al. (2012) have analyzed the censorship program of the Chinese government, dubbed by the authors the “most extensive effort to selectively censor human expression ever implemented.” They show that Chinese authorities do not react to strong negative criticism of the state; rather, they target and silence “comments that represent, reinforce, or spur social mobilization, regardless of content.” Chinese censors do not fear dissent: they fear mobilization. An experiment of mobilization via social influence conducted by Bond et al. (2012) claims to have changed the participation behavior of 340,000 American citizens in the 2010 congressional elections by implementing simple Facebook banners on election day. The fact that mobilization has a far-reaching impact helps explain the massive endeavors of civic organizations and campaigns. For example, in the 2012 US general election, the National Council of La Raza (the largest national Hispanic civil rights organization in the US) conducted the Mobilize to Vote campaign, a multi-state door-to-door effort to register 180,000 Latino voters.² On the partisan side, the Faith and Freedom Coalition pledged to spend approximately \$12 million to get religious voters to the polls in support of the Republican candidate, Mitt Romney. The organization had files with cell phone, e-mail or other contact information on 17.3 million potential voters in several states that were seen as key in that election. All those voters were supposedly contacted, many of them multiple times, and two million were visited by volunteers.³

²Source: National Council of La Raza, <http://goo.gl/gfIHjJ> (retrieved August 29, 2013).

³Source: Wall Street Journal, <http://goo.gl/cAU3rx> (retrieved August 29, 2013).

The field experimental framework pioneered by Gerber and Green (1999) has been used to test a number of psychological theories with respect to voter participation. Nickerson (2007) tested for conversational network peer-to-peer effects with respect to turnout decisions; contrary to the significant effect found in canvassing by strangers, the evidence shows that an outreach from friends and neighbors fails to mobilize voters. However, it seems that there is voting behavior contagion within households (Nickerson 2008), i.e. an appeal delivered to an household member has spillover effects on the participation behavior of other members of the household.

Several papers test the effectiveness of field interventions in generating participation. In general, these works change the incentives faced by citizens by manipulating the salience of, and feedback on, the social norm of voting. Gerber et al. (2008) vary the scrutiny on one's turnout decision (ranging from ex-post scrutiny by the experimenter to publicly revealing one's participation or abstention to the neighborhood). As expected, moving from a mail-delivered civic duty appeal to the threat of 'public shaming' leads to a considerable increase in participation, showing that social pressure does work. Feelings of personal pride and shame also seem to influence participation decisions: receiving a mail that encourages turnout and shows either a past abstention (shame) or participation (pride) has an effect, with the latter bearing a stronger impact (Gerber et al. 2010). In a similar spirit, Panagopoulos (2010) provides evidence that there is a significant impact of publicly disclosing past participation or abstention on turnout decisions. The same author (Panagopoulos 2011) shows that a gratitude message for past participation increases turnout in future elections. The evidence presented in Gerber and Rogers (2009) suggests that communicating a high turnout social norm (urging citizens to vote because many will do so) is more effective than communicating a low turnout one (many will stay home) on reported intention to vote for occasional and infrequent voters.

3.1.1 Candidate Explanations

The evidence presented thus far sheds light on what social norm interventions can boost voter turnout, but tells us little on the mechanism of mobilization. In other words, why do people participate in politics when a complete stranger asks them to do so? Three candidate explanations can be conceived: a reduction in information costs, reciprocation from the reached citizen to the mobilizer, or adherence to a salient social norm. Regarding the first, mobilization usually conveys practical information on how to participate: how to register, when and where to show up, what one needs to do, etc. In theory, mobilization reduces the costs of participating by providing useful information along these lines. Yet the evidence seems to

suggest that endowing subjects with information on election details (like the location of polling stations) does not significantly increase voter turnout (e.g. Bond et al. 2012). It is also possible that citizens tend to forget upcoming elections and mobilization efforts provide an effective reminder; however, in the words of Green and Gerber (2004), “low voter turnout reflects low motivation, not amnesia”, as evidenced by the fact that prerecorded messages reminding people to vote just before an election do little to increase voter turnout.

The second candidate explanation hinges on the observation that those being reached by civic associations or campaigns might want to return the favor by performing the act they are urged to. A citizen receiving a mobilization attempt, be it a leaflet or a knock on the door, observes an investment of time and money into persuading her to adopt a certain behavior. If this gesture is perceived as kind, the citizen being reached might want to reciprocate by doing what she was asked to. A long literature in psychology and economics has shown that reciprocity is a pervasive behavioral motive (Fehr and Gächter 2000). Online tools of grassroots mass mobilization, like MoveOn.org, harness the power of reciprocity and trust in social networks. The extant evidence points to a potential role for reciprocity. Arceneaux (2007) shows that canvassing by the running candidate of a local election seems to be more effective than by paid canvassers, even though the effect is not statistically distinguishable. The reciprocity hypothesis suggests that people would be more responsive to an appeal by someone who is sacrificing more to mobilize citizens, which in this particular case would be the candidate (whose value of time is particularly high during a campaign). In their meta-study, Green et al. (2013) show that phone calls from volunteer phone banks have an impact on turnout that is roughly double the increase in turnout from commercial phone banks, and more than ten times the impact of pre-recorded messages. This is a hint of the influence of reciprocity, as a volunteer is seen as investing his time in a disinterested manner.⁴

The third candidate explanation for why mobilization works is social norms. This norm is transmitted by the mobilizing agent to the reached citizen and is commonly accepted within the relevant social group. Voting in national or local elections or showing up at a rally can be seen as duties expected from members of groups in society. In this sense they reflect social norms (Elster 1989): they are “shared by other people and partly sustained by their approval and disapproval.” Members of all societies respect a plethora of social norms in their daily lives. In Western societies examples include waiting in queues, covering one’s mouth when yawning, or

⁴In principle, nothing prevents a reached citizen from reciprocating institutionalized or monetized efforts like the ones of commercial phone banks; however, it seems plausible that reciprocal actions should be increasing in non-monetary inputs.

tipping in restaurants. Social norms are “sustained by the feelings of embarrassment, anxiety, guilt and shame that a person suffers at the prospect of violating them” (Elster, 1989). Participating in politics is often regarded as fulfilling a social norm. When behavior is observable, external sanctioning of the social norm of participation by third parties is possible. For example, a union member not showing up at a small rally can be sanctioned by her fellow members. However, behavior is not observable in many instances of political participation like voting or writing letters to officials. In such cases the sanction from disrespecting the norm will tend to be internal rather than external. It is reasonable to conjecture that the sanction from failing to comply with the norm is increasing in its salience. A mobilization appeal does precisely this, i.e. it makes a norm salient. Many people feel compelled to stick to the norm, and most feel bad for not doing so, as evidenced by self-reported turnout 15 percentage points in excess of actual turnout rates (Holbrook and Krosnick 2010).

In this study, I investigate both the reciprocity and the social norm mechanisms as determinants of mobilization’s effectiveness. If a reciprocity rationale explains the effectiveness of mobilization, then who mobilizes makes a difference, as reciprocity behavior may differ depending on the person at which it is directed. The evidence on the impact of reciprocity in political behavior is scarce and mixed in this respect. Whereas there is scant evidence that canvassers who match the ethnic profile of the neighborhood have more success than those who do not (Green and Gerber 2004), the work of Nickerson (2007) testifies to the low impact of friends and neighbors’ attempts at mobilizing each other. The evidence concerning message content is also mixed. On the one hand, Matland and Murray (2012), among others, have shown that message content does not matter; on the other hand, Bryan et al. (2011) find that subtle linguistic clues can have a profound impact on political behavior.⁵ If the salience of a social norm is responsible for the increase in participation then it is worth investigating more systematically the way in which this impact depends on specific messengers and messages.

Testing the role of reciprocity in participation behavior is best done in a controlled laboratory environment, as reciprocity preferences are hard to measure in the field.⁶ A laboratory experiment allows us to isolate the impact of reciprocity by

⁵Asking “How important is it to you to be a voter in the upcoming election?” increases turnout by 10 percentage points more than “How important is it to you to vote in the upcoming election?”

⁶One could devise field experiments where the mobilizing agents differ in characteristics that are hypothesized to influence reciprocity. For example, if the ethnicity or other observable characteristic of a canvasser matches the racial profile of the neighborhood, a higher boost to participation might be expected. However, this requires strong assumptions concerning the preferences of the reached citizens, namely that someone of their own ethnicity commands higher reciprocity. Measuring such preferences in an incentivized way, e.g. by asking citizens to materially reward the mobilizing agent, may crowd out the intrinsic motivation inherent to the act. Measuring such

comparing a treatment where the mobilization decisions are the product of a human subject's effort with one where this is not the case. The laboratory further allows us to observe the specific impact of transmitting a normative appeal by selectively allowing the mobilizing agents to convey a message to the target group. Other questions that can be addressed by an experiment include whether mobilization efforts are directed at subjects contingent on their costs of participation, and whether a voter's reaction to the mobilization effort depends on these costs. In particular, high cost subjects might reciprocate to a greater extent the higher effort of mobilizing them relative to low cost subjects.

Such hypotheses can be tested using a game theoretic model that extends the participation game of the previous chapters. A group member decides how many of the other members of his group to mobilize; the members who have been mobilized then have to decide whether to participate. Mobilization and participation are both costly. The group where more members participate are better off than those where fewer do. Complete information is assumed in order to prevent beliefs from clouding the predictive power of the hypotheses. Applying the concept of quantal response equilibrium (QRE), equilibrium probabilities of participation for all players in all possible electorates can be derived. Given the sequential structure of the game, the subjects mobilizing others can solve for the equilibrium levels of mobilization (sub-game perfect equilibrium).

The experimental design manipulates two experimental features: mobilization and one-way communication. In half of the sessions mobilization is decided by a human subject, while in the other half this observed mobilization pattern is replicated mechanically. This removes the intentionality from mobilization decisions, and therefore should preclude reciprocity-induced boosts to participation. In all sessions, at a pre-determined point in the experiment, the agents who mobilize have the opportunity to transmit a message to their group – which, as we will see, takes the form of an appeal regarding what the group is normatively expected to do.

The results show that mobilization alone is not sufficient to generate higher participation. Furthermore, there is no significant correlation between subjects' reciprocity preferences and participation behavior. The experimental data thus suggests that mobilization's success is not driven by reciprocity concerns per se. However, mobilization is effective when coupled with a normative appeal delivered by the mobilizing subject, leading to an increase in participation of approximately 10 per-

preferences in a non-incentivized way, e.g. using survey questions, may also crowd out intrinsic motivations and further provide unreliable measurements due to an experimenter demand effect (citizens might overstate their reciprocity to individuals of the canvasser's ethnicity as a token of appreciation).

centage points, which is in line with what is typically found in the field.

This chapter is structured as follows: Section 3.2 presents the model, Section 3.3 presents the experimental design and formulates the hypotheses, and Section 3.4 presents the results. A final section concludes.

3.2 The Model

The model put forward in this section is an extension of the participation game of Appendix 1.A. A number of differences are introduced. First, there are two player roles: Alpha and Beta. There is one Alpha player in each group, the remaining are Beta players, i.e. each group is composed of one Alpha and M Betas.⁷ Second, the game has two stages. Alphas move at the first stage while Betas move at the second stage. Alphas decide simultaneously how many Betas to activate (mobilize), i.e. the Alpha of G_i decides how many $m_i \in [0, M_i]$ Betas to activate. Note the change in notation from Chapter 1: $m_i \leq M_i$ now stands for the number of active Betas out of a total M_i Betas. In comparison, the model presented in Chapter 2 corresponds to a situation where there are only Beta players, all of them are active, and $M_i = M_j$.

If a Beta has been activated by the Alpha in her group she is said to be ‘active’; otherwise she is ‘inactive’. Active Betas play a standard participation game. In order to prevent Alphas from sending a leadership or commitment signal to active Betas, Alphas cannot participate.⁸ All players on the winning group (Alpha and Betas, both active and inactive) receive the high material payoff (B^W), while all players on the losing group (Alpha and Betas, both active and inactive) receive the low material payoff (B^L).

There are two types of Beta player in each group: Low and High, denoted by $\Sigma = \{l, h\}$. Activating a Low and High Beta costs c_l^A and $c_h^A > c_l^A$, respectively. Each Alpha has a budget A_i at his disposal, which is exactly enough to activate all the Betas in his group; therefore $A_i = M_i^l c_l^A + (M_i - M_i^l) c_h^A$, where M_i^l denotes the number of Low Betas in group i . What is spent on activating Betas is lost, but what is left (A_i net of activation costs) is kept by the Alpha member. For the sake of tractability I assume that Alphas always activate Low Betas before High Betas. This assumption imposes a minimal degree of rationality on Alphas, as Low Betas

⁷Henceforth, “Alpha(s)” and “Beta(s)” stand for “Alpha player(s)” and “Beta player(s)”, respectively. For the most part, the exposition adopts the viewpoint of the players in group G_i , without loss of generality.

⁸Such a leadership signal could cloud the role of reciprocity in Betas’ choices, in the sense that an Alpha who participated could also reveal a commitment to future activation decisions. Leadership signals have been shown to help overcome coordination failure in situations of group cooperation (e.g. Brandts et al. 2007).

cost less to activate and, as will become clear below, participate at higher rates.

I define an ‘active electorate’ as a pair (m_i, m_j) , which is the number of Betas in each group who are active and can therefore participate. For each (m_i, m_j) , a standard participation game with complete information is played between the active Betas. Participation is costly for (active) Betas: participation costs c_h for a High Beta and $c_l < c_h$ for a Low Beta. The structure and parameters of the game are common knowledge. Therefore, for each (m_i, m_j) there exist equilibrium participation probabilities for each Beta type in each group.⁹ From these participation probabilities the Alphas can derive, for each active electorate, a probability distribution over the possible outcomes: G_i wins and G_j loses, G_i and G_j tie, or G_j wins and G_i loses. For the Alpha in G_i , define a vector containing this probability distribution as:

$$v_i(m_i, m_j) = \left[\Pr_i(\text{win}|m_i, m_j) \quad \Pr_i(\text{tie}|m_i, m_j) \quad \Pr_i(\text{lose}|m_i, m_j) \right]' \quad (3.1)$$

The derivation of each event’s probability is carried out in detail in Appendix 3.A.1; the result for a particular parameterization is provided later in this section. With this information in hand and knowledge of the model’s parameters, the Alpha decides how much of the budget A_i to spend on activating Betas, i.e. she solves:

$$\max_{m_i} A_i - a(m_i) + v_i(m_i, m_j)' b \quad (3.2)$$

where $b = \left[B^W \quad \frac{B^W + B^L}{2} \quad B^L \right]'$ and $a(m_i)$ describes the costs to an Alpha of activating m_i Betas. The Alpha in G_j solves a symmetric problem.

The Alphas choose simultaneously and are aware that the Betas will know how many others (and of which type) have been activated in the two groups. The equilibrium is a pair (m_i^*, m_j^*) . For certain parameter values an interior equilibrium (in pure strategies) exists. This simple model does not have a closed form solution for most group sizes (M_i and M_j larger than 2), as the probability terms underlying $v_i(m_i, m_j)$ are non-linear. I solve the model using numerical methods in order to obtain theoretical predictions.

Before we move on to the parameterization that allows us to characterize the solution that was implemented in the laboratory, a few words about how the model relates to mobilization are in order. The model is quite stylized in order to allow for a straightforward implementation. In particular, the fact that only Betas who have been activated by Alphas are able to participate does not closely represent all

⁹I carry out the analysis for the case of selfish preferences, i.e. the expected utility of participation and abstention are equations 1.1 and 1.2. An extension to social and group-discriminating preferences could be obtained in the fashion of Chapter 2.

forms of mobilization. However, we can think of the activation as voter registration or transportation to the polls provided by a political activist, in the absence of which participation would not be possible. Braconnier et al. (2013) show that home registration visits are successful in increasing both voter registration and voter turnout among the newly registered voters. Of course, ‘real world Betas’ can provide registration or transportation for themselves. However, allowing for an intermediate stage where Betas who had not been activated could choose to activate themselves would add another stage to the game and complicate the model considerably. The gain would be marginal as the purpose of the model is to provide a framework for the empirical test of the behavioral response of subjects who have been activated by an Alpha member, and compare it to the response of those who have been activated by the computer. A further feature that might not seem to be present in all modes of mobilization is the fact that Alphas can keep the money that they do not spend on activation. In the world of campaigning and activism the person mobilizing other citizens is often not handed money that she can ultimately keep. However, there exist organizations behind the people delivering the mobilization efforts, and these organizations have competing ends for their money. In a sense, organizations can use their money to mobilize citizens or allocate it to competing ends. The choice faced by the Alpha member incorporates this trade-off.

3.2.1 Implemented Parameterization and Equilibrium Predictions

The parameters of the model, which are the same for the two groups, are presented in Table 3.1. In solving the model I apply the concept of QRE, for reasons discussed in the Appendix of Chapter 1.

Group size, M	6	Activation cost: High, c_h^A	1
Number of Low Betas, M^l	2	Activation cost: Low, c_l^A	0.5
Benefit for the winning group, B^W	4	Participation cost: High, c_h	1
Benefit for the losing group, B^L	1	Participation cost: Low, c_l	0.5
Alpha’s budget, A	4		

Table 3.1: Parameter values.

In order to obtain point predictions we must choose a value for the μ assigned to Betas, μ^{Beta} .¹⁰ A value of $\mu^{Beta} = 0.4$ was chosen for two reasons. First, it is

¹⁰A detailed derivation of QRE can be found in Appendix 3.A.2. In Appendix 3.A.2 I provide graphical displays of participation probabilities of each Beta type in all possible active electorates, for a large range of μ .

consistent with estimates from experiments which implement participation games. Goeree and Holt (2005) rationalize the data of Schram and Sonnemans (1996) using μ ranging from 0.8 to 0.4; Cason and Mui (2005) estimate μ to be between 0.4 and 0.6; Levine and Palfrey (2007) find $\mu = 0.17$; and Grosser and Schram (2010) base their theoretical predictions on $\mu \geq 0.3$. In a strategic voting setting, Tyszler (2008) estimates $\mu = 0.55$. Second, $\mu = 0.4$ is partially consistent with a maximum likelihood estimation performed on pilot sessions of my experimental design.¹¹ Table 3.2 presents the participation probabilities and outcome distribution for this value of μ^{Beta} .

Active Electorate	G_i				Outcome Distribution			Expected Participation		
	m_i, m_j	p_l	p_h	q_l	q_h	G_i wins	Tie	G_j wins	G_i	G_j
1,0	0.92	-	-	-	-	0.92	0.08	0	0.92	0
2,0	0.58	-	-	-	-	0.82	0.18	0	0.58	0
3,0	0.55	0.15	-	-	-	0.83	0.17	0	0.42	0
4,0	0.52	0.15	-	-	-	0.83	0.17	0	0.33	0
5,0	0.49	0.14	-	-	-	0.84	0.16	0	0.28	0
1,1	0.92	-	0.92	-	-	0.07	0.86	0.07	0.92	0.92
2,1	0.77	-	0.56	-	-	0.75	0.22	0.03	0.77	0.56
3,1	0.72	0.21	0.57	-	-	0.74	0.22	0.04	0.55	0.57
4,1	0.66	0.20	0.56	-	-	0.75	0.21	0.04	0.43	0.56
5,1	0.61	0.20	0.55	-	-	0.76	0.20	0.04	0.36	0.55
2,2	0.90	-	0.90	-	-	0.16	0.68	0.16	0.90	0.90
3,2	0.80	0.45	0.69	-	-	0.57	0.32	0.11	0.68	0.69
4,2	0.74	0.33	0.64	-	-	0.61	0.27	0.11	0.53	0.64
5,2	0.68	0.28	0.61	-	-	0.65	0.24	0.11	0.44	0.61
3,3	0.78	0.54	0.78	0.54	-	0.32	0.36	0.32	0.70	0.70
4,3	0.73	0.41	0.69	0.44	-	0.49	0.29	0.22	0.57	0.60
5,3	0.69	0.33	0.64	0.38	-	0.56	0.26	0.18	0.47	0.55
4,4	0.70	0.42	0.70	0.42	-	0.36	0.29	0.36	0.56	0.56
5,4	0.68	0.36	0.66	0.38	-	0.46	0.27	0.27	0.49	0.52
5,5	0.66	0.36	0.66	0.36	-	0.37	0.26	0.37	0.48	0.48

Table 3.2: Point predictions

Notes: This table presents participation probabilities for each Beta type in each active electorate. Players in G_i (G_j) participate with probability p (q).

The equilibrium probabilities and the resulting probability distribution over out-

¹¹A maximum likelihood estimation using data from the pilot sessions yields $\mu = 0.2$ for the Low Betas, and $\mu \rightarrow \infty$ for High Betas. The reason is that, as we will see below, High Betas over-participate relative to any admissible prediction, and therefore bias the estimate in the direction of random behavior. A value of $\mu = 0.4$ reconciles the high participation observed in the pilot sessions with the typical values found by previous authors. However, using a value of $\mu = 0.8$, as was used in Chapter 2, would only change point predictions and not comparative statics.

comes allows us to close the model. The Alpha members of each group use the information in the 6th-8th columns of Table 3.2 to solve equation 3.2. The two Alphas know the expected payoffs that will result in each of the sub-games that correspond to an active electorate (m_i, m_j) . Using backward induction, they choose the optimal number of Betas to activate, (m_i^*, m_j^*) , which constitutes a sub-game perfect equilibrium.

As with Betas, I derive a QRE of the game played by the Alphas. Figure 3.1 shows the QRE predictions for the number of Betas to be activated by an Alpha, assuming $\mu^{Beta} = 0.4$. Note that the equilibrium is always conditional on the μ assigned to both Alphas (denoted μ^{Alpha}) and Betas. However, this equilibrium is robust to a wide range of μ^{Beta} and μ^{Alpha} . For the depicted values of μ^{Alpha} , we observe that QRE predicts $m_i = m_j = 2$ should be chosen with highest probability. As μ^{Alpha} increases, the difference between the probability of choosing $m_i = m_j = 2$ and other values of m_i and m_j decreases, but the prediction is never overturned, regardless of μ^{Beta} . The most natural assumption is to choose $\mu^{Alpha} = 0.4$, in which case Alphas are expected to activate two (Low) Beta players with probability 0.484, in equilibrium. Figure 3.1 also depicts the NE prediction ($m_i^* = m_j^* = 2$), which is obtained for all values of μ^{Beta} .

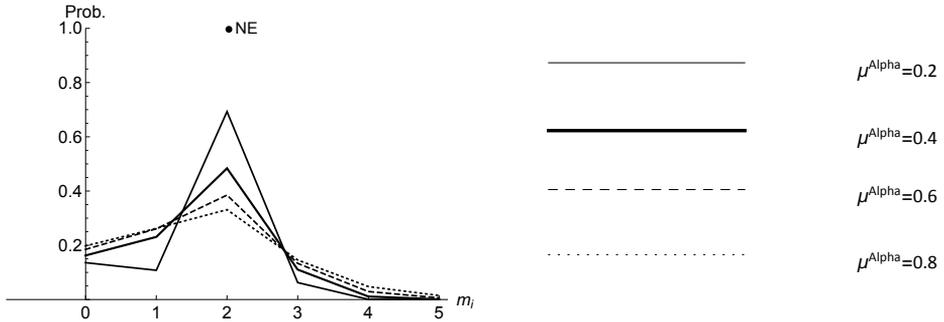


Figure 3.1: QRE and NE predictions for Alphas

Notes: All lines are drawn assuming $\mu^{Beta} = 0.4$.

3.2.2 Comparative Statics

The information presented in Table 3.2 and Figure 3.8 in Appendix 3.A allow us to derive a number of comparative statics predictions. First of all, we note that, within the same group, Low Betas always participate at higher rates than High

Betas. I call this the cost effect. This is an intuitive result: players with higher participation costs will participate less often in equilibrium. Another robust pattern, documented for example by Levine and Palfrey (2007), is the underdog effect: for $m_i + m_j > 3$, expected aggregate participation in the minority is always higher than expected aggregate participation in the majority (i.e. the number in the 9th column of Table 3.2 is always smaller than that in the 10th). The underdog effect reflects both the fact that minorities must participate at higher rates if they are to have a chance of upsetting the majority and the more complicated coordination problem faced by majorities. The competition effect is a related and commonly observed theoretical regularity: holding the total size of the active electorate constant, the participation probability of Betas is decreasing in $m_i - m_j$, $m_i > m_j$. The intuition is that as groups become more similar in size, races become more disputed and individuals respond by participating more. In our example, this translates into the following inequalities: $p_l(3, 0) < p_l(2, 1)$, $p_\sigma(3, 1) < p_\sigma(2, 2)$, $p_\sigma(4, 1) < p_\sigma(3, 2)$, and $p_\sigma(5, 1) < p_\sigma(4, 2) < p_\sigma(3, 3)$, $\sigma = \{l, h\}$; the same inequalities hold for q_l . Finally, in active electorates where $m_i = m_j$, aggregate participation is decreasing in $m_i + m_j$. I refer to this as the weak size effect, which reflects the fact that pivotality is decreasing in electorate size.¹² These predictions will be confronted with the data, as they can help us structure our interpretation of behavioral patterns, ascribing strategic and behavioral effects where they are due. Moreover, they bear interest in themselves as a further test of the pivotal model of turnout vindicated by the evidence of Levine and Palfrey (2007) and Herrera et al. (2014).

3.3 Experimental Design and Hypotheses

The experimental design investigates the role of reciprocity and normative appeals by manipulating two treatment variables: activation by a human subject and a one-way appeal sent by the Alpha to the Betas. Activation by a human subject corresponds to a mobilization effort, and an appeal sent by a subject who mobilizes others can serve the purpose of creating a norm for the group. I will first describe the features of the experimental framework and then discuss the implemented treatments. Earnings in the experiment are expressed in tokens, which are converted into Euro at the rate of 15 Euro cents (approximately 20 Dollar cents) per token.

Each experimental session has two parts. The first consists of a simplified trust game and the second one of what I call the ‘activation-participation game’ (APG),

¹²Levine and Palfrey (2007) test for the size effect, which is a more structured hypothesis that does not hold in this model due to the presence of two cost type subjects.

which is based on the model presented in the previous section. All subjects go through the two parts in this order. The trust game is the same for all subjects, as the treatments are implemented in the APG. A transcript of instructions and comprehension questions can be found in Appendix 3.B.

The trust game (Berg et al. 1995) is a widely used game in the study of trust and reciprocity, allowing us to measure individual attitudes/preferences concerning each of the two. In the trust game, one subject, called the ‘sender’, is endowed with 8 experimental tokens. The sender has to choose how many tokens to keep and how many to send to another anonymous subject, the ‘receiver’. Money sent to the receiver is multiplied by 3. The receiver faces a similar decision: she has to decide how many tokens to keep and how many to send back to the sender.¹³ The amount returned by the receiver is increasing in how reciprocal she is, i.e. the extent to which she rewards and punishes kind or unkind actions. A meta-study by Johnson and Mislin (2011) shows that, on average, senders hand over 50.2% of their endowment to receivers. Receivers send back 37.2% of the multiplied amount, on average. This leads to an approximate equalization of payoffs.

The receiver’s decision is more important for the purpose of this chapter’s research questions, and therefore I introduce three changes to the typical protocol of the trust game. First, subjects play both roles and do so sequentially. This allows us to obtain both trust and reciprocity measurements from all subjects.¹⁴ Second, the receivers input their decision using the strategy method, which is adopted to capture subjects’ reciprocity attitudes towards different levels of trust. Third, senders must send either 0, 4 or 8 tokens to the receiver. This restriction was implemented in order to allow for the use of the strategy method while keeping this part of the experiment parsimonious. Note that 8 tokens induces round numbers when calculating the amount to send back according to many heuristics (the typical fairness heuristics of returning $1/2$ or $2/3$ of the multiplied amount are easy to compute, as well as the less frequent $1/4$ or $3/4$).

Each subject is informed that she is paired with two other subjects: the one she sends 0, 4 or 8 tokens to, and the one to whom she may send a share of the 12 or 24 tokens back. The pairing of a subject with two others prevents interdependencies

¹³Assuming selfish preferences, the sub-game perfect equilibrium is for the sender to send 0 tokens to the receiver and for the receiver to keep whatever she is handed by the sender. The amount sent to the receiver is increasing in the extent to which the sender ‘trusts’ the receiver to share the multiplied amount. In addition, the amount sent to the receiver is increasing in the degree of altruism and the preference for efficiency of the sender.

¹⁴Burks et al. (2003) report evidence that playing both roles decreases the amount sent if subjects are aware of this beforehand, and decreases the amount returned, regardless of whether playing both roles was known beforehand. This pattern does not show up in my data.

between the amount sent and the amount returned.¹⁵ In order not to induce different individual histories, subjects are informed that results from this part of the experiment would only be known at the end. Looking at the results we conclude that the data from this experiment’s trust game is very similar to the typical pattern found in the literature.¹⁶

The second part of the experiment (APG) is based on a parameterization of the model of Section 3.2. The values in Table 3.1 correspond to the token amounts used in the experiment. In each period, the stage game corresponding to the model of Section 3.2 is played. Each subject goes through three blocks of 27 periods each. Each block is further divided into three sub-blocks of 9 periods each. In the first period of each sub-block Alphas make their activation decisions. In the remaining 8 periods of a sub-block, active Betas decide whether or not to participate. Active Betas have no period-to-period feedback, both in order to avoid reaction to period-to-period results and to allow them to implement mixed strategies (recall that the model’s predictions are participation probabilities).¹⁷ Alphas and inactive Betas observed their electorate’s results every period. At the end of a sub-block all members of an electorate were shown the results of the past 8 periods. Results consist of aggregate participation in each period and cumulative earnings in the block. All subsequent sub-blocks had the same structure. Groups, roles (Alpha and Beta) and Beta types (High or Low) were kept constant within each block; this means that from a sub-block to the next, within a block, only the activation status of Beta players could change. At the end of each block, groups, roles and Beta types were drawn anew. Figure 3.2 presents a diagram exemplifying the sequence in block 1 (sub-blocks were numbered from 1 to 9, so block 3 comprised sub-blocks 7 to 9, for example).

In the periods in which the active Betas face the participation decision, Alphas and inactive Betas are asked to provide a rating of the previous period’s results using a three-step rating scale.¹⁸ This information is useful as a proxy for the

¹⁵For example, a subject A who sends her full endowment to subject B might feel less inclined to reciprocate what subject B sends her. If subject A receives tokens from subject C, this kind of interdependency is avoided.

¹⁶Appendix 3.C provides a concise analysis of the trust game data.

¹⁷In participation games, previous period results (whether a subject was pivotal, how many others participated, etc.) are strong predictors of current period behavior. For example, Duffy and Tavits (2008) show that individual participation is increasing in pivotality events and previous period participation, whereas it is decreasing in the own group’s aggregate participation. Grosser and Schram (2006) also show that previous period participation in the own group correlates significantly with participation decisions. This is undesirable to the extent that treatment effects become statistically clouded. Regarding the implementation of mixed strategies, subjects employed them in a third of all sub-blocks.

¹⁸There were three available ratings: “satisfied”, “neutral” and “dissatisfied”.

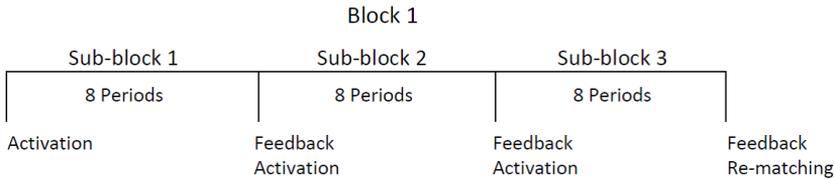


Figure 3.2: Sequence in one of the blocks of the experiment.

beliefs of Alphas and inactive Betas. For example, we can assess whether an Alpha member who activates more Betas than the payoff-maximizing number becomes disappointed when participation is below the winning threshold. Asking Alphas and inactive Betas for input also helps them remain engaged with the experiment. Comparing this feature with instances of political participation, we can conceive of inactive Betas as those citizens who are not able to participate themselves but who observe the (aggregate) outcome. Subjects are paid for the trust game and for one randomly picked block of the APG, both of which are added to the show-up fee.

3.3.1 Treatments

In order to investigate the role of reciprocity, sessions where activation is carried out by human subjects are compared (between subjects) to sessions where activation happens in a pseudo-random way. In the former type of session, mobilization is carried out by the Alpha of each group and thus constitutes an intentional process; I refer to this treatment as *Mob*. The activation patterns observed in *Mob* sessions were subsequently implemented in sessions where the Alphas were present but were not in control of activation decisions. These sessions serve as a control treatment (*Ctr*), as no intentionality in activation decisions is present and therefore reciprocity concerns should be absent. The Alpha members are still affected by the observed activation pattern, but all subjects take it as given and not as the product of the Alphas' choices. In *Mob*, active Betas are activated by their group's Alpha, and may thus incorporate reciprocity considerations in their decision to participate. In *Ctr*, active Betas are simply informed that they did or did not become active, and therefore reciprocity concerns cannot play a role in their decision to participate.

The second manipulated experimental variable was free form one-way communication from the Alpha to the Betas in his group. This was varied within subjects. In both treatments, at the beginning of block 2 (after activation but before Betas could make any decision), subjects are told that the Alphas would have 90 seconds to send written text to the Betas in their group. The experiment only proceeded

after all subjects have read this message and give their permission for the one-way communication to begin, i.e. Alphas have enough time to plan what they want to transmit to the Betas in their group. This window of free form communication allows Alphas to convey an appeal to the group, which is expected to take the form of a group norm. The participation behavior from block 2 (the ‘appeal block’) can be compared to the corresponding behavior in blocks 1 and 3, where an appeal is not present.

3.3.2 Hypotheses

In this sub-section I formulate a number of testable hypotheses on the questions under study. The first hypothesis concerns the role of mobilization on participation, which is assumed to hinge on a reciprocity rationale. To be sure, a considerable literature in experimental economics has studied the importance of reciprocity in shaping individual behavior within the framework of social preferences. The influential outcome-based social preference models (e.g. Fehr and Schmidt 1999 and Bolton and Ockenfels 2000) cannot account for many regularities in games where intentions and reciprocity play a role. Using a simple trust game framework, McCabe et al. (2003) show that reciprocity plays a major part in subject’s reciprocation decisions. Falk et al. (2008) provide further evidence on the role of intentions for reciprocal behavior. Their experimental design bears resemblance to the one proposed in this chapter as they compare a treatment where trust behavior is the product of a subject’s choice with one where trust behavior follows a random draw from the observed distribution of choices. In a game with distributional consequences, subjects are more generous towards kind distributional choices made by human subjects than to choices (with the same distributional consequences) dictated by chance. The authors show that fairness intentions are crucial for both positive and negative reciprocal behavior, and largely supersede outcome inequality considerations in terms of magnitude.

Bearing in mind the extant field and experimental evidence, I conjecture that subjects who have been intentionally mobilized by another subject will reciprocate the mobilization effort by participating at an increased frequency. The underlying rationale is reciprocity: Betas who have been activated by an Alpha perceive mobilization as a kind act, since they become able to participate and pursue their group’s interests. They should reciprocate by participating more. The first hypothesis follows:

Hypothesis 1: *Mobilization (activation) leads to higher participation.*

Participation in Mob is higher than in Ctr.

This hypothesis will be evaluated in light of the theoretical model proposed in Section 3.2. We expect subjects in *Ctr* to participate in line with the predictions presented in Table 3.2, whereas subjects in *Mob* are expected to participate at higher rates. This hypothesis does not follow directly from theory, as players' preferences do not explicitly account for reciprocity. There exist game theoretical models that incorporate reciprocity (e.g. Dufwenberg and Kirchsteiger 2004 and Falk and Fischbacher 2006), but their application would pose significant problems. First, the theory of Dufwenberg and Kirchsteiger (2004) makes reciprocity exclusively dependent on beliefs about how kind other players are. As the authors note, a good set of predictions requires a proper measurement of first- and second-order beliefs, which is possible but would complicate further the experimental design. A problem that is common to both theories is that the parameter that governs reciprocity preferences is exogenous and therefore would have to be estimated from data for the predictions not to be *ad hoc*. The calibration of the model would require more than a simple trust game. To be sure, the authors apply their models to games that are substantially simpler than the APG. In any case, Falk and Fischbacher's (2006) model predicts that in a sequential public goods game there can be positive contributions to the public good in situations in which the Nash equilibrium is to contribute zero. This hints at the direction in which reciprocity might work in a framework like the present one, where free-riding and group benefits are also present.

Casual evidence suggests that reciprocity is not unconditional: kinder acts will tend to generate kinder responses than less kind acts. Put differently, reciprocity is increasing in an act's kindness. In the realm of mobilization efforts, I conjecture that someone who requires a larger mobilization effort will reciprocate to a greater extent. In other words, those who have higher mobilization costs will return the act of being mobilized with a higher boost in participation than those who have lower mobilization costs. It is well known that campaigns and activists target those who are more likely to participate, i.e. those who have low costs of participation (Rosenstone and Hansen, 1993). Often, those with the lowest costs of participation (the more educated, the more mobile, etc.) are also the easiest to mobilize. However, to the best of my knowledge, there is no evidence on how citizens with different mobilization costs respond to the mobilization effort. For example, inhabitants of a remote region have both high costs of participation and high costs of mobilization; however, if a campaign or an activist reaches them, the increase in participation may be higher than observed for easily reached citizens. The proposed setting allows

for such a test since we can observe deviations from the predicted participation probabilities for both high and low cost subjects. If the first hypothesis is confirmed, a corollary hypothesis follows:

Hypothesis 1b: *High cost subjects will deviate from the predicted participation levels more than low cost subjects.*

A normative appeal to the group is expected to guide its members actions in the desired direction. My conjecture is that Alphas will appeal to Betas to participate. Conditional on this being observed, I expect participation to increase. As discussed in the Introduction, a normative appeal to the group makes an implicit social norm salient. The participation game is a situation in which subjects would generally be better off if all others in their group participated while they abstained themselves. This creates strong free-rider incentives. Drawing a parallel with similar situations outside of the laboratory (an election, a sports competition, a faculty meeting, etc.), participation is prescribed by social norms: a guiding principle they adhere to in order to make the group better off, despite the fact that this can be detrimental from a rational and selfish perspective. The appeal block summons obedience to a social norm insofar as an expectation on how subjects ought to act is presented, and some form of sanctions or rewards is involved. Whereas material sanctions are undoubtedly important for the observance of the conditional cooperation social norm (Fehr and Fischbacher 2004), the fact that aggregate results are made public conjures a “still, small voice of censure or approbation from within”, which has been shown to be crucial for norm adherence (Kerr et al. 1997). The second hypotheses is:

Hypothesis 2: *An appeal that makes an implicit group norm salient leads to higher participation. Participation in the appeal block will be higher than in the other two blocks.*

I refer to these three hypotheses as ‘behavioral’ hypotheses since they imply a departure from standard self-interested preferences (reciprocity must matter for 1a and 1b, and normative behavior for 2). All three hypotheses make clear comparative statics predictions, which are tested by cross-treatment comparisons. In addition to the three behavioral hypotheses, the strategic comparative statics of the game enumerated in Sub-section 3.2.2 will also be tested. This can be done by pooling the data across treatments together and focusing on subject types, subject roles and electorate configurations.

3.4 Results

The experimental sessions were run at the CREED laboratory of the University of Amsterdam in April and June 2013. The experiment was programmed and conducted in z-Tree (Fischbacher 2007). A total of 144 subjects participated in six sessions, three for each treatment. Participants were recruited online from a subject pool of students at the University of Amsterdam.¹⁹ Forty-four per cent of the participants were female and 58% were Economics (and related subjects) majors. The typical session lasted 1 hour and 30 minutes, with average earnings of 18 Euro (which includes a show-up fee of 7 Euro).

I start the analysis of results by looking at the activation decisions of Alphas, and then proceed to test the behavioral hypotheses. At this point I will ignore how the data relates to the point predictions of the model, which will be carried out in a subsequent sub-section. A final sub-section assesses the robustness of the results using regression analysis.

In equilibrium, Alphas should activate two Low Betas and zero High Betas each (see Figure 3.1). Table 3.3 presents activation decision frequencies. We observe that Alphas activate two Low Betas 86.1% of the times, but at least one High Beta is also activated 82.4% of the times. The median activation decision is 2 Low Betas and 3 High Betas, i.e. the Alphas activate all the Betas in their group more than half of the times. This over-activation contradicts the equilibrium predictions. However, we do observe that Alphas activate Low Betas before High Betas; this so-called minimum rationality requirement is violated 0.9% of the times only, which lends credibility to the assumption used in the theoretical analysis. The distribution of activation decisions is remarkably stable throughout time, as can be observed in Figure 3.3. The observed over-activation is thus not the product of an escalation (or de-escalation) of competition, in which activation levels would be low in the beginning of a block and high towards the end (or vice-versa). The reasons underlying the observed over-activation lie probably somewhere between social preferences and joy of winning. An Alpha member facing the decision of allocating a budget between his or her own pocket and activating his or her group's members, tends to choose the latter. This maximizes the chances of his or her group winning at the cost of reduced earnings for himself or herself. For example, the equilibrium active electorate, $(m_i^*, m_j^*) = (2, 2)$, entails an expected payoff of 5.5 for each Alpha, while the modal electorate,

¹⁹Two pilot sessions were run (48 subjects) prior to those reported here. These sessions implemented a different parameter configuration, and served the purpose of testing subject comprehension of the experimental protocol, assessing the pace of subjects in the experiment and providing data for model calibration.

$(m_i, m_j) = (5, 5)$, implies an expected payoff of 2.5. An Alpha activating 2 Low Betas makes his or her group an ‘easy prey’ for a competing group with more than 2 active Betas, despite his or her high material prospects. For moderate levels of social preferences the material losses imposed on her group’s Betas might not compensate the personal gain. In the same direction, an Alpha who wants to simply win the game will activate all the Betas since the probability of winning is strictly increasing in the own group’s size, all else equal.

The actions of an Alpha who activates many Betas at a personal material cost, such that Betas can participate and determine the outcome through their actions, are probably regarded as kind acts. The fact that Alphas deviate from the equilibrium predictions, both in the case of full (NE) and bounded (QRE) rationality, disconfirms the model’s predictions under the assumption of selfish preferences on the part of Alphas. However, this should not interfere with the test of the listed hypotheses; if anything, Alphas that over-activate are more likely to be regarded as kind and command more reciprocity as a consequence.

		Low Betas			
High Betas	0	1	2	Total	
0	8.3	4.6	4.6	17.6	
1	0	0	3.7	3.7	
2	0	0	20.4	20.4	
3	0	0.9	57.4	58.3	
Total	8.3	5.6	86.1	100	

Table 3.3: Activation decisions

Notes: Frequency of each active electorate in the data (in percentage).

3.4.1 The Impact of Mobilization

Before proceeding with the analysis, a previous point on the data’s dependence structure is in order. I argue that there is a substantial degree of independence of observations across blocks. To be sure, the independence of observations across blocks can only be assessed statistically, something that is done in detail in Appendix 3.E. That analysis shows that the results of one block do not influence behavior in subsequent blocks. This is not to deny that experience accumulates and that learning can take place, but rather that the results experienced in one block do not influence subsequent block’s observed behavior. For the non-parametric tests I will use a pair of electorates in a given block (i.e. all subjects in the same session) as an independent observation. All tests are two-sided, unless otherwise noted.

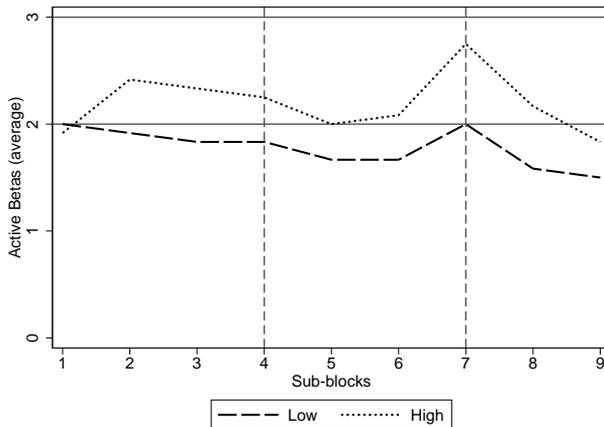


Figure 3.3: Activation decisions of Alphas through sub-blocks.

The first hypothesis posits that participation should be higher when subjects are mobilized by a human subject (*Mob*) relative to when they are not (*Ctr*). The rationale underlying this conjecture is that reciprocity should lead to an increase in participation. To test this hypothesis we look at the aggregate levels of participation of *Mob* and *Ctr*, which are depicted in Figure 3.4. Aggregate participation is measured conditional on being active, i.e. it is the ratio of the number of active Betas who participate in an active electorate divided by the size of the active electorate, $m_i + m_j$.

The data shown in Figure 3.4 provides no evidence of an overall impact of mobilization on aggregate participation levels. The participation levels in the two treatments are remarkably close to each other: 75.2% in *Mob* and 75.8% in *Ctr*. Participation is higher in *Ctr* in the first block, whereas in the second block the pattern is reversed. In the last block participation levels in the two treatments are close to each other. Since the experimental design imposes the same pattern of activation in the two treatments, I carry out Wilcoxon matched-pairs signed-ranks tests (W-MP) on the data. Using the average participation in a block as an observation, a comparison between the 9 pairs of treatment and control observations shows no statistical evidence of a mobilization effect (W-MP $p = 0.76$). Restricting attention to the three pairs of observations from the two blocks where a difference seems to be present, we obtain no statistical evidence for block 1 (W-MP $p = 1.00$) and marginal insignificance for block 2 (W-MP $p = 0.11$). These results point to the irrelevance of mobilization for aggregate participation. The null of no difference between *Mob* and

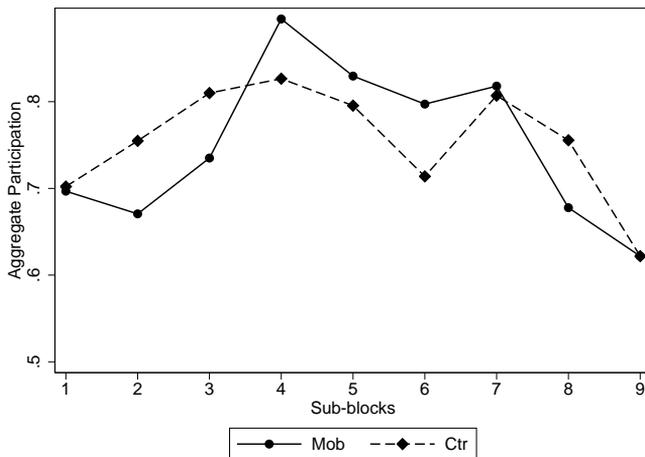


Figure 3.4: Aggregate participation across treatments

Notes: Each point represents average aggregate participation in the corresponding sub-block.

Ctr cannot be rejected, which means that reciprocity from the Betas to the Alphas does not seem to influence participation decisions. However, there is a hint of an interaction effect between mobilization and the appeal. The participation levels in the appeal block of *Ctr* are in line with the other two blocks, whereas in *Mob* that is not the case. In *Mob* the appeal seems to induce a boost in participation. Given the low number of observations of the presented non-parametric test, a more definite answer to this question will be given in the next sub-section.

Hypothesis 1b, a corollary of Hypothesis 1, states that if reciprocity lead to higher participation, High cost subjects would participate at higher levels as their activation should be perceived as a kinder act. The observed null effect of reciprocation can in principle be due to a differential response from High and Low cost subjects, e.g. Low cost subjects are irresponsive to activation while High cost subjects are responsive, but since the former are activated more often the pattern does not show up in the data. In order to investigate this possibility, Figure 3.5 presents the participation patterns of the two cost level groups in each treatment. Whereas High cost subjects participate at lower rates than Low cost subjects (80% versus 72%), the response to treatments is not more pronounced for High cost subjects. We fail to reject Hypothesis 1b (W-MP $p > 0.50$).

A final check on the potential role of reciprocity can be carried out by investi-

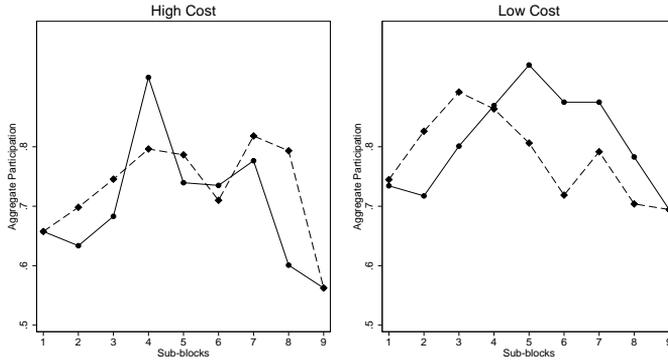


Figure 3.5: Participation and individual costs

Notes: This graphic presents the data of Figure 3.4 broken down by cost levels (High and Low).

gating whether reciprocal attitudes correlate with participation behavior. For this purpose, I define two types of reciprocator: strong and weak. A strong reciprocator is a subject who, as a receiver in the trust game, returns more than what the sender sent her. A strong reciprocator sends back simultaneously more than 4 and more than 8 tokens in case the sender sent her 4 or 8 tokens, respectively; recall that tokens are multiplied by 3 by the experimenter, which means that the receiver gets 12 or 24 tokens in each of these cases. In other words, a strong reciprocator leaves the sender better off than he would be had he not trusted the receiver. All other subjects are weak reciprocators. The share of strong reciprocators in the data is 43%. Figure 3.6 presents the data of Figure 3.4 broken down by reciprocator type. What we observe is that, contrary to our conjecture, strong reciprocators are the least sensitive to treatment manipulations (W-MP $p = 0.37$ and $p = 0.29$ for strong and weak reciprocators, respectively). Instead, it seems to be weak reciprocators who are most susceptible to the manipulation implemented in the appeal block, exhibiting higher participation levels (the participation gap between the appeal block and the two other ones is 14.1 percentage points for weak reciprocators, and 0.9 percentage points for strong reciprocators). The next sub-section investigates the response to the Alpha's appeal in more detail.

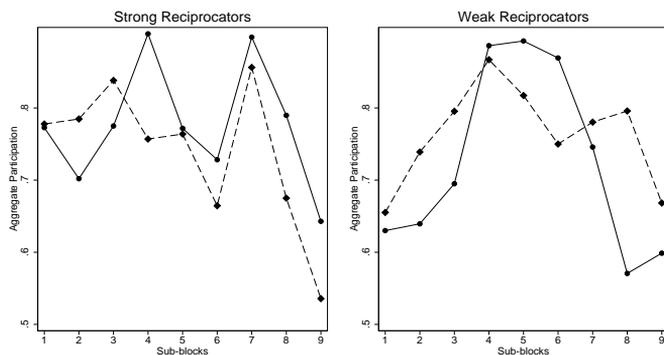


Figure 3.6: Participation by reciprocator type

Notes: This graphic presents the data of Figure 3.4 broken down by reciprocator type.

3.4.2 Participation Appeals and Group Norms

Hypothesis 2 puts forward that an appeal by the Alpha to the Betas further increases participation, provided it proposes a desirable course of action - i.e. a norm for the group to follow. Recall that Alphas have the possibility of delivering written messages to their group's Betas for 90 seconds after the activation decision of sub-block 4. It has been shown that within-group discussion in participation games leads to increased participation (Bornstein et al. 1989, Schram and Sonnemans 1996b, Goren and Bornstein 2000). I expect a one-way appeal to have the same effect, but it is important to stress the differences between the open discussion implemented by the cited works and this experiment's manipulation. Open discussion allows subjects to propose, discuss and commit to a given strategy. The fact that these promises amount to cheap talk has not prevented intra- and inter-group cooperation strategies to arise. Furthermore, open discussion allows subjects to infer others' intentions as well as to create a sense of group identity. The messages sent by the Alpha can propose a strategy and motivate it, but no commitment for the participation game is present as Alphas cannot participate themselves. Alphas could potentially convey intentions about future activation decisions (sub-blocks 5 and 6), but that was never the case. The absence of communication from the Betas to the Alpha also attenuates the case of group identity enhancement.

The appeals sent by the Alphas to the Betas can be found in Appendix 3.F. Alphas sent an average of 15 words to the Betas when they were responsible for activation decisions (*Mob*), and 12 when they were not (*Ctr*). This difference is

not statistically significant according to a Wilcoxon-Mann-Whitney rank-sum test (MW; $p = 0.13$). Two Alphas in *Ctr* chose not to send an appeal to the Betas, while all Alphas in *Mob* sent an appeal. Two typical appeals are reproduced here:²⁰

Alpha₁ (Mob): I want everybody to buy a disc each round

Alpha₂ (Ctr): lets try buying 4 discs. everyone buy one and lets see what they do.

The great majority of appeals request participation (disc buying) from Betas in one form or another: conditional on a message being sent, the share is 86.4% (79.1% if we consider all appeals). This number is similar across treatments: 91.6% in *Mob*, 80% in *Ctr*. In sum, the Alphas transmit an appeal which puts forward a norm to be followed by the group (e.g. ‘I want everybody to buy a disc each round’, i.e. ‘I want everybody to participate’; ‘everyone buy one and lets see what they do’, i.e. ‘everyone participate and let’s see what they do’), often accompanied by a rationale relating to group payoffs or how participation will leave the group better off despite implying a personal sacrifice. This configures a manipulation that comes as close as possible to making an implicit social norm (costly contribution to the group’s effort) explicit in a laboratory environment, without literally invoking external norms.

Figure 3.7 shows aggregate participation in the appeal block compared to the other two blocks. We observe that participation is higher when subjects receive an appeal: 81.1% versus 72.9%. Comparing the aggregate participation rate of the appeal block of with the pooled aggregate participation rate of the other two blocks of each session (i.e. each session contributing two observations, one for the appeal block and one for the non-appeal blocks), the difference proves to be statistically significant (MW $p = 0.02$). The null of no difference between appeal blocks and non-appeal blocks is rejected in favor of the alternative Hypothesis 2. In other words, the Alphas’ appeals induce an increase in participation. However, a closer observation of Figure 3.7 suggest that this effect is mainly driven by the mobilization treatment. In fact, the appeal induces an increase of 13.3 percentage points of aggregate participation in *Mob*, compared to an increase of 3.2 percentage points in *Ctr*. Performing the same statistical test on each treatment separately confirms this observation: whereas the difference is still statistically significant in *Mob* (MW $p = 0.05$), in *Ctr* it becomes insignificant (MW $p = 0.28$).

The more pronounced effect of a group norm by those who are in charge of mobilization can be due to a number of reasons. On the one hand, the Alphas

²⁰Recall that the participation decision is framed as buying a disc. An appeal to buying a disc is therefore equivalent to an appeal to participation.

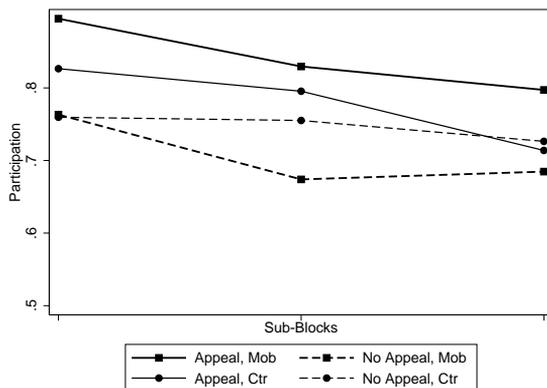


Figure 3.7: Participation and normative appeals

Notes: Participation in Appeal versus non-Appeal blocks, across treatments.

in *Mob* might feel more entitled to demand a particular behavior from Betas, and deliver the norm in a more assertive or motivational way. On the other hand, the group norms might be perceived differently by Betas in the two treatments. For example, a Beta in *Mob* may perceive the Alpha's request as more legitimate than if the exact same request was delivered in *Ctr*. Investigating these points further is interesting, but an eventual difference in the motivational content of messages is endogenous to the mobilization process, and therefore part of the effect we are trying to observe. We conclude that a social or group norm is especially effective when the person delivering it is also the one responsible for mobilizing others, possibly because the message or the context in which it is delivered is, or is perceived as, more motivational, legitimate, etc. The norm becomes less effective when delivered by someone who was not responsible for mobilization.

3.4.3 Test of Equilibrium Predictions

In this sub-section I analyze the experimental data in light of the predictions presented in Table 3.2 and the four comparative statics results put forward in Sub-section 3.2.2. The participation probabilities of Table 3.2 correspond to what we should observe when individuals are self-interested, exhibit a moderate degree of bounded rationality and are not subject to influences that fail to be captured by the strategic structure of the game. In other words, these predictions correspond to a situation where reciprocity concerns and other influences, like a normative appeal,

are absent. Table 3.4 presents the observed participation rates for each Beta type and the resulting aggregate participation rates for each active electorate. As mentioned at the beginning of the section, the activation levels were in general high, which leads to a concentration of data points on the larger electorates.

A noteworthy feature of the data is the high participation levels (above 70%). However, the deviation was driven for the most part by high cost subjects. Subjects seem to be less responsive to cost than the theory predicts and other experimental results have shown (Levine and Palfrey 2007). It is possible that the circumstances of the APG, in which not all subjects can participate and determine the group outcome induces a lower sensitivity to cost.

Roughly 2/3 of the estimated p and q are above the model's predictions. However, on average, the difference is not substantial: the estimated participation probabilities lie 9 percentage points above the predicted probabilities. This deviation is mostly driven by high cost types. The model predicts that high cost types should never participate at a probability higher than 0.5 (see Figure 3.8). In the data this is only supported in lopsided electorates, i.e. when $m_i - m_j \geq 3$. In the remaining cases the participation of high cost types is invariably and substantially above 50%. In fact, high cost subjects participate 25 percentage points above the predicted levels, on average, both when they are in the majority and in the minority. Low cost subjects in the majority over-participate slightly (6 percentage points on average), while those in the minority under-participate slightly (6 percentage points on average, if we exclude the outlier $\hat{q}_l(4, 1)$). We conclude that the calibration of the model captures the behavior of Low Betas to a considerable extent, but fails to predict the behavior of High Betas by a large margin. In fact, participation rates of high cost Betas above 50% cannot be rationalized by any calibration of the model. It appears that subjects are less sensitive to participation costs than our model predicts.

However, the cost effect, which hypothesizes that $\hat{p}_l(m_i, m_j) > \hat{p}_h(m_i, m_j)$ and $\hat{q}_l(m_i, m_j) > \hat{q}_h(m_i, m_j)$ for all (m_i, m_j) , is largely observed in the data. For Betas in the majority, this regularity is observed in all 13 cases. In the minority it is verified in 40% of the cases, but the values are close in terms of magnitude. Restricting attention to balanced electorates, we observe a difference of 17 percentage points between Low and High cost for (4, 4) and 10 percentage points for (5, 5). Comparing the average participation of cost types in all electorates, the statistical evidence supports the existence of a cost effect (MW $p = 0.00$).²¹ The effect is mostly driven by majorities: using a matched pairs signed rank test on the average participation

²¹The unit of observation is the average participation of a given type in each electorate, i.e. a sub-block.

Active Electorate m_i, m_j	Model						Data				n	
	G_i		G_j		Participation		G_j		Participation			
	p_i	p_h	q_i	q_h	G_i	G_j	\hat{p}_i	\hat{p}_h	\hat{q}_i	\hat{q}_h	\hat{G}_i	\hat{G}_j
3,0	0.55	0.15	-	-	0.42	0	0.66 (0.04)	0.13 (0.18)	-	-	0.48 (0.09)	0
4,0	0.52	0.15	-	-	0.33	0	0.45 (0.13)	0.19 (0.26)	-	-	0.32 (0.17)	0
5,0	0.49	0.14	-	-	0.28	0	0.53 (0.32)	0.38 (0.16)	-	-	0.44 (0.13)	0
1,1	0.92	-	0.92	-	0.92	0.92	0.78 (0.15)	-	0.78 (0.15)	-	0.78 (0.30)	0.78 (0.30)
4,1	0.66	0.20	0.56	-	0.43	0.56	0.25 (0.35)	0.19 (0.09)	0	-	0.22 (0.22)	0
5,1	0.61	0.20	0.55	-	0.36	0.55	0.61 (0.27)	0.39 (0.24)	0.22 (0.26)	-	0.47 (0.17)	0.22 (0.26)
2,2	0.90	-	0.90	-	0.90	0.90	0.98 (0.03)	-	0.98 (0.03)	-	0.98 (0.03)	0.98 (0.03)
4,2	0.74	0.33	0.64	-	0.53	0.64	0.84 (0.22)	0.72 (0.13)	0.81 (0.09)	-	0.78 (0.04)	0.81 (0.09)
5,2	0.68	0.28	0.61	-	0.44	0.61	0.72 (0.23)	0.46 (0.14)	0.19 (0.22)	-	0.56 (0.10)	0.19 (0.22)
4,3	0.73	0.41	0.69	0.44	0.57	0.60	0.84 (0.19)	0.73 (0.06)	0.50 (0.45)	0.53 (0.37)	0.79 (0.12)	0.51 (0.34)
5,3	0.69	0.33	0.64	0.38	0.47	0.55	0.91 (0.04)	0.75 (0.06)	0.375 (0.09)	0.44 (0.44)	0.81 (0.05)	0.40 (0.09)
4,4	0.70	0.42	0.70	0.42	0.56	0.56	0.88 (0.20)	0.71 (0.26)	0.88 (0.20)	0.71 (0.26)	0.80 (0.14)	0.80 (0.14)
5,4	0.68	0.36	0.66	0.38	0.49	0.52	0.97 (0.09)	0.78 (0.22)	0.66 (0.27)	0.71 (0.23)	0.86 (0.13)	0.69 (0.14)
5,5	0.66	0.36	0.66	0.36	0.48	0.48	0.92 (0.15)	0.82 (0.17)	0.92 (0.15)	0.82 (0.17)	0.86 (0.13)	0.86 (0.13)

Table 3.4: Model predictions and observed results

Notes: This table presents predicted (columns 3-8) and estimated (columns 4-10) participation probabilities for all observed active electorates. There are 108 active electorates in total: two per sub-block for each session. The two electorates in which Low Betas were not activated before High Betas were excluded. Standard deviations in parentheses.

of low and high cost subjects (the 8th to 11th columns of Table 3.4), the difference is significant for the majority (W-MP $p = 0.00$) but not for the minority (W-MP $p = 0.69$). In general, low cost subjects participate more than high cost ones, but the gap between the two is smaller than predicted. Moreover, it seems that the relative gap between high and low cost subjects' participation rates becomes smaller as the electorates become larger. Focusing on the majority, the ratio $\pi = [\hat{p}_h(m_i, m_j) - \hat{p}_l(m_i, m_j)] / [p_h(m_i, m_j) - p_l(m_i, m_j)]$ is significantly decreasing in $m_i + m_j$, while controlling for $m_i - m_j$.²² As groups become larger and therefore more inclusive, cost differences seem to cease to matter; this is a pattern deserving more investigation.

Another regularity that emerges from the theoretical analysis of the game is the underdog effect, which states that the minority should participate at higher rates than the majority: we should observe $G_i(m_i, m_j) < G_j(m_i, m_j)$ for all electorates where $m_j > 0$, $m_i \neq m_j$ and $m_i + m_j > 3$. With the exception of (4, 2), the underdog effect does not show up in the data, which is in contrast with the results of Levine and Palfrey (2007) and Herrera et al. (2014). In the relevant electorates, average participation is 71% in the majority (G_i) and 48% in the minority (G_j); this difference is statistically significant (MW $p = 0.00$, using the average participation in each side of a sub-block electorate as the unit of observation). The minority seems to 'give up' in very lopsided contests, something that is also present in the data of Cason and Mui (2005). In my framework this tendency might be aggravated by the fact that Betas might protest with abstention whenever their group's Alpha didn't mobilize enough Betas (preventing the two groups from competing on a level playing field).

Balanced electorates are more competitive. Hence, we expect participation probabilities to be decreasing in $m_i - m_j$ while controlling for $m_i + m_j$ (see Section 3.2 for all inequalities). This is a result documented in Levine and Palfrey (2007) and Herrera et al. (2014), for example. Due to the absence of many active electorate configurations in the data, the only comparison we can make is between $(m_i, m_j) = (5, 1)$ and $(m_i, m_j) = (4, 2)$. The competition effect posits that $p_\sigma(5, 1) < p_\sigma(4, 2)$, $\sigma = \{l, h\}$ and $q_l(5, 1) < q_l(4, 2)$. The data produces an ordering that conforms to the hypothesis: $\hat{p}_l(5, 1) = 0.61 < 0.84 = \hat{p}_l(4, 2)$; $\hat{p}_h(5, 1) = 0.39 < 0.72 = \hat{p}_h(4, 2)$ and $\hat{q}_l(5, 1) = 0.22 < 0.81 = \hat{q}_l(4, 2)$. The latter comparison achieves marginal statistical significance (MW $p = 0.06$; $p > 0.16$ for the other two cases). Given the low number of observations, there is moderate support in favor of the competition

²²A simple regression of π on electorate size, $m_i + m_j$, and electorate difference, $m_i - m_j$, yields coefficients (and standard errors) of -0.09 (0.03) and -0.05 (0.05), respectively; $N = 14$ and $R^2 = 0.38$. The data for the minority is less conclusive

effect.

As the size of the electorate increases, aggregate participation should decrease. For the sake of comparison, let us restrict attention to cases in which $m_i = m_j$. The model predicts $G(1, 1) > G(2, 2) > G(3, 3) > G(4, 4) > G(5, 5)$. This is largely disproved by the data: $\hat{G}(2, 2) > \hat{G}(5, 5) > \hat{G}(4, 4) > \hat{G}(1, 1)$. The weak size effect found in the literature is not observed in the data.

In sum, the aggregate participation levels observed in the experiment are substantially above the ones predicted by theory. The explanation lies in the over-participation of high cost subjects, which shows that subjects might be less cost sensitive than the available evidence suggests. Despite the mismatch between the theory and the data in terms of magnitude, the comparative statics of cost are mostly observed in the data. The cost effect follows the predicted pattern, despite the lower gap between high and low cost types. The underdog effect is not observed, as the minority seems to give up in lopsided electorates. The competition effect is reasonably supported by the data. However, participation does not seem to decrease as electorate size increases, i.e. no weak cost effect is present.

3.4.4 Regression Analysis

The results presented up to this point show that reciprocity does not seem to drive a response to mobilization efforts and that subjects with different reciprocity preferences do not seem to participate differently. However, we have seen that there is a boost to participation when the mobilization effort is accompanied by a group norm appeal. From a strategic point of view, participation costs seem to matter but less than predicted by theory. In what follows I will use regression analysis to test the robustness of these findings. Table 3.5 presents two models in which individual participation is regressed on the treatment variables, participation costs and trust and reciprocity preferences. Marginal effects are estimated using the method of Ai and Norton (2003) and Norton et al. (2004), and evaluated at the average sample value of all variables. Model 1 presents no interaction effect between the treatment variables whereas Model 2 does.

Regarding reciprocity preferences, the pattern observed so far is confirmed: the variable Reciprocator Type is insignificant, attesting that reciprocal behavior is largely irrelevant for individual participation behavior. The amount trusted (Trust) is equally irrelevant. Participation costs, on the contrary, matter. For the average subject in the sample, a high participation cost leads to a 13.73 percentage point decrease in participation relative to a low participation cost. In line with what was concluded in sub-section 3.4.3, the magnitude of the cost effect is smaller than

predicted by the theoretical analysis, but is nevertheless substantial.

	Model 1		Model 2	
	Coeff.	Mg.Effect (%)	Coeff.	Mg. Effect (%)
Activation	-0.21 (0.29)	-3.58 (0.05)	-0.35 (0.29)	-3.19 (0.06)
Appeal	0.41*** (0.08)	6.95*** (0.02)	0.14 (0.11)	7.23*** (0.02)
Activation*Appeal	-	-	0.54*** (0.15)	9.76*** (0.04)
Cost	-0.79*** (0.08)	-13.73*** (0.04)	-0.80*** (0.08)	-14.57*** (0.04)
Trust	0.04 (0.05)	0.69 (0.01)	0.04 (0.05)	0.70 (0.01)
Reciprocator Type	0.30 (0.31)	6.03 (0.06)	0.30 (0.30)	6.23 (0.06)
Constant	1.49*** (0.47)		1.59*** (0.47)	
Log Likelihood		-3212.09		-3205.95

Table 3.5: Panel regression results

Notes: This table presents model coefficients and marginal effects estimated using a panel data logit specification with random effects. The marginal effects are reported as percentage values and are evaluated at sample average values. Standard errors in parentheses, N=142 in both models. *** indicates significance at the 1% level.

As suggested by the non-parametric analysis, mobilization alone (i.e. activation) does not seem to influence participation. Both the coefficient and the marginal effects are insignificant. In contrast, the effect of a normative appeal is positive and significant according to Model 1, but it becomes insignificant when the interaction between the appeal and activation is included in Model 2. This interaction effect is highly significant, which leads us to conclude that an appeal per se does not lead to a significant increase in participation. However, when the subject delivering the appeal is responsible for mobilization the effect is considerable and highly significant. The observed boost in participation is 9.76 percentage points, which is in line with the typical effect of canvassing found in field studies.²³

²³The marginal effect of an appeal alone is significant because it is mostly driven by the interaction effect. Restricting the sample to *Ctr* sessions, the marginal effect drops to 3.08% and is insignificant.

3.5 Conclusion

This work provides a laboratory test of the psychological mechanisms underlying the effectiveness of political mobilization. Political mobilization is a massive endeavor in many modern electoral democracies, with a considerable impact on participation behavior. The fact that simple gestures and appeals can foster participation has fascinated social scientists for a long time. In the words of Cox (1999), “one may ask why it is rational for voters to participate merely because they are urged to vote.” The question is particularly puzzling given the evidence discussed in the Introduction, showing that strangers seem to be at least as good as friends when trying to persuade others to participate in politics. According to Cox (*idem*), “if the person urging a citizen to vote is a stranger”, explaining why mobilization works “is harder.” He conjectures that “perhaps there exists a norm enjoining people to vote and they are reluctant publicly to violate this norm.” The evidence presented in this chapter lends support to this conjecture: personalized normative appeals are crucial for mobilization to work, even if delivered by absolute strangers. This finding also helps explain the fact that some mobilization efforts are more effective than others, e.g. why door-to-door canvassing is more effective than direct mailings.

In this chapter I tested two channels through which mobilization could work: reciprocity and group-level normative appeals. Citizens reached by an activist or a campaign might perceive the mobilization act as kind, and wish to return the kindness by participating. Concurrently, mobilization efforts invariably include an appeal akin to a watchword or slogan. These appeals often harness the compliance to a social or group norm, be it participating in a national election or a local demonstration. The social or group norm is made salient by the mobilizers, which renders its violation more prone to both internal and external sanctions.

The experimental design tests these two non-mutually-exclusive forces by having subjects play a game in which two groups compete for benefits on the basis of participation. Participation is an individual and costly decision. The novelty of the proposed framework is that subjects have to be activated (mobilized) by another subject in the same group, who faces a trade-off between keeping the budget handed to him or her by the experimenter and activating subjects in his or her group. I compared treatments in which the mobilization process is in control of a human subject, and is therefore intentional, to treatments where this process is not in the hands of a subject and thus perceived as non-intentional. If reciprocity considerations drove the response to mobilization, we should observe those activated by a human subject participating at higher rates than those who were told they simply became active. In addition, in a part of the experiment the subject responsible for activation

decisions could send an appeal to the subjects in his or her group. This appeal usually provided a norm for the group - a request for others to participate.

The results show that mobilization does not lead to higher participation when taken separately, hinting at the irrelevance of reciprocity concerns for participation decisions. Furthermore, a measurement of individual trust and reciprocity preferences (administered at the beginning of the experiment) does not correlate with participation behavior. However, appeals to participation seem to be effective in raising participation levels, especially when this appeal is sent by the subject responsible for mobilizing his or her peers. The appeal typically provides a norm for the group (e.g. 'I want everybody to participate'). An appeal by a subject who was not responsible for this process is of similar nature, but it does not lead to a significant increase in participation.

The overall evidence sheds light on the psychological mechanisms underlying the mobilization process. The experimental analysis presented in this chapter allowed for a decomposition of the mobilization process into the material act of mobilizing others and the normative appeal simultaneously conveyed. It seems that the simple investment of material resources that make participation possible do not garner enough gratitude as to increase participation. However, normative appeals delivered by the individual responsible for mobilization are extremely effective. This suggests that violating the normative appeal conveyed by the subject responsible for mobilization is psychologically costly, and thus primes subjects into participating more.

Whether this evidence describes the processes taking place in instances of political participation outside of the laboratory is (always) an open question, but it suggests a direction for future field work. In particular, we should try to explore what is hinted at by this experiment, in particular by decoupling the material resources invested in mobilization efforts from the strength of their normative content. For theoretical work in the rational choice tradition, this chapter provides a starting point for incorporating mobilization in a pivotal voter framework. It also suggests that modeling the mobilization process can perhaps eschew preferences for reciprocity aspects, but should take seriously into account the intricacies of norm transmission.

Appendix

3.A The Model - Details

3.A.1 Participation Probabilities for Beta Players

This sub-section presents the derivation of the probability vector in equation 3.1. For an active electorate (m_i, m_j) , let $m'_h \leq m_h$, $h = i, j$, be the number of active Betas who choose to participate. From the perspective of Beta i , define m''_i and m''_j as the number of other players who participate in G_i and G_j , respectively. The mechanics of the model are identical to the basic case presented in Chapter 1, with the only exception that now $m''_i < m'_i \leq m_i \leq M_i$ (players other than i who participate, players who participate including i , active Betas in G_i , size of G_i), where before we had $m'_i \leq m_i \leq M_i$ (players other than i who participate, players who participate including i , size of G_i).

For group G_i , equation 3.1 can be expanded as follows:

$$\begin{aligned}
 v_i(m_i, m_j) &= \left[\Pr_i(\text{win}|m_i, m_j) \quad \Pr_i(\text{tie}|m_i, m_j) \quad \Pr_i(\text{lose}|m_i, m_j) \right]' \\
 &= \left[\Pr[m'_i > m'_j | (m_i, m_j)] \quad \Pr[m'_i = m'_j | (m_i, m_j)] \quad \Pr[m'_i < m'_j | (m_i, m_j)] \right]'
 \end{aligned} \tag{3.A.1}$$

and analogously for G_j .

Determining the value of each term in this vector requires knowledge of the participation probabilities of each Beta type in the corresponding active electorate, (m_i, m_j) . These equilibrium probabilities are obtained following the procedure described in Chapter 1, equations 1.1 and 1.3. The direct equivalent of equation 1.3 is:

$$\Pr[m''_i = m''_j] + \Pr[m''_i = m''_j - 1] = \frac{2c_\sigma}{(B^W - B^L)}, \sigma = l, h. \tag{3.A.2}$$

To proceed with the analysis we only need to specify the probability terms in these equations. Let $M^l \leq M$ be the number of Low Betas in a group. Define the probability of participating as p_l and p_h for Low and High Betas in G_i , respectively; and q_l and q_h for Low and High Betas in G_j , respectively. In other words, we base our analysis on a quasi-symmetric mixed equilibrium, in which the same type within a group implements the same strategy.

The formulas that follow refer to a Low Beta in G_i ; formulas for High Betas

and Betas in G_j can be written in an analogous way. Recall that we have assumed that Alphas activate Low Betas before High Betas. If $m_i = 0$ and $m_j \leq M_j^l$, the probability terms in equation 3.A.2 are simply $\Pr [m_i'' = m_j''] = (1 - q_l)^{m_j}$ and $\Pr [m_i'' = m_j'' - 1] = m_j q_l (1 - q_l)^{m_j - 1}$. If $0 < m_i \leq M_i^l$ and $m_j \leq M_j^l$ the probability terms in equation 3.A.2 are:

$$\begin{aligned} & \Pr [m_i'' = m_j''] & (3.A.3) \\ & = \sum_{k=0}^{\min[m_i-1, m_j]} \binom{m_i-1}{k} \binom{m_j}{k} p_l^k (1-p_l)^{m_i-1-k} q_l^k (1-q_l)^{m_j-k} \end{aligned}$$

$$\begin{aligned} & \Pr [m_i'' = m_j'' - 1] & (3.A.4) \\ & = \sum_{k=0}^{\min[m_i-1, m_j-1]} \binom{m_i-1}{k} \binom{m_j}{k+1} p_l^k (1-p_l)^{m_i-1-k} q_l^{k+1} (1-q_l)^{m_j-k-1} \end{aligned}$$

If $m_i > M_i^l$ and $m_j \leq M_j^l$, the probability terms in equation 3.A.2 are:

$$\begin{aligned} & \Pr [m_i'' = m_j''] & (3.A.5) \\ & = \sum_{k=0}^{m_j} \left\{ \begin{array}{l} \binom{m_j}{k} q_l^k (1-q_l)^{m_j-k} \\ \sum_{s=0}^k \binom{M_i^l-1}{s} \binom{m_i-M_i^l}{k-s}^* \\ p_l^s (1-p_l)^{M_i^l-1-s} p_h^{k-s} (1-p_h)^{(m_i-M_i^l)-(k-s)} \end{array} \right\} \end{aligned}$$

$$\begin{aligned} & \Pr [m_i'' = m_j'' - 1] & (3.A.6) \\ & = \sum_{k=0}^{m_j-1} \left\{ \begin{array}{l} \binom{m_j}{k+1} q_l^{k+1} (1-q_l)^{m_j-(k+1)} \\ \sum_{s=0}^k \binom{M_i^l-1}{s} \binom{m_i-M_i^l}{k-s}^* \\ p_l^s (1-p_l)^{M_i^l-1-s} p_h^{k-s} (1-p_h)^{(m_i-M_i^l)-(k-s)} \end{array} \right\} \end{aligned}$$

Finally, if $m_i > M_i^l, i = 1, 2$, the probability terms in equation 3.A.2 are:

$$\Pr \left[m_i'' = m_j'' \right] \tag{3.A.7}$$

$$= \sum_{k=0}^{\min[m_i-1, m_j]} \left\{ \begin{array}{l} \sum_{s=\max[0, k-(m_i-M_i^l)]}^{\min[k, M_i^l]} \binom{M_i^l}{s} \binom{m_i - M_i^l}{k-s} * \\ p_l^s (1-p_l)^{M_i^l-s} p_h^{k-s} (1-p_h)^{(m_i-M_i^l)-(k-s)} \\ * \\ \sum_{s=\max[0, k-(m_j-M_j^l)]}^{\min[k, M_j^l]} \binom{M_j^l}{s} \binom{m_j - M_j^l}{k-s} * \\ q_l^s (1-q_l)^{M_j^l-s} q_h^{k-s} (1-q_h)^{(m_j-M_j^l)-(k-s)} \end{array} \right\}$$

$$\Pr \left[m_i'' = m_j'' - 1 \right] \tag{3.A.8}$$

$$= \sum_{k=0}^{\min[m_i-1, m_j-1]} \left\{ \begin{array}{l} \sum_{s=\max[0, k-(m_i-M_i^l)]}^{\min[k, M_i^l]} \binom{M_i^l}{s} \binom{m_i - M_i^l}{k-s} * \\ p_l^s (1-p_l)^{M_i^l-s} p_h^{k-s} (1-p_h)^{(m_i-M_i^l)-(k-s)} \\ * \\ \sum_{s=\max[0, k-(m_j-M_j^l)]}^{\min[k+1, M_j^l]} \binom{M_j^l}{s} \binom{m_j - M_j^l}{k+1-s} * \\ q_l^s (1-q_l)^{M_j^l-s} q_h^{k+1-s} (1-q_h)^{(m_j-M_j^l)-(k+1-s)} \end{array} \right\}$$

The probability terms presented in equations 3.A.3-3.A.8 can be plugged into the equilibrium conditions (equation 3.A.2) to derive both Nash equilibria (pure and mixed) and QRE. For the reasons discussed in Chapter 1, this chapter implements only QRE.

3.A.2 Quantal Response Equilibrium

Beta Players

Following the procedure outlined in Appendix 1.B of Chapter 1, we can proceed with the derivation of QRE for each player type in a given active electorate. In equilibrium, the participation probabilities of the two types in both groups are determined simultaneously, and represent a probabilistic best response to all other Betas' probabilistic best responses. For each value of μ^{Beta} there exists a logit

equilibrium, but uniqueness is not guaranteed. However, it is possible to identify a branch that pins down an equilibrium for each value of $\mu^{Beta} \in (0, \infty)$; when $\mu^{Beta} \rightarrow 0$ the branch converges to a Nash equilibrium (McKelvey and Palfrey, 1995).

Having determined p_σ and q_σ for every (m_i, m_j) , it is straightforward to derive the probability distribution of events in each active electorate. This is done by plugging the equilibrium values of p_σ and q_σ in equations 3.A.3-3.A.8. Define the resulting probability vector as $v_i(m_i, m_j, \mu^{Beta})'$, which is identical to the probability term presented in equation 3.A.1 except that we explicitly acknowledge that the equilibrium probability distribution depends on μ^{Beta} and the ensuing equilibrium values of p_σ and q_σ . Figure 3.8 depicts the participation probabilities for each type of player in each possible electorate. Table 3.2 in Sub-Section 3.2.1 presents the values that result from this procedure for $\mu^{Beta} = 0.4$.

Alpha Players

A similar procedure can be applied to the Alpha's decision in order to obtain a logit QRE, with the difference that we must implement a multinomial logistic specification, as Alphas have 6 actions at their disposal ($m_i \in [0, 5]$). Define $\sigma_i(m_i, m_j, \mu^{Beta})$ as the payoff accruing to Alpha i when the active electorate is (m_i, m_j) and μ^{Beta} is assumed. The payoff is thus:

$$\sigma_i(m_i, m_j, \mu^{Beta}) = A_i - a(m_i) + v_i(m_i, m_j, \mu^{Beta})' b \quad (3.A.9)$$

which is analogous to equation 3.2.

As it was the case with the Betas, we need to extend Alpha's payoffs by a stochastic component: $\sigma_i(m_i, m_j, \mu^{Beta}) + \mu^{Alpha} \varepsilon_i^{m_i}$. Define $s_i := \Pr[m_i]$ as the probability assigned by Alpha i to activating m_i Betas. Imposing the assumptions on $\varepsilon_i^{m_i}$ discussed in Chapter 1, we can write:

$$s_i = \frac{\exp\left[\frac{\sigma_i(m_i, m_j, \mu^{Beta})}{\mu^{Alpha}}\right]}{\sum_{m_i=0}^5 \exp\left[\frac{\sigma_i(m_i, m_j, \mu^{Beta})}{\mu^{Alpha}}\right]}, i = 0, \dots, 5 \quad (3.A.10)$$

Fixing μ^{Beta} , Alphas know $v_i(m_i, m_j, \mu^{Beta})$ and therefore also the payoffs of the activation game, $\sigma_i(m_i, m_j, \mu^{Beta})$. The payoffs for the case $\mu^{Beta} = 0.4$ are reproduced in Table 3.6.

These payoffs constitute the input for calculating s_i . For a given μ^{Alpha} , the QRE consists of a probability distribution over the m_i for each Alpha, which is a best

(m_i, m_j)	0	1	2	3	4	5
0	6.50,6.50	5.11,7.39	5.26,6.74	5.26,5.74	5.25,4.75	5.25,3.75
1	7.39,5.11	6.00,6.00	4.92,6.58	4.94,5.56	4.94,4.56	4.93,3.57
2	6.74,5.26	6.58,4.92	5.50,5.50	4.81,5.19	4.75,4.25	4.69,3.31
3	5.74,5.26	5.56,4.94	5.19,4.81	4.50,4.50	4.09,3.91	3.94,3.06
4	4.75,5.25	4.56,4.94	4.25,4.75	3.91,4.09	3.50,3.50	3.22,2.78
5	3.75,5.25	3.57,4.93	3.31,4.69	3.06,3.94	2.78,3.22	2.50,2.50

Table 3.6: Payoff matrix of Alphas for $\mu^{Beta}=0.4$.

response to the other Alpha's probability distribution. Figure 3.1 in the main text depicts QRE for different values of μ^{Alpha} . The Nash Equilibrium, $(m_i^*, m_j^*) = (2, 2)$, can be easily computed from the payoff matrix presented in Table 3.6.

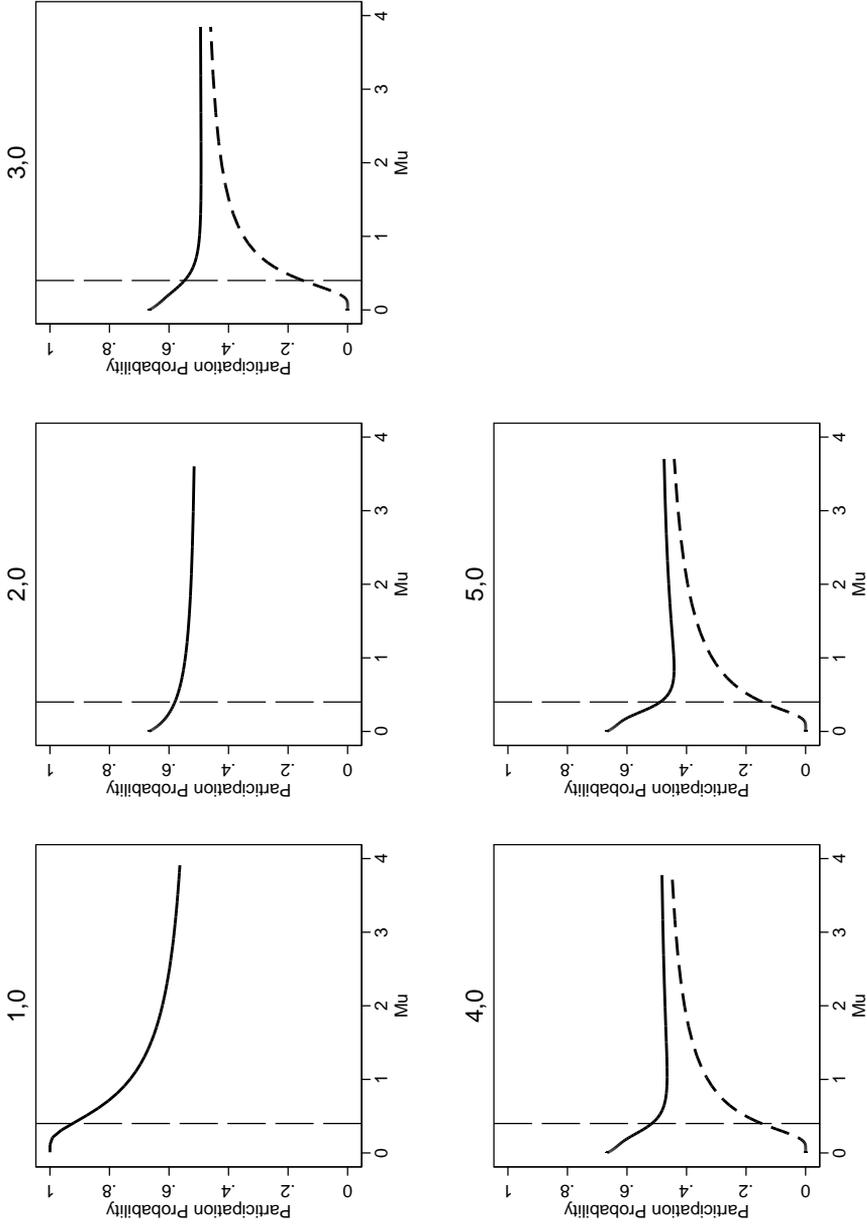


Figure 3.8: QRE equilibrium participation probabilities for different values of μ .

Notes: The numbers above each graph indicate the corresponding active electorate (m_i, m_j) . Thick (thin) lines denote G_i (G_j) and dashed (solid) lines denote High (Low) Betas. The vertical bar corresponds to $\mu = 0.4$.

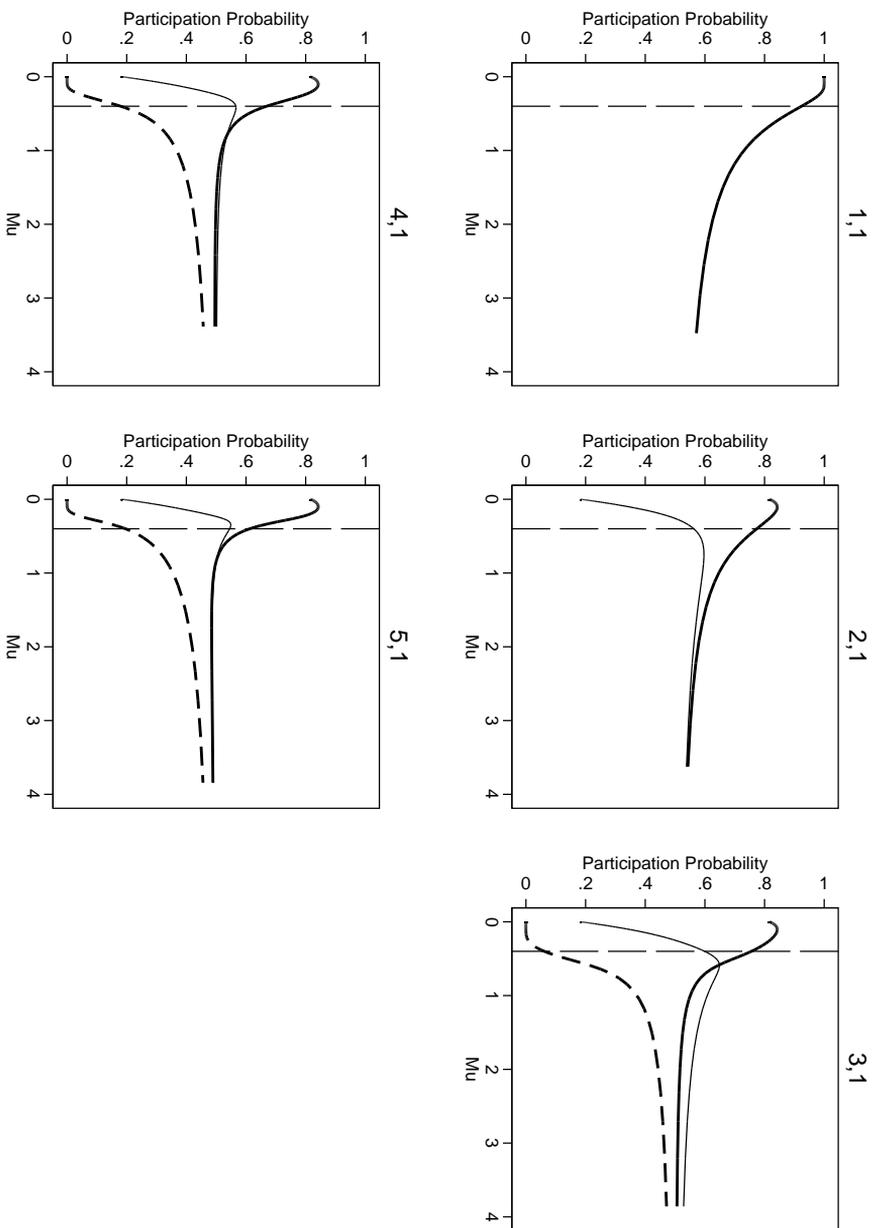


Figure 3.8 (continued).

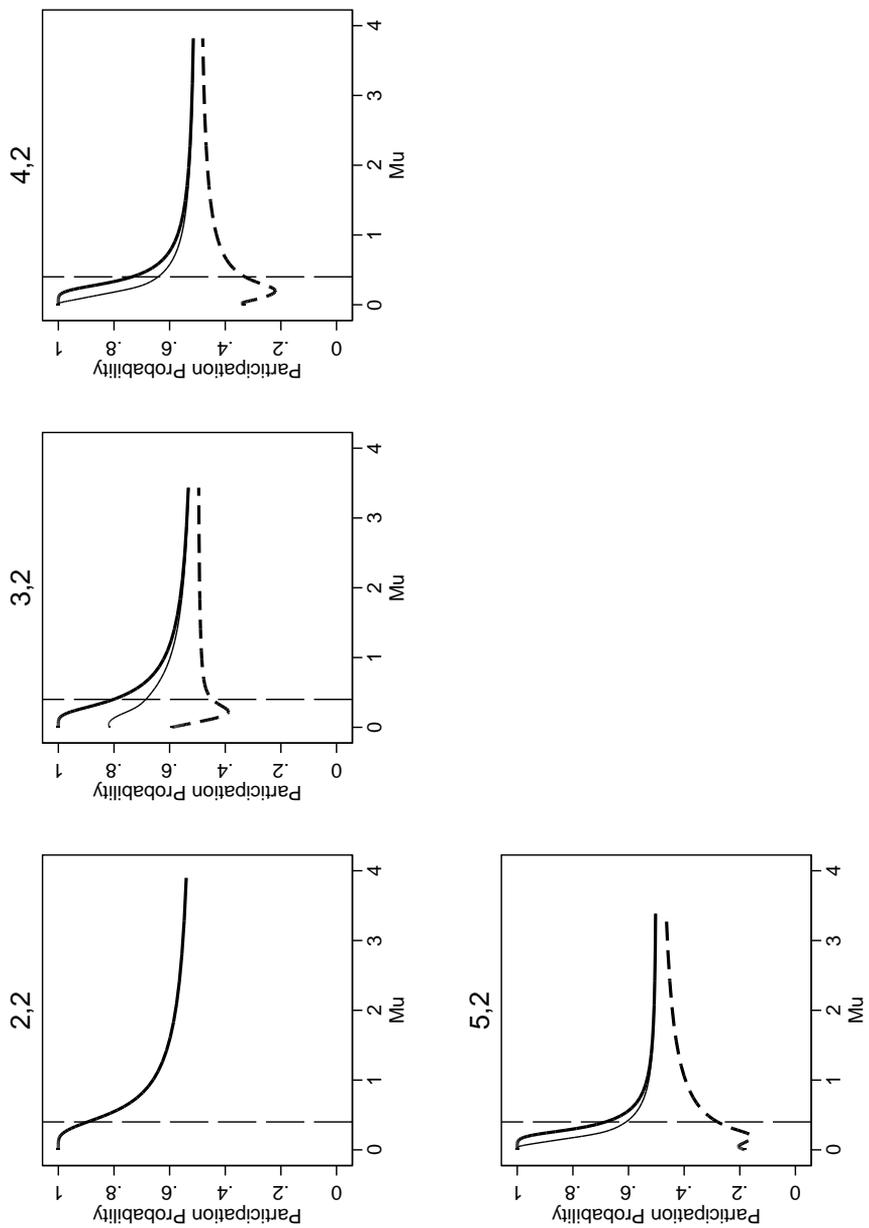


Figure 3.8 (continued).

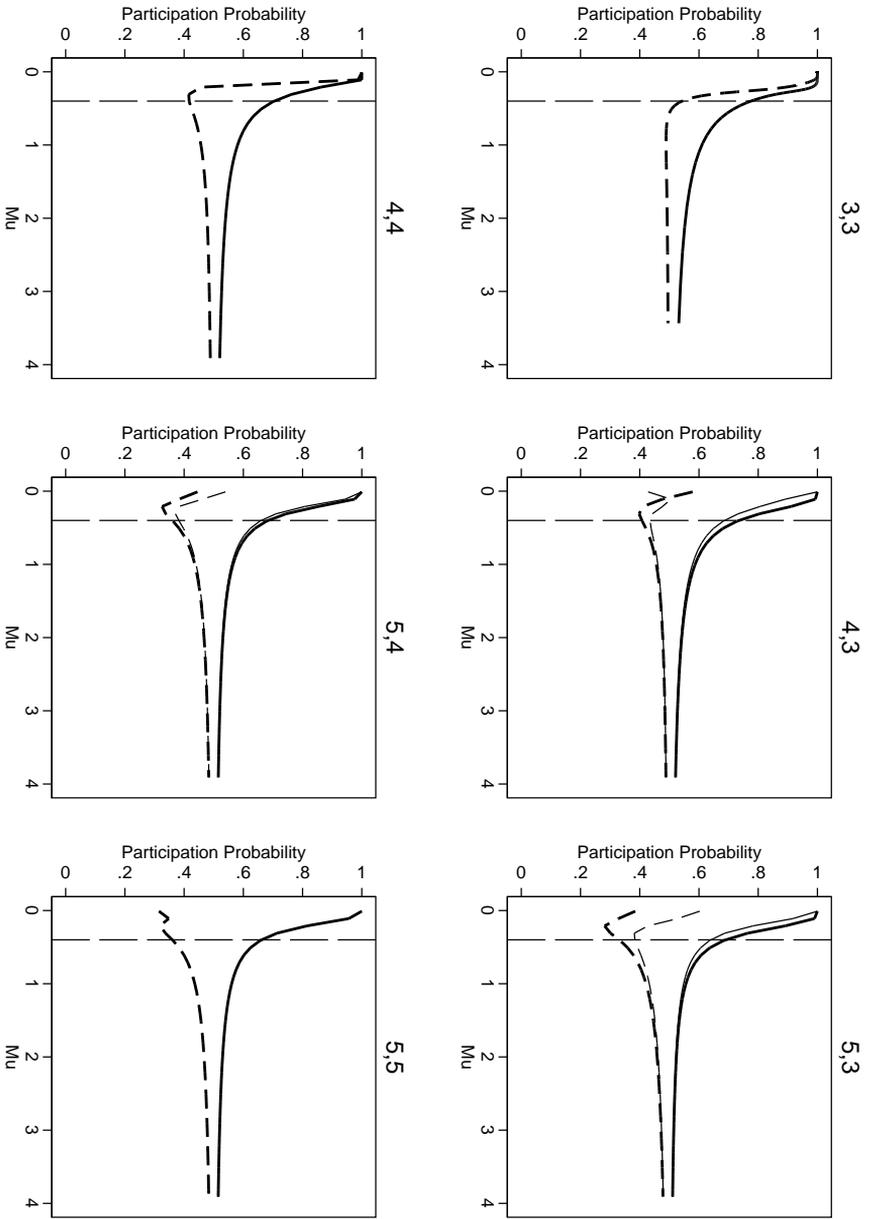


Figure 3.8 (continued).

3.B Experiment Instructions

What follows is an abridged transcript of the experiment's instructions. The changes from *Mob* to *Ctr* are shown within square brackets.

This experiment is composed of two tasks: Task 1 and Task 2. You will receive instructions for Task 2 after Task 1 has been completed. Tasks are made up of rounds. Task 1 has one round. Task 2 is divided into 3 blocks of 27 rounds each, which makes a total of 81 rounds.

In Task 1 you have to make two decisions. You will be paired with two different participants of this experiment. You will not know their identity and they will not know yours. All results from this task will only be revealed at the end of the experiment. For the sake of explanation let us refer to the two participants with whom you are paired as Participant A and Participant B.

For this task you have been given 8 tokens by the experimenter, which you can keep for yourself or send to Participant A. You can send 0, 4 or 8 tokens to Participant A. The tokens that you send to Participant A will be multiplied by 3. Participant A then has to decide how many of them to keep for himself or herself and how many to send back to you.

After subjects submit their decision.

Participant B, who is not the same person as Participant A, has been asked how many tokens he or she wants to send to you. Basically, Participant B was asked to make the same decision which you had to make regarding Participant A. Participant B could choose to send you 0, 4 or 8 tokens out of the 8 tokens that he or she received from the experimenter. These tokens, if any, have been multiplied by 3. This means that if Participant B decided to send you 0 tokens there will be 0 tokens at your disposal, if Participant B decided to send you 4 tokens there will be 12 tokens at your disposal and if Participant B decided to send you 8 tokens there will be 24 tokens at your disposal. The tokens that you do not send back to Participant B are yours and will be converted into earnings.

Task 2 consists of 3 blocks of 27 rounds each, which makes a total of 81 rounds. What follows is a complete description of Task 2. At all times during Task 2 there will be a summary of relevant instructions on your screen.

You are part of a group of 6 participants, you and 5 others. Your group will interact with another group that is identical in every respect and which will face the same decisions. In each group there is 1 Alpha member and 5 Beta members. The Alpha member is appointed randomly for the duration of a block (27 rounds). This means that any participant can be either an Alpha member or a Beta member, and will remain so for the duration of a block of 27 rounds. At the beginning of a new block, new Alpha and Beta members will be randomly chosen. This means that group composition remains constant

for the duration of a block. Groups change when a new block starts.

Each block has the same structure: 1 Activation round followed by 8 Decision rounds, repeated 3 times. In the Activation round only the Alpha member makes a decision. In the Decision rounds all members have to make a decision, even though the kind of decision might differ per member. At the end of each 8th decision round a summary of decisions and payoffs will be shown to all members.

Task 2

In the Activation round only the Alpha member makes a decision. [In the Activation round it is announced how many Beta members have been activated for the following 8 Decision rounds. The activation status of each Beta member will change every 8 rounds and is determined randomly. No decision has to be made by any member in the Activation round.] In the Decision rounds all members have to make a decision, even though the kind of decision might differ per member.

Alpha Member

The Alpha member of each group has to decide how many Beta members to activate. The activation decision is made every 8 rounds and applies to the 8 rounds between activation decisions. Alpha members get an amount of tokens (a budget), which they can either keep for themselves or use to activate Beta members. Activating Beta members has a cost, to be incurred per each round that a Beta member is active. [The Alpha member is affected by the activation state of the Beta members: active or inactive. Alpha members get an amount of tokens (a budget) in each round. Active Beta members represent a cost to the Alpha member. For each round that a Beta member is active a cost is deducted from the budget of the Alpha member.] The budget and costs will be explained in more detail below.

During the 8 Decision rounds the Alpha member will observe and rate the decisions made by the active Beta members. The ratings will not be revealed to any participant of the experiment.

Beta Member

The Beta members of a group can be in one of two states: active and inactive. This depends on whether they have been activated by the Alpha member or not. [The Beta members of a group can be in one of two states: active and inactive, which is determined randomly.]

The Beta members that are active will have to make the following decision. In every round, each active Beta member of a group will have to decide on whether to buy a “disc” or not. A “disc” has a cost, to be explained in more detail below. The members of the group with more “discs” receive a higher reward in that round: 4 tokens. The members of the group with fewer “discs” receive a lower reward in that round: 1 token. If the number

of discs in the two groups is the same, the group who gets the higher reward in that round is picked with equal probability. In other words, in case of a tie each group has a 50% chance of getting the high reward. Note that if one of the groups gets the high reward the other necessarily gets the low reward.

Important: these rewards apply equally to all the members of a group in a given round, be it an Alpha member, an active Beta member or an inactive Beta member.

The Beta members that are inactive observe and rate the decisions of the active Beta members. The ratings will not be revealed to any participant of the experiment.

Budget, Activation Costs and Disc Buying Costs

Both Activation Costs and Disc Buying Costs depend on the type of Beta member. There are 3 High Cost Beta members and 2 Low Cost Beta members in each group. You will be informed which one you are, in case you are a Beta member.

In each round the Alpha member gets a budget of 4 tokens that he can keep for himself or herself or use to activate Beta members. The tokens spent on activating Beta members cannot be recovered. The budget is just enough to activate all Beta members in his or her group, which means that nothing remains of the budget if all Beta members are activated. The cost of activating a High Cost Beta member and a Low Cost Beta member are 1 and 0.5 per round, respectively. Note that the amounts we refer to are per round. However, the Alpha member makes a decision that is valid for 8 rounds. [In each round the Alpha member gets a budget of 4 tokens. The tokens deducted from the Alpha member's budget due to active Beta members cannot be recovered. The budget is just enough to pay for all Beta member activation in his or her group, which means that nothing remains of the budget if all Beta members are active. The costs of an active High Cost Beta member and a Low Cost Beta member are 1 and 0.5 per round, respectively. Note that the amounts we refer to are per round. However, the activation state of Beta members is valid for 8 rounds.]

The cost of buying a disc is also different for High Cost Beta members and Low Cost Beta members: 1 and 0.5 per round, respectively.

In sum, the earnings per round of each member are as follows:

Alpha member = Reward + Budget - Eventual Activation Costs

Active Beta member = Reward - Eventual Disc Buying Cost

Inactive Beta member = Reward

Before we head to the 3 blocks there will be some practice questions to make sure you understand Task 2.

3.B.1 Comprehension Questions

The following questions were administered after the APG's instructions. Subjects could only proceed after a correct answer to all questions.

Q1: What is the maximum number of Beta members that can be activated by the Alpha member? A1: 5.

Q2: If an Alpha member activates all the Beta members of his group, how much is left of the budget? Recall that an Alpha member has a budget of 4 tokens at his disposal, that there are 2 Low Cost Beta members and 3 High Cost Beta members in each group, and that activating a Low Cost Beta member and a High Cost Beta member costs 0.5 and 1, respectively. A2: 0.

Q3: Suppose that the Alpha member activates 2 Low Cost Beta members and 1 High Cost Beta member. Suppose further that, in a given round, these Beta members all buy a disc whereas the Beta members of the other group only buy 1 disc. What is the payoff of the Alpha member in this round? Recall that activating a Low Cost Beta member and a High Cost Beta member costs 0.5 and 1 tokens per round, respectively, and that a member of a group receives 4 tokens if his or her group buys more discs than the other group. A3: 6.

Q4: At the end of a block of 27 rounds, all participants are assigned new (random) member roles and new groups are formed. True or false? A4: True.

Q5: Suppose that the other group buys 2 discs in total. What is the minimum number of discs that your group has to buy such that your group gets the high reward of 4 tokens for sure? A5: 3.

Q6: Suppose that your group bought 2 discs and the other group bought 2 discs as well. What is the probability that your group gets the high reward in this round of the game? Note: Express probability in decimal terms; for example, if you want to answer "20% probability" please type "0.2". A6: 0.5.

Q7: Suppose you are a High Cost Beta member and that all Beta members have been activated in both groups. Suppose further that, in a given round, your group buys 5 discs and the other group buys 4 discs. What are your earnings (in tokens) in this round of the game? A7: 3.

3.C Trust Game Results

The data from the trust game is presented in this Appendix. Approximately half of the subjects chose to send the intermediate amount (4), roughly 30% sent the high amount (8), and 20% sent 0 (see Figure 3.9). Averaging, subjects sent 55.5% of their

endowment (standard deviation: 35.5%), which is consistent with the typical 50% result found in trust games with a finer action space (Johnson and Mislin 2011).

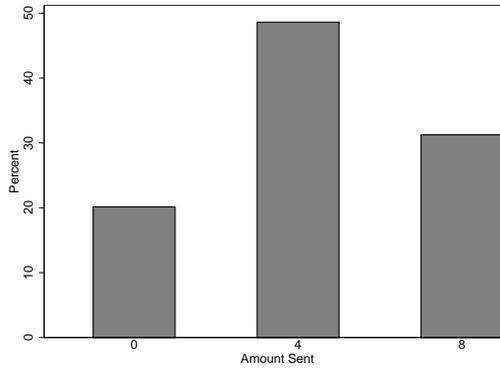


Figure 3.9: Amount sent in the trust game.

The amount returned is also in line with typical results. The distribution of the returned amount is remarkably close for the intermediate and high received amounts: subjects chose to return 34.4% and 35.3%, respectively (standard deviations of 21.2% and 21.0%; see Figure 3.10).

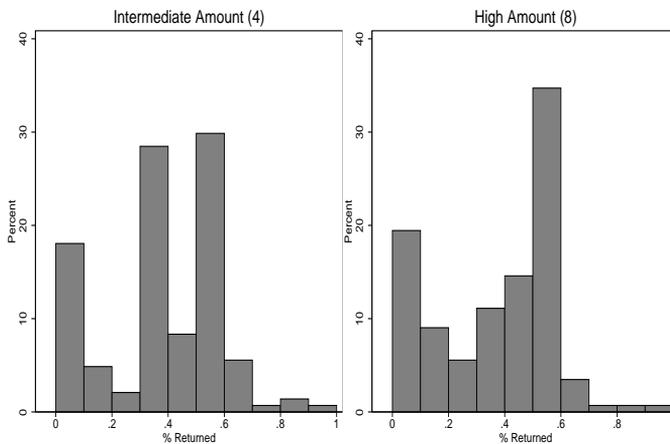


Figure 3.10: Amount returned in the trust game.

3.D Subject Ratings

This Appendix provides a short description of the ratings data. During the APG, Alphas and inactive Betas observe the outcome of every period, which they are asked to rate. Subjects choose between “dissatisfied”, “neutral” and “satisfied” (coded as 1, 2 and 3, respectively). This information provides important clues on what both Alphas and inactive Betas expect from the active Betas (in terms of participation) and what Betas expect from the Alphas (in terms of activation).

Figure 3.11 presents the ratings distribution of Alphas and inactive Betas. The main conclusion to be drawn is that inactive Betas tend to be more dissatisfied than Alphas, which is to be expected if we assume that they would prefer to be active. Alphas, on the other hand, are mostly satisfied with the observed results.

Table 3.7 presents regression results on how Betas perceive activation decisions. The evidence shows that Betas rate high activation decisions favorably, while activation in the other group assumes the expected sign, but is statistically insignificant. Ratings are not significantly higher in the presence of an appeal, and high cost Betas don't rate differently than low cost ones.

Ratings also convey important information on subjects' expectations regarding participation. In particular, we want to know what Alphas and inactive Betas consider satisfactory with respect to the participation behavior of active Betas. This can shed light on the extent to which participation is expected and can therefore be considered norm for group members. Table 3.8 presents regression results, where we can observe that high participation is rated favorably by both Alphas and Betas, while controlling for victory, tying and margin of victory. This means that participation is valued beyond its instrumental consequences, which provides grounding to participation being regarded as normative behavior for the group.

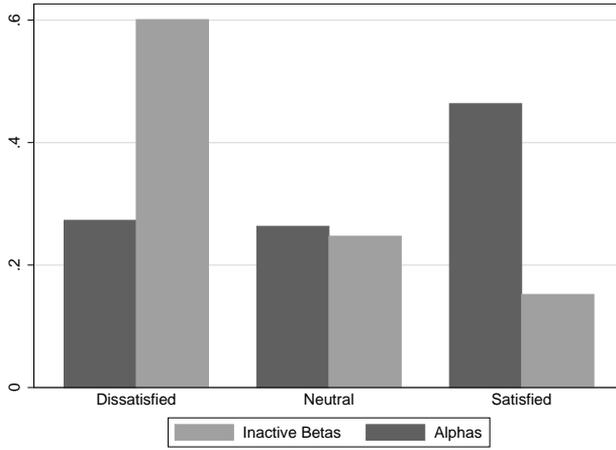


Figure 3.11: Ratings distribution.

	<i>Mob</i>	<i>Ctr</i>
Activation Own Group	0.21*** (0.06)	0.16*** (0.05)
Activation Other Group	-0.03 (0.06)	-0.06 (0.06)
Appeal	0.19 (0.23)	-0.09 (0.20)
High Cost	-0.10 (0.28)	0.31 (0.24)
Constant	1.31*** (0.36)	1.46*** (0.31)
R ² (overall)	0.24	0.13

Table 3.7: Panel regression results

Notes: Panel regression with fixed effects; the cross section variable is subject and the time variable is sub-block. The dependent variable is an inactive Beta's rating of the activation decisions (once per sub-block). Activation Own (Other) Group is the total number of active subjects in the subject's own (other) group, and both Appeal and High Cost are dummy variables. Standard errors in parentheses, N=111. *** indicates significance at the 1% level.

	Alphas	Betas
Participation Own Group _(t-1)	0.33** (0.14)	0.35*** (0.10)
Participation Other Group _(t-1)	-0.29** (0.14)	-0.10 (0.11)
Victory _(t-1)	1.89*** (0.09)	1.51*** (0.13)
Tie _(t-1)	1.10*** (0.06)	0.63*** (0.08)
Margin of Victory _(t-1)	-0.10*** (0.04)	-0.09** (0.04)
Appeal	-0.10 (0.08)	0.18 (0.12)
High Cost		-0.63*** (0.20)
Constant	1.16*** (0.09)	1.61*** (0.08)
R ² (overall)	0.49	0.32

Table 3.8: Panel regression results

Notes: Panel regression with fixed effects; the cross-section variable is subject and the time variable is period. The dependent variable is the Alpha's and the inactive Betas' rating of the results in their electorate, respectively. Participation Own (Other) Group is the participation rate in the subject's own (other) group. Victory is a dummy variable which takes value 1 if the own group achieved outright victory (victories from coin tosses excluded), Margin of Victory is the difference in participation between the own group and the other group, Appeal and High Cost are dummy variables. Standard errors in parentheses. N=1260 and N=744, respectively. **/** indicates significance at the 5%/1% level.

3.E Independence of Observations Across Blocks

Part of the non-parametric analysis presented in this chapter relied on the assumption of observation independence across blocks. The crucial aspect is that there was a complete re-matching of groups and re-assignment of both roles and types from one block to the next. A subject thus faced a probability of 2.97×10^{-5} of being part of the exact same group, while the probability that his or her group was composed exclusively of subjects who were not part of his or her previous group is 0.25. These two aspects dampen the influence of experienced history for future interac-

tion. Furthermore, new roles and types were randomly assigned at the beginning of each block. A High Beta, a Low Beta and an Alpha faced a 1/2, 1/3 and 1/6 probability of being assigned the same role from one block to the next. In addition, earnings were set to 0 at the beginning of each block, and subjects knew that they would only be paid for one randomly picked block.

In order to statistically assess the independence claim, I present regression results on how the determinants of participation correlate across blocks. The procedure consists of two steps: first, I investigate which aspects influence participation behavior from one sub-block to the next; second, I test whether these aspects have an influence across blocks.

Given that feedback is given at the end of each sub-block, it is conceivable that elements of the game in a sub-block influence behavior in future sub-blocks. Such elements include participation levels in the subject's own group and in the other group, whether the subject was pivotal, and whether a tie was observed (e.g. Duffy and Tavits 2008 or Grosser and Schram 2006 use some of these variables as controls in parametric statistical analysis). Table 3.9 presents regression results on the relationship between these variables and participation decisions.

The results show that individual participation in a given sub-block is significantly influenced by having been active in the previous sub-block and by the previous sub-block's participation level in the subject's own group provided he or she was active. Participation in the other group, the number of pivotal events or the number of ties do not seem to have an effect.

Nevertheless, all of these elements are included in regression models that investigate the impact of a given block's play on subsequent ones (Table 3.10). Since subjects are part of multiple groups in the experiment and that lagged dependent variables must be included, I opt for a simple linear regression specification. The dependent variable is average individual participation in block 2 or block 3. The independent variables are average individual participation (p), average group participation (in the own group, P^G , and in the other group, $P^{G'}$), the number of pivotal (Piv) and tie (Tie) events, the share of the block in which the subject's group achieved victory (W), and the share of the block in which the subject was active (A). Since historical effects might depend on whether a subject was active or not, interaction effects are included in some specifications.

Most variables are insignificant, be it in parsimonious or more comprehensive specifications. The results suggest that the history of one block does not affect individual participation in the next block systematically. There exist three statistically significant patterns worth discussing though. First, aggregate participation

in the first block seems to affect individual participation in the second block differently for active and inactive subjects (model 2). Whereas a subject who was inactive tends to participate significantly more in the second block, for active subjects (the large majority) this effect is insignificant ($H_0 : \beta_{P_1^G} + \beta_{A_1 * P_1^G} = 0$, $F(1, 84) = 1.10$, $p = 0.30$). An identical pattern is observed for block 3, but with respect to individual participation (p_2) (models 5 and 6). Again it seems that subjects who were seldom active will tend to participate more in the following block, but the result is not statistically significant for subjects who were often active ($H_0 : \beta_{p_2} + \beta_{A_2 * p_2} = 0$; $F(1, 86) = 0.17$, $p = 0.68$; $F(1, 53) = 0.42$, $p = 0.52$, respectively). In any case, this result would be hardly surprising, as each subject might have an idiosyncratic tendency to participate. A third set of significant variables are observed in model 4, but they are not part of any systematic pattern. All in all, the independence of observations across blocks does not appear to be a problematic assumption.

Individual Participation			
Part. Own Group $_{t-1}$	-0.01 (0.02)	Active $_{t-1}$ *Part. Own Group $_{t-1}$	0.05* (0.03)
Part. Other Group $_{t-1}$	0.03 (0.02)	Active $_{t-1}$ *Part. Other Group $_{t-1}$	-0.02 (0.02)
Tie $_{t-1}$	-0.01 (0.01)	Active $_{t-1}$ *Tie $_{t-1}$	0.02 (0.01)
Piv $_{t-1}$	0.02 (0.02)	Active $_{t-1}$ *Piv $_{t-1}$	-0.00 (0.02)
Active $_{t-1}$	0.16** (0.07)	Constant	0.24*** (0.06)

Table 3.9: Determinants of participation across sub-blocks

Notes: Panel regression with fixed-effects, where the cross-section variable is subject and the time variable is sub-block. The dependent variable is the rate of individual participation (in a sub-block), while the independent variables are the rate of participation in the subject's group (Part. Own Group), the rate of participation in the other group (Part. Other Group), the number of pivotal events (Piv), the number of tie events (Tie) and whether the subject was active (Active). N=1152. ***/**/* indicates significance level at the 1%/5%/10%. level.

	P_2			P_3			P_2			P_3		
	(1)	(2)	(3)	(4)	(5)	(6)	(1)	(2)	(3)	(4)	(5)	(6)
P_1	0.08 (0.17)	-0.69 (0.91)		0.08 (0.17)		-0.01 (0.92)	$A_1^*p_1$	0.89 (0.92)				0.11 (2393)
P_1^G	0.05 (0.08)	0.58** (0.24)		-0.15* (0.08)		-0.19 (0.26)	$A_1^*P_1^G$	-0.69** (0.27)				-0.01 (0.30)
$P_1^{G'}$	-0.01 (0.09)	-0.34 (0.32)		0.07 (0.09)		-0.06 (0.33)	$A_1^*P_1^{G'}$	0.41 (0.35)				0.26 (0.37)
Piv_1	-0.01 (0.02)	-0.02 (0.09)		-0.06*** (0.02)		0.01 (0.11)	$A_1^*Piv_1$	0.01 (0.10)				-0.07 (0.11)
Tie_1	0.01 (0.01)	-0.04 (0.04)		0.01 (0.01)		0.03 (0.04)	$A_1^*Tie_1$	0.05 (0.04)				-0.04 (0.04)
W_1	-0.13 (0.40)	-0.94 (1.59)		1.08** (0.43)		-0.07 (1.66)	$A_1^*W_1$	1.16 (1.72)				1.41 (1.76)
A_1	-0.19 (0.16)	-0.60 (0.79)		-0.16 (0.16)		-0.94 (0.81)						
P_2			0.15 (0.20)	-0.05 (0.23)	1.24* (0.65)	2.11** (0.96)	$A_2^*p_2$				-1.16* (0.66)	-2.36** (1.03)
P_2^G			0.10 (0.11)	0.21 (0.13)	0.24 (0.37)	0.10 (0.47)	$A_2^*P_2^G$				-0.31 (0.42)	-0.24 (0.54)
$P_2^{G'}$			-0.10 (0.10)	-0.13 (0.13)	-0.23 (0.26)	-0.33 (0.29)	$A_2^*P_2^{G'}$				0.32 (0.32)	0.64 (0.36)
Piv_2			0.00 (0.02)	0.01 (0.02)	0.01 (0.11)	0.07 (0.11)	$A_2^*Piv_2$				-0.02 (0.11)	-0.08 (0.12)
Tie_2			-0.01 (0.01)	-0.02* (0.01)	-0.00 (0.05)	0.00 (0.05)	$A_2^*Tie_2$				-0.01 (0.05)	-0.03 (0.06)
W_2			-0.20 (0.43)	-0.46 (0.52)	-0.30 (1.22)	-1.12 (1.32)	$A_2^*W_2$				0.71 (1.44)	2.17 (1.66)
A_2			-0.15 (0.20)	-0.19 (0.23)	-0.22 (0.67)	-1.21 (0.89)	Const.	0.64*** (0.24)	1.08*** (0.71)	0.77*** (0.21)	0.76*** (0.38)	1.59 (1.03)

Table 3.10: Determinants of participation across blocks

Notes: Least squares regression where the dependent variable is the rate of individual participation (in either block 2 or block 3), Subscripts indicate the block to which the variable is reported. N=98 for models (1) and (2), N=100 for models (3) and (5) and N=80 for models (4) and (6). ***/**/* indicates significance level at the 1%/5%/10% level.

3.F Transcript of Alphas' Messages

The transcript of the messages sent by the Alphas to the Betas in their group is presented below. Alpha_i stands for the Alpha of group i ; Group 1 competed with Group 2, and Group 3 with Group 4. The numbers in parentheses after 'Session' indicate the active electorate in sub-block 4 (the one for which choices had been made). The number in parentheses after ' Alpha_i ' indicates the time in seconds at which the message was sent: the clock started at 90 and counted down to 0. Sessions A, B and C implemented *Mob*, the remaining implemented *Ctr*.

Session A (5-5, 4-5):

$\text{Alpha}_1(73)$: Hi there, I think we should all buy discs

$\text{Alpha}_1(37)$: In that way we can make the most tokens

$\text{Alpha}_2(50)$: I want everybody to buy a disc each round

$\text{Alpha}_3(71)$: all choose to buy a disc

$\text{Alpha}_3(40)$: we are with four, but one of the other does not choose to buy. if we all buy we can get 4

$\text{Alpha}_3(34)$: payoff

$\text{Alpha}_4(24)$: Let's work together and get the most tokens for sure!

Session B (5-4, 1-1):

$\text{Alpha}_1(51)$: Please, always buy a disc in every round. We have got 5 members, they've got 4, so we could earn a lot of points here. Thanks in advance.

$\text{Alpha}_2(40)$: okay guys, if you're low cost always buy the disc, payoff will still be 0.5 if you lose, high cost at least 1 per round, payoff could be 4-1

$\text{Alpha}_2(10)$: scorched earth

$\text{Alpha}_3(41)$: i hope you can buy discs, because we have higher chance to win

$\text{Alpha}_4(19)$: the beta member always has to buy a disc! the expected payoff is higher because the other group also has one beta. the expected payoff is 2

Session C (5-5, 5-4):

$\text{Alpha}_1(44)$: please all buy a disk, if we do it together we can make more profit. Otherwise I'd rather make a lot of profit on my own.

$\text{Alpha}_1(2)$: I think the best way is to stick together :)

Alpha₂(74): please all buy a disc!

Alpha₂(7): thats all i have to say lol

Alpha₃(12): I've activated you all. If you all buy discs, we all will earn the maximum to my opinion.

Alpha₄(54): hey, Ok so there is one more beta activated in the other group. here there are 4 and in the other group 5. So in my opinion you should

Alpha₄(33): all buy discs, so we have a higher probability of getting the reward

Alpha₄(24): lets see how it goes :)

Session D(5-5, 4-5):

Alpha₂(58): hi

Alpha₃(41): think critically what to choose, last round the other group had 5 (we only 4) members just as in this case. They all bought disk!

Alpha₄(84): He all :-)

Alpha₄(55): Let's try to buy all the discs we can. So that we always go for the highest reward

Alpha₄(41): It will cost you a little bit, but makes it more likely that you earn more money

Alpha₄(22): 4-1 or 4-0.5 is still more than.. 1-1 or 1-0.5

Alpha₄(16): Let's do it!!!!

Alpha₄(13): And earn some money

Session E (5-4, 1-1):

Alpha₁(80): buy

Alpha₁(60): buy them all

Alpha₁(12): "good luck ;)"

Alpha₂(48): lets try buying 4 discs. everyone buy one and lets see what they do.

Alpha₃(69): In the first 8 rounds we have one active member

Alpha₃(63): and the other group also one active member

Alpha₃(39): so if the active member buys all the time the disc we have 50% chance to get the high reward which is good

Alpha₃(17): I think it would be good to try and have the higher rewards in the other rounds as well

Alpha₄(70): everybody buy discs

Session F (5-5, 5-4):

Alpha₁(29): both groups have 5 members, please all buy discs, it has shown in the past that in round there are always some who dont, we buy all!!!

Alpha₂(54): Ok, we can get them if we play disciplined. Which means every one buy discs every round!

Alpha₃(88): Hi

Alpha₃(77): We should buy all discs each round\

Alpha₃(68): Especially now we have more active beta members

Alpha₃(56): That way we're certain we get 4 points each round

Alpha₃(29): So, everyone. Don't skip.

Chapter 4

Paying is Believing: The Effect of Costly Information on Bayesian Updating¹

The quality of collective decisions often hinges on efficient information aggregation. When information is free to decision makers, increasing the number of decision makers should lead to better decisions. However, when information is costly free-rider incentives in information acquisition arise and collective decisions might become sub-optimal. A common assumption underlying both results is that the quality of decision making does not depend on the cost of information, as any information costs are sunk at decision time. This chapter challenges this assumption and explores whether individual decision making under risk is affected by the cost of information. To do so one must distinguish the effect of cost from self-selection by individuals who value information the most. Outside of the laboratory it is difficult to disentangle these two effects. We design an experimental environment where subjects are offered additional, useful and identical information on the state of the world across treatments. We find a systematic effect of sunk costs on the manner in which subjects update their beliefs. Subjects over-weigh costly information relatively to free information, which results in a ‘push’ of beliefs towards the extremes. This shift does not necessarily lead to behavior more attuned with Bayesian updating. We find that an intensification of representativeness bias is the most likely explanation of our results.

¹This chapter is based on Robalo and Sayag (2014).

4.1 Introduction

Information is a crucial determinant of collective decision quality, be it a large scale election or a committee of experts deciding on the adoption of a certain policy. However, in collective action situations the incentive to acquire costly information is low: one's decision is unlikely to change the collective choice, and therefore there is little reason to invest in information. Information becomes an under-provided public good. This is, in a nutshell, the rational ignorance phenomenon (Downs 1957). Formal models of large elections (Martinelli 2006, 2007) have confirmed Downs conjecture: voters will only acquire information when its cost is very close to or equal to zero. The electorate will be poorly informed and likely to make sub-optimal choices.

The situation is somewhat different when information is free. In a two candidate race (or any binary decision subject to vote), if individuals have identical preferences over candidates but face uncertainty as to which candidate is best in each state of the world, voting is an efficient way of aggregating each individual's pieces of information. This result is known as the Condorcet jury theorem: if information is costless, increasing the size of the electorate will lead to better decisions via efficient information aggregation (Young, 1988).

The fundamental conclusions of this literature assume that the way information is used does not depend on its cost, i.e. the way individuals take into account information does not depend on how much they had to pay for it. In practice, this means that individuals are expected to perform Bayesian updating in the same fashion at all cost levels. However, if the cost of information influences the way it is incorporated in decision making, these results might no longer hold. In particular, a scenario with costly information and an ensuing sub-optimal level of aggregate information acquisition might be preferable if information is better incorporated in decisions. Compare this with a situation in which information is costless, leading to a high level of aggregate information, but in which it is not fully incorporated in decision making because of deficient Bayesian updating.

There exist a few recent experimental tests of voting with endogenous costly information acquisition: Bhattacharya et al. 2014, Elbittar et al. 2014, and Grosser and Seebauer 2014. The former two vary the cost of information but assume, as is customary, that the extent of Bayesian updating does not depend on its cost (the latter work does not change the cost of information). To be sure, the underlying theoretical predictions are Bayesian Nash equilibria, which assume Bayes rule is applied when updating beliefs. They leave open the question of whether, and how, the cost of information interacts with voting outcomes (efficiency and correct choices).

This chapter investigates whether the cost paid for information influences the way it is used in decision making. Since this constitutes a first approach to the question we restrict our attention to an individual decision-making environment. However, there exist several similarities between our set-up and the cited works: there is uncertainty regarding (two) possible states of the world, the optimal choices are state-dependent, and imperfect but informative signals can be used to reduce uncertainty. Absent are the strategic complexity of a voting situation and the potential free-riding incentives in information acquisition.

4.1.1 Sunk Information Costs

Conditional on receiving information, the behavior of a rational individual should not depend on the price paid for it. However, we conjecture otherwise: decision makers might put a higher weight on information they had to pay for, and paying for information can interact with optimization behavior. Underlying our conjecture is the possibility that individuals fall prey to a variant of the sunk cost effect (Thaler 1980), and “use” information relatively more when it comes at a cost.²

We contribute to the accumulated evidence on the sunk cost fallacy by investigating the existence of sunk costs in a scenario of decision making under risk. If a relationship exists between information costs and decision making, a follow-up question is whether it leads to better decisions. We set out to investigate these matters using a laboratory experiment. Field data are likely to be contaminated by serious selection issues: individuals who choose to acquire information in the field are likely to differ along several dimensions from individuals who choose not to do so. The laboratory allows us to correct for these selection issues through carefully constructed procedures. One way in which we disentangle selection from sunk cost effects is by imposing the cost of information on subjects. This is something that is easily done in the laboratory, but arguably difficult to implement in the field.³ Moreover, the laboratory allows us to assess the extent to which individuals value information and are able to use it; in other words, we can identify different types of individuals from their revealed demand for information. To the best of our knowledge, this work provides the first experimental test of how information costs affect

²According to Thaler (1980): “paying for the right to use a good or service will increase the rate at which the good will be utilized, *ceteris paribus*. This hypothesis will be referred to as the sunk cost effect.”

³Field tests of sunk cost effects in product use have been carried out (Arkes and Blumer 1985, Ashraf et al. 2010 and Cohen and Dupas 2010, for example), but doing so with respect to information is arguably more complicated. In particular, measuring product usage (a theater season ticket, a bottle of water disinfectant and bed nets, respectively for the cited works) is easier than measuring information usage.

updating behavior.

The sunk cost fallacy's main prescription is that only marginal costs and benefits should matter for decision-making. The vintage normative prescriptions (e.g. "don't push yourself through a movie which you are not enjoying") are one of the first lessons that business and economics students are exposed to. And indeed the sunk cost fallacy still seems to plague many courses of action, be it continuing a failed relationship because one has already invested many years in it or a failure to withdraw from a lost war because of an extensive death toll. Thaler put forward a compelling loss aversion-based rationale for why people fall prey to the sunk cost fallacy. Given the convexity of the utility of losses, a decision-maker facing a risky investment has an incentive to recover an incurred loss because the increase in utility of a gain will be larger than what a further comparable loss would entail.⁴ Despite the abundance of casual and anecdotal evidence, the literature's verdict on the sunk cost fallacy is still mixed. The pioneering field experiment of Arkes and Blumer (1985) found that granting a random discount for a theater season ticket significantly decreases attendance. Drawing inspiration from this study, Ashraf et al. (2010) and Cohen and Dupas (2010) test for selection and sunk cost effects in the pricing of health products in the developing world; they find weak evidence of sunk cost effects. Other tests with field data have also produced mixed evidence: Staw and Hoang (1995) find considerable sunk cost effects in the drafting of National Basketball Association players (a result later corroborated by Camerer and Weber 1999), while Borland et al. (2011) find no such effects for the Australian Football League.

The experimental laboratory evidence is slightly more supportive of the sunk cost fallacy. On the one hand, using a search environment specifically designed to observe sunk cost effects, Friedman et al. (2007) find that experimental subjects are surprisingly consistent with optimal behavior, falling prey to the sunk cost fallacy occasionally at most. On the other hand, in an Industrial Organization setting both Offerman and Potters (2006) and Buchheit and Feltovich (2011) find that sunk costs influence pricing decisions. Cunha and Caldieraro (2009) show that sunk costs not only affect decisions over material investments, but also purely behavioral ones, i.e. those which stem from the cognitive effort invested in a task. They show that subjects are more likely to switch to a slightly better alternative if the sunk level of effort was low. However, an attempt at replicating these findings was not successful

⁴Eyster (2002) puts forward a taste-for-consistency-based explanation of the sunk cost fallacy. In face of sequential decisions under risk, decision makers trade off revenue-maximizing choices for consistency-maximizing ones, i.e. a decision maker gives up revenue today if this choice makes yesterday's decision seem more optimal.

(Otto 2010). Gino’s contribution (2008) is methodologically close to our work, but focuses on the role of the cost of advice. Subjects in her experiment answer trivia questions, for which they could use free or paid advice. Subjects who paid for advice incorporated it significantly more often in their decisions than those who obtained advice for free. From the competing explanations, the author shows that sunk cost effects drive the results. However, her results are limited in application to standard models of decision making under risk as subjects’ prior and posterior beliefs are not known to the experimenter. Our experimental design allows for these measurements and therefore uncovers the conditions under which costly information leads to better decision making. Furthermore, we are able to investigate the channel through which sunk costs operate and test for potential selection effects, both within- and between-subjects.

Our study investigates the impact of the cost of information in a setting where subjects have to make a decision under risk. Information is provided in a way that can help them reduce uncertainty in a Bayesian fashion, and therefore our work relates to a long literature in economics and psychology that deals with optimal decision making under risk, as well as the associated heuristics and biases (see DellaVigna 2009 for an overview). In particular, we are interested in knowing whether the cost of information can play a role in dampening some of the traditional biases or interact with some popular heuristics. To be sure, the verdict on whether “man is a Bayesian” is still out. When combining information on prior probabilities of possible states of the world with informative state-dependent signals, three main inter-related phenomena have been observed (see Camerer 1995 for a detailed overview). First, individuals often exhibit conservatism in their choices, failing to use the signal to the extent normatively prescribed by Bayes’ formula (e.g. Eger and Dickhaut 1982). Second, there is a systematic tendency for individuals to neglect the prior probabilities in their judgment, often referred to as the “base rate neglect” (see Koehler 1996 for an appraisal of the literature). Third, when the signal is representative of one of the states, the tendency to overweigh the signal’s information content is exacerbated. This heuristic is known as “representativeness”.⁵ For example, if a decision maker draws a sample which *exactly* matches the distribution of the signal in a given state, he will tend to overweigh the probability that this state will occur (in which case it is often referred to as “exact representativeness”).

Early evidence (e.g. Kahneman and Tversky, 1972 and 1973) showed that representativeness was a serious and systematic bias, leading these authors to claim that

⁵According to Kahneman and Tversky (1972): “this heuristic evaluates the probability of an uncertain event, or a sample, by the degree to which it is: (i) similar in essential properties to its parent population; and (ii) reflects the salient features of the process by which it is generated.”

“man is not a Bayesian at all” (1973). A number of experiments by David Grether (1980, 1992; El-Gamal and Grether 1995) produced more optimistic evidence: subjects do use representativeness (especially when it is “exact”), but behavior is not always far from Bayesian. Even though experimental subjects prove not to be perfect Bayesians, the “most likely rule that people use is Bayes’s rule.” (El-Gamal and Grether 1995), followed by representativeness. Experimental market tests of this heuristic (Duh and Sunder 1986 and Camerer 1987) have shown that behavior converges to Bayesian over time and that the observed deviation is mostly explained by representativeness. In sum, with respect to conservatism, base rate neglect and representativeness, the accumulated evidence seems to show that “base rates are underweighted in some settings but sample information is underweighted in others. Base rates are incorporated when they are salient or interpreted causally.” (Camerer 1995). Not only that, base rates’ “degree of use depends on task representation and structure” (Koehler 1996).

Building upon these conclusions, we ask a natural question: can the cost of information influence the extent to which conservatism, base rate neglect and representativeness prey on decision makers? In other words, can the cost of information mediate the difficulties posed by Bayesian updating (as emphasized by economists) and a tendency to disregard underlying prior probabilities (as documented by psychologists)? If that is the case, the cost of information can be used to dampen some of the shortcomings associated with decision making under risk. We seek to establish an existence result which would allow for further investigations into context-specific fine-tuning where information cost is the control variable.

In our design each subject has to make a number of discrete decisions with state-dependent payoff consequences. There are two states with known and constant priors. Subjects sometimes have the opportunity of reducing uncertainty by drawing a sample (a “ball”) from a state-dependent lottery (an “urn” with balls). Our treatments change the way in which this information is made available: in the Free treatment it is made available at no cost, while in the Costly treatment it is only accessible if purchased. A treatment where the cost is imposed on subjects (Forced) corrects for selection while leaving the role of cost intact. Moreover, and for all treatments, subjects subsequently go through a reduced version of the three treatments. Observing subjects’ revealed demand for information allows us to classify them by types and further analyze the role of selection.

Our results show that individual decisions are in line with the described biases, with deviations from the Bayesian normative model explained both by under- and over-updating. Paying for information in the Costly and Forced treatments leads to

an over-weighting of newly obtained information, which leads to moves in the posterior that are more extreme than in the Free treatment. This pattern is explained by a sunk cost effect, as the only difference between the two former treatments and the latter is the cost charged for information. These results cannot be explained by selection, as the data shows no significant differences between Forced and Costly. Moreover, subject types do not explain the overall pattern, which reinforces the sunk cost explanation. Regarding decision optimality, more extreme choices can lead to better or worse decisions. Most subjects benefit from having access to information (regardless of the cost) as it allows them to reduce uncertainty. However, some subjects do not benefit from information as the return derived from reduced uncertainty does not compensate for the cost paid for it.

Our main conclusion is that costly information is weighed more heavily than free information. However, this does not necessarily lead to more optimal decision making. From a policy perspective, charging for information is beneficial if the decisions made using free information correspond to a situation of Bayesian under-updating, since costly information leads subjects to put a higher weight on newly obtained information. These results suggest that individuals who have to vote on a binary issue will incorporate the information signal in their decision when this information is costly. A test of this conjecture in a voting context proper is left for future research. The remainder of the chapter is organized as follows. Section 2 presents the experimental design, Section 3 presents our results and Section 4 presents a small exercise on information pricing. A final section concludes.

4.2 Experimental Design

Each subject has to make decisions in two blocks: a “Decision block” and an “Identification block”, comprising 40 and 30 periods, respectively. Our analysis focuses on the data obtained from the Decision block, while data from the Identification block is used to account for the discussed selection issues. The decision is identical across the two blocks except for parameterization. Paper instructions are distributed in the beginning of each block, which subjects are asked to read silently.⁶ Each block starts after all subjects have finished reading the instructions. A set of practice questions to test understanding of the experiment is administered before the Decision block starts. In the experiment all values are expressed in tokens, which are converted at an exchange rate of 0.75 Euro per token. Subjects were paid for six randomly determined periods, three from each block.

⁶A transcript of the instructions can be found in Appendix 4.B.

4.2.1 Choice Framework

We presented subjects with an intuitive, yet non-trivial individual choice task in which information can be used in a Bayesian fashion. In each period the decision maker faces one of two states of the world (**Left** and **Right**), for which probabilities are known: $p \equiv \Pr(L)$. In the Decision block $p = 0.4$.⁷ The payoffs are determined by a state-dependent scoring function (see Figure 4.1).⁸ The parameterizations were chosen such that the loss domain was restricted while still providing substantial incentives to perform Bayesian updating.

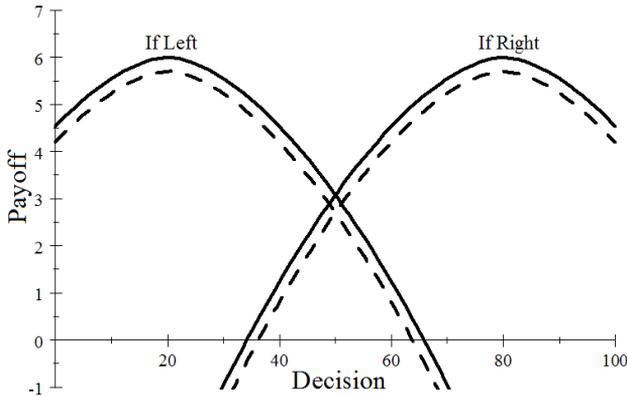


Figure 4.1: The scoring function

Notes: the solid (dashed) line corresponds to the Decision (Identification) block.

Subjects choose a number between 0 and 100 in steps of 0.5. If the state is L (R) the optimal decision is 20 (80). Note that decisions below 20 and above 80 are strictly dominated. The information on the two state-dependent payoffs is made available to subjects in three distinct ways: on the screen (subjects can interactively learn the state-dependent payoffs that result from any particular decision at all times using a slider bar), in graphical format and in table format (both in the paper instructions). Our state-dependent payoff function is an adjusted quadratic scoring rule, which was

⁷We chose not to implement symmetric priors for two reasons. First, the task could become trivial (Camerer 1987) or invite the usage of “obvious” (but possibly wrong) heuristics. Second, two scenarios with symmetric priors make the alignment of incentives and moves in the posterior across scenarios impossible to achieve (the urns would have to be different, rendering moves in the posterior not comparable). This aspect is important because we want to identify types in an environment (the Identification block) that is as similar as possible to the Decision block.

⁸See Appendix 4.A.1 for a detailed description of the choice environment and derivation of optimal decisions according to the normative model. Visit <http://www.reisayag.com/EmulationIntro.html> for an emulation of a decision round.

preferred to other proper scoring rules (e.g. Offerman et al. 2009). The fundamental reason behind our choice is the fact that many proper scoring rules do not provide a substantial incentive to update beliefs unless radical moves in the posterior are observed. In other words, we need a scoring rule function that is steep enough in the region where probability updating takes place. A common problem with proper scoring rules is that risk attitudes may play a role in the observed choices. However, this is not problematic in our setting as risk attitudes influence subjects’ decisions identically across treatments. Nonetheless, we statistically control for risk attitudes in our analysis.

The information signal we provide to subjects is a lottery, for which we use an “urn” filled with balls. There are five balls in the urn, some black and some white. The distribution of balls is itself state-dependent, but does not change across periods within a block and is visible to subjects before every draw. In our design, drawing a ball from the urn is informative of the state of the world, i.e. the probability of the realized state being L or R should be updated after drawing a ball from the urn. In the Decision block, the urn contains one (three) black balls if the state is L (R) (see Table 4.1).

	Left	Right
Prior	0.4	0.6
Urn	■□□□□	■ ■ ■ □ □

Table 4.1: Decision block: priors and lottery distributions (urns).

4.2.2 Treatments

We implement our treatments by varying the way in which the information is made available to subjects in the Decision block. The Identification block is identical across treatments. In Free the ball can be drawn from the urn at no cost. In Costly a ball can be drawn at a cost. In Forced the price of drawing a ball is imposed upon subjects (subjects are told that “a ball has been drawn for you”). In Costly a subject buying a ball observes it automatically while in Free and Forced subjects can choose whether to see the drawn ball or not. In the Decision block there is a 50% chance that subjects can draw a ball from the urn in each period of Free and Costly, or that a ball is drawn for them in the case of Forced. If information were available in all periods we would run the risk of subjects automatically discounting the costs of information to be incurred at the beginning of the experiment, which would dissolve

the psychological impact of an imposed cost. This also forces subjects to experience decisions without information, which provides us with individual decisions made without a ball draw - a likely anchor for decisions when information is available. In the Costly and Forced treatments the information is priced at $c = 0.3$ tokens, which is roughly 60% of the expected gain if expected utility maximization with Bayesian updating is performed by a risk- and loss-neutral decision maker. Note that the quality of the information is the same regardless of cost.

The Identification block uses the same framework with a slightly different parameterization. The idea is to create a decision environment that is equivalent in terms of incentives but looks sufficiently different for it not to be trivial nor invite the application of the decision rules employed in the Decision block. In particular, the ratio of the expected gain from using costly information to the expected gain from not using information is similar across blocks (see Appendix 4.A.1 for details).

The Identification block consists of three sequences of ten periods each. When available, information is provided in every period. In the first sequence (I_1), information is available for free. In the second sequence (I_2) information is available at a cost ($c = 0.25$ tokens, which is again roughly 60% of the expected gain). The first two sequences are akin to the Free and Costly treatments with a 100% probability of getting information. In the third sequence subjects have to choose between ten periods where they always have to pay for information (which is identical to Forced with a 100% probability of having information) and ten periods where information is never available. See Figure 4.2 for a time-flow diagram of the experiment.

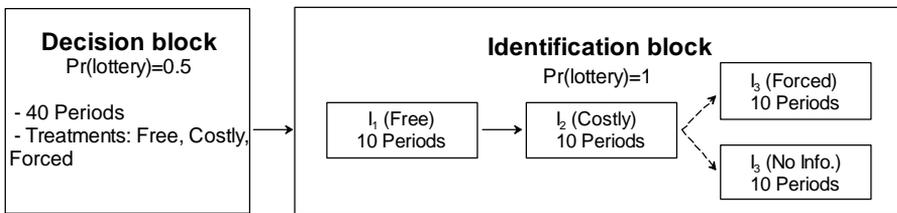


Figure 4.2: Outline of the experimental design.

The Identification block allows us to measure the value of information to subjects, i.e. how their expected benefits compare to the costs they have to incur. We can distinguish between two types of cost: monetary and cognitive. We present a classification that takes both into account. Accordingly, a subject buys information

if:

$$V_i(\text{Draw}) - C_{1,i} - C_2(\theta) \geq V_i(\text{No Draw})$$

where $\theta \in \{\text{Free}, \text{Costly}, \text{Forced}\}$, $V_i(\cdot)$ is the expected payoff of a subject (which depends on many cognitive factors like aptitude, mathematical training, confidence, etc.), $C_{1,i}$ is the cognitive cost of processing information, and $C_2(\cdot)$ is the monetary cost of acquiring information (equal to 0 in I_1 and equal to c in I_2 and I_3).⁹ In this sense, subjects incur C_1 in I_1 and $C_1 + C_2$ in I_2 in exchange for information. In the Identification block subjects make this choice in every period of I_1 and I_2 . In I_3 subjects also choose whether they want to incur $C_1 + C_2$ or not, but their choice is binding for ten periods. This stylized framework allows us to create an intuitive classification of types.

A subject who chooses not to see information in I_1 considers the cognitive cost of processing it higher than the benefits. A subject who chooses not to buy information in I_2 finds the sum of the cognitive and material costs of information higher than the benefits. Sequence I_3 measures the same relationship, but the choice is presented in a dichotomous way. The first two sequences not only provide useful measurements in themselves, they also allow all subjects to experience what it is like to use information for free and at a cost, especially considering that they faced different treatment conditions in the Decision block. Combining data from sequences I_1 and I_3 allows us to classify subjects in a way that improves our understanding of the major selection issues at hand. In particular, we classify subjects into four types:¹⁰

	$\Delta V_i(\cdot) \leq C_{1,i} + c$	$\Delta V_i(\cdot) > C_{1,i} + c$
$\Delta V_i(\cdot) \leq C_{1,i}$	Type 1	Type 0
$\Delta V_i(\cdot) > C_{1,i}$	Type 2	Type 3

Table 4.2: Subject types.

Type 3 individuals are those whose expected gain from using information exceeds not only the cognitive cost of using it but also the monetary cost charged for it. Type 2 individuals expect a net gain from using information but are not willing to buy it at price c . Type 1, on the other hand, do not expect a net gain from using information, even if there is no material cost involved. Type 0 are inconsistent types and they are considered for completeness (as we will see, they are a residual category in the data).

⁹See Appendix 4.A.2 for further details.

¹⁰Where:

$$\Delta V_i(\cdot) = V_i(\text{Draw}) - V_i(\text{No Draw})$$

We assume that $\Delta V_i(\cdot) \leq C_{1,i}$ if a subject observes information less than 9 out of 10 times in I_1 , and $\Delta V_i(\cdot) \leq C_{1,i} + c$ if a subject chooses to have no information in I_3 (in Section 4.3.3 we analyze the distribution of types that we obtain in light of these criteria).¹¹

In order to control for risk attitudes and demographic characteristics in our statistical treatment of the data, we end the experiment with the Charness-Gneezy-Potters task for risk attitude elicitation (Gneezy and Potters 1997, Charness and Gneezy 2010) and a questionnaire.¹²

4.3 Experimental Results

The experimental sessions were run at the CREED laboratory of the University of Amsterdam between February and May 2012; they were programmed and conducted with the experiment software z-Tree (Fischbacher 2007). A total of 166 subjects participated in 8 sessions, recruited online from a subject pool of students at the University of Amsterdam. No subject participated in more than one session. Fifty-five per cent of the participants were male and 57% were Business or Economics majors. The typical session took 1 hour and 20 minutes with average earnings of 24 Euro (which includes a show-up fee of 7 Euro). Two of the sessions (47 participants, 22 in Free and 25 in Costly) had a different Identification block.¹³ Unless mentioned otherwise, all data discussed in this section pertains to decision making in the Decision block. Subsection 4.3.1 analyzes the difference in decision making across treatments. Subsection 4.3.2 investigates possible channels through which cost affects decision making. Subsection 4.3.3 expands the analysis to include subject type data.

To start, table 4.3 provides a summary of descriptive statistics for the collected data. Differences in individual traits are not statistically significant across treatments (Pearson's chi-square test $p > 0.17$). Average period payoff refers to the gross payoff, i.e. not including differences in costs of information across treatments. There is no significant differences in the average payoff between treatments. Per-

¹¹The first criterion is employed as we are looking for subjects who would buy information whenever it is free (which is 10 times in I_1) while allowing for one mistake.

¹²The risk attitude elicitation task consists in asking subjects how they wish to allocate an endowment of three tokens between a safe account and an account that multiplies the invested amount by a factor of 2.5 with 50% probability and destroys the money with 50% probability. The questionnaire asked whether subjects had had Math in high school, how many Math courses they had completed at university, their gender, age, and major.

¹³The Decision Block was identical across all sessions. The Identification Block was changed after the first two sessions in order to enhance the validity of the type classification. For this reason no data from these two sessions is used in analyses containing type variables.

centage of information seen refers to the fraction of times subjects chose to observe information when it was available. Naturally, when information was costly and optional, fewer subjects chose to observe it. The Costly treatment thus shows significantly fewer information views than the Free and Forced treatments. We observe that subjects chose not to see information (draw a ball) sometimes, even when it was free or already paid for. This is possible as in all treatments we let subjects have the option of not drawing a ball, reasoning in terms of the stylized model discussed in Section 4.2.2. That is, some subjects, denoted as Type 0 and Type 1, found the cognitive costs of using information higher than the benefits. Additionally, many subjects experimented with drawing and not drawing a ball and thus did not observe information in some of the periods.¹⁴ We further observe no difference in average information use across treatments in I_1 and I_3 . In I_2 , subjects in the Free treatment are slightly less likely to pay for information than in the Costly and Forced treatments. This however is not significant for both the difference between Free and Costly and Free and Forced (Mann-Whitney-Wilcoxon rank-sum test, $p \geq 0.12$).

	Free	Costly	Forced
N	65	65	36
Risk	2.06	1.94	2.10
Math courses	2.38	2.95	2.56
% Female	37%	48%	56%
Average period payoff	3.10	3.08	3.10
% Information seen	79%	55%	74%

Table 4.3: Summary statistics

4.3.1 Treatment effects

We begin with the analysis of aggregate treatment outcomes. For ease of exposition and later analysis we define benchmark decision making as the optimal decisions made by a rational, risk- and loss-neutral individual (hereafter, the Bayesian benchmark). The Bayesian benchmark decisions are 74, 44 and 58.5 while the average decisions are 69.4, 39 and 55.6 for a Black draw, a White draw and No draw, respectively. Figure 4.3 presents five-period average decisions over the duration of the Decision block by treatment and information condition (Black draw, No draw and White draw). Decision averages visibly differ across treatments and in all periods after a ball draw. No such difference is discernible after No draw. Table 4.4

¹⁴Ignoring type 0 and 1 subjects, the percentage of information seen changes to 92% in the Free treatment and 87% in the Forced treatment. Of all subjects, only 6% in Free and 11% in Forced chose never to draw a ball (this figure is 21% in Costly).

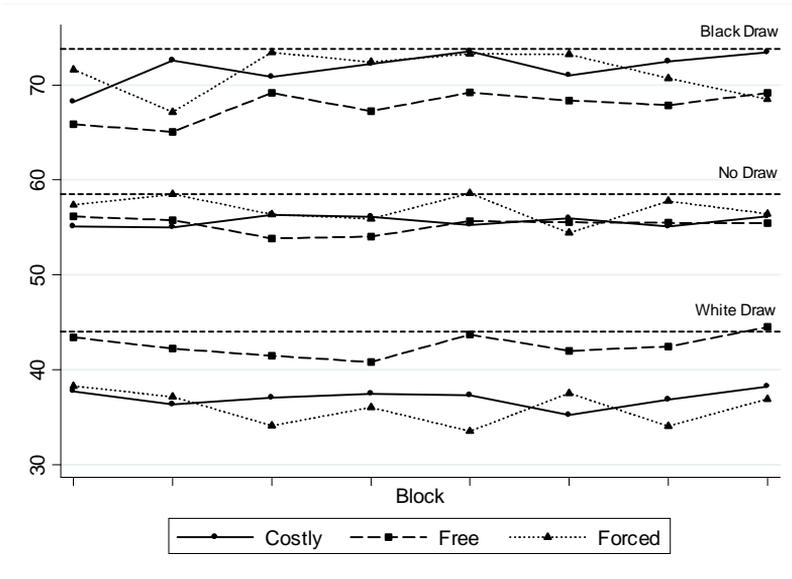


Figure 4.3: Average decisions by blocks of 5 periods.

Notes: Decision averages for each block are calculated using all decisions made in a five-period interval (8 blocks of 5 periods) by all subjects within an information condition. For example, the top solid line oscillating around 70 represents the five-period decision average of subjects who observed a Black ball in the Costly treatment. The dashed straight lines denote the Bayesian benchmark.

presents benchmark decisions and treatment averages by information condition and treatment. After both a White draw and a Black draw there are significant differences between the Costly and Free treatments (two-sided Mann-Whitney-Wilcoxon rank-sum test, MW hereafter: $p = 0.01$ and $p = 0.02$, respectively), and between the Forced and Free treatments (MW: $p = 0.00$ and $p = 0.02$, respectively). No significant differences are found between the Costly and Forced treatments. Without a draw from the urn there are no significant pair-wise differences between any of the treatments. The shift in average decision making between both the Costly and Forced treatments and the Free treatment is thus significant only after a ball draw. It is an upward shift after a Black draw and downward one following a White draw.

Figure 4.4 presents the cumulative distribution functions of individual decision by information condition. We define individual decision $d_{i\phi t}$ as the decision of subject i in information condition $\phi \in \{Black, White, No Info\}$ and in period t . Average decision $\bar{d}_{i\phi}$ is defined as the sum of subject i 's decisions when in information condition ϕ divided by $n_{i\phi}$, the number of periods in which subject i is in information

	A Black draw	A White draw	No draw
Free	66.05 (1.30)	42.57 (1.26)	55.08 (0.71)
Costly	70.26 (1.38)	38.62 (1.51)	55.96 (0.81)
Forced	70.6 (1.89)	36.55 (1.97)	56.47 (1.38)
Benchmark	74	44	58.5

Table 4.4: Decision averages by treatment and information condition.

Notes: Standard errors in brackets.

condition ϕ :¹⁵

$$\bar{d}_{i\phi} = \frac{\sum_{t=1}^{40} d_{i\phi t}}{n_{i\phi}} \quad (4.1)$$

As with Figure 4.3, a shift in decision making between the Free treatment and the Costly and Forced treatments is clearly observable after a ball has been drawn. Decision distributions in Costly and Forced first-order stochastically dominate the decision distribution in Free after a Black draw, and are first-order stochastically dominated after a White draw. The differences in distributions are all significant at the 10% level (two-sample Kolmogorov-Smirnov test).¹⁶

From Figure 4.3 and Figure 4.4 it is clear that the availability of costly information “pushes” subjects’ decisions more to the extremes. Subjects who incur a cost tend to make more extreme decisions vis-a-vis those who did not. Thus, as subjects who pay for information behave differently than those who do not, we conclude that a sunk cost effect exists in our experiment. Thaler’s (1980) description of the mechanism behind the sunk cost phenomenon provides an explanation for the observed data. A subject paying for information, found in a loss state, has a higher marginal benefit of payoff than a subject receiving information at no cost. If effort is costly, a paying subject should exert higher effort than a non-paying subject. Higher effort may change the way in which information is used. Another manner in which sunk costs can lead to this change is by increasing the relative salience of a ball draw to the subject who just paid for it. Following Koehler (1996), the change in the saliency of information may affect the use of information.

Result 1 *Paying for information alters individual decision making in a systematic*

¹⁵A small minority of individuals (6%) have a strictly dominated average decision (i.e. outside the range [20, 80]). Only 1% are outside the range [19, 81].

¹⁶See Table 4.11 in Appendix 4.C.

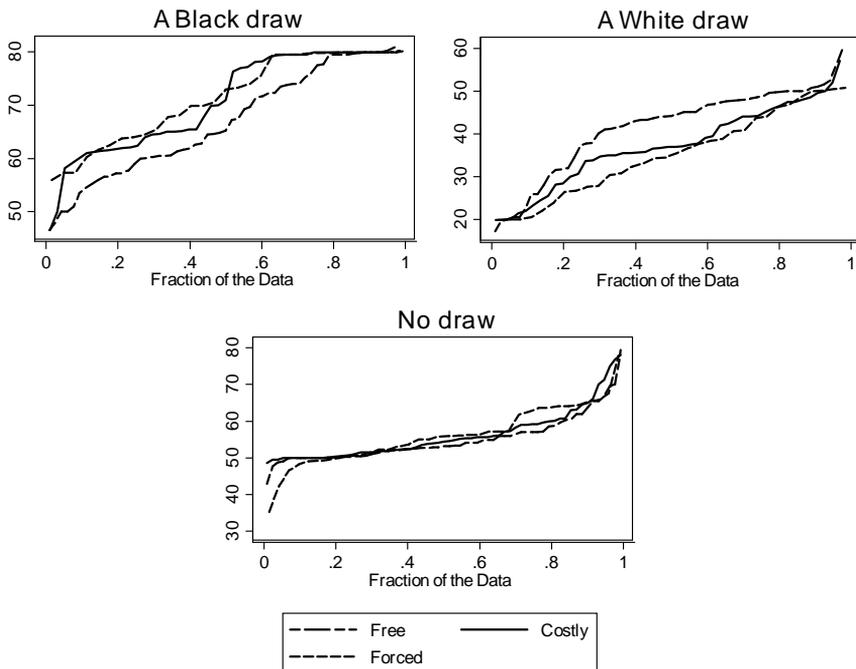


Figure 4.4: Distributions of individual decision making.

Notes: Outliers were discarded (one observation after a black draw and one after a white draw).

manner. Individuals who incur a cost weigh the newly acquired information more heavily than individuals who do not. This behavior can be explained by a sunk cost effect.

A natural concern is that selection is driving the difference between the Costly and Free treatments. That is, the two samples are not identical as subjects who choose to pay for information may differ in many dimensions from those who are only willing to observe free information. The average decision in the Free treatment can then be perceived as a weighted average of two sub-groups: those who are willing to pay for information and those who are not. The average decision in the Costly treatment after a ball draw is the outcome of the paying subjects' decisions alone. However, selection can not play a part in the differences between the Free and the Forced treatments as evidenced by the similar information acquisition rates in Table 4.3. Any difference in decision making should thus be attributed to the difference

in cost incurred by the subjects between these two treatments. Selection effects are further discussed in Subsection 4.3.3.

We now turn to decision optimality across treatments, where we define optimality as the absolute distance from the Bayesian benchmark. Table 4.4 presents the treatment means by information condition. We observe that the average decision in the Free treatment is the most optimal after a White draw, but the least optimal after a Black draw.

Result 2 *Costly information does not necessarily lead to a more optimal use of information.*

We conclude that, depending on the state of the world, costly information may improve or worsen a subject’s performance. This is a result of the “push” towards the extremes that costly information induces. After a Black draw, as subjects tend to under-update new information in the Free treatment, costly information leads to a more optimal decision. In case of a White draw, where subjects over-update, costly information leads to a less optimal decision.¹⁷

4.3.2 Interpreting the effect of cost

Our experimental design is devised to detect the existence of cost effects on the use of information, which we have shown to exist. Nonetheless, an exploration of the underlying channels through which the cost of information affects decision making is of interest. This subsection investigates whether our results can be explained by the biases discussed in the Introduction. Though our experiment does not single out one interpretation unequivocally, it allows us to point at a likely candidate.

Representativeness refers to an individual’s tendency to overweigh new information, biasing her choice away from the prior. When observing a White ball draw, this bias leads an individual to overweigh the likelihood that this all draw signals the state of the world is Left. A similar overweighing of the Right state of the world occurs if the ball is Black. Representativeness thus results in decisions which are closer to the extremes of 20 and 80 than the benchmark decision.

Base rate neglect pertains to the tendency of an individual to underweigh the prior when receiving new information, i.e. the individual has a prior belief that is closer to 0.5 than it actually is. We follow Camerer (1987) in our analysis, and

¹⁷Considering Result 2 in light of Thaler’s (1980) explanation of sunk cost, it might be surprising at first look that increased effort can lead to a less optimal result. Still, similar observations in the literature exist which demonstrate that more effort can lead to lesser performance (e.g. Camerer and Hogarth 1999, Ariely et al. 2009, Leuven et al. 2011).

equate full base rate neglect with Bayesian updating with erroneous and equal priors $\Pr(Left) = \Pr(Black) = 0.5$. Note that this is *not* the same as representativeness. Since the prior in our design indicates that Left is less likely than Right, this bias brings subjects to over-update the probability of Left after any ball draw. In case an individual exhibits base rate neglect she thus deviates towards 20 after any ball draw. See Appendix 4.D for a graphical illustration of the effect of these biases.

An increase in the strength of base rate neglect due to costly information could be a natural channel for the effect of cost on decision making. A subject who focuses more on the new information she just paid for discounts the underlying prior. This explanation fits the observed behavior after a White draw but not after a Black draw. If base rate neglect increases with a costly ball draw, average decision after a Black draw should shift towards 50 relatively to a free ball draw, while we observe the opposite.

Representativeness bias describes the data well. If paying for information engenders or intensifies representativeness, a subject observing a White ball perceives it to be more representative of Left (out of 5 balls, 4 are white if Left and 2 are white if Right), and a Black ball to be more representative of Right. This would directly lead to more extreme decisions.¹⁸

We can investigate the representativeness explanation by using a parametric estimation of a simple model of information updating. To do so we build upon a model used by Grether (1980, 1992) to assess representativeness and conservatism, and by Gonzales and Wu (1999) and Holt and Smith (2009) to evaluate probability weighting in Bayesian updating tasks. The individual belief that the state of the world is Left after a Black ball draw can be written as:

$$\Pr(Left|Black) = \frac{(\Pr(Black|Left))^\eta \cdot (\Pr(Left))^\gamma}{(\Pr(Black|Left))^\eta \cdot (\Pr(Left))^\gamma + (\Pr(Black|Right))^\eta \cdot (\Pr(Right))^\gamma} \quad (4.2)$$

If both $\gamma = 1$ and $\eta = 1$ the individual is Bayesian. The parameter $\gamma \in [0, 1]$ represents the strength of base rate neglect. If $\gamma = 1$ there is no base rate neglect, while as γ goes to zero the strength of base rate neglect increases, i.e. the prior probabilities are perceived to become more equal. Similarly to Camerer (1987), if $\gamma = 0$ the individual perceives the prior probabilities of the two states of the world as equal. The representativeness/conservatism bias is brought about by the

¹⁸Risk loving behavior resulting from convex utility function in a loss-frame can also explain the observed shift in decision making. We find this explanation unlikely though as the cost of information is much too small to reasonably explain the observed change in behavior.

parameter $\eta \in [0, \infty]$. If $\eta > 1$ the individual overweighs new information such that representativeness bias exists. If $\eta < 1$ the opposite occurs and conservatism bias is observed. We can write the individual updated odds ratio after a Black ball draw as:

$$\frac{\Pr(Left|Black)}{\Pr(Right|Black)} = \left(\frac{\Pr(Black|Left)}{\Pr(Black|Right)} \right)^\eta \cdot \left(\frac{\Pr(Left)}{\Pr(Right)} \right)^\gamma \quad (4.3)$$

A similar expression can be written in case of a White draw. To estimate this model while allowing for treatment effect we use a log-linear least squares specification:¹⁹

$$\ln \frac{r}{1-r} = \alpha + (\beta_1 + \beta_2 \cdot D_{Costly} + \beta_3 \cdot D_{Forced}) \cdot \ln \left(\frac{\Pr(Draw|L)}{\Pr(Draw|R)} \right) \quad (4.4)$$

where r is the inferred probability of Left in each information condition (Black ball, White ball or No Draw) and derived from individual average decision making $\bar{d}_{i\phi}$, D_{Costly} and D_{Forced} are treatment dummies, $\eta = \beta_1 + \beta_2 \cdot D_{Costly} + \beta_3 \cdot D_{Forced}$ and $\alpha = \tau + \gamma \cdot \ln(\Pr(Left) / \Pr(Right))$. The parameter η is decomposed to allow the capturing of treatment effects via the likelihood ratios. The parameter τ is a constant suggested by Gonzales and Wu (1999) in their two-parameter estimation. The prior odds are included in the constant term since we cannot estimate γ due to the lack of variation in the prior odds of the state of the world in the Decision block.

Column (1) in Table 4.5 presents the estimation results of expression (4.4). We cannot reject that β_1 is equal to one (Wald, $p = 0.86$), suggesting that there is no representativeness or conservatism bias in the Free treatment, i.e. $\eta = 1$. The interaction coefficients β_2 and β_3 are both positive and significantly different than zero. This suggests that the effect of costly information is to increase representativeness bias. Column (2) presents an estimation with independent treatment dummy variables.²⁰ This allows for the possibility of the Costly and Forced treatment to have an effect on individual decisions through channels other than representativeness (i.e. through the constant τ or the scope of base rate neglect γ). We do not find a significant effect of treatment dummies on inferred probability odds. The coefficient β_1 remains statistically indistinguishable from 1 (Wald, $p = 0.80$). Our

¹⁹For decision with no ball draw equal likelihood are assumed $\Pr(Black|Left) = \Pr(Black|Right) = 0.5$. Following Grether (1992) and Holt and Smith (1999), we change inferred probabilities of 0 and 1 to 0.01 and 0.99, respectively. Additionally, average subject decision outside $x \in [20, 80]$ are also converted to 0.01 if $x < 20$ and 0.99 if $x > 80$. Results are robust to the exclusion of these observations (See Appendix 4.C).

²⁰See Appendix 4.C for the estimated model in column (2).

	Model	
	(1)	(2)
β_1 (Likelihood ratio)	1.02*** (0.11)	1.03*** (0.11)
β_2 (Likelihood ratio x Costly)	0.51*** (0.16)	0.49*** (0.17)
β_3 (Likelihood ratio x Forced)	0.71*** (0.19)	0.71*** (0.19)
Costly	—	-0.09 (0.12)
Forced	—	0.01 (0.14)
Constant	-0.24***	-0.21***
N	461	461

Table 4.5: Representativeness bias and treatment effects.

Notes: Each observation corresponds to the average individual decision in a specific information condition. The number of observation is lower than $166 \times 3 = 498$ since for some subjects there are no observations for all information conditions. Clustered standard errors used. ***/**/* indicates significance level at the 1%/5%/10%.

estimation results support the preceding qualitative discussion on the possible channels through which sunk costs affect individual decision making. An intensification of representativeness bias following costly information is the most likely explanation of our results.

4.3.3 Selection effects

This sub-section extends the analysis to include the type variables described in Sub-section 4.2.2.²¹ Table 4.6 presents summary statistics for subjects for which Identification block data is available. Type 1 subjects are defined as those who choose neither to draw a ball when it is free nor when it is costly. Type 2 subjects are defined as those who choose to draw a ball when it is free but rather not draw a ball when it is costly. Type 3 subjects are defined as those who always choose to draw a ball. Type 0 subjects are inconsistent: they choose to draw a ball when it is costly but not when it is free. There are no significant differences in the proportion of subject types across treatments. Type also does not correlate strongly with observables like risk preferences (Pearson's $r=0.18$), math courses taken at the university level ($r = 0.12$), math taken in high school ($r = 0.05$) or age ($r = 0.02$).

²¹In the analysis using subject type variables we only include data for the 119 subjects who participated in sessions with the Identification block. Moreover, some subjects never drew a ball and thus the number of subjects used is $N = 106$ for a Black draw and $N = 103$ for a White draw.

There is no statistical evidence for a difference of observables across types (Pearson's chi-square test $p > 0.13$). This suggests that our type classification measures the ability and willingness to use (costly) information, and not simply mathematical proficiency or willingness to take risk. The only variable that correlates significantly with type is gender: males tend to be of higher type (Pearson's chi square test $p = 0.02$), albeit moderately ($r = 0.27$). See Table 4.10 in Appendix 4.A.2 for percentage of information seen in the decision block by type.

	Free	Costly	Forced
N	43	40	36
% Type 0	2%	0%	3%
% Type 1	14%	15%	17%
% Type 2	47%	43%	42%
% Type 3	37%	43%	39%

Table 4.6: Type distribution by treatment

Figure 4.5 presents aggregate average decisions by subject type, treatment and ball draw.²² Note that this stratification lowers our sample sizes and results in a lower power of statistical tests. As a result, in this sub-section we use the 10% level as the threshold for qualitative statements about statistical significance. Average decisions of Type 2 and Type 3 subjects are significantly different in the Free and Costly treatments after a White draw (MW: $p = 0.05$ and $p = 0.04$, respectively). A borderline insignificant difference is found in the Costly treatment after a Black draw (MW: $p = 0.13$). All other differences are highly insignificant.

In order to further understand how decision making differs across types we conduct a parametric analysis. The following model is estimated:

$$\begin{aligned} \bar{d}_{i\phi} = & \alpha + \beta_1 \cdot D_{i\phi}^{Costly} + \beta_2 \cdot D_{i\phi}^{Costly} \cdot D_{i\phi}^{Type\ 2} + \beta_3 \cdot D_{i\phi}^{Forced} \\ & + \beta_4 \cdot D_{i\phi}^{Type\ 1} + \beta_5 \cdot D_{i\phi}^{Type\ 2} + X_{i\phi}' \cdot \gamma + \epsilon_{i\phi} \end{aligned} \quad (4.5)$$

where the dependent variable $\bar{d}_{i\phi}$ and the subscripts i and ϕ are as defined in equation (4.1). D^{Costly} and D^{Forced} are treatment dummies and $D^{Type\ 1}$ and $D^{Type\ 2}$ are subject type dummies. The vector of control variables X includes gender, mathematical knowledge and risk aversion. Table 4.7 presents OLS regression coefficients over decisions made after a Black draw and after a White draw for different combinations of independent variables. The baseline is Type 3 subjects in the Free

²²Type 0 and Type 1 are not presented since our sample size for these subject types is very small. See Table 4.13 in the Appendix 4.C for aggregate means of all subject types with standard errors.

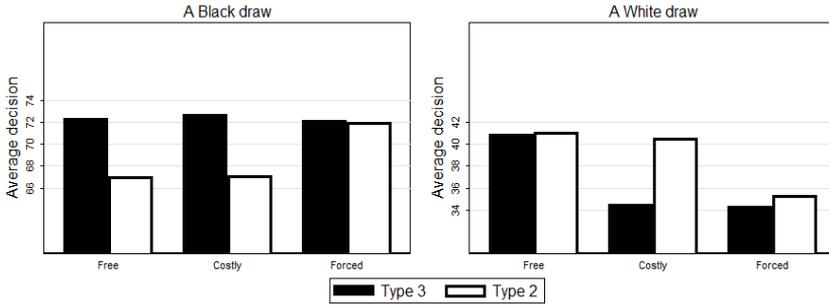


Figure 4.5: Average decision by subject type.

treatment. Columns (4) and (8) present results with the addition of an interaction term between Costly and Type 2 (CT2). This interaction term is added in order to account for the difference shown in Figure 4.5 between the average decisions made by Type 2 and Type 3. The results are consistent with the analysis done in Subsection 4.3.1 for the effect the Costly and Forced treatments. Costly and Forced treatments shift decisions upwards after a Black draw and downwards after a White one, even after controlling for subject types.²³

The difference in the effect of the Costly and Forced treatment is not significant after both a Black and a White draw (Wald test, $p = 0.72$ and $p = 0.76$, respectively). If the difference in average decision making between Free and Costly was driven by those subjects who choose to purchase costly information then any change between these treatments must be larger than the difference between the Free and Forced treatments. The reason is that the average decision making in the Forced treatment is the weighted average of those subjects who would voluntarily buy costly information and those who would not. It is thus the cost of information itself that directly affects subject behavior.

Result 3 *The observed differences in decision making across treatments are not driven by selection.*

Using Figure 4.5 we can attempt to explain why a stronger deviation is observed in Forced than in Costly. To do so, we focus our attention on the behavior of Type 2 subjects.²⁴ Type 2 subjects' average decision after both a Black and a White draw

²³See Table 4.14 and Table 4.15 in the Appendix 4.C for regression results including controls and for decisions made after No draw.

²⁴Remember that though we define Type 2 subjects as those who are willing to draw a ball only

	A Black draw				A White draw			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Costly	4.22** (1.89)	4.84*** (1.81)	2.13 (2.19)	5.42** (2.20)	-4.21** (1.97)	-4.95*** (1.90)	-4.89** (2.33)	-8.18*** (2.69)
CT2	-	-	-	-7.19* (3.84)	-	-	-	6.57 (4.45)
Forced	5.09*** (2.33)	5.73** (2.34)	4.33* (2.40)	4.40* (2.38)	-6.44*** (2.37)	-7.12*** (2.39)	-7.27*** (2.41)	-7.37*** (2.43)
Type 1	-	-	-13.68*** (4.04)	-13.42*** (3.93)	-	-	14.16*** (3.45)	13.91*** (3.43)
Type 2	-	-	-2.97* (1.91)	-0.74 (2.18)	-	-	1.66 (2.12)	-0.33 (2.48)
Constant	65.74	66.35	71.42	70.62	42.84	40.95	39.19	39.88
Controls	No	Yes	Yes	Yes	No	Yes	Yes	Yes
N	146	146	106	106	145	145	103	103

Table 4.7: Regression results by ball draw.

Notes: The dependent variable is the subject average decision $\bar{d}_{i\phi}$. Type 0 subject decisions and outliers are not included in the analysis. Robust standard errors used. ***/**/* indicates significance level at the 1%/5%/10%.

are similar in the Free and Costly treatments and only shift in the Forced treatment. This is consistent with our type dichotomy if we consider Type 2 subjects' behavior in the Costly treatment as an exploratory one, investigating the benefits and costs of opting for a ball draw. Such "testing of the waters" may be less affected by the cost of drawing a ball. Average decisions for Type 2 and Type 3 subjects should then be equal in Free and Forced while being different in Costly. This is indeed the case after a White draw but not entirely after a Black draw.²⁵ The weak effect of Costly in Table 4.7 is thus the average effect of Type 2 and Type 3 subjects. This also explains the weakly significant effect of the Type 2 dummy seen in Table 4.7. When adding the interaction term CT2 the shift in decision making after both Costly and Forced becomes highly similar. Additionally, after a Black draw, the Type 2 dummy coefficient is no longer significant. Table 4.7 also exhibits significant differences between Type 3 and Type 1 subjects' performance. Using a Wald test we also find significant differences between Type 2 and Type 1 subjects (p-values: 0.09 after a White draw and 0.01 after a Black draw).

Result 4 *Subjects who use freely available information (Type 2 and Type 3 subjects) perform similarly in the Free and Forced treatments and differently from subjects who do not (Type 1). Type 1 subjects substantially under-weigh new information.*

The experimental results discussed in Sub-section 4.3.1 and Sub-section 4.3.3 are closely related. We show that the variance in the cost of information is the driver of our result using both the Forced treatment and by using the subject type data. We are thus able to evaluate the effect that subject heterogeneity, with respect to information purchasing decisions, has on our results. The change in behavior due to the cost of information is significant and systematic. Sunk cost effects are thus shown to have an effect on decision making in our experimental setting.

4.4 Pricing Information

In general, reliable and useful information is a lever towards better results. Our choice framework presents subjects with the opportunity of increasing expected gains through the incorporation of new information. In Sub-section 4.3.1 we observed that subjects put relatively more weight on information they had to pay for, but this does not always lead to better decisions. However, even if information always

at no cost, we still observe those subjects drawing a ball in the Costly treatment occasionally.

²⁵After a Black draw Type 3 subjects average decision is significantly higher than Type 2 subjects in the Free treatment (MW: $p = 0.07$).

pushed subjects closer to the optimum, this would tell us little about the efficiency of using information. In other words, the gains realized from using information (because of higher-payoff choices due to reduced uncertainty) might not compensate the cost paid for it.²⁶ In this section we investigate this implicit trade-off using a small exercise that demonstrates how information should be priced (or subsidized) in order for it to be profitable for subjects. Hence, in this section we move our focus away from the effect of sunk costs to the efficiency gains obtained from using information.

To answer this question we compute the cost levels that would make subjects indifferent between paying for information and having no information. We use data from the Decision block in this analysis, and only from the Free and Forced treatments (the selection present in Costly complicates data interpretation as many subjects chose never to observe information - cf. Table 4.3). We want to calculate the individual cost level c_i that makes the following equality hold:

$$\sum_{\phi \in \{Black, White\}} \Pr(\phi) V(\bar{d}_{i,\phi}) - c_i = V(\bar{d}_{i,No Draw})$$

where $V(\cdot)$ is the payoff that would result from implementing the average decision $\bar{d}_{i\phi}$ (as defined in equation 4.1) of each subject in the respective information condition.

Figure 4.6 presents a plot of the implied individual cost levels in ascending order. To be more precise, c_i is the cost level which would make subject i indifferent between facing one decision with paid information and one decision without information. This value is conditional on i 's average decisions in each information condition. We observe that the great majority of subjects should be willing to pay for information ($c_i \geq 0$): only 10% would have to be subsidized ($c_i < 0$). Moreover, 60% of our subjects have an implied cost level above 0.3, which means that information was priced in a beneficial way for the majority of them. A noteworthy aspect is the fact that Forced did not lead to a better overall use of information, as the distribution of c_i in this treatment follows the one in Free quite closely. The reason is that, as mentioned, subjects in Forced got very close to the optimum in case of a Black ball but overshot in case of a White ball (see Figure 4.3).

The main message from this exercise is that information pricing is far from trivial from a policy perspective due to the underlying trade-off. On the one hand, it can provide (the right) incentives if it leads to a better incorporation of information in decision making. On the other hand, individuals might end up worse off if the price

²⁶Recall that we priced information at roughly 60% of the expected gain of the Bayesian benchmark.

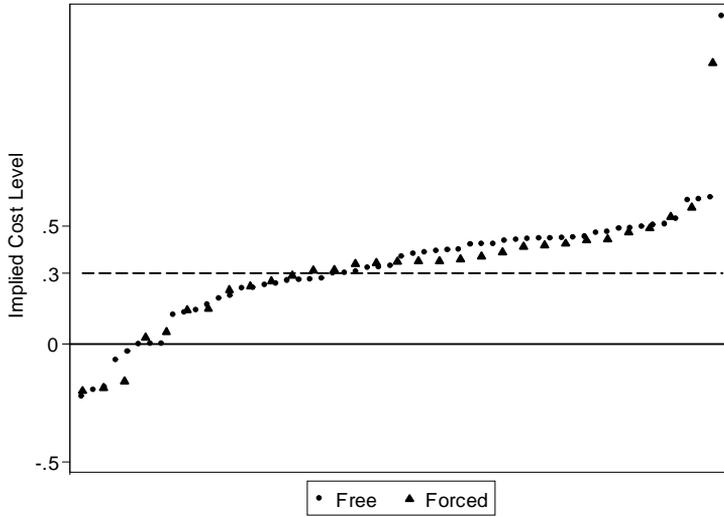


Figure 4.6: Implied cost levels

Notes: $N_{Free}=57$ and $N_{Forced}=31$ as some subjects did not see a Black or a White ball. The observations of each treatment are equally spaced over the axis. The dashed line (0.3) is the price of information in the Decision block.

paid for information cancels the benefits derived from reducing uncertainty.

4.5 Conclusion

The work presented in this chapter sets out to explore how individuals' use of information in a situation of decision making under risk is affected by its cost. Within the scope of political economy, this question has relevance for two of the standard results of the rational choice approach - rational ignorance and the Condorcet jury theorem - which assume that an individual's demand for information is decreasing in price but that its incorporation in decision making does not depend on it. To be sure, standard economic theory posits that the cost of a given piece of information should not influence the way it is incorporated in updating beliefs, all else equal.

This chapter challenges the established view. We thus touch upon two known issues concerning individual decision making: the sunk cost fallacy and Bayesian updating. Individuals who are prone to sunk cost effects may behave differently after receiving information for free than after paying for it. Consequentially, sunk costs may have an effect on individual deviations from Bayesian updating. Biases in the updating of beliefs may thus be exacerbated or alleviated by costs.

To examine these issues we use a laboratory experiment which enables us to control for the selection problem which is often found in field data. That is, we control for the possibility that variations in the data are engendered by the difference in behavior between subjects who have high valuation of information and those with low valuation of information, rather than by the cost of information itself. We do so by using two independent procedures. First, we implement a treatment in which we force subjects to pay for information regardless of their willingness to use it. Second, we identify subjects' demand for costly information in all treatments. We can then compare those subjects who choose to buy information when it is costly with those who do not.

We find a significant sunk cost effect on individual decision making. Subjects who pay for information put higher weight on it relative to subjects who receive identical information at no cost. This effect leads to a shift of updated beliefs towards the extremes. Decision making with costly information can be closer to or further from the correct Bayesian updating compared to decision making with free information. If subjects under-update their beliefs using free information, then costly information "pushes" their decision closer to the optimum. In case the opposite occurs, i.e. subjects over-update with free information, then costly information "pushes" them further away from the optimal outcome.

In a voting situation, our results suggest that individuals who paid for more expensive signals are more likely to vote in accordance with the received information. An interesting extension is to know whether more expensive signals reduce

abstention when this is allowed. More expensive information might result in a more extreme posterior probability of each state of the world, and produce less indifference and more certain votes. These extrapolations are subject to obvious caveats, which can only be addressed by future work.

Appendix

4.A Details on the Choice Framework and Type Classification

4.A.1 Choice Framework

In this Appendix we provide details on the choice framework and the normative prescriptions of the implemented parameterizations (the Bayesian benchmark). The two-part payoff function that we employed is:

$$F(x, \sigma) = \alpha - \beta \|x - s(\sigma)\|^\gamma$$

where $\sigma \in \Sigma = \{L, R\}$ is the state of the world; $s : \Sigma \rightarrow \{l, r\}$ is a state-dependent function such that $s(L) = l$ and $s(R) = r$; $\alpha, \beta, \gamma \gg 0$ are parameters; x is the decision maker’s decision variable. For $p \equiv \Pr(L)$, expected value maximization yields:

$$x^* = \frac{1}{\left(1 + \left(\frac{p}{1-p}\right)^{\frac{1}{\gamma-1}}\right)} \left(r + l \left(\frac{p}{1-p}\right)^{\frac{1}{\gamma-1}}\right) \quad (4.A.1)$$

In our experiment $x \in [0, 100]$, $l = 20$ and $r = 80$. As explained in Sub-section 4.2.1 there are three possible information conditions, $\phi \in \{Black, White, No Info\}$, which induce different distributions of the lottery. We define a “Draw” as a “Black” or a “White” ball draw. Table 4.8 presents the two parameterizations that were implemented: *A* was used in the Decision block and *B* was used in the Identification block.

	α	β	γ	$\Pr(L)$	$\Pr(Black L)$	$\Pr(Black R)$	Cost	Exch. Rate
<i>A</i>	6	0.009	1.7	0.4	0.2	0.6	0.3	0.75
<i>B</i>	5.7	0.00925	1.7	0.7	0.8	0.4	0.25	0.75

Table 4.8: Parameterizations A and B.

The posterior probabilities $\Pr(L|\phi)$, optimal decisions $x^*|\phi$, and the expected values in different information conditions, $E[F(x^*, \sigma)|\phi]$ and $E[F(x^*, \sigma)|Draw]$, are provided for both parameterizations in Table 4.9.

Note that the scenarios induce similar expected values across parameterizations, which makes the incentive to optimize and acquire information similar in *A* and *B*. Notwithstanding, the scenarios look sufficiently different from each other such that the rules employed in one are not easily translated to the other. Note that the prior

	Pr(L ϕ)		$x^* \phi$			$E[F(x^*, \sigma) \phi]$			Pr(Black)	$E[F(x^*, \sigma) Draw]$
	B	W	ND	B	W	ND	B	W		
A	0.18	0.57	58.5	74	44	3.22	4.40	3.15	0.44	3.70
B	0.82	0.44	34	26	55	3.26	4.10	2.76	0.68	3.67

Table 4.9: Values for parameterizations A and B.

Notes: ND, B and W stand for No Draw, Black draw and White draw respectively.

probability changes while the urn composition remains unchanged (there is a mere relabeling of colors and states).

4.A.2 Type Classification

We define the value derived from individual decision $x_i(\phi)$ given the state of the world σ as $F(x_i(\phi), \sigma)$. Individual utility is thus given by $U(F(x_i(\phi), \sigma))$ where $u(\cdot)$ is a general utility function. Given that individual utility only depends on the primitives ϕ and σ , we simplify our notation and write individual utility as $U_i(\phi, \sigma)$.

We define our types according to subjects' willingness to buy information. A subject buys information if:

$$V_i(Draw) - C_{1,i} - C_2(\theta) \geq V_i(No\ Info) \quad (4.A.2)$$

where

$$\begin{aligned} V_i(Draw) &= E(U_i(Draw, \sigma)) \\ &= \Pr(L) \cdot [\Pr(White|L) U_i(White, L) + \Pr(Black|L) U_i(Black, L)] \\ &\quad + \Pr(R) \cdot [\Pr(White|R) U_i(White, R) + \Pr(Black|R) U_i(Black, R)] \\ V_i(No\ Info) &= E(U_i(No\ Info, \sigma)) \\ &= \Pr(L) \cdot U_i(No\ Info, L) + \Pr(R) \cdot U_i(No\ Info, R) \end{aligned}$$

We can re-write equation 4.A.2 as:

$$\Pr(L) [\Pr(White|L) U_i(White, L) + \Pr(Black|L) U_i(Black, L) - U_i(No\ Info, L)]$$

$$\begin{aligned}
& + \Pr(R) [\Pr(White|R) U_i(White, R) + \Pr(Black|R) U_i(Black, R) - U_i(No\ Info, R)] \\
& \geq C_{1,i} + C_2(\theta)
\end{aligned}$$

which tells us that a subject acquires information if her estimate of the expected gain of having information available is sufficiently higher than the expected gain when no information is available. Our specification makes use of a couple of assumptions. First, both types of cost are separable from benefits in the utility function and they are linearly additive. Second, cognitive costs for No Draw are normalized in such a way that we can write $C_{1,i} = C_{1,i}|Draw - C_{1,i}|NoDraw$.

Table 4.10 presents percentage of information seen in the decision block by type and treatment. As can be observed, Type 2 and Type 3 subjects observe information more often than Type 1, particularly in the Free and Forced treatments. In the costly treatment the difference between Type 2 and Type 3 subjects goes in the expected direction but falls short of statistical significance (MW: $p = 0.16$).

	Free	Costly	Forced
Type 0	0.23	–	0.86
	–		–
Type 1	0.17 (0.09)	0.06 (0.04)	0.11 (0.08)
Type 2	0.89 (0.05)	0.5 (0.09)	0.81 (0.05)
Type 3	0.96 (0.02)	0.68 (0.08)	0.94 (0.04)

Table 4.10: Information acquisition rates of types across treatments.

Notes: We drop the first 10 decision periods for these calculations as these are likely to be noisy with respect to information acquisition.

4.B Experiment Instructions

Below we provide an abridged transcript of the instructions. Square parentheses indicate changes in the sessions with a different Identification block.

In this experiment you will be asked to make decisions in 70 [80] periods, with one decision per period. The 70 periods are divided in 2 blocks of 40 decisions each. The first block has 40 periods, and the second block has 30 periods. [The 80 periods are divided in 2 blocks of 40 decisions each.] The type of decision is similar, but not identical, across the two blocks. The second block will only start when every participant in this room has finished the first block. You will receive instructions for the second block after the first

one is finished. The periods are not timed, which means that you can make decisions at your own pace. We estimate that each block should not take more than 40 minutes to complete.

Your earnings will be determined according to your performance in the experiment. Out of each block, 3 periods will be randomly selected to be paid (that is, 6 periods in total). All payoffs in the experiment are expressed in tokens. Each token in the experiment is worth 0.75 Euro.

FIRST BLOCK:In each period you can be in one of two States, Left and Right. There is some probability that you are in Left and some probability that you are in Right. Think of this as tomorrow's weather in Sydney: with a certain probability tomorrow will be cloudy and with a certain probability tomorrow will be sunny, but we don't know for sure what the weather in Sydney will be tomorrow. The same applies to the States in this experiment. The probability that the state is Left is 40% and the probability that the state is Right is 60%. As you can see, the two probabilities sum to 100%. These probabilities will be shown on your screen at all times.

Your decision in each period is to pick a number from 1 to 100. You can pick numbers in steps of 0.5, which means that 24 and 24.5 are possible, but 24.4 and 24.6 are not. Your payoff in each period will depend on your decision (the number you choose) and the actual State (Left or Right). Below you can see two graphs showing how the payoffs depend on your decision and the State:

(a graph similar to the one in Figure 4.1 was shown here)

These graphs show that if the State is Left, choosing 20 yields the highest payoff, and if the State is Right choosing 80 yields the highest payoff. However, if 20 is chosen and the State is Right, a negative payoff results. The same is true if 80 is chosen and the state is Left. Given that the actual state is not known when you must make your decision, choosing other values can make sense.

You can find a Table with the payoffs for all possible combinations of decisions and States in the last sheet. You will also be able to see those payoffs on the computer screen before making your decision.

In each period, a basket with 5 balls is presented. Some balls are black and some are white. The composition of the basket depends on the State. If the state is Left then there is one black ball and four white balls in the basket. If the state is Right then there are three black balls and two white balls in the basket.

(a graph depicting the distribution presented in Table 4.1 was shown here, see a graphical representation of the urns in Appendix D)

In each period, there is a 50% chance that you can see a ball drawn from the basket. Note that when the ball is drawn you still do not know what the State is, which means

that you don't know from which basket composition you are drawing the ball.

To summarize, the events in each period of the first block occur in the following order:

1. The State is randomly determined. You do not know what the State is at this point.

2. With a 50% chance you have the option of seeing a ball drawn from the basket.

3. You make your decision.

4. The State is revealed and your payoff is known.

SECOND BLOCK: You will now begin the second block of the experiment. Note that the State probabilities and the payoffs have changed from the first block you have just finished.

In this block the probability that the state is Left is 70% and the probability that the state is Right is 30%. As you can see, the two probabilities sum to 100%. These probabilities will be shown on your screen at all times.

Your payoff in each period will depend on your decision (the number you choose) and the actual State (Left or Right). Below you can see two graphs showing how the payoffs depend on your decision and the State:

(a graph similar to the one in Figure 4.1 was shown here)

These graphs show that if the State is Left, choosing 20 yields the highest payoff, and if the State is Right choosing 80 yields the highest payoff. However, if 20 is chosen and the State is Right, a negative payoff results. The same is true if 80 is chosen and the state is Left. Given that the actual state is not known when you must make your decision, choosing other values can make sense.

You can find a Table with the payoffs for all possible combinations of decisions and States in the last sheet. You will also be able to see those payoffs on the computer screen before making your decision.

In each period, a basket with 5 balls is presented. Some balls are black and some are white. The composition of the basket depends on the State. If the state is Left then there are four black balls and one white ball in the basket. If the state is Right then there are two black balls and three white balls in the basket.

(a graph depicting the distribution presented in Table 4.1 was shown here, see a graphical representation of the urns in Appendix D)

Note that when the ball is drawn you still do not know what the State is, which means that you don't know from which basket composition you are drawing the ball.

This block is composed of three sets of 10 decisions. Each set differs in the manner in which a ball can be drawn from the basket. Further instructions will be given on the computer screen before each set of 10 decisions. [In each period, there is a 50% chance that you can see a ball drawn from the basket.]

4.C Additional results

Table 4.11 presents the p-values of the Kolmogorov-Smirnov test for equality of distribution. It is applied to the differences in individual average decision distribution across treatment.

	A Black draw	A White draw	No draw
Free vs. Costly	0.08	0.01	0.83
Free vs. Forced	0.1	0.00	0.08
Costly vs. Forced	0.89	0.69	0.28

Table 4.11: p-values for Kolmogorov-Smirnov tests of decision distributions

In Table 4.5 two different models are estimated. The model of column (2), an extension of the model presented in equation 4.4, is:

$$\ln \frac{r}{1-r} = \tau + \gamma \cdot \ln \left(\frac{\Pr(\textit{Left})}{\Pr(\textit{Right})} \right) + \alpha_1 \cdot D_{\textit{Costly}} + \alpha_2 \cdot D_{\textit{Forced}} \\ + (\beta_1 + \beta_2 \cdot D_{\textit{Costly}} + \beta_3 \cdot D_{\textit{Forced}}) \cdot \ln \left(\frac{\Pr(\textit{Draw}|L)}{\Pr(\textit{Draw}|R)} \right)$$

The addition of the dummy variables in model 2 explicitly allows the possibility that the Costly and Forced treatments have an effect on individual decisions through channels other than representativeness (i.e. through the constant τ or the scope of base rate neglect γ). Table 4.12 presents additional estimations of equation 4.4 presented in Section 4.3.1. Columns (1) and (2) are identical to the columns presented in Table 4.5. Columns (3) and (4) include the same independent variables as Columns (1) and (2), respectively, but discard observations of subjects who make decisions outside $x \in [20, 80]$. As expected, the results are slightly less pronounced but remain significant.

Table 4.13 presents average aggregate decision by subjects types and ball draw. Table 4.14 is identical to Table 4.7, with the control variables explicitly shown. Math has a significant effect after both draws. A larger number of math courses leads to a shift towards less extreme decision making. Gender has a significant effect only after a Black draw. Female subjects tend towards less extreme decisions. Risk does not have a significant effect. It is worth mentioning that our measure of risk is truncated at risk neutrality ($Risk \in [0, 3]$ where $Risk = 3$ is risk neutrality). Our measure can thus not detect risk-loving behavior. Table 4.15 expands the analysis shown in Table 4.14 to *No draw* data. No variable is significant but the gender dummy: female subjects tend to make decisions closer to 50 than male subjects.

	Model			
	(1)	(2)	(3)	(4)
β_1 (Likelihood ratio)	1.02*** (0.11)	1.03*** (0.11)	0.96*** (0.10)	0.96*** (0.10)
β_2 (Likelihood ratio x Costly)	0.51*** (0.16)	0.49*** (0.17)	0.36** (0.15)	0.35** (0.16)
β_3 (Likelihood ratio x Forced)	0.71*** (0.19)	0.71** (0.19)	0.68*** (0.17)	0.68*** (0.18)
Costly	– –	–0.09 (0.12)		–0.08 (0.11)
Forced	– –	0.01 (0.14)		0.00 (0.13)
Constant	–0.24***	–0.21***	–0.26***	–0.23***
N	461	461	451	451

Table 4.12: Representativeness bias and treatment effects

Notes: Clustered standard errors used. ***/**/* indicates significance level at the 1%/5%/10%.

	Black draw			White draw			No draw		
	Type 1	Type 2	Type 3	Type 1	Type 2	Type 3	Type 1	Type 2	Type 3
Free	53.21 (4.67)	65.35 (1.48)	71.78 (1.91)	54.92 (2.21)	42.53 (1.46)	40.61 (2.55)	57.26 (2.87)	54.00 (0.83)	56.68 (1.49)
Costly	70.00 (7.42)	69.01 (1.80)	73.16 (2.08)	48.5 (9.5)	39.63 (1.86)	34.62 (2.48)	52.04 (1.13)	55.81 (1.15)	57.69 (1.08)
Forced	58.36 (15.44)	72.27 (2.05)	71.96 (2.21)	49.18 (9.43)	34.87 (3.12)	35.30 (2.28)	63.19 (3.78)	58.05 (1.56)	52.20 (2.18)

Table 4.13: Decision averages by treatment, information and type.

Notes: Standard errors in brackets.

	A Black draw			A White draw				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Costly	4.22** (1.89)	4.84** (1.81)	2.13 (2.19)	5.42** (2.20)	-4.21** (1.97)	-4.95** (1.90)	-4.89** (2.33)	-8.18*** (2.69)
CT2	-	-	-	-7.19* (3.84)	-	-	-	6.57 (4.45)
Forced	5.09** (2.33)	5.73** (2.34)	4.33* (2.40)	4.40* (2.38)	-6.44** (2.37)	-7.12*** (2.39)	-7.27*** (2.41)	-7.37*** (2.43)
Type 1	-	-	-13.68*** (4.04)	-13.42*** (3.93)	-	-	14.16*** (3.45)	13.91*** (3.43)
Type 2	-	-	-2.97* (1.91)	-0.74 (2.18)	-	-	1.66 (2.12)	-0.33 (2.48)
Math	-	-0.80** (0.39)	-0.77** (0.39)	-0.83** (0.38)	-	1.06*** (0.37)	1.32*** (0.41)	1.41*** (0.42)
Female	-	-3.68** (1.75)	-4.50** (2.03)	-4.67** (1.99)	-	2.55 (1.79)	2.71 (2.03)	2.73 (2.05)
Risk	-	1.29 (0.98)	0.98 (1.14)	0.94 (1.10)	-	-0.75 (0.9)	-1.22 (1.09)	-1.16 (1.08)
Constant	65.74	66.35	71.42	70.62	42.84	40.95	39.19	39.88
N	146	146	106	106	145	145	103	103

Table 4.14: Regression results by information condition (ball draw)

Notes: Type 0 subject decisions and outliers are not included in the analysis. Robust standard errors used. ***/**/* indicates significance level at the 1%/5%/10%.

	No draw			
	(1)	(2)	(3)	(4)
Costly	0.93 (1.09)	1.42 (1.03)	0.41 (1.26)	0.92 (1.53)
CT2	—	—	—	-1.12 (2.56)
Forced	1.61 (1.58)	2.04 (1.58)	1.31 (1.61)	1.34 (1.60)
Type 1	—	—	2.58 (2.07)	2.59 (2.10)
Type 2	—	—	0.89 (1.28)	1.27 (1.62)
Math	—	-0.40 (0.26)	-0.26 (0.26)	-0.27 (0.26)
Female	—	-1.97 (1.09)	-2.68* (1.37)	-2.70* (1.38)
Risk	—	0.33 (0.56)	0.67 (0.72)	0.65 (0.73)
Constant	54.98	55.97	55.17	55.06
N	164	164	117	117

Table 4.15: Regression results by information condition (no draw)

Notes: Type 0 subject decisions are not included in the analysis. Robust standard errors used. ***/**/* indicates significance level at the 1%/5%/10%..

4.D Biases in Bayesian updating

Figure 4.7 displays average decision in our data (DA), the Bayesian benchmark (BBM) and the ‘No Base Rate’ (NBR) decisions (the choice that would result from complete base-rate neglect, i.e. taking $p = 0.5$). The arrows in the bottom of the figure display the direction in which biases and risk attitudes affect individual decisions.

- Representativeness: according to this bias, posterior probabilities are assessed by the extent to which the signal is representative of the states. In our design a Black (White) ball is more representative of Right (Left), and therefore choices after a ball draw are shifted away from the decision when no ball is drawn. Conservatism (not displayed), in our context, is the opposite of representativeness. This bias consists in a failure to use the signal to the extent prescribed by Bayes rule.
- Base rate neglect: prior probabilities are underweighted and therefore perceived to be equal. In our experiment this shifts decisions in the direction of the optimal value in the less likely state (20). See NBR for the case of complete base rate neglect ($p = 0.5$ instead of $p = 0.6$).
- Risk aversion: a risk averse individual shifts her decision towards less risky options, which is translated in our design to a deviation towards 50. Such a deviation ensures a lower gap in the state dependent payoff than that entailed by BBM.

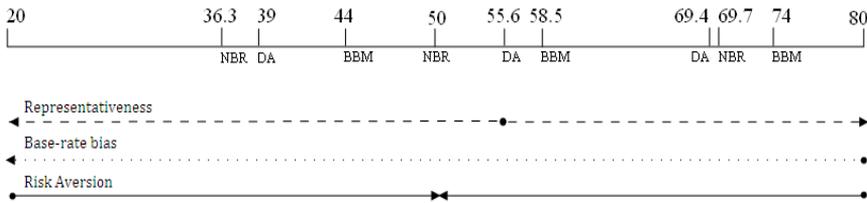


Figure 4.7: Decisions and the effect of biases on decision making.

Bibliography

- Ai, C. and Norton, E. C. (2003). Interaction terms in logit and probit models. Economics letters, 80(1):123–129.
- Arceneaux, K. (2007). I’m asking for your support: The effects of personally delivered campaign messages on voting decisions and opinion formation. Quarterly Journal of Political Science, 2:43–65.
- Ariely, D., Gneezy, U., Loewenstein, G., and Mazar, N. (2009). Large stakes and big mistakes. Review of Economic Studies, 76(2):451–469.
- Arkes, H. R. and Blumer, C. (1985). The psychology of sunk cost. Organizational Behavior and Human Decision Processes, 35(1):124–140.
- Ashraf, N., Berry, J., and Shapiro, J. M. (2010). Can higher prices stimulate product use? Evidence from a field experiment in zambia. The American Economic Review, 100(5):2383–2413.
- Berg, J., Dickhaut, J., and McCabe, K. (1995). Trust, reciprocity, and social history. Games and Economic Behavior, 10(1):122–142.
- Bhattacharya, S., Duffy, J., and Kim, S.-T. (2014). Information acquisition and voting mechanisms: Theory and evidence. Working Paper.
- Bolton, G. E. and Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. The American Economic Review, 90(1):166–193.
- Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., and Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. Nature, 489(7415):295–298.
- Borland, J., Lee, L., and Macdonald, R. D. (2011). Escalation effects and the player draft in the AFL. Labour Economics, 18(3):371–380.
- Bornstein, G. (1992). The free-rider problem in intergroup conflicts over step-level and continuous public goods. Journal of Personality and Social Psychology, 62(4):597–606.

- Bornstein, G., Rapoport, A., Kerpel, L., and Katz, T. (1989). Within- and between-group communication in intergroup competition for public goods. Journal of Experimental Social Psychology, 25(5):422–436.
- Braconnier, C., Dormagen, J.-Y., and Pons, V. (2014). Voter registration costs and disenfranchisement: experimental evidence from France. Working Paper.
- Brandts, J., Cooper, D. J., and Fatas, E. (2007). Leadership and overcoming coordination failure with asymmetric costs. Experimental Economics, 10(3):269–284.
- Brewer, M. B. (2007). The social psychology of intergroup relations: Social categorization, ingroup bias, and outgroup prejudice. In Kruglanski, A. and Higgins, E. T., editors, Social Psychology: Handbook of Basic Principles. Guilford Press, New York.
- Bryan, C. J., Walton, G. M., Rogers, T., and Dweck, C. S. (2011). Motivating voter turnout by invoking the self. Proceedings of the National Academy of Sciences, 108(31):12653–12656.
- Buchheit, S. and Feltovich, N. (2011). Experimental evidence of a sunk-cost paradox: a study of pricing behavior in Bertrand-Edgeworth duopoly. International Economic Review, 52(2):317–347.
- Burks, S. V., Carpenter, J. P., and Verhoogen, E. (2003). Playing both roles in the trust game. Journal of Economic Behavior & Organization, 51(2):195–216.
- Camerer, C. F. (1987). Do biases in probability judgment matter in markets? Experimental evidence. The American Economic Review, 77(5):981.
- Camerer, C. F. (1995). Individual decision making. In Kagel, J. H. and Roth, A. E., editors, The Handbook of Experimental Economics, pages 587–703. Princeton University Press.
- Camerer, C. F. and Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. Journal of Risk and Uncertainty, 19:7–42.
- Camerer, C. F., Loewenstein, G., and Rabin, M. (2004). Advances in Behavioral Economics. Princeton University Press.
- Camerer, C. F. and Weber, R. A. (1999). The econometrics and behavioral economics of escalation of commitment: a re-examination of Staw and Hoang’s NBA data. Journal of Economic Behavior & Organization, 39(1):59–82.
- Caplan, B. (2007). The myth of the rational voter: Why democracies choose bad policies. Princeton University Press.

- Carney, D. R., Jost, J. T., Gosling, S. D., and Potter, J. (2008). The secret lives of liberals and conservatives: Personality profiles, interaction styles, and the things they leave behind. Political Psychology, 29(6):807–840.
- Cason, T. N. and Mui, V.-L. (2005). Uncertainty and resistance to reform in laboratory participation games. European Journal of Political Economy, 21(3):708–737.
- Charness, G. and Gneezy, U. (2010). Portfolio choice and risk attitudes: An experiment. Economic Inquiry, 48(1):133–146.
- Charness, G., Rigotti, L., and Rustichini, A. (2007). Individual behavior and group membership. The American Economic Review, 91(4):1340–1352.
- Chen, Y. and Li, S. X. (2009). Group identity and social preferences. The American Economic Review, 99(1):431–457.
- Cohen, J. and Dupas, P. (2010). Free distribution or cost-sharing? Evidence from a randomized Malaria prevention experiment. The Quarterly Journal of Economics, 125(1):1–45.
- Cox, G. W. (1999). Electoral rules and the calculus of mobilization. Legislative Studies Quarterly, 24(3):387–419.
- Cunha, Jr, M. and Caldieraro, F. (2009). Sunk-cost effects on purely behavioral investments. Cognitive Science, 33(1):105–113.
- Dawes, C. T., Loewen, P. J., and Fowler, J. H. (2011). Social preferences and political participation. Journal of Politics, 73(03):845–856.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. Journal of Economic Literature, 47(2):315–72.
- Diehl, M. (1990). The minimal group paradigm: Theoretical explanations and empirical findings. European Review of Social Psychology, 1(1):263–292.
- Downs, A. (1957). An economic theory of democracy. New York: Harper and Row.
- Duffy, J. and Tavits, M. (2008). Beliefs and voting decisions: A test of the pivotal voter model. American Journal of Political Science, 52(3):603–618.
- Dufwenberg, M. and Kirchsteiger, G. (2004). A theory of sequential reciprocity. Games and Economic Behavior, 47(2):268–298.
- Duh, R. R. and Sunder, S. (1986). Incentives, Learning and Processing of Information in a Market Environment: An Examination of the Base-Rate

Fallacy. Center for Economic and Management Research, University of Oklahoma.

- Eckel, C. C. and Grossman, P. J. (2005). Managing diversity by creating team identity. Journal of Economic Behavior & Organization, 58(3):371–392.
- Edlin, A. S., Gelman, A., and Kaplan, N. (2007). Voting as a rational choice: Why and how people vote to improve the well-being of others. Rationality and Society, 1:293–314.
- Eger, C. and Dickhaut, J. (1982). An examination of the conservative information processing bias in an accounting framework. Journal of Accounting Research, 20(2):711–723.
- El-Gamal, M. A. and Grether, D. M. (1995). Are people Bayesian? Uncovering behavioral strategies. Journal of the American Statistical Association, 90(432):1137–1145.
- Elbittar, A., Gomberg, A., Martinelli, C., and Palfrey, T. R. (2014). Ignorance and bias in collective decisions. Working Paper.
- Ellemers, N. (2012). The group self. Science, 336(6083):848–852.
- Elster, J. (1989). Social norms and economic theory. The Journal of Economic Perspectives, 3(4):99–117.
- Engel, C. (2011). Dictator games: A meta study. Experimental Economics, 14(4):583–610.
- Evren, O. (2012). Altruism and voting: A large-turnout result that does not rely on civic duty or cooperative behavior. Journal of Economic Theory, 147(6):2124–2157.
- Eyster, E. (2002). Rationalizing the past: a taste for consistency. Working Paper.
- Falk, A., Fehr, E., and Fischbacher, U. (2008). Testing theories of fairness - intentions matter. Games and Economic Behavior, 62(1):287–303.
- Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. Games and Economic Behavior, 54(2):293–315.
- Falk, A. and Heckman, J. J. (2009). Lab experiments are a major source of knowledge in the social sciences. Science, 326(5952):535–538.
- Feddersen, T., Gailmard, S., and Sandroni, A. (2009). Moral bias in large elections: Theory and experimental evidence. American Political Science Review, 103(2):175–192.

- Feddersen, T. J. (2004). Rational choice theory and the paradox of not voting. The Journal of Economic Perspectives, 18(1):99–112.
- Fehr, E. and Fischbacher, U. (2004). Social norms and human cooperation. Trends in Cognitive Sciences, 8(4):185–190.
- Fehr, E. and Gächter, S. (2000). Fairness and retaliation: The economics of reciprocity. The Journal of Economic Perspectives, 14(3):159–181.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. The Quarterly Journal of Economics, 114(3):817–868.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. Experimental Economics, 10:171–178.
- Fowler, J. H. (2006). Altruism and turnout. Journal of Politics, 68(3):674–683.
- Fowler, J. H. and Kam, C. D. (2007). Beyond the self: Social identity, altruism, and political participation. Journal of Politics, 69(3):813–827.
- Friedman, D., Pommerenke, K., Lukose, R., Milam, G., and Huberman, B. (2007). Searching for the sunk cost fallacy. Experimental Economics, 10:79–104.
- Friedman, J. (1996). The rational choice controversy: Economic models of politics reconsidered. Yale University Press.
- Gerber, A. S. and Green, D. P. (1999). Does canvassing increase voter turnout? A field experiment. Proceedings of the National Academy of Sciences, 96(19):10939–10942.
- Gerber, A. S. and Green, D. P. (2000). The effects of canvassing, telephone calls, and direct mail on voter turnout: A field experiment. American Political Science Review, 94(3):653–663.
- Gerber, A. S., Green, D. P., and Larimer, C. W. (2008). Social pressure and voter turnout: Evidence from a large-scale field experiment. American Political Science Review, 102(1):33–48.
- Gerber, A. S., Green, D. P., and Larimer, C. W. (2010). An experiment testing the relative effectiveness of encouraging voter participation by inducing feelings of pride or shame. Political Behavior, 32(3):409–422.
- Gerber, A. S., Huber, G. A., Doherty, D., Dowling, C. M., Raso, C., and Ha, S. E. (2011). Personality traits and participation in political processes. Journal of Politics, 73(3):692–706.

- Gerber, A. S. and Rogers, T. (2009). Descriptive social norms and motivation to vote: Everybody's voting and so should you. Journal of Politics, 71(1):178–191.
- Gino, F. (2008). Do we listen to advice just because we paid for it? the impact of advice cost on its use. Organizational Behavior and Human Decision Processes, 107(2):234–245.
- Gneezy, U. and Potters, J. (1997). An experiment on risk taking and evaluation periods. The Quarterly Journal of Economics, 112(2):631–645.
- Goeree, J. K. and Holt, C. A. (2005). An explanation of anomalous behavior in models of political participation. American Political Science Review, 99(2):201–213.
- Gonzalez, R. and Wu, G. (1999). On the shape of the probability weighting function. Cognitive Psychology, 38(1):129–166.
- Goren, H. and Bornstein, G. (2000). The effects of intragroup communication on intergroup cooperation in the repeated intergroup prisoner's dilemma (IPD) game. The Journal of Conflict Resolution, 44(5):700–719.
- Green, D. P. and Gerber, A. S. (2004). Get Out the Vote!: How to Increase Voter Turnout. Brookings Institution Press.
- Green, D. P., McGrath, M. C., and Aronow, P. M. (2013). Field experiments and the study of voter turnout. Journal of Elections, Public Opinion & Parties, 23(1):27–48.
- Green, D. P. and Shapiro, I. (1994). Pathologies of rational choice theory: A critique of applications in political science. Cambridge University Press.
- Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. The Quarterly Journal of Economics, 95(3):537–557.
- Grether, D. M. (1992). Testing bayes rule and the representativeness heuristic: Some experimental evidence. Journal of Economic Behavior & Organization, 17(1):31–57.
- Großer, J. and Schram, A. (2006). Neighborhood information exchange and voter participation: an experimental study. American Political Science Review, 100(2):235–248.
- Großer, J. and Schram, A. (2010). Public opinion polls, voter turnout, and welfare: an experimental study. American Journal of Political Science, 54(3):700–717.

- Großer, J. and Seebauer, M. (2014). The curse of uninformed voting: An experimental study. Working Paper.
- Herrera, H., Morelli, M., and Palfrey, T. (2014). Turnout and power sharing. The Economic Journal, 124(154):131–162.
- Hofmann, K. (1991). Aus Gesprächen mit Thomas Bernhard. DTV Deutscher Taschenbuch Verlag.
- Hogg, M. A. and Abrams, D. (1988). Social identifications: A social psychology of intergroup relations and group processes. Routledge.
- Holbrook, A. L. and Krosnick, J. A. (2010). Social desirability bias in voter turnout reports tests using the item count technique. Public Opinion Quarterly, 74(1):37–67.
- Holt, C. and Smith, A. (2009). An update on bayesian updating. Journal of Economic Behavior & Organization, 69(2):125–134.
- Isaac, R. M. and Walker, J. M. (1988). Group size effects in public goods provision: The voluntary contributions mechanism. The Quarterly Journal of Economics, 103(1):pp. 179–199.
- Jankowski, R. (2002). Buying a lottery ticket to help the poor: Altruism, civic duty, and self-interest in the decision to vote. Rationality and Society, 14(1):55–77.
- Jankowski, R. (2007). Altruism and the decision to vote: Explaining and testing high voter turnout. Rationality and Society, 19(1):5–34.
- John, O. P., Naumann, L. P., and Soto, C. J. (2008). Paradigm shift to the integrative Big-Five trait taxonomy: History, measurement, and conceptual issues. In Handbook of Personality: Theory and Research, pages 114–158. Guilford Press.
- Johnson, N. D. and Mislin, A. A. (2011). Trust games: A meta-analysis. Journal of Economic Psychology, 32(5):865–889.
- Jost, J. T. (2006). The end of the end of ideology. American Psychologist, 61(7):651–670.
- Jost, J. T., Glaser, J., Kruglanski, A. W., and Sulloway, F. J. (2003). Political conservatism as motivated social cognition. Psychological Bulletin, 129(3):339–375.
- Kahneman, D. and Tversky, A. (1972). Subjective probability: a judgment of representativeness. Cognitive Psychology, 3(3):430–454.

- Kahneman, D. and Tversky, A. (1973). On the psychology of prediction. Psychological Review, 80(4):237–251.
- Kenny, C. B. (1992). Political participation and effects from the social environment. American Journal of Political Science, 36(1):259–267.
- Kerr, N. L., Garst, J., Lewandowski, D. A., and Harris, S. E. (1997). That still, small voice: Commitment to cooperate as an internalized versus a social norm. Personality and Social Psychology Bulletin, 23(12):1300–1311.
- King, G., Pan, J., and Roberts, M. E. (2013). How censorship in China allows government criticism but silences collective expression. American Political Science Review, 107(02):326–343.
- Knack, S. (1992). Social altruism and voter turnout: Evidence from the 1991 NES pilot study. Technical Report nes002294, ANES Pilot Study Report.
- Koehler, J. J. (1996). The base rate fallacy reconsidered: Descriptive, normative and methodological challenges. Behavioral and Brain Sciences, 19:1–53.
- Leighley, J. E. and Vedlitz, A. (1999). Race, ethnicity, and political participation: Competing models and contrasting explanations. Journal of Politics, 61(4):1092–1114.
- Leuven, E., Oosterbeek, H., Sonnemans, J., and Van Der Klaauw, B. (2011). Incentives versus sorting in tournaments: Evidence from a field experiment. Journal of Labor Economics, 29(3):637–658.
- Levine, D. K. and Palfrey, T. R. (2007). The paradox of voter participation? A laboratory study. American Political Science Review, 101(1):143–158.
- Liebrand, W. B. (1984). The effect of social motives, communication and group size on behaviour in an n-person multi-stage mixed-motive game. European Journal of Social Psychology, 14(3):239–264.
- Lijphart, A. (1997). Unequal participation: democracy’s unresolved dilemma. American Political Science Review, 91(91):1–14.
- Martinelli, C. (2006). Would rational voters acquire costly information? Journal of Economic Theory, 129(1):225–251.
- Martinelli, C. (2007). Rational ignorance and voting behavior. International Journal of Game Theory, 35(3):315–335.
- Matland, R. E. and Murray, G. R. (2012). An experimental test of mobilization effects in a Latino community. Political Research Quarterly, 65(1):192–205.

- McCabe, K. A., Rigdon, M. L., and Smith, V. L. (2003). Positive reciprocity and intentions in trust games. Journal of Economic Behavior & Organization, 52(2):267–275.
- McFadden, D. (2001). Economic choices. The American Economic Review, 91(3):351–378.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. Games and Economic Behavior, 10(1):6–38.
- Mondak, J. J., Hibbing, M. V., Canache, D., Seligson, M. A., and Anderson, M. R. (2010). Personality and civic engagement: An integrative framework for the study of trait effects on political behavior. American Political Science Review, 104(1):85–110.
- Mutz, D. C. (2002). The consequences of cross-cutting networks for political participation. American Journal of Political Science, 46(4):838–855.
- Nickerson, D. W. (2007). Friends dont make friends vote: Selection and reputation in voter mobilization. Annual Meeting of the American Political Science Association, Chicago, IL.
- Nickerson, D. W. (2008). Is voting contagious? Evidence from two field experiments. American Political Science Review, 102(1):49–57.
- Norton, E. C., Wang, H., and Ai, C. (2004). Computing interaction effects and standard errors in logit and probit models. Stata Journal, 4:154–167.
- Offerman, T. and Potters, J. (2006). Does auctioning of entry licences induce collusion? An experimental study. Review of Economic Studies, 73(3):769–791.
- Offerman, T., Sonnemans, J., and Schram, A. (1996). Value orientations, expectations and voluntary contributions in public goods. The Economic Journal, 106(437):817–845.
- Offerman, T., Sonnemans, J., Van De Kuilen, G., and Wakker, P. (2009). A truth serum for non-Bayesians: Correcting proper scoring rules for risk attitudes. Review of Economic Studies, 76(4):1461–1489.
- Olson, M. (1965). The Logic of Collective Action: Public Goods and the Theory of Groups. Harvard University Press.
- Otto, A. R. (2010). Three attempts to replicate the behavioral sunk-cost effect: A note on Cunha and Caldieraro (2009). Cognitive Science, 34(8):1379–1383.

- Palfrey, T. R. (2009). Laboratory experiments in political economy. Annual Review of Political Science, 12:379–388.
- Panagopoulos, C. (2010). Affect, social pressure and prosocial motivation: Field experimental evidence of the mobilizing effects of pride, shame and publicizing voting behavior. Political Behavior, 32(3):369–386.
- Panagopoulos, C. (2011). Thank you for voting: Gratitude expression and voter mobilization. Journal of Politics, 73(03):707–717.
- Pollock III, P. H. (1982). Organizations as agents of mobilization: How does group activity affect political participation? American Journal of Political Science, 26(3):485–503.
- Rabin, M. (2013a). An approach to incorporating psychology into economics. The American Economic Review, 103(3):617–622.
- Rabin, M. (2013b). Incorporating limited rationality into economics. Journal of Economic Literature, 51(2):528–543.
- Robalo, P. and Sayag, R. (2014). Paying is believing: The effect of costly information on bayesian updating. Working Paper.
- Robalo, P., Schram, A., and Sonnemans, J. (2013). Other-regarding preferences, group identity and political participation: an experiment. Timbergen Institute Discussion Paper, 13-079(I).
- Rosenstone, S. and Hansen, J. M. (1993). Mobilization, participation and democracy in America. MacMillan Publishing.
- Rowley, C. K., Schneider, F., and Tollison, R. D. (1993). The next twenty-five years of Public Choice. Springer.
- Schram, A. and Sonnemans, J. (1996a). Voter turnout as a participation game: An experimental investigation. International Journal of Game Theory, 25(3):385–406.
- Schram, A. and Sonnemans, J. (1996b). Why people vote: Experimental evidence. Journal of Economic Psychology, 17(4):417–442.
- Staw, B. M. and Hoang, H. (1995). Sunk costs in the NBA: Why draft order affects playing time and survival in professional basketball. Administrative Science Quarterly, 40(3):474–494.
- Tajfel, H. (1982). Social psychology of intergroup relations. Annual Review of Psychology, 33(1):1–39.

- Thaler, R. (1980). Toward a positive theory of consumer choice. Journal of Economic Behavior & Organization, 1(1):39–60.
- Tyszler, M. (2008). Information and strategic voting. Master's thesis, University of Amsterdam, Tinbergen Institute.
- Van Dijk, F., Sonnemans, J., and van Winden, F. (2002). Social ties in a public good experiment. Journal of Public Economics, 85(2):275–299.
- Verba, S. and Nie, N. H. (1972). Participation in America: Social equality and political democracy. Harper and Row.
- Young, H. P. (1988). Condorcet's theory of voting. American Political Science Review, 82(04):1231–1244.

Summary

Understanding political participation is one of the big challenges in the social sciences. Whenever they are called upon to intervene in political affairs, individuals face a dilemma: on the one hand, a desire to participate - to vote, protest, campaign, or speak out; on the other hand, a propensity to be “left alone” and let others take responsibility and bear the costs of the political process.

This thesis explores how non-standard preferences (Chapters 2 and 3) and decision making biases (Chapter 4) might influence the decision to participate and the optimality of political choices, respectively. The standard models in political economy have mirrored the central tenets of the neoclassical or rational choice view, namely self-interested preference orderings and full rationality (Rowley et al. 1993). This paradigm has yielded substantial insights into the political process. However, some phenomena seem to elude the paradigm, chiefly among them the inability to fully explain individual political participation. This thesis aims at extending the rational choice model by providing directions in which its microfoundations can be enriched. The final goal is to increase its explanatory content and predictive power. Despite looking for new directions to improve the microfoundations of political economy models, the work in this thesis retains the economic approach’s conventional techniques. All chapters use a combination of theoretical models and laboratory experiments. Primacy is always given to what can be learned from the experimental data - the models are developed to the extent that they can make predictions regarding what to expect in the laboratory.

Chapter 2 deals with the effect of group identity and altruism on the decision to participate and is based on joint work with Arthur Schram and Joep Sonnemans. When groups with diverging interests settle their disputes via democratic politics, which is the case in *any* election, the allegiances each individual has to the group should matter for his decision to participate. More broadly, how much more she cares about an individual of her group than an individual of the other group should determine her willingness to endure the costs of participation. In fact, group identity seems to be a driving force of participation: in the United States, African-americans

participate at higher rates than their socioeconomic status would predict (Leighley and Vedlitz 1999). This constitutes a puzzle because socioeconomic status is typically the best predictor of individual political participation. One of the candidate explanations for this puzzle is a heightened sense of group identification. However, field data makes it extremely hard to identify the causal effect of group identification per se on participation, as it evolves concomitantly with mobilization and socialization processes. A laboratory experiment grants us the control necessary for this investigation.

In the laboratory, each participant is part of a group, which interacts with another group in simulated elections (the participation game of Palfrey and Rosenthal 1983). Using a novel procedure, we manage to induce different levels of group identification in the experimental treatments. Based on the theoretical results and the extant evidence, we hypothesize that both individual and aggregate participation should be increasing in the level of individual group identification and treatment-level group identity, respectively. Concurrently, our experiment also allows us to test whether individuals with more altruistic preferences tend to participate more often. At the aggregate level, no pronounced differences in participation are observed between high group identity and low group identity environments. At the individual level, there is a modest effect of group identification on participation. A more robust effect is found for non-group-specific altruistic concerns. This is in line with the altruism theories of voter turnout (e.g. Evren 2012). A by-product contribution of this study is methodological, as we propose a procedure that manages to induce different levels of group identity in the laboratory without resorting to natural groups.

Just as much a concern for others is a reason to participate, having been asked by others do to so also seems to work. A large field experimental literature (spanned by the seminal work of Gerber and Green 1999), has showed that certain mobilization tactics (like door-to-door canvassing) are successful in increasing turnout. However, other mobilization tactics, like mass mailings, are largely ineffective. The question remains as to what drives people to respond to mobilization efforts. Chapter 3 investigates the psychological mechanism underlying this phenomenon. The starting point is the observation that all mobilization efforts involve both a material effort and a normative appeal (to participate). The effect of mobilization on participation could then work either via reciprocity concerns, i.e. participation as a token of appreciation for the material effort of mobilization; and/or via compliance with normative appeals, i.e. participation in order to avoid the disutility associated with violating the participation norm that mobilization makes salient.

The laboratory election framework of Chapter 2 is extended to allow for both participation and mobilization. The experimental treatments consist in varying the mobilization method (human-driven, by group leaders, or automated) and the normative appeal conveyed by the mobilizing subject to others in the same group (present or absent). The main results show that the normative appeal is successful in increasing participation, in particular when it is coupled with a mobilization effort. Mobilization alone is not enough to increase participation, which disconfirms the reciprocity conjecture. I also carry out an assessment of the model's point and comparative statics predictions, and conclude that the behavior of group leaders is not in line with the point predictions, while most comparative statics results of the model seem to hold.

Information is a crucial determinant of good decisions, be they individual or collective. Chapter 4, which is based on joint work with Rei Sayag, proposes a first approach to the question of whether the costs of information affect the way it is incorporated in decision making. The findings are potentially relevant for two well-known and related results in political economy: rational ignorance and the Condorcet jury theorem.

The literature dealing with these topics, as most of the literature on individual decision making, assumes that information is incorporated in individuals' judgments via Bayes Rule. It is further assumed that Bayes rule is applied irrespective of the cost of information. In our study we ask whether this presumption is legitimate. We construct an individual decision making task under risk and vary the way in which information is made available to subjects: for free, optionally at a cost or imposed at a cost. The laboratory allows us to circumvent the problematic selection issues present in the field, where the subjects who acquire information are the ones most likely to benefit from it. We find that the assumption that Bayesian updating does not depend on information's cost should be questioned: subjects weigh more heavily both the signals they choose to acquire and the signals that they were forced to acquire. In sum, costly information is given a higher weight than free information, which leads to more extreme moves in posterior beliefs. Whether this results in more optimal decision making depends on how far the posterior under free information lies from the normative optimum. In other words, if a decision maker is not using information to the extent prescribed by the normative model, letting her buy it or make her pay for it is likely to lead to a better outcome.

Samenvatting (Summary in Dutch)¹

Het begrijpen van politieke participatie is een van de grote uitdagingen in de sociale wetenschappen. Wanneer men wordt opgeroepen om deel te nemen in politieke zaken, worden individuen geconfronteerd met een dilemma. Aan de ene kant een verlangen om te participeren - om te stemmen, te protesteren, campagne te voeren of om zich uit te spreken; aan de andere kant is er een neiging om 'alleen gelaten' te willen worden en anderen de verantwoordelijkheid te laten nemen en de kosten van het politieke proces te laten dragen.

Dit proefschrift onderzoekt hoe niet-standaard voorkeuren en afwijkende ('biased') besluitvorming participatie en de optimaliteit van politieke keuzes kunnen beïnvloeden. De standaard modellen in de politieke economie weerspiegelen de grondbeginselen van het neoklassieke uitgangspunt, namelijk eigenbelang en volledige rationaliteit (Rowley et al. 1993). Dit paradigma heeft veel inzicht gegeven in het politieke proces. Maar sommige verschijnselen lijken niet binnen het paradigma te passen. In het bijzonder heeft dit paradigma het moeilijk om individuele politieke participatie te verklaren. Dit proefschrift richt zich op het uitbreiden van het rationele keuzemodel door nieuwe richtingen aan te geven waarmee de microgrondslagen van het politieke gedrag kunnen worden verrijkt. Het uiteindelijke doel is om de verklarende en voorspellende kracht van het model te vergroten. Bij de zoektocht naar nieuwe manieren om de microgrondslagen van de politiek economische modellen te verbeteren, blijft het werk in dit proefschrift gebruik maken van technieken die binnen de economische benadering gebruikelijk zijn. Alle hoofdstukken gebruiken een combinatie van theoretische modellen en laboratorium experimenten. Voorrang wordt steeds verleend aan wat geleerd kan worden van de experimentele data - de modellen zijn ontwikkeld om zo goed mogelijk te voorspellen wat er verwacht kan worden in het laboratorium.

¹Deze samenvatting is tot stand gekomen met behulp van Boris van Leeuwen, Jozefina Milanovski en Arthur Schram.

Hoofdstuk 2 gaat over het effect van groepsidentiteit en altruïsme op de beslissing om te participeren en is gebaseerd op gezamenlijk werk met Arthur Schram en Joep Sonnemans. Wanneer groepen met uiteenlopende belangen hun geschillen beslechten via de democratische politiek, zoals bijvoorbeeld het geval is bij verkiezingen, dan zou de loyaliteit van een individu aan een groep van invloed kunnen zijn op de beslissing om te participeren. Meer in het algemeen zal de bereidheid om de kosten van participatie te dragen mede afhangen van de mate waarin men meer geeft om iemand uit de eigen groep dan om iemand uit een andere groep. In de praktijk lijkt groepsidentiteit inderdaad een drijvende kracht achter participatie te zijn; in de Verenigde Staten, bijvoorbeeld, is participatie onder de zwarte bevolkingsgroep hoger dan op basis van sociaal-economische status te verwachten zou zijn (Leighley en Vedlitz 1999). Dit wordt gezien als een raadsel, omdat sociaal-economische status normaal gesproken de beste voorspeller van individuele politieke participatie is. Een van de mogelijke verklaringen voor dit verschijnsel is een verhoogd gevoel van groepsidentificatie onder de zwarte bevolking. Het is echter uiterst moeilijk om het causale effect van groepsidentificatie op participatie te identificeren op basis van data uit het veld, omdat het gelijktijdig evolueert met mobilisatie en socialisatie processen. Een laboratoriumexperiment geeft ons de noodzakelijke controle over de situatie om dit te onderzoeken.

In het laboratorium is elke deelnemer lid van een groep, die samen met een andere groep deelneemt aan gesimuleerde verkiezingen (het participatie-spel' van Palfrey and Rosenthal 1983). Aan deze laboratorium verkiezingen voegen wij de mogelijkheid toe dat men preferenties heeft die waarde toekennen aan hoe het met de anderen in de eigen groep gaat. In het experiment slagen we erin om verschillende niveaus van groepsidentificatie te induceren. Op basis van theoretische analyses en eerder empirisch onderzoek, veronderstellen we dat zowel individuele als geaggregeerde participatie toeneemt met de mate van respectievelijk de individuele groepsidentificatie en het niveau van de genduceerde groepsidentiteit in de verschillende experimenten. Tegelijkertijd biedt ons experiment de mogelijkheid om te testen of individuen met altruïstische voorkeuren vaker participeren. Op geaggregeerd niveau vinden we geen verschil in participatie tussen omgevingen met een sterke of zwakke groepsidentiteit. Op individueel niveau vinden we een bescheiden effect van groepsidentiteit op de mate van participatie. We vinden een duidelijker effect van niet-groepsgebonden-altruïsme. Dit is in overeenstemming met theorieën over altruïsme en de opkomst bij verkiezingen (bijv. Evren 2012). Daarnaast levert deze studie een methodologische bijdrage: we introduceren een nieuwe procedure om verschillende niveaus van groepsidentiteit in het lab te induceren, zonder een

toevlucht te nemen tot natuurlijke groepen.

Net zoals groepsidentiteit en altruïsme een reden kunnen zijn om te participeren, kan een verzoek van een ander ook leiden tot participatie. Een uitgebreide experimentele literatuur die gebruik maakt van veld-data (zie het baanbrekende werk van Gerber en Green 1999), heeft aangetoond dat bepaalde mobilisatie tactieken (zoals deur-tot-deur colportage) succesvol zijn in het verhogen van de opkomst bij verkiezingen. Echter, andere mobilisatie tactieken, zoals massa-mailings, zijn grotendeels ineffectief. De vraag is, wat mensen drijft om te reageren op mobilisatie inspanningen. Hoofdstuk 3 onderzoekt het onderliggende psychologische mechanisme achter dit fenomeen. Het uitgangspunt is dat alle mobilisatie inspanningen zowel een materile inspanning als een normatief beroep op de kiezer omvatten. Mobilisatie kan dan werken via wederkerigheid, dat wil zeggen als een blijk van waardering voor de materile inspanning; en/of via de naleving van het normatieve beroep dat gedaan wordt, dat wil zeggen dat men het schenden van de norm wil voorkomen.

Het kader van hoofdstuk 3 is uitgebreid ten opzichte van de oorspronkelijke participatiespelen om het bestuderen van mobilisatie en participatie mogelijk te maken. In de experimenten varieer ik systematisch tussen de gebruikte mobilisatiemethode (een oproep door een persoon, de groepsleider, of een geautomatiseerde oproep) en of er (wel of niet) een normatief beroep wordt gedaan op het individu. De belangrijkste resultaten tonen aan dat een normatief beroep op de kiezer leidt tot verhoogde participatie, met name wanneer deze is gekoppeld aan mobilisatie. Mobilisatie op zichzelf is niet genoeg om participatie te verhogen, wat in tegenspraak is met de wederkerigheidshypothese. Daarnaast voer ik een analyse van de (punt)voorspellingen en de comparatieve statica van het model uit. Hierin vind ik dat het gedrag van leiders niet in overeenstemming is met de puntvoorspellingen maar we met de meeste comparatieve statica.

Informatie is een cruciale factor in het nemen van de juiste beslissingen, zowel voor individuele als voor collectieve beslissingen. Hoofdstuk 4, dat gebaseerd is op gezamenlijke werk met Rei Sayag, biedt een eerste benadering van de vraag of de kosten van de informatie invloed hebben op de manier waarop deze informatie wordt verwerkt in besluitvorming. De bevindingen zijn van belang voor de twee bekende en verwante resultaten in de politieke economie: rationele onwetendheid en het Condorcet jury theorema.

De literatuur over deze onderwerpen, zoals het merendeel van de literatuur met betrekking tot individuele besluitvorming, gaat ervan uit dat individuen informatie verwerken via de stelling van Bayes. Verder wordt aangenomen dat de stelling van Bayes wordt toegepast ongeacht de kosten van de informatie. In ons onderzoek

stellen we de vraag of deze aanname legitiem is. Wij construeren een individuele besluitvormingstaak waar risico een rol speelt, en varieren hoe de informatie aan proefpersonen beschikbaar wordt gemaakt: gratis, tegen een prijs of opgelegd tegen een prijs. Het laboratorium stelt ons in staat om selectieproblemen die optreden met veld-data te omzeilen: in het veld is het waarschijnlijk dat degenen die informatie verwerven degenen zijn die er het meeste baat bij hebben. Wij vinden dat de kosten van informatie invloed hebben op het toepassen van de stelling van Bayes: deelnemers aan het experiment hechten teveel waarde aan zowel de informatie die ze kiezen om te kopen als de informatie die ze gedwongen kopen. Kortom, kostbare informatie wordt zwaarder gewogen dan gratis informatie, wat leidt tot meer extreme bewegingen in de posterieure verwachtingen. Of dit leidt tot betere besluitvorming hangt af van hoever de posterieure verwachting met gratis informatie af ligt van het normatieve optimum.

Sumário (Summary in Portuguese)

Compreender a participação política é um dos grandes desafios das ciências sociais. Sempre que são chamados a intervir em assuntos políticos, os indivíduos enfrentam um dilema: por um lado, o desejo de participar - de votar, protestar, fazer campanha, ou fazer-se ouvir; por outro lado, e aludindo à citação que abre o Capítulo 1, um desejo de ser 'deixado em paz' para que os outros assumam a responsabilidade do processo político e acarretem com os seus custos.

Esta tese explora como preferências não-*standard* e enviesamentos cognitivos sistemáticos podem influenciar a decisão de participar bem como a optimalidade das escolhas políticas. Os modelos tradicionais de economia política espelham os princípios centrais da visão 'neoclássica' ou da 'escolha racional', nomeadamente preferências guiadas pelo interesse próprio e racionalidade pura dos agentes (Rowley et al., 1993). Este paradigma tem rendido uma compreensão mais aprofundada do processo político. No entanto, alguns fenómenos parecem eludir o paradigma, com particular relevância a incapacidade de explicar a participação política individual. Esta tese tem como objetivo estender o modelo de escolha racional, providenciando direcções ao longo das quais as suas fundações *micro* podem ser enriquecidas. O objectivo final é aumentar o seu conteúdo explicativo e poder preditivo. Apesar de procurar novos caminhos para melhorar as fundações *micro* dos modelos de economia política, o trabalho contido nesta tese retém as técnicas convencionais da abordagem económica. Todos os capítulos usam uma combinação de modelos teóricos e experiências de laboratório. A primazia é dada ao que pode ser deduzido dos dados experimentais - os modelos são desenvolvidos na medida em que podem ajudar a prever e interpretar o que poderá ser observado no laboratório.

O Capítulo 2 aborda o efeito da identidade de grupo e do altruísmo na decisão de participar, e é baseado em trabalho conjunto com Arthur Schram e Joep Sonnemans. Quando grupos com interesses divergentes resolvem os seus diferendos através do processo democrático, o que é o caso em qualquer eleição, os laços que cada indivíduo tem para com o grupo são passíveis de influenciar a sua decisão de participação. Em termos mais gerais, a preocupação que um indivíduo tem para

com um elemento do seu grupo relativamente a um indivíduo de outro grupo deve determinar a sua disponibilidade para suportar os custos de participação. Na verdade, a identidade de grupo parece ser uma das forças motrizes da participação: nos Estados Unidos, a população Afro-americana participa a taxas mais elevadas do que o seu estatuto socioeconómico poderia prever (Leighley e Vedlitz 1999). Trata-se de um *puzzle*, uma vez que o nível socioeconómico é geralmente a variável que melhor prevê a participação política individual. Uma das potenciais explicações para este *puzzle* é o sentimento de identidade de grupo. No entanto, os dados de campo tornam extremamente difícil identificar o efeito causal da identidade de grupo *per se* na participação, uma vez que esta evolui concomitantemente com os processos de mobilização e de socialização. Uma experiência de laboratório dá-nos o controlo necessário a esta investigação.

No laboratório, cada participante é parte de um grupo, o qual interage com outro grupo em eleições simuladas (o jogo de participação de Palfrey e Rosenthal 1983). Usando um novo procedimento experimental, conseguimos induzir diferentes níveis de identificação de grupo nos tratamentos. Com base nos resultados teóricos e na evidência existente, tomamos como hipóteses que tanto a participação individual como a participação agregada estarão positivamente correlacionadas com o nível individual de identificação com o grupo e com o nível de identidade de grupo induzido nos diferentes tratamentos, respectivamente. Ao mesmo tempo, a nossa experiência também nos permite testar se os indivíduos com preferências altruístas tendem a participar mais vezes. Ao nível agregado, não são observadas diferenças substanciais na participação consoante o nível de identidade de grupo. A nível individual, há um efeito modesto da identidade de grupo sobre a participação. Um efeito mais forte é encontrado em indivíduos com preferências altruístas. Este resultado está em consonância com as teorias de participação eleitoral altruísta (por exemplo, Evren 2012). Um sub-produto importante deste estudo é de cariz metodológico, uma vez que propomos um procedimento que consegue induzir diferentes níveis de identidade de grupo no laboratório sem recorrer a categorias existentes fora dele, como a etnia ou o estrato socio-económico.

Assim como uma preocupação pelo bem-estar de outros é uma razão para participar, tendo sido incentivado por outros a fazê-lo parece também surtir efeito. Uma extensa literatura de experiências de campo (iniciada pelo trabalho seminal de Gerber e Green 1999), mostra que certas tácticas de mobilização (como campanhas porta-a-porta) são bem-sucedidas em aumentar a afluência às urnas. No entanto, outras tácticas de mobilização, como o envio de cartas em massa, redundam na sua maioria ineficazes. A questão permanece relativamente ao que leva as pessoas a re-

sponder aos esforços de mobilização. O Capítulo 3 investiga o mecanismo psicológico subjacente a este fenómeno. O ponto de partida é a observação de que todos os esforços de mobilização envolvem tanto um esforço material como um apelo normativo à participação. A mobilização poderia então surtir o seu efeito na participação devido a uma atitude de reciprocidade, ou seja, a participação enquanto gesto de agradecimento pelo esforço material envolvido no acto de mobilização; e/ou através da adesão aos apelos normativos veiculados pela mobilização, de forma a evitar a perda de bem-estar associada à violação da norma de participação.

O modelo teórico do Capítulo 2 é estendido de forma a permitir mobilização e participação. Os tratamentos experimentais diferem consoante o método de mobilização empregue (por seres humanos – os líderes do grupo – ou automatizado), e consoante a presença ou ausência de um apelo normativo veiculado por um indivíduo a outros do seu grupo. Os principais resultados mostram que o apelo normativo é bem-sucedido em aumentar a participação, em particular quando é combinado com o esforço de mobilização. A mobilização por si só não é suficiente para aumentar a participação, o que rejeita a conjectura da reciprocidade. No Capítulo 3 é também levada a cabo uma avaliação das previsões do modelo, o que nos permite concluir que o comportamento dos líderes não está em linha com as previsões numéricas, embora a maioria dos resultados de estática comparada sejam observados.

A informação é um determinante crucial de boas decisões, sejam elas individuais ou coletivas. O Capítulo 4, que é baseado em trabalho conjunto com Rei Sayag, propõe uma primeira abordagem à questão de saber se o custo da informação afecta a forma como ela é incorporada no processo decisório. Os resultados são potencialmente relevantes para dois conhecidos resultados em economia política: a ignorância racional e o teorema do júri de Condorcet.

A literatura que trata estes temas, assim como a maior parte da literatura sobre a tomada de decisão individual, assume que a informação é incorporada no julgamento individual através da regra de Bayes. Assume-se ainda que a regra de Bayes é aplicada independentemente do custo da informação. Este estudo questiona a legitimidade desta presunção. Para tal, delineamos um exercício de tomada de decisão individual sob risco e variamos a forma como a informação é disponibilizada aos participantes: de graça, opcionalmente com um custo ou imposta com um custo. O laboratório permite-nos contornar os problemas de seleção presentes no campo, onde os indivíduos que adquirem informação são os mais susceptíveis de beneficiar dela. Concluímos que a presunção de que o uso da regra de Bayes não depende do custo da informação deve ser questionada: os indivíduos tendem a sobrevalorizar tanto a informação que escolhem adquirir como a informação que são forçados a

adquirir. Em suma, a informação com custos é alvo de uma ponderação mais forte na tomada de decisões do que a informação sem custos. Se tal leva a melhores ou piores decisões depende de quão longe os indivíduos se encontram do ideal normativo quando a informação não tem custos.

The Tinbergen Institute is the Institute for Economic Research, which was founded in 1987 by the Faculties of Economics and Econometrics of the Erasmus University Rotterdam, University of Amsterdam and VU University Amsterdam. The Institute is named after the late Professor Jan Tinbergen, Dutch Nobel Prize laureate in economics in 1969. The Tinbergen Institute is located in Amsterdam and Rotterdam. The following books recently appeared in the Tinbergen Institute Research Series:

548. B. PELLOUX, the Role of Emotions and Social Ties in Public On Good Games: Behavioral and Neuroeconomic Studies
549. N. YANG, Markov-Perfect Industry Dynamics: Theory, Computation, and Applications
550. R.R. VAN VELDHUIZEN, Essays in Experimental Economics
551. X. ZHANG, Modeling Time Variation in Systemic Risk
552. H.R.A. KOSTER, The internal structure of cities: the economics of agglomeration, amenities and accessibility.
553. S.P.T. GROOT, Agglomeration, globalization and regional labor markets: micro evidence for the Netherlands.
554. J.L. MHLMANN, Globalization and Productivity Micro-Evidence on Heterogeneous Firms, Workers and Products
555. S.M. HOOGENDOORN, Diversity and Team Performance: A Series of Field Experiments
556. C.L. BEHRENS, Product differentiation in aviation passenger markets: The impact of demand heterogeneity on competition
557. G. SMRKOLJ, Dynamic Models of Research and Development
558. S. PEER, The economics of trip scheduling, travel time variability and traffic information
559. V. SPINU, Nonadditive Beliefs: From Measurement to Extensions
560. S.P. KASTORYANO, Essays in Applied Dynamic Microeconometrics
561. M. VAN DUIJN, Location, choice, cultural heritage and house prices
562. T. SALIMANS, Essays in Likelihood-Based Computational Econometrics
563. P. SUN, Tail Risk of Equidity Returns
564. C.G.J. KARSTEN, The Law and Finance of M&A Contracts
565. C. OZGEN, Impacts of Immigration and Cultural Diversity on Innovation and Economic Growth
566. R.S. SCHOLTE, The interplay between early-life conditions, major events and health later in life
567. B.N.KRAMER, Why dont they take a card? Essays on the demand for micro health insurance
568. M. KILI, Fundamental Insights in Power Futures Prices
569. A.G.B. DE VRIES, Venture Capital: Relations with the Economy and Intellectual Property
570. E.M.F. VAN DEN BROEK, Keeping up Appearances
571. K.T. MOORE, A Tale of Risk: Essays on Financial Extremes
572. F.T. ZOUTMAN, A Symphony of Redistributive Instruments

573. M.J. GERRITSE, Policy Competition and the Spatial Economy
574. A. OPSCHOOR, Understanding Financial Market Volatility
575. R.R. VAN LOON, Tourism and the Economic Valuation of Cultural Heritage
576. I.L. LYUBIMOV, Essays on Political Economy and Economic Development
577. A.A.F. GERRITSEN, Essays in Optimal Government Policy
578. M.L. SCHOLTUS, The Impact of High-Frequency Trading on Financial Markets
579. E. RAVIV, Forecasting Financial and Macroeconomic Variables: Shrinkage, Dimension reduction, and Aggregation
580. J. TICHEM, Altruism, Conformism, and Incentives in the Workplace
581. E.S. HENDRIKS, Essays in Law and Economics
582. X. SHEN, Essays on Empirical Asset Pricing
583. L.T. GATAREK, Econometric Contributions to Financial Trading, Hedging and Risk Measurement
584. X. LI, Temporary Price Deviation, Limited Attention and Information Acquisition in the Stock Market
585. Y. DAI, Efficiency in Corporate Takeovers
586. S.L. VAN DER STER, Approximate feasibility in real-time scheduling: Speeding up in order to meet deadlines
587. A. SELIM, An Examination of Uncertainty from a Psychological and Economic Viewpoint
588. B.Z. YUESHEN, Frictions in Modern Financial Markets and the Implications for Market Quality
589. D. VAN DOLDER, Game Shows, Gambles, and Economic Behavior
590. S.P. CEYHAN, Essays on Bayesian Analysis of Time Varying Economic Patterns
591. S. RENES, Never the Single Measure
592. D.L. IN T VELD, Complex Systems in Financial Economics: Applications to Interbank and Stock Markets
593. Y. YANG, Laboratory Tests of Theories of Strategic Interaction
594. M.P. WOJTOWICZ, Pricing Credits Derivatives and Credit Securitization
595. R.S. SAYAG, Communication and Learning in Decision Making
596. S.L. BLAUW, Well-to-do or doing well? Empirical studies of wellbeing and development
597. T.A. MAKAREWICZ, Learning to Forecast: Genetic Algorithms and Experiments