

---

SOCIAL PREFERENCES,  
BEHAVIORAL RULES,  
AND RESPONSE TIMES

---

Inauguraldissertation  
zur  
Erlangung des Doktorgrades  
der  
Wirtschafts- und Sozialwissenschaftlichen Fakultät  
der  
Universität zu Köln

**2019**

vorgelegt  
von

Alexander Ritschel, M.Sc.

aus  
Hennigsdorf

Referent: Prof. Dr. Carlos Alós-Ferrer  
Korreferent: Prof. Dr. Erik Hölzl  
Tag der Promotion: 02. April 2019

---

## TABLE OF CONTENTS

---

List of Tables	vi
List of Figures	viii
Introduction and Summary	viii
1 Performance Curiosity	5
1.1 Introduction . . . . .	5
1.2 Related Literature . . . . .	8
1.3 Experiment 1 . . . . .	12
1.3.1 Design and Procedures . . . . .	12
1.3.2 Distribution choices . . . . .	15
1.3.3 Expected Performance . . . . .	17
1.3.4 Regression Analysis . . . . .	19
1.4 Experiment 2 . . . . .	21
1.4.1 Design and Procedures . . . . .	22
1.4.2 Distribution Choices . . . . .	23
1.4.3 Expected Performance . . . . .	26
1.4.4 Regression Analysis . . . . .	28
1.5 Discussion . . . . .	30
Appendix 1.A: Robustness Analysis . . . . .	33
Appendix 1.A.1: Experiment 1 . . . . .	33
Appendix 1.A.2: Experiment 2 . . . . .	36

2	The Big Robber Game	39
2.1	Introduction . . . . .	39
2.2	Experimental Design . . . . .	48
2.2.1	Design and Procedures . . . . .	48
2.2.2	Power Analysis . . . . .	54
2.2.3	Data . . . . .	55
2.3	The Big Robber Question . . . . .	55
2.3.1	Robbers' Choices: To Rob or Not to Rob . . . . .	55
2.3.2	Decision Times . . . . .	59
2.3.3	Beliefs . . . . .	61
2.3.4	The Selfishness of Big Robbers . . . . .	62
2.4	Does the Big Robber Question Affect Behavior? . . . . .	67
2.4.1	The more selfish robbers . . . . .	68
2.4.2	The less selfish robbers . . . . .	71
2.4.3	The victims . . . . .	72
2.5	Donations . . . . .	73
2.6	Models of Social Preferences . . . . .	76
2.7	Discussion . . . . .	85
3	The Reinforcement Heuristic in Normal Form Games	91
3.1	Introduction . . . . .	91
3.2	Experimental Design and Procedures . . . . .	96
3.3	Results . . . . .	98
3.3.1	Classification of the Decision Situations . . . . .	98
3.3.2	Choice Data . . . . .	101
3.3.3	Cumulative Reinforcement . . . . .	103
3.3.4	Response Times . . . . .	104
3.4	Discussion . . . . .	109
	Appendix 3.A: Frequency and Response Times by Game . . . . .	111
	Appendix 3.B: Additional Regressions . . . . .	112

4	Multiple Behavioral Rules in Cournot Oligopolies	117
4.1	Introduction . . . . .	117
4.2	Predictions for Multiple Behavioral Rules . . . . .	121
4.3	Experiment 1 . . . . .	125
4.3.1	Experimental Design and Procedures . . . . .	125
4.3.2	Classification of Decisions . . . . .	128
4.3.3	Results . . . . .	129
4.4	Experiment 2 . . . . .	137
4.4.1	Cognitive Load . . . . .	137
4.4.2	Experimental Design and Procedures . . . . .	138
4.4.3	Cognitive Load and Behavior . . . . .	140
4.4.4	Replication and Robustness . . . . .	142
4.5	Reinforcement . . . . .	148
4.6	Inertia . . . . .	151
4.7	Conclusion . . . . .	157
5	Cognitive Load in Economic Decisions:	
	How To Tell Whether It Works	161
5.1	Introduction . . . . .	161
5.2	Working Memory and Cognitive Load . . . . .	166
5.3	Experiment 1: Cournot Oligopoly Under Light Cognitive Load	168
5.3.1	Experimental Design and Procedures . . . . .	168
5.3.2	Results . . . . .	170
5.4	Experiment 2: Cournot Oligopoly Under Heavy Cognitive Load	172
5.4.1	Experimental Design and Procedures . . . . .	172
5.4.2	Results . . . . .	173
5.5	Experiment 3: Voting . . . . .	175
5.5.1	Experimental Design and Procedures . . . . .	175
5.5.2	Results . . . . .	177
5.6	Experiment 4: Bayesian Updating . . . . .	179
5.6.1	Experimental Design and Procedures . . . . .	179
5.6.2	Results . . . . .	182
5.7	Discussion . . . . .	185



---

## List of Tables

---

1.1	Treatment Overview: Distribution and Information Revelation.	14
1.2	Experiment 1, Probit Regressions on Performance-based Choices. . . . .	20
1.3	Only Winner Treatment in Experiment 2. . . . .	23
1.4	Experiment 2, Probit Regressions on Performance-based Choices. . . . .	29
1.5	Experiment 1, Additional Probit Regression on Performance-based Choices. . . . .	37
1.6	Experiment 2, Additional Probit Regression: Epistemic Curiosity (EC) and Self-Motives (SM). . . . .	38
2.1	Big Robber Choices. . . . .	51
2.2	Out-of-sample analysis for the models of Charness and Rabin (2002), Bolton and Ockenfels (2000) and Alger and Weibull (2013). . . . .	84
3.1	The three games used in the experiment. . . . .	98
3.2	Frequency of the different experienced regret levels in lose situations. . . . .	99
3.3	Random-effects panel probit regressions for choices. . . . .	102
3.4	Panel regressions for log-transformed response times. . . . .	108
3.5	Random-effects panel probit regressions with additional controls.	112
3.6	Regression models while bootstrapping the standard errors. . .	113
3.7	Regressions models for cumulative proportional reinforcement.	113
3.8	Panel regression for log-transformed response times including an additional dummy. . . . .	114
3.9	Panel regression for ln(response time) with additional controls.	115

3.10	Panel regression for log-transformed response times while bootstrapping the standard errors. . . . .	116
4.1	Overview of prescribed actions. . . . .	129
4.2	Experiment 1, Overview of observations in conflict and alignment situations. . . . .	132
4.3	Experiment 1, Random effects panel regressions for $\ln(\text{ResponseTimes})$ . . . . .	135
4.4	Experiment 2, Overview of observations in conflict and alignment situations. . . . .	141
4.5	Experiment 2, Random effects panel regressions for $\ln(\text{ResponseTimes})$ . . . . .	147



---

## List of Figures

---

1.1	Experiment 1, Proportion of Performance-Based Splits. . . . .	16
1.2	Experiment 1, Proportion of Performance-Based Splits and Expected Performance. . . . .	18
1.3	Experiment 2, Proportion of Performance-Based Splits. . . . .	24
1.4	Experiment 2, Proportion of Performance-Based Splits and Expected Performance. . . . .	27
2.1	Overview of the Experimental Design. . . . .	50
2.2	Absolute Frequency of Robbing Choices. . . . .	56
2.3	Robbing Choices by Treatment. . . . .	57
2.4	Decision Times for the Big Robber Question. . . . .	60
2.5	Robbers' Decisions in the Intermediate Games. . . . .	63
2.6	Behavior of Robbers in the Dictator Game. . . . .	64
2.7	Behavior of Robbers-50% in the Dictator Game. . . . .	69
2.8	Behavior of Robbers-50% in the Ultimatum Game. . . . .	70
2.9	Average donations in percentage and expected donations in Euros. . . . .	76
3.1	Average individual choice frequencies. . . . .	100
3.2	Average individual response times. . . . .	105
3.3	Average individual choice frequencies and response times by games. . . . .	111
4.1	Classification of actions. . . . .	130
4.2	Experiment 1, Overview of observations and their classifica- tion according to different behavioral rules. . . . .	132
4.3	Experiment 1, Average response times of correct responses and errors. . . . .	133

4.4	Experiment 2, Overview of observations and their classification according to different behavioral rules. . . . .	140
4.5	Experiment 2, Relative frequency of errors and correct responses for NoLoad and Load treatments. . . . .	143
4.6	Experiment 2, Average response times of correct responses and errors. . . . .	144
4.7	Experiment 2, Average response times for the NoLoad and Load treatments. . . . .	145
4.8	Average response times of reinforcement and imitating-others decisions. . . . .	149
4.9	Average response times of of stay and shift correct decisions and errors. . . . .	152
4.10	Average response times of correct decisions and errors excluding inertia decisions. . . . .	155
5.1	Experiment 1, Average response times and payoff. . . . .	172
5.2	Experiment 2, Average response times and payoff. . . . .	174
5.3	Experiment 3, Average response times of voting decisions. . . .	178
5.4	Experiment 3, Relative frequency of sincere votes and number of votes. . . . .	179
5.5	Experiment 4, Average response times of the second draw. . .	183
5.6	Experiment 4, Average error rates of the second draw. . . . .	184

---

## Introduction and Summary

---

This dissertation consists of five self-contained research papers that cover experimental studies on different economic topics. My research interests are mainly focused on two interrelated areas within economics. One area is concerned with inequality aversion and social preferences. The other area is more process-orientated and concerns behavioral rules and heuristics. Especially, the use of response times as process data is among my research interests. Chapters 1 and 2 cover topics from the former area, whereas the remaining three chapters contribute to the latter area. Chapter 1 concerns the role of performance curiosity which can trump inequality aversion. Chapter 2 presents a novel experiment analyzing prosocial behavior in a high-stakes setting. Chapter 3 shows that the reinforcement process can work as cognitive shortcut to apparent myopic best reply behavior. Chapter 4 demonstrates that behavior in Cournot oligopolies is codetermined by multiple behavioral rules. Chapter 5 establishes a consistent reduction in response times for complex decisions under cognitive load. In the remainder of this section I present a brief introduction for each chapter and summarize the main findings.

Chapter 1 is the result of joint work with Carlos Alós-Ferrer (University of Zurich) and Jaume García-Segarra (University of Cologne) and has been published under the title “Performance Curiosity” in the *Journal of Economic Psychology*. We show that performance curiosity – the desire to know one’s own (relative) performance – can trump inequality aversion. In two experiments (combined  $N = 450$ ), participants chose between an equal allocation and a performance-based one after generating surplus in a real-effort task. In the experimental treatment, choosing an equal allocation came at the cost of not knowing the own performance, which led to a substantial increase of performance-based choices in comparison with the control treatment. The effect seems especially pronounced for women, but the gender effect is due to a difference in expectations regarding performance. Interestingly, the manipulation equalized the proportion of equal allocation choices between males

and females compensating for their difference in expectations. Work on this paper was shared among the authors as follows: Carlos Alós-Ferrer 33%, Jaume García-Segarra 33%, Alexander Ritschel 33%.

Chapter 2, entitled “The Big Robber Game,” is the result of joint work with Carlos Alós-Ferrer (University of Zurich) and Jaume García-Segarra (University of Cologne). We present a novel design measuring a correlate of social preferences in a high-stakes setting. In the Big Robber Game, a “robber” can obtain large personal gains by appropriating the gains of a large group of “victims” as seen in recent corporate scandals. We observed that more than half of all robbers take as much as possible. At the same time, participants displayed standard, prosocial behavior in the Dictator, Ultimatum, and Trust games. That is, prosocial behavior in the small is compatible with highly selfish actions in the large, and the essence of corporate scandals can be reproduced in the laboratory even with a standard student sample. We show that this apparent contradiction is actually consistent with received social-preference models. In agreement with this view, in the experiment more selfish robbers also behaved more selfishly in other games and in a donation question. We conclude that social preferences are compatible with rampant selfishness in high-impact decisions affecting a large group. Work on this paper was shared among the authors as follows: Carlos Alós-Ferrer 33%, Jaume García-Segarra 33%, Alexander Ritschel 33%.

Chapter 3 is the result of joint work with Carlos Alós-Ferrer (University of Zurich) and has been published under the title “The Reinforcement Heuristic in Normal Form Games” in the *Journal of Economic Behavior and Organization*. We analyze simple reinforcement-based behavioral rules in  $3 \times 3$  games through choice data and response times. We argue that there is a large overlap between reinforcement-based heuristics (win-stay, lose-shift) and the more “rational” behavioral rule of myopic best reply. However, evidence from response times shows that choices in agreement with the common prescription of those rules are comparatively fast, and choices of the form “lose-shift” occur more frequently for larger differences with bygone payoffs. Both observations speak in favor of reinforcement processes as a cognitive shortcut for apparent myopic best reply, and advise caution when interpreting behav-

ioral results in favor of optimizing behavior. Work on this paper was shared among the authors as follows: Carlos Alós-Ferrer 50%, Alexander Ritschel 50%.

Chapter 4, entitled “Multiple Behavioral Rules in Cournot Oligopolies,” is the result of joint work with Carlos Alós-Ferrer (University of Zurich). We show that in a Cournot oligopoly, multiple behavioral rules codetermine behavior. While imitation of successful behavior is an intuitive behavioral rule, myopic best reply is a deliberative rule. Previous literature has deduced the relevance of imitation from convergence, but we analyze individual decisions and derive testable predictions from a formal drift-diffusion model. The model predicts a non-trivial asymmetry which was readily found in two laboratory experiments. Our results suggest that a dual view of behavior incorporating rational behavior and evolutionary ideas on bounded rationality might provide a useful synthesis and improve our understanding of how economic decisions are made. Work on this paper was shared among authors as follows: Carlos Alós-Ferrer 50%, Alexander Ritschel 50%.

Chapter 5, entitled “Cognitive Load in Economic Decisions: How To Tell Whether It Works,” is the result of joint work with Anja Achtziger (Zeppelin University, Friedrichshafen) and Carlos Alós-Ferrer (University of Zurich). Cognitive load manipulations are routinely used in psychology to causally manipulate the reliance on deliberative or intuitive decision modes or processes, hence uncovering performance tradeoffs (e.g., between speed and accuracy) and default behavioral tendencies and motivations. The success of cognitive load manipulations in economic tasks, which are typically far more complex and time-consuming, is comparatively moderate, with many studies finding null effects. Those, however, might be due to either a failure of cognitive load to influence behavior or performance, or to a failure to induce relevant levels of cognitive load. A fundamental difficulty is that there exist no simple manipulation check to disentangle both possibilities. We argue that response times provide exactly such a manipulation check, but that, contrary to intuition and to received results for tasks from psychology, a successful cognitive load manipulation in typical economic tasks will *decrease* response times as decision makers reallocate resources to faster, less deliber-

ative processes. We test this hypothesis in four separate experiments using different, purely economic tasks (belief updating, voting, and competition in a dynamic Cournot oligopoly) and different cognitive load manipulations. Our evidence shows that response times are a simple test to check whether cognitive load was successfully induced in economic experiments, and should be systematically reported in such studies to allow for a proper interpretation of the results. Work on this paper was shared among the authors as follows: Anja Achtziger 33%, Carlos Alós-Ferrer 33%, Alexander Ritschel 33%.

---

## CHAPTER 1

### Performance Curiosity

---

#### 1.1 Introduction

When different people exert different levels of effort, an egalitarian distribution of jointly-generated proceeds is not necessarily “fair.” In this sense, egalitarianism might not always be socially desirable. If agents anticipate the results of their effort, social conventions dictating egalitarian allocations can greatly diminish incentives to exert effort to achieve a high performance.

This simple observation points to a conflict arising from different human motivations. On the one hand, most human beings strive to perform well. Indeed, psychological research has identified *achievement motivation* as a basic (intrinsic) motive (e.g., McClelland et al., 1989; Brunstein and Heckhausen, 2008). This goes hand-in-glove with the view that individual effort should be rewarded. On the other hand, research in behavioral economics has identified *inequality aversion*, which implies a preference in favor of egalitarian outcomes even if they result in a reduction of the own payoffs, as an important factor in decisions concerning (re)distribution (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002). When confronted with the decision to allocate the proceeds from individual effort, these two motivations result in opposite tendencies, which are at the heart of discussions on many socioeconomic issues, ranging from performance pay within firms to redistribution of income through tax systems.

In this work, we contribute to the investigation into the motivations underlying preferences among distributional allocations. We aim to show that preferences for egalitarian redistribution, as opposed to rewarding individual performance, hang in a fragile balance and can be significantly reduced with subtle interventions. In particular, we focus on a manipulation derived from *performance curiosity*, defined as the desire to know the own performance

(especially in relative terms). The rationale is that the motivation to achieve and perform well is linked to self-reputation and self-image concerns, which are themselves closely linked to the notion of *identity* (Mazar et al., 2008; Bénabou and Tirole, 2011). Obviously, information on the own performance has a crucial impact on the self-image. We hence postulate that human beings might be willing to give up the benefits of their own effort to appease their inequality aversion, but this willingness only lasts as long as they at least receive the information on how well they have performed. Specifically, if social conventions, on top of imposing an egalitarian redistribution, go to the extreme of eliminating (relative) performance feedback, we postulate that preferences for egalitarian distributions will be greatly reduced.

In this work, we report on two experiments with an innovative design which pits inequality aversion against curiosity regarding the own performance in a real-effort task. The task is used to generate surplus, and is followed by a distributional allocation decision within a group. The design creates a tradeoff between an egalitarian allocation and receiving information on performance in the real-effort task. We show that this subtle manipulation can shift preferences away from egalitarian allocations.

In our first experiment ( $N = 180$ ) participants generated income by working on a real-effort task and subsequently decided on the allocation of the joint proceedings within a small group. In the control treatment, participants chose between an egalitarian allocation and a performance-based one, leading to the standard observation of inequality aversion. In the experimental (“No Info”) treatment, the only change was that opting for the egalitarian distribution came at the cost of not knowing the actual performance and ranking in the real-effort task. The design highlights *relative performance*, because the information, if revealed, includes the performance and rank of all group members, enabling social comparisons. This change had a large effect, with participants choosing the performance-based distribution rather than the egalitarian one. In the second experiment ( $N = 270$ ) we replicated the results of the first experiment and also added a third treatment and additional questionnaires to test for further explanations of the basic effect.



Our tasks were explicitly incentivized. Hence, given the own expected performance, participants could easily compute which of the two allocations would maximize their (expected) monetary rewards. The allocations were designed in such a way that only those expecting to be strictly above average would be better off under the performance-based allocation, with all others being better off under the egalitarian allocation. Hence, we elicited expected performance in the real-effort task and controlled for it by testing the basic hypotheses also within different groups (in terms of expectations) and by including the corresponding variable as a control in our regressions. As was to be expected, a higher expectation led to a higher percentage of performance-based choices, but the basic effect remains clearly significant.

In view of evidence on gender differences and competition (e.g., Gneezy et al., 2003; Niederle and Vesterlund, 2007), we were also interested in gender effects. For this reason, we were led to a design with perfectly balanced participation across gender and a sample size large enough to examine gender differences. Specific hypotheses regarding those can also be derived from research on inequality aversion. Some studies in this field have found gender differences, as reflected, e.g., in the proportion of egalitarian allocations in Dictator and Ultimatum games (see, e.g., Fehr et al., 2006). However, the evidence is mixed. Croson and Gneezy (2009) argue that there is no difference in social preferences, and that the observed effects arise due to gender differences in the sensitivity to cues in experimental contexts. The latter hypothesis is motivated by research in psychology establishing that women are more sensitive to social cues and feedback than men (Gilligan, 1982; Roberts and Nolen-Hoeksema, 1989). Since our manipulation involves the provision of feedback, it is natural to expect that the basic effect should be especially pronounced for women, compared to men. Indeed, the treatment effect in our experiments was clearly stronger among women, with behavior in the control treatment exhibiting clear gender differences which vanished in the No Info treatment. Controlling for expected performance allows to uncover the roots of this gender effect: women had a lower expected performance than men, leading to a lower percentage of performance-based choices in the

control treatment (due to extrinsic, monetary rewards) which was overcome in the experimental treatment.

We conclude that performance curiosity, i.e. wanting to know one's own (relative) performance, can counteract inequality aversion as a motivation and tilt decisions away from egalitarian distributions and toward performance-based ones. Additionally, creating a tradeoff between information and egalitarian allocations substantially reduces gender differences in behavior, with women becoming more willing to accept performance-based allocations. This is a potentially important insight for incentive design and the reduction of gender differences.

The Chapter is organized as follows. Section 1.2 discusses some related literature strands and helps place our experiments in a broader context. Section 1.3 describes the design of Experiment 1 and presents the analysis of the data. Section 1.4 does the same for Experiment 2. Section 1.5 discusses the results. Appendix A analyzes some closely related constructs which were measured through questionnaires in our experiments and discards them as possible alternative explanations. Appendix B contains the experimental instructions.

## 1.2 Related Literature

Our results are in agreement with the literature showing that preferences for egalitarian allocations, or, more generally, reducing inequality, seem not to be stable. For instance, dictator-game giving is reduced if the perception of anonymity and social distance is increased (Hoffman et al., 1994; Charness and Gneezy, 2008; Franzen and Pointner, 2012). Also, subtle content-free psychological manipulations cause radical shifts of behavior (Achtziger et al., 2015b, 2016, 2018). Further, part of the motivation for reducing inequality might be related to the desire to show others that one is not selfish (Bardsley, 2008; Cappelen et al., 2013). Further, unequal payoffs are often deemed acceptable, provided some justification is given, e.g. merit, entitlement, or needs (Forsythe et al., 1994; Hoffman et al., 1994; Konow, 2003; Fershtman et al., 2012). This already points out that the balance between rewarding

performance and distributing resources equitably can be shifted if the link to performance is emphasized (e.g., through merit or entitlement).

The studies reported here are also related to the extensive social-psychological literature on social comparisons, which goes back to Festinger (1954). This literature stresses the human desire to evaluate oneself. The very first hypothesis is that people have a basic need to evaluate themselves and their performance. The second hypothesis is that social comparisons are used to fulfill this desire since objective evaluations are often unavailable or difficult to obtain (see also Moore and Klein, 2008). For instance, Trope (1980) asked experimental participants to rate computer games requiring eye-hand coordination, but which differed in how accurately skill was linked to outcomes. Participants preferred games with accurate feedback, bearing the potential negative consequences for their self-esteem, over games which were not diagnostic for the participant's skill.

Recent studies in economics have also clearly shown that information on the own (relative) performance has a strong impact on motivation. Several studies have shown that the provision of feedback has a positive effect on performance in the laboratory and in the field (Hannan et al., 2008; Azmat and Iriberri, 2010; Blanes i Vidal and Nossol, 2011; Kuhnen and Tymula, 2012). Azmat and Iriberri (2016) have recently shown that information on relative performance increases effort and affects satisfaction when payoff depends on performance (piece-rate wage), but has no effect in either when pay is independent of performance (flat rate). Schoenberg and Haruvy (2012) gave participants in an experimental asset market information on either the highest- or lowest-earning trader in the market, period by period. Market prices were significantly higher when the information of the highest-performing trader was presented, and the satisfaction of the traders was higher if they had a good relative position relative to the presented information.

In agreement with developments in research on social comparisons as quoted above, research in economics has also argued that people may care about relative performance, as captured by rankings, even when rankings have no financial consequences, because of its impact on self-image (Bén-

abou and Tirole, 2006, 2011; Köszegi, 2006).<sup>1</sup> For instance, Berger et al. (2013) show that ratings on performance are more effective for increasing productivity if there is a high level of differentiation in the feedback.<sup>2</sup> However, this strand of the literature has typically concentrated on the link between information on performance and performance itself, while our aim is to establish the link between information on performance and *preferences* on distributional allocations.

Of course, the concept of curiosity (in a general sense) has received a great deal of attention in psychology. An influential curiosity-related construct is *lay epistemic theory* (Kruglanski, 1989), which describes the factors involved in the general knowledge formation process. It describes four epistemic motivations related to the need for specific or nonspecific closure, and the need to avoid them. Those determine the length of hypothesis generation and testing sequence to evaluate evidence and form a belief. Although similarities between social comparison processes (Festinger, 1954) and lay epistemic theory seem apparent, the latter does not suggest a general drive towards social comparison and weighs informational sources by their relevance, not by their similarity towards oneself.

Another important construct is *epistemic curiosity* (Loewenstein, 1994; Litman et al., 2005), defined as “the desire for knowledge that motivates individuals to learn new ideas, eliminate information-gaps, and solve intellectual problems” (Litman, 2008, p. 1586). Epistemic curiosity is subdivided in two broad categories. The first is *interest-type curiosity*, which involves the anticipated pleasure of new discoveries. The second is *deprivation-type*

---

<sup>1</sup>Work in evolutionary game theory has shown that relative-payoff concerns, which go hand-in-glove with imitating behavior, can tilt long-run predictions away from Nash equilibria (Vega-Redondo, 1997; Alós-Ferrer and Ania, 2005) and towards more competitive ones in experimental markets. This prediction has been confirmed in several Cournot-oligopoly experiments, e.g. Huck et al. (1999) and Offerman et al. (2002). In alignment with those, Fatas et al. (2015) found that the availability of a relative-performance measure led to more competitive outcomes.

<sup>2</sup>Payment schemes based on *relative* performance can sometimes decrease productivity. In a field experiment, Bandiera et al. (2005) found that productivity under piece rates was substantially higher than under relative incentives, because workers took into account the negative externality their effort imposed on others. However, an altruistic motivation was ruled out because this effect was only present when monitoring was possible.

*curiosity*, which is the need to reduce uncertainty and eliminate undesirable knowledge gaps. To check for the possible relation between these concepts and performance curiosity, we included the epistemic curiosity questionnaire by Litman and Mussel (2013) in Experiment 2 (see Appendix A).

Our explanation for the relevance of performance curiosity hinges on the value of feedback for the own self-image. Psychological research (Sedikides and Strube, 1995, 1997) has identified four cardinal *self-motives* as relevant to the development, maintenance, and modification of self-views. These are *self-enhancement* (the desire to see oneself positively), *self-verification* (the desire to confirm a preexisting view of oneself), *self-assessment* (the desire to know the truth about oneself), and *self-improvement* (the desire to improve oneself). To check the possible relation between these concepts and performance curiosity, we included the questionnaire by Gregg et al. (2011) in Experiment 2.

A study by Van de Ven et al. (2005) investigated how curiosity regarding an uncertain outcome contributes to the *endowment effect*, which causes a disparity between selling and buying prices. In an experiment involving lotteries, sellers were more curious about the outcome than buyers and the curiosity positively correlated with the minimum selling price. A second experiment involved tokens initially held by sellers whose actual value (exchange rate) was uncertain. Post-trade information about that value differed across conditions. Minimum selling prices were significantly higher for sellers who would not learn about the exchange rate of a token they sold, compared to sellers who would learn about the exchange rate independently of whether they sold it or not. This showed a willingness to pay for the satisfaction of curiosity, which added to the endowment effect.

Last, our experiments are related to the literature on gender effects in competition. In our case, the decision on how to redistribute the joint earnings of the group came after performing the task. Therefore, we do not deal with performance differences or with preferences for competition, but rather with preferential choices in distributions. Still, our work is obviously related to those strands of the literature at a conceptual level. There are well-established gender differences in the attitudes toward competition. Gneezy

et al. (2003) and Niederle and Vesterlund (2007) showed that women tend to shy away from competition in real-effort tasks, and Gneezy et al. (2009) showed that matrilineal and patriarchal societies differ with regard to competitive behavior.<sup>3</sup> A recent field experiment by Azmat et al. (2016) has shown that women perform worse than men when stakes are high, while the opposite is true when the stakes are low.

### 1.3 Experiment 1

#### 1.3.1 Design and Procedures

The experiment was conducted at the Cologne Laboratory for Economic Research and programmed in z-Tree (Fischbacher, 2007).<sup>4</sup> Participants were recruited from the University of Cologne’s pool, excluding psychology students, using ORSEE (Greiner, 2015). The experiment was a between-subject design with 180 participants, 45 females and 45 males in each of the two treatments. Average payoff was EUR 13.25 (USD 17.80) for a session that lasted around 60 minutes. Subjects were assigned to groups of five players, but they did not know the identities of the other group members.<sup>5</sup> Further, both the task and the subsequent allocation decision were made individually and there was no interaction of any kind except in the determination of payoffs at the end of the experiment. Initially, subjects worked individually in a real-effort task requiring to add up sets of five two-digit numbers (e.g.  $30+87+19+16+38=\square$ ) for a predetermined time. Each correct answer generated 10 points, with an exchange rate of 5 Euro cents per point, but no feedback on the correctness of the responses was provided until the end of the

---

<sup>3</sup>Our design bears procedural similarities to Task 4 in Niederle and Vesterlund (2007) in the sense that participants perform the task without knowing the incentive rule and choose the compensation scheme later. Niederle and Vesterlund (2007) offered a choice between piece-rates or competitive tournaments, while our design involves neither, but rather different distributions of the joint income generated by the group.

<sup>4</sup>The data and the codebooks for both experiments can be downloaded from <http://osf.io/5hkj6>.

<sup>5</sup>In particular, they had no information on the gender composition of their group. This is particularly important, because gender composition might affect decision making and performance within a group (Gneezy et al., 2003; Apesteguía et al., 2012).

experiment. Subjects were not allowed to use calculators or other electronic devices, but they were allowed to use scratch paper. This real-effort task was proposed in Niederle and Vesterlund (2007) and has been often used in experiments (e.g., Azmat and Iriberry, 2016).<sup>6</sup>

The real-effort task lasted for eight minutes, which is relatively long. This duration was selected to put enough weight on the effort side rather than on the skills of the subjects. Also, a shorter task duration would have increased the probability of ties inside each group. After performing the task, the experiment moved to the decision-making part without any feedback being provided.

Participants were asked to make a decision on how they would prefer to distribute the joint amount of points generated by the group. They were informed that one of the group members would be selected at random and his or her decision would be implemented as stated.<sup>7</sup> Two distributions were available. The first was an *equal split*, where all group members earned the same amount of points. The second was a *performance-based split*, leading to a different amount of points depending on the relative rank within the group (see Table 1.1). Specifically, the performance-based split made the three worst-placed participants earn less than under the equal split, in order to create a clear tradeoff between monetary rewards and feedback provision for the median participant.<sup>8</sup>

---

<sup>6</sup>Women often think that their performance in math will be worse than that of men, which sometimes leads to actually worse performance due to “stereotype threat” (Spencer et al., 1999; Gneezy et al., 2003). However, a meta-analysis of 100 studies on gender differences in math performance showed that there is no gender difference in arithmetic or algebraic performance (Hyde et al., 1990).

<sup>7</sup>Subjects were informed that they were part of a group at the beginning of the experiment, and they were aware that after finishing the real-effort task, they would decide among two distribution rules for sharing the joint surplus, with one decision implemented at random. However, they were not informed about the specifics of the possible distribution rules until the decision had to be made.

<sup>8</sup>The performance-based distribution allocates 15%, 10%, and 5% of the joint proceeds to the 3rd, 4th, and 5th group member, respectively. Hence choosing the performance-based distribution entails a clear monetary cost for participants who expect to be average or worse. The participants ranked first and second are rewarded with 40% and 30% of the joint proceeds, that is, by a factor of 2 and 1.5, respectively, compared to the equal-split distribution. This makes clear that the distribution indeed rewards performance.

Table 1.1: Treatment Overview: Distribution and Information Revelation.

Control Treatment				No Info Treatment			
Equal Split		Performance-Based		Equal Split		Performance-Based	
Rank	Share (%)	Rank	Share (%)	Rank	Share (%)	Rank	Share (%)
1	20	1	40	-	20	1	40
2	20	2	30	-	20	2	30
3	20	3	15	-	20	3	15
4	20	4	10	-	20	4	10
5	20	5	5	-	20	5	5

The key manipulation between treatments was the amount of information that was revealed with each possible distribution. In the *Control Treatment* participants chose one of the two distributions, but later the ranking and the number of correctly solved calculations was revealed independently of which distribution was finally implemented. However, in the *No Info Treatment* participants were told that the ranking and the number of correctly solved calculations would only be revealed if the performance-based split was actually the one chosen and implemented. The alternatives are detailed in Table 1.1.

In the absence of self-image concerns, the difference across treatments should not affect behavior. In the presence of performance curiosity, however, we should observe a higher proportion of performance-based choices in the No Info treatment than in the control treatment.

To control for various other possible explanations for differences in behavior across treatments we also elicited *self-efficacy* before and after the real-effort task. That is, participants were asked how many additions they thought they could solve correctly in eight minutes. Further, they were asked to report their expected relative performance in a 7 seven-point scale ranging from “Far below average” to “Far above average.” To control for individual differences in the regression analysis in attitudes and motivation, we also



elicited risk aversion attitudes using the lottery questionnaire of Holt and Laury (2002) and the reduced Achievement Motive Scale of Lang and Fries (2006), using the German version of Dahme et al. (1993) (the analysis of those is relegated to Appendix A). At the end of the experiment participants received feedback (if any) on a screen where their rank and the number of correctly solved additions were highlighted.

### 1.3.2 Distribution choices

The percentage of performance-based choices in both distributions is illustrated in Figure 1.1(a). In the control treatment, there is no tradeoff since the same information is revealed independently of which distribution is implemented. Hence, this treatment is a pure test of inequality aversion. In agreement with standard results in the social preferences literature, only 35 of the 90 participants (38.89%) in this treatment chose the performance-based distribution. In contrast, this proportion rose to 70.00% in the No Info treatment (63 out of 90). The difference is highly significant according to a test of proportions ( $z = -4.191$ ,  $p < 0.0001$ ).<sup>9</sup> Thus, subjects chose the performance-based distribution in the No Info treatment more often than in the control treatment. Since the only difference between treatments is the information that is revealed when the equal split is implemented, the difference in results is due to the tradeoff regarding information on performance. We hence conclude that performance curiosity overcomes inequality aversion in our data.

Since we expected a gender effect, we analyzed the behavior of the women and men separately. Figure 1.1(b) shows the proportions of participants choosing the performance-based split across treatments by gender. The overall effect observed in Figure 1.1(a) seems to be driven by female behavior. While only 10 women out of 45 (22.22%) chose the performance-based split in the control treatment, the proportion rose to 68.89% (31 out of 45) in the No Info treatment (test of proportions,  $z = -4.445$ ,  $p < 0.0001$ ). The difference goes in the same direction for men, with 55.56% of performance-based

---

<sup>9</sup>A Fisher's exact test delivers the same conclusion.

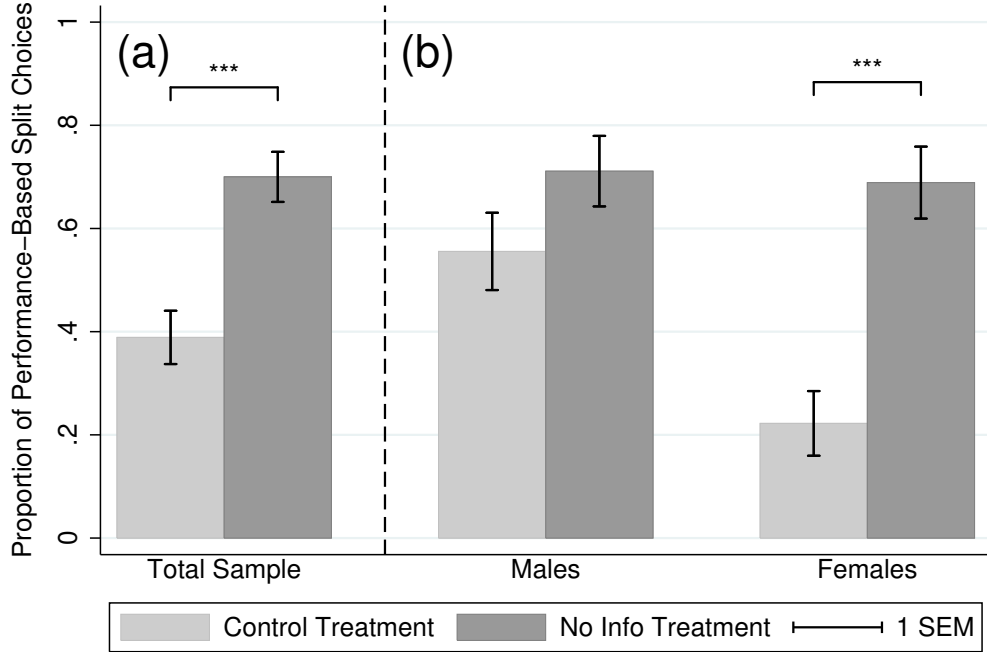


Figure 1.1: Experiment 1, Proportion of Performance-Based Splits.

*Notes.* (a) Proportion of participants choosing the performance-based split across treatments. (b) Proportion of participants choosing the performance-based split by treatment and gender. Bars represent 1 Standard Error of the Mean. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , test of proportions.

choices in the control treatment vs. 71.11% in the No Info treatment, but the difference misses significance ( $z = -1.531$ ,  $p = 0.1257$ ). As a result, while females chose the equal split more often than males in control treatment (test of proportions,  $z = 3.2434$ ,  $p = 0.0012$ ), the difference disappeared in the No Info treatment ( $z = 0.2300$ ,  $p = 0.8181$ ). We delay the analysis of the interaction between treatment and gender to the regression analysis below (Table 1.2, Model 2; Section 1.3.4).

That is, while men choose the performance-based distribution more often, as soon as a tradeoff involving performance curiosity is introduced, women behave exactly as men.

### 1.3.3 Expected Performance

There is one obvious, rational reason for some participants to choose the performance-based split. If the participant believes that he or she is likely to be strictly above the group's average, the individual payoff under the performance-based split is higher than under the egalitarian allocation. Reciprocally, if the participant believes to be at or below average, the individual payoff is higher under the egalitarian allocation (recall Table 1.1). Hence, we need to control for expected performance.

The median number of correctly solved additions was 15 (ranging from 5 to 43), and the median error rate was 0.19 (ranging from 0.03 to 0.69).<sup>10</sup> Crucially, participants were asked their expected relative performance (after completing the real-effort task) in a seven-point scale. The expected relative performance did not differ across treatments according to a Mann-Whitney-Wilcoxon (MWW) test (control treatment, mean 4.089; No Info treatment, 4.300;  $z = -1.055$ ,  $p = 0.2913$ ). We classified participants into three groups according to their beliefs, i.e. below, exactly, and above average.<sup>11</sup>

Figure 1.2(a) disentangles the proportions of performance-based choices within the three expectation groups. The performance-based split was chosen less often in the control treatment than in the No Info treatment across all three groups. For participants who expected to be above average, the performance-based split was chosen 24 out of 36 times (66.67%) in the control treatment, compared to 39 out of 43 (90.70%) in the No Info treatment (test of proportions,  $z = -2.647$ ,  $p = 0.0081$ ). For participants who expected to perform exactly on average, the performance-based split was chosen 6 out of 23 times (26.09%) in the control treatment, compared to 14 out of 21 (66.67%) in the No Info treatment ( $z = -2.700$ ,  $p = 0.0069$ ). For partic-

<sup>10</sup>After completing the real-effort task but before making the distribution choice, participants were asked about their expected performance, that is, how many additions they thought they had solved correctly. The average expectation was 15.2 additions in the control treatment and 17.3 in the No Info treatment. The difference was not significant according to a Mann-Whitney-Wilcoxon test ( $z = -1.457$ ,  $p = 0.1452$ ).

<sup>11</sup>For simplicity, the question on expected relative performance did not refer to a group. Since allocation to groups was random, the participants' expectations translate into expected ranking within the group.

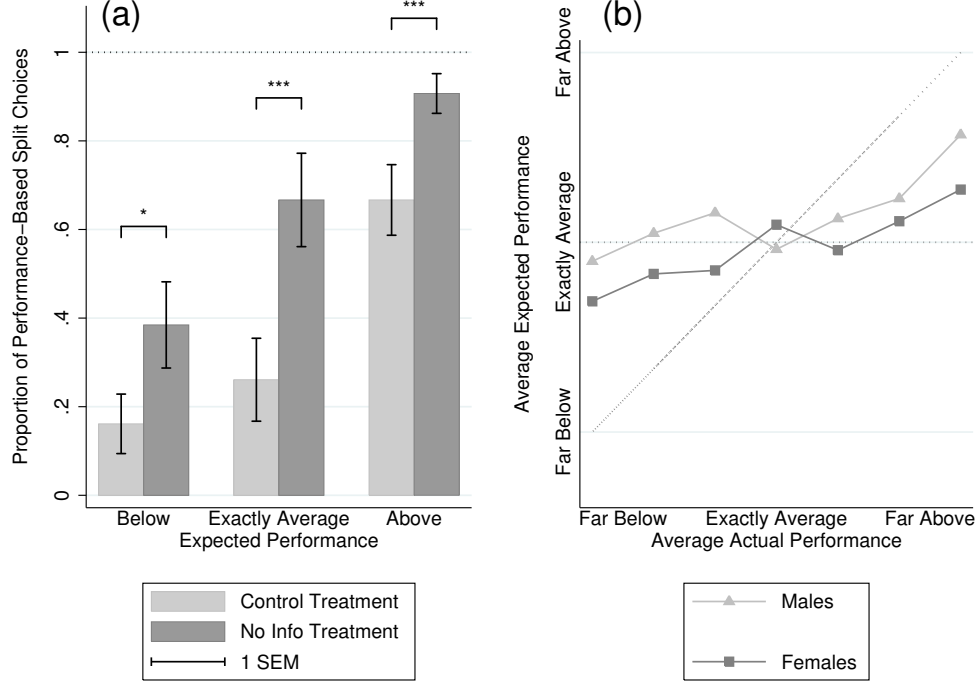


Figure 1.2: Experiment 1, Proportion of Performance-Based Splits and Expected Performance.

*Notes.* (a) Proportion of participants choosing the performance-based split by treatment and expected ranking. (b) Expected versus actual performance by gender. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , test of proportions.

ipants who expected to be below average, the performance-based split was also chosen less often in the control treatment (16.13%; 5 of 31) than in the No Info treatment (38.46%; 10 of 26), but the difference narrowly missed significance ( $z = -1.907$ ,  $p = 0.0565$ ). Additionally, comparing the proportions of choices across groups, one sees that, as was to be expected, a higher expected performance resulted in a higher proportion of performance-based choices (see regression analysis below).<sup>12</sup>

We did not find any difference in beliefs across treatments. However, women believed that they had performed worse (average 3.800 in the 7-point scale), compared to men's expectations (average 4.589 in the 7-point scale).

<sup>12</sup>A separate probit regression showed no significant interaction between treatment and expected performance.

The difference is highly significant (MWW test,  $z = 3.924$ ,  $p = 0.0001$ ). Figure 1.2(b) plots the expected performance (average report in the 7-point Likert scale) for males and females, conditional on the actual population septiles in the experiment, as computed from the number of correctly solved additions. This figure makes apparent that females tend to underestimate their performance compared to males, as previously pointed out, e.g., by Gneezy et al. (2003).<sup>13</sup>

#### 1.3.4 Regression Analysis

Recapitulating, in the previous subsections we have seen that performance curiosity shifts choices toward the performance-based split and away from the egalitarian allocation, and that this effect is far stronger for women. However, we have also seen that choices are influenced by the expected ranking, and that there are differences in beliefs across genders. In view of these results, our next step is to conduct regression analyses controlling for expected ranking. This will also allow us to investigate the determinants of the gender effect.

We hence turn to probit regressions on the choice of the performance-based split accounting for treatment, gender, and expected ranking. We report on three different model specifications in Table 1.2.

Model 1 investigates the basic treatment effect without further controls. The treatment dummy is highly significant and positive, reproducing in the regression the basic message of the test of proportions in Section 1.3.2, as illustrated in Figure 1.1(a). Model 2 adds gender and its interaction with the treatment dummy. The reference group is hence made of males in the control treatment. The treatment dummy becomes non-significant, reflecting that the treatment effect misses significance for males as found in Section 1.3.2 (recall Figure 1.1(b)). In contrast, the interaction of the treatment dummy and the female dummy is significant, showing that women react to

---

<sup>13</sup>Although the task has been determined to be gender-neutral (see Niederle and Vesterlund, 2007), in our sample women actually performed worse than men, with an average 18.4 correct additions for males and 15.0 for females (MWW test,  $z = 3.349$ ,  $p = 0.0008$ ). Causality, of course, might go in the other direction because of stereotype threat (recall Footnote 6).

Table 1.2: Experiment 1, Probit Regressions on Performance-based Choices.

	Model 1	Model 2	Model 3
No Info Treatment	0.807*** (0.1930)	0.417 (0.2725)	0.679** (0.3031)
Female		-0.904*** (0.2801)	-0.395 (0.3137)
No Info X Female		0.840** (0.3946)	0.440 (0.4350)
Expected Performance in Septiles			0.529*** (0.0930)
Constant	-0.282** (0.1340)	0.140 (0.1875)	-0.224 (0.2126)
LogLikelihood	-115.120	-109.701	-90.676
Wald Test	17.458***	26.590***	50.785***
Linear combination tests:			
No Info + No Info X Female		1.257*** (0.2854)	1.119*** (0.3102)
Female + No Info X Female		-0.064 (0.278)	0.045 (0.303)

*Notes.* Standard errors in brackets, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

the manipulation more strongly than men. A post-hoc linear combination test (No Info + No Info  $\times$  Female, bottom of Table 1.2) reveals that the treatment effect for women is highly significant ( $p < 0.0001$ ). The female dummy is negative and highly significant, showing that women choose the egalitarian allocation more often than men in the control treatment. However, this gender difference does not exist in the No Info treatment, as shown by the corresponding linear combination test (Female + No Info  $\times$  Female, bottom of Table 1.2;  $p = 0.8181$ ).

Model 3 is the key regression, which introduces the expected ranking of participants and sheds light into the determinants of the gender differences. First, the independent variable “Expected Performance in Septiles”, centered at 0, is highly significant and positive, confirming that the higher the expectations the more likely it is that the participant chooses the performance-based

split, as higher expectations make the latter split more attractive. Once we control for this natural effect, the treatment dummy becomes highly significant and positive, indicating that performance curiosity is a significant driver of behavior for males. The linear combination test confirms that the treatment effect is also highly significant for females. That is, after controlling for the expectations we see that both male and female participants are more likely to choose the performance-based split when inequality aversion conflicts with performance curiosity.

In contrast, in Model 3 the female dummy and its interaction with the treatment dummy are not significant anymore. A linear combination test also shows that there is no gender difference in the No Info treatment ( $p = 0.8822$ ). That is, the entire gender effect observed in previous regressions and tests originates exclusively in differences in expectations. In other words, since females have lower expectations, the performance-based split is less attractive, and hence chosen less often by females than by males. Once we control for expectations, we see that the effects of performance curiosity are present for both genders.

## 1.4 Experiment 2

The purpose of our second experiment was threefold. First, the experiment was conceived as a replication in order to establish the reliability of the effects. Accordingly, the first two treatments of Experiment 2 were identical to those of Experiment 1. Second, we added a third treatment to the design, the *Only Winner Treatment*, in order to test whether revealing only partial information, namely whether one is the winner or not, is enough to produce the same effect we found in Experiment 1. Human beings have intrinsic preferences for winning, perceived as a reward in itself. That is, they choose to exert effort to win (Fershtman et al., 2012) regardless of the payoff distribution resulting from their eventual victory. Hence, one could speculate that the effects we uncovered in Experiment 1 are actually derived from this motivation only, that is, that the desire to know whether one is the winner is the actual driver behind the effect. In the new treatment, the information

regarding whether one is the winner or not is revealed independently of the distribution choice. Hence, if the desire to know whether one is the winner were enough to satisfy the need for information about performance, in the new treatment the effect should disappear.

Third, we added additional controls (personality scales) to the design in order to test for the robustness of the effects with respect to individual differences. The new controls were of two kinds, related, first, to general curiosity, and, second, to self-motives. Specifically, we focused on epistemic curiosity (Loewenstein, 1994; Litman et al., 2005), using the epistemic curiosity questionnaire by Litman and Mussel (2013), and self-motives (Sedikides and Strube, 1995, 1997), using the questionnaire by Gregg et al. (2011). Both constructs are described in Section 1.2 and their analysis is relegated to Appendix 1.A.

#### 1.4.1 Design and Procedures

The experiment was a between-subject design with three treatments, again conducted at the Cologne Laboratory for Economic Research, and programmed in z-Tree. We recruited 270 participants through ORSEE from the University of Cologne’s pool, excluding psychology students. The sample was perfectly balanced by gender, leaving 45 females and 45 males for each of the three treatments. The average payoff was EUR 10.56 (USD 14.21) for a session that lasted around 55 minutes.

The experiment differed from Experiment 1 in two respects. First, since after Experiment 1 we concluded that neither risk attitudes nor achievement motivation were determinant for the treatment effect, we replaced those questionnaires by the ones measuring epistemic curiosity and self-motives as explained above. Since questionnaires were placed at the end of the experiment, the change could not affect behavior and the experiments remain fully comparable. Second, we added a third treatment, but the other two treatments were identical to the control and No Info treatments of Experiment 1 (recall Table 1.1), hence the new experiment contains a pure replication.



Table 1.3: Only Winner Treatment in Experiment 2.

Only Winner Treatment			
Equal Split		Performance-Based	
Rank	Share (%)	Rank	Share (%)
1	20	1	40
-	20	2	30
-	20	3	15
-	20	4	10
-	20	5	5

The third treatment is the *Only Winner Treatment*. In this treatment, the egalitarian and performance-based allocations are identical to those of other treatments, but the associated information is different. If the performance-based split is chosen and implemented, as in the other two treatments all information regarding how many additions have been correctly solved by each member of the group and the resulting ranking is revealed. The difference is that, if the equal split is the one chosen and implemented, then the only information revealed is how many additions were correctly solved *by the winner*, plus a signal indicating whether the participant is the winner or not. See Table 1.3 for a summary of the treatment.

The rest of Experiment 2 was identical to Experiment 1 regarding the real-effort task, the procedure for implementing the chosen allocation, the order of tasks and decisions, self-efficacy elicitation, and the final possibility to privately receive feedback.

#### 1.4.2 Distribution Choices

Figure 1.3 shows the proportions of participants choosing the performance-based split across the three treatments, for the whole sample (a) and split by gender (b). We successfully replicated the main result of Experiment 1. There were significantly more performance-based choices in the No Info treatment (75.56% of 90) than in the control treatment (43.33% of 90) (test of proportions,  $z = -4.402$ ,  $p_{\text{adj}} < 0.0001$ , adjusted for multiple compar-

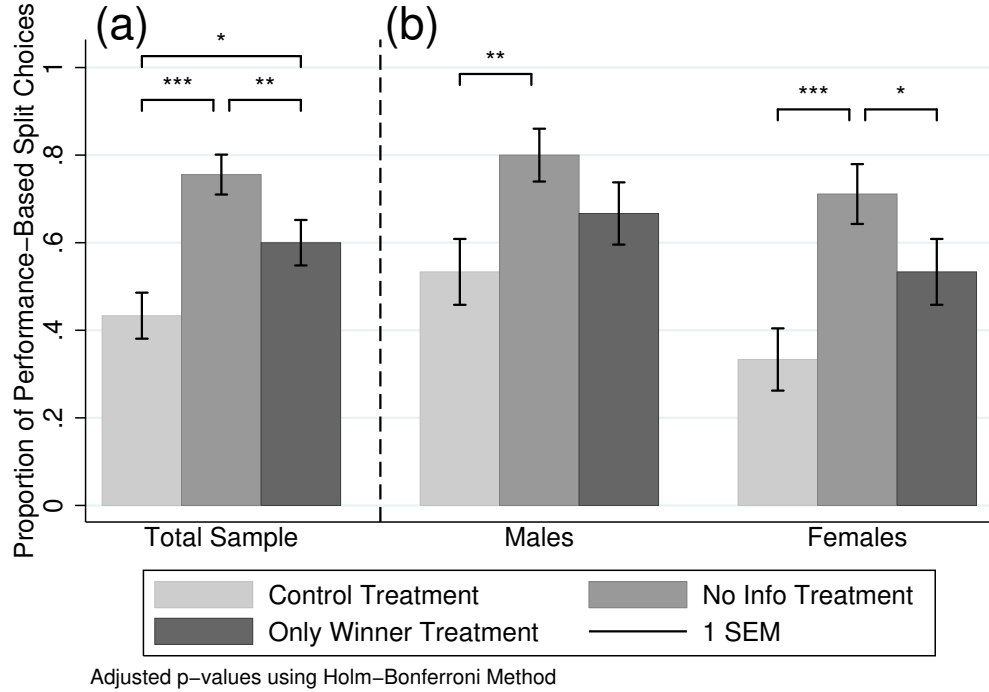


Figure 1.3: Experiment 2, Proportion of Performance-Based Splits.

*Notes.* (a) Proportion of participants choosing the performance-based split across treatments. (b) Proportion of participants choosing the performance-based split by treatment and gender. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , test of proportions.

isons following the Holm-Bonferroni method). Further, the proportion of performance-based choices in the Only Winner treatment (60.00% of 90) was significantly smaller than in the No Info treatment ( $z = 2.233$ ,  $p_{\text{adj}} = 0.0256$ ), and larger than in the control treatment ( $z = -2.237$ ,  $p_{\text{adj}} = 0.0505$ ), although the latter difference narrowly misses significance. That is, the new treatment lies clearly “in the middle” between the two previous ones.

Females chose the performance-based split less often than males in the control treatment (females, 33.33%; males, 53.33%;  $z = 1.915$ ,  $p = 0.0556$ ), although the difference narrowly misses significance, whereas there were no gender differences in the No Info treatment (females, 71.11%; males, 80.00%;  $z = 0.981$ ,  $p = 0.3265$ ) or in the Only Winner treatment (females, 53.33%; males, 66.67%;  $z = 1.291$ ,  $p = 0.1967$ ). That is, the apparently higher degree

of inequality aversion of women in the control treatment disappears in both of the treatments with an informational tradeoff.

As in Experiment 1, when looking at females only, we also observe significantly more performance-based choices in the No Info treatment (71.11%) than in the control treatment (33.33%) (test of proportions,  $z = -3.588$ ,  $p_{\text{adj}} = 0.0010$ ). Further, the effect is also significant for males (No Info treatment, 80.00%; control treatment, 53.33%;  $z = -2.683$ ,  $p_{\text{adj}} = 0.0219$ ). That is, although the effect for males missed significance in Experiment 1, it is significant in the replication, and we conclude that both men and women chose the performance-based split significantly more often when information was not revealed in the equal split. Looking at females only, the proportion of performance-based choices in the Only Winner treatment (53.33%) again lies between the control and the No Info treatments, although the effects miss significance (No Info vs. Only Winner,  $z = 1.739$ ,  $p_{\text{adj}} = 0.0820$ ; control vs. Only Winner,  $z = -1.915$ ,  $p_{\text{adj}} = 0.1111$ ). The proportion of performance-based choices for males in the Only Winner treatment (66.7%) was also between those of the other treatments, but the effects are not significant (No Info vs. Only Winner,  $z = 1.430$ ,  $p_{\text{adj}} = 0.3053$ ; control vs. Only Winner,  $z = -1.291$ ,  $p_{\text{adj}} = 0.1967$ ).<sup>14</sup>

If partial information, namely whether one was or not the winner, was enough to satisfy performance curiosity, the effect seen in Experiment 1 should disappear in the Only Winner treatment, because the information regarding the winner was available in both allocations. Quite to the contrary, we observe that the effect is already present in the Only Winner treatment. Hence, we discard that revealing partial information about the winner is enough to satisfy performance curiosity. The general picture that arises so far is that Experiment 2 successfully replicated the results of Experiment 1, but additionally the new treatment, where the information tradeoff also exists but is less severe than in the No Info treatment (because less information is hidden under the equal split) is simply midway between the control treatment and the No Info one. That is, the stronger the informational tradeoff,

---

<sup>14</sup>The interaction between gender and treatments was not significant in an additional probit regression.

the stronger the influence of performance curiosity, leading the participants to choose the fully informative performance-based split.

### 1.4.3 Expected Performance

The median number of correctly solved additions in Experiment 2 was 15 (ranging from 3 to 36), and the median error rate was 0.15 (ranging from 0 to 0.73).<sup>15</sup> As in Experiment 1, we control for the individual beliefs regarding the expected ranking, elicited after the completion of the real-effort task but before the distribution choice. Figure 1.4(a) illustrates the expected own ranking of the participants. In contrast to Experiment 1, there were some minor but significant differences in expectations across treatments. Specifically, the average expected septile for participants in the No Info treatment (4.378) was significantly higher than that of the participants in the Only Winner treatment (3.889;  $z = 2.487$ ,  $p_{\text{adj}} = 0.0387$ ), and also higher than that of the participants in the control treatment (3.978; MWW test,  $z = 2.100$ ,  $p_{\text{adj}} = 0.0714$ ) although the difference missed significance. There were no significant differences in expectations between the control and Only Winner treatments ( $z = 0.430$ ,  $p_{\text{adj}} = 0.6674$ ).

As in Experiment 1, the proportion of performance-based choices was significantly higher in the No Info treatment than in the control treatment, both for participants who expected to perform above average (control, 73.53% of  $N = 34$ ; No Info, 97.67% of  $N = 43$ ; test of proportions,  $z = -3.130$ ,  $p = 0.0017$ ) and for those who expected to perform on average (control, 23.81% of  $N = 21$ ; No Info, 72.00% of  $N = 25$ ;  $z = -3.256$ ,  $p = 0.0011$ ). The difference for participants who expected to perform below the average was not significant (control, 25.71% of  $N = 35$ ; No Info, 36.36% of  $N = 22$ ;  $z = -0.856$ ,  $p = 0.3922$ ). However, we do not observe “last-place aversion”

---

<sup>15</sup>As in Experiment 1, participants were asked how many additions they thought they had solved correctly (after completing the real-effort task but before making the distribution choice). The average expectation was 16.2 additions in the control treatment, 17.3 additions in the No Info treatment, and 15.0 additions in the Only Winner treatment. The differences were generally not significant according Mann-Whitney-Wilcoxon tests, adjusted for multiple testing (control vs. No Info,  $z = -1.177$ ,  $p_{\text{adj}} = 0.2393$ ; control vs. Only Winner,  $z = 1.344$ ,  $p_{\text{adj}} = 0.3577$ ; No Info vs. Only Winner,  $z = 2.217$ ,  $p_{\text{adj}} = 0.0799$ ).

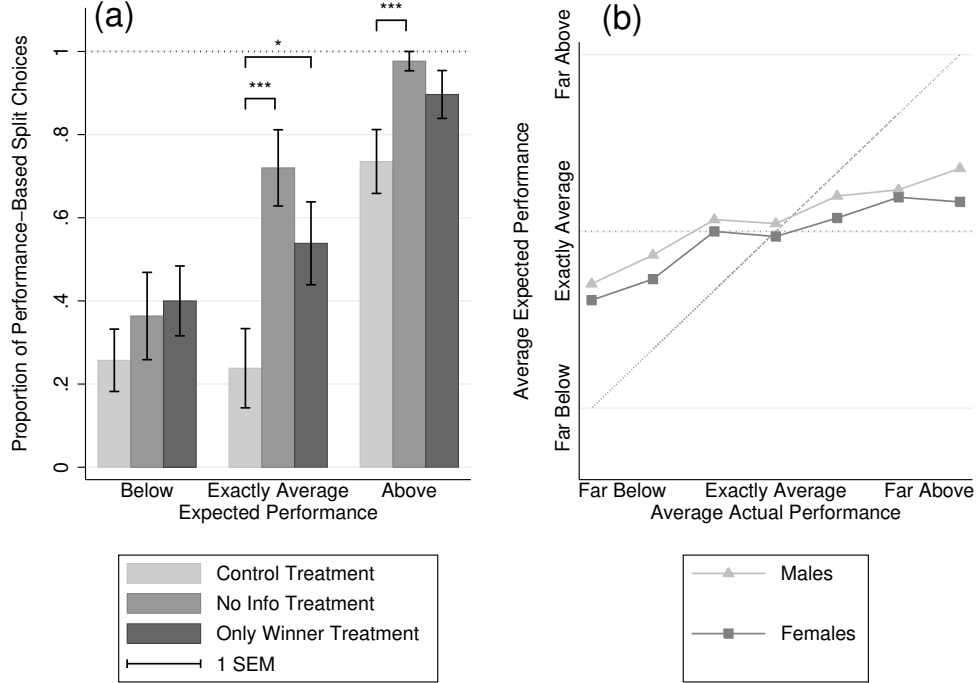


Figure 1.4: Experiment 2, Proportion of Performance-Based Splits and Expected Performance.

*Notes.* (a) Proportion of participants choosing the performance-based split by treatment and belief. (b) Expected versus actual performance by gender. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , test of proportions.

(Kuziemko et al., 2014) in the sense that for those subjects the equal split distribution is not chosen significantly more often in the No Info treatment, where they could hide their expected poor placement.

As commented above, the proportion of performance-based choices in the Only Winner treatment lies between the proportion of performance-based split in the control and No Info treatments. Of course, this points to weaker effect sizes when comparing with either of the other two treatments. As a consequence, when splitting the data into the three expected performance ranges, the differences are in general not significant.

Figure 1.4(b) plots the expected performance (average report in the 7-point Likert scale) for males and females, conditional on the actual population

septiles in the experiment, as computed from the number of correctly solved additions (compare with Figure 1.2(b)). As in Experiment 1, we observe that females tend to underestimate performance compared to males. This is confirmed by a Mann-Whitney-Wilcoxon test on the whole sample (males,  $N = 135$ , average expected septile 4.304; females,  $N = 135$ , average 3.859;  $z = 2.843$ ,  $p = 0.0045$ ).<sup>16</sup>

#### 1.4.4 Regression Analysis

We now turn to a regression analysis to confirm the results of our non-parametric tests while controlling for expected ranking and for the various other individual variables measured in the experiment. In view of the results above, instead of treatment dummies we introduce two dummies representing what kind of information is revealed. The first dummy, *Hidden Info*, captures all cases when some information is hidden in the equal split, that is, both the No Info and the Only Winner treatment. The second dummy, *Only First*, further differentiates the cases where only information regarding the winner is revealed when the equal split is implemented, that is, identifies the Only Winner treatment. This way we can better analyze the differential effects of the informational treatments and emphasize the comparisons between the Only Winner and the No Info treatments (captured directly by the Only First dummy) and between the No Info and the control treatments (captured by the Hidden Info dummy).

Model 1 (see Table 1.4) captures the basic treatment effects without additional controls and reproduces the insights from the non-parametric tests as illustrated in Figure 1.3(a). The Hidden Information dummy is positive and highly significant, showing that choosing the performance-based split was more likely in the No Info treatment than in the control treatment. The Only First dummy is significant and negative, showing that performance-based choices were less likely in the Only Winner treatment than in the No Info treatment. A linear combination test (bottom of Table 1.4) shows that

---

<sup>16</sup>As in Experiment 1, women actually performed worse than men, with an average 17.2 correct additions for males and 15.1 for females (MWW test,  $z = 2.294$ ,  $p = 0.0218$ ). Recall Footnote 13.

Table 1.4: Experiment 2, Probit Regressions on Performance-based Choices.

	Model 1	Model 2	Model 3
Hidden Info	0.860*** (0.1961)	0.875*** (0.1978)	0.874*** (0.2208)
Only First	-0.439** (0.1967)	-0.446** (0.1980)	-0.305 (0.2216)
Female		-0.387** (0.1592)	-0.219 (0.1759)
Expected Performance in Septiles			0.595*** (0.0809)
Constant	-0.168 (0.1328)	0.0220 (0.1547)	-0.0634 (0.1694)
LogLikelihood	-172.205	-169.224	-136.021
Wald Test	19.261***	24.540***	68.681***
Linear Combination Test Hidden Info + Only First	0.421** (0.1884)	0.430** (0.1899)	0.569*** (0.2077)

*Notes.* Standard errors in brackets, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

performance-based choices were more likely in the Only Winner treatment than in the control treatment.

In Model 2, we further control for gender. Females are less likely to choose the performance-based split than men, as reflected by the significantly negative female dummy. Adding interactions between gender and treatments does not change the results qualitatively. Otherwise, the results of Model 1 are not affected by controlling for gender.

Model 3 incorporates the expected performance (in septiles and centered at 0), which, as in Experiment 1, is shown to have a positive and highly significant effect on the likelihood of choosing the performance-based split. That is, subjects are naturally more likely to choose the performance-based split when they expect to be highly ranked. As expected, the female dummy becomes insignificant controlling for expectations, confirming that the apparent gender effect is fully driven by the fact that females have lower expectations than males. The key treatment effect remains significant, showing that

even when controlling for expectations the performance-based choice is more likely in the No Info treatment than in the control treatment (as in Experiment 1). However, the difference between the No Info and the Only Winner treatments becomes insignificant ( $p = 0.1686$ ) when we control for expectations. Accordingly, the difference between the control and the Only Winner treatments becomes *more* pronounced and highly significant ( $p = 0.0061$ ), suggesting that, after controlling for expectations, both the Only Winner and the No Info treatments actually have similar effects. Recall that if partial information were already enough to satisfy performance curiosity, we would expect no effect at all in the Only Winner treatment, as this information is available independently of the distribution choice. Hence, it seems that the driver of behavior is the actual curiosity to know the own performance, independently of whether one is the winner or not.

## 1.5 Discussion

We have shown that performance curiosity, a form of intrinsic motivation related to self-image, can overcome preferences for egalitarian allocations of jointly generated proceeds. In our experiments, a subtle manipulation in the structure of information regarding how well each participant performed in a real-effort task was enough to produce large shifts away from equal-split choices and toward performance-based schemes. This effect subsists even when one controls for expected relative performance, which of course correlates with the choice of the performance-based allocation, as the latter maximizes expected income if one expects to perform better than average.

As we expected, the effect of performance curiosity appears to be much stronger for females than for males, with the consequence that the manipulation greatly helps reduce gender differences in egalitarian choices. The gender difference in our results is fully explained by differences in performance beliefs. Apparent gender differences regarding the proportion of equal-split choices (and hence, inequality aversion) are caused by differences in beliefs regarding the own performance, namely the fact that, compared to women, men are typically more overconfident. The difference in beliefs creates an



apparent gender effect in the control treatment, with females choosing the egalitarian allocation more often than males. This difference (and the effect of beliefs) disappears in the presence of performance curiosity. Our analysis hence identifies a factor which helps overcome gender differences arising from relative underconfidence.

Appendix A analyzes and discards a number of alternative explanations for our results, on the basis of supplementary data (questionnaires) collected in the two experiments. Some other alternative interpretations of our results and possible extensions are as follows. First, it is well-known that different factors as entitlement, merits, etc, can be used as justifications in order to depart from “fair” allocations and switch to more favorable ones (in terms of income). It could be argued that the presence of additional information attached to the performance-based split could have been used as precisely such a justification even if the participants did not care about the information itself. However, this explanation is unlikely because our data shows a larger proportion of performance-based distribution choices in the No Info treatment than in the control treatment even for subjects who expected to perform exactly on or below average. For those participants, by design the income-maximizing decision would have been the equal split. Hence, the additional information has no value as an excuse to pursue an income-maximizing allocation.

Second, a possible alternative interpretation would be that the effect is simply due to the fact that, in the No info treatment, the performance-based distribution comes with an additional informational attribute than the equal split distribution. This explanation would imply that general curiosity is far stronger than assumed, in the sense that the particular content of the additional information is not relevant. By its very nature, a performance-based split will always reveal some information on performance, hence it is not possible to shut down that informational attribute and replace it with an alternative, less informative one. However, the explanation could be tested with an additional design where the equal split is endowed with some non-informative attribute (with respect to performance; e.g., the weather forecast of a remote country) but the performance-based split is still endowed with

performance information. Hence, the number of informational attributes would be kept constant, and any difference would be due to the actual nature of the information. This design goes beyond the scope of the present work.

A related point is that in the No Info treatment it was made salient that information on performance was being left out, which might have increased the focus on performance. This is unavoidable, since the point of our design is to make the tradeoff between information and an equal-split allocation apparent. However, the real-effort task was run before the actual choices were presented. There was no information whatsoever about the fact that performance information might or might not be revealed depending on later choices (in the No Info treatment, the initial instructions mentioned that choices might differ in attached information, but the nature of such information was not explained until after the real-effort task). That is, the fact that there was a tradeoff with respect to information on performance was only apparent right before the actual choice of distribution, and hence any possible salience effects would have been kept to a minimum. Also, we remark that, as shown in the regression analysis, the effect subsists when controlling for expected performance.

Some extensions of the present research would be natural. We mention here just two of them. First, Loewenstein (1994) argued that curiosity may be an impulsive hedonic drive that easily wanes after being satisfied. If this is the case, prosocial decision makers might regret having chosen feedback over equality, leading to different decisions in the future. Since our objective was to show that performance curiosity can trump inequality aversion, we adopted a purely one-shot design, and hence we cannot test this additional possibility. It would be possible to expand the design incorporating repeated decisions in order to test for this hypothesis, but if performance was revealed between decisions, it would be necessary to appropriately control for it.

A second natural avenue for further research would be to elicit the actual willingness to pay for information on performance. In view of well-known behavioral phenomena pointing at discrepancies between valuations and actual choices and the noisy character of evaluations (e.g., Delquié, 1993; Alós-Ferrer et al., 2016a), we aimed to establish that the postulated effect had

consequences on actual behavior (choice data), but once this effect has been established, evaluating the willingness to pay for information is a natural second step. Although we leave this task for future research, we can offer some preliminary evidence because, as mentioned above, participants who believed to be exactly on average or below would have maximized expected payoffs by choosing the equal-split allocation, but did not choose that allocation more often. Hence they actually revealed the willingness to incur monetary costs in order to obtain the information.

We view our research as a contribution to the general investigation into the causes of and motivations behind preferences among different distributional allocations, especially in frameworks where the resources that are distributed can be traced back to individual contributions. This is potentially important for reward schemes in firms and organizations, and for a deeper understanding of attitudes with respect to fairness and redistribution at the societal level. Our results, however, can be read in two different ways. If egalitarian allocations are viewed as fair and hence as a worthy policy objective, the results point out that support for egalitarian principles might be diminished in society if information on individual performance is restricted, for instance in order to protect privacy or to avoid making interpersonal comparisons prominent. While people might be willing to accept egalitarian redistributions independently of whether they profit from them themselves or not, they still have a strong preference for receiving information on their relative individual contribution. On the other hand, if rewarding performance through incentives is viewed as fair and desirable, linking information to rewards might result in increased social support of the corresponding distribution schemes.

## Appendix 1.A: Robustness Analysis

### Appendix 1.A.1: Experiment 1

In this subsection, we report on a few additional controls which help discard possible alternative interpretations. First, it is natural to speculate that some

participants might choose the egalitarian distribution in the No Info treatment in order to protect their self-image, especially if they fear being below average, because in this way they avoid being exposed to negative feedback. However, this would create a tendency leading to less performance-based splits, which we do not observe, even among participants who expected to be below average. In any case, we anticipated this alternative explanation and we introduced a final question for participants in the No Info treatment, after all other decisions have been made but before the distribution was implemented. In this question, we asked them whether they wanted to be privately informed about their own performance (both absolute and relative, that is, number of correctly solved additions and ranking within the group), independently of which distribution was finally implemented. Recall that even if a participant chose the performance-based split, it was not guaranteed that the information would be revealed, since the random mechanism could select a different participant, which might have chosen the egalitarian allocation instead. Hence the additional question made sense to participants. All participants but one (that is, 89 out of 90) chose to see their own performance, indicating that the possible desire to protect their self-image against feedback was not a factor.

Second, we measured risk attitudes by means of the lottery questionnaire of Holt and Laury (2002). There were no differences in risk attitudes by gender (males: mean relative risk aversion parameter 0.497; females: mean 0.558; MWW test,  $z = -0.405$ ,  $p = 0.6851$ ), treatment (control treatment: mean 0.526; No Info treatment: mean 0.530;  $z = -0.203$ ,  $p = 0.8388$ ), or choice (egalitarian choice,  $N = 82$ , mean 0.530; performance-based choice,  $N = 98$ , mean 0.526;  $z = -0.062$ ,  $p = 0.9508$ ).<sup>17</sup>

Third, we also used the reduced Achievement Motive Scale of Lang and Fries (2006) to elicit achievement motives. That scale is divided in questions eliciting two components, aptly named “hope of success” and “fear of failure.” Each subscale is made out of 5 items on a 4-point Likert scale, and

---

<sup>17</sup>Pooling together a large number of experimental studies, Filippin and Crosetto (2016) find significant but very small gender differences in risk attitudes, as measured by the lottery-questionnaire method.

the (added) scores of the subscale range from 4 to 20. We found the gender differences typically described in the literature, namely that men score significantly higher on the hope of success scale (males: mean 17.122; females: mean 15.944; MWW test,  $z = 3.518$ ,  $p = 0.0004$ ) and women score significantly higher on the fear of failure scale (males: mean 10.344; females: mean 12.400;  $z = -4.346$ ,  $p < 0.0001$ ). There were no differences across treatments, neither for hope of success (control treatment, mean 16.300; No Info treatment, mean 16.767;  $z = -1.379$ ,  $p = 0.1680$ ) nor for fear of failure (control treatment, mean 11.400; No Info treatment, mean 11.344;  $z = 0.200$ ,  $p = 0.8417$ ).

There was no difference by choices in the scores of the fear of failure scale (egalitarian choice:  $N = 82$ , mean 11.780; performance-based choice,  $N = 98$ , mean 11.031;  $z = 1.470$ ,  $p = 0.1415$ ). However, participants who chose the performance-based split scored higher on the hope of success scale (egalitarian choice:  $N = 82$ , mean 15.841; performance-based choice,  $N = 98$ , mean 17.112;  $z = -3.689$ ,  $p = 0.0002$ ). This difference is also significant if we look separately at the control treatment (egalitarian choice:  $N = 55$ , mean 15.872; performance-based choice,  $N = 35$ , mean 16.971;  $z = -2.914$ ,  $p = 0.0283$ ) and the No Info treatment (egalitarian choice,  $N = 27$ , mean 15.778; performance-based choice,  $N = 63$ , mean 17.190,  $z = -2.776$ ,  $p = 0.0055$ ). This makes sense as hope of success should be associated with higher expectations, hence a higher likelihood of choosing the performance-based split. In other words, we do not gain any new insights from the achievement motivation scales.

To further examine the effects of achievement motivation and risk attitudes, and also to study the robustness of the basic effects, we added them as controls in a further regression model (see Table 1.5). This model also includes interactions with the No Info treatment, which should reflect any possible explanation of the basic effect based on the individual differences captured by the new measures. This has the problem that the interpretation of the treatment dummy becomes more complex. To test for the treatment effect of the manipulation, we calculated the average marginal effects of the

treatment dummy on the performance-based split choice.<sup>18</sup> As shown in the bottom line of Table 1.5, the average marginal effects of the treatment dummy are highly significant and quantify a 25% increase in the probability of choosing the performance-based split after including the manipulation in the structure of the information.

The regression model, however, did not detect any significant effects of risk attitudes, hope of success, or fear of failure. Hence, we conclude that these additional variables do not hide any alternative explanation and our basic effects, as captured in Model 3 (Table 1.2), are robust.

#### Appendix 1.A.2: Experiment 2

As in Experiment 1, subjects in the No Info and Only Winner treatments were asked whether they wanted to be informed about the own absolute and relative performance after making the distribution choice. Only two subjects (of 180) declined to receive feedback. Hence, we conclude again that self-image protection does not appear to be a factor in our experiments.

In Experiment 2, we further investigated the possible relation of performance curiosity to existing constructs capturing curiosity. We focused on *epistemic curiosity* (Loewenstein, 1994; Litman et al., 2005; recall Section 1.2) and included the epistemic curiosity questionnaire by Litman and Mussel (2013). Further, since our explanation for the relevance of performance curiosity hinges on the value of feedback for the own self-image, we explored whether this link was captured by differences in standard *self-motives* (Sedikides and Strube, 1995, 1997; recall Section 1.2) and included the questionnaire by Gregg et al. (2011).

Model 4 (see Table 1.6) includes the two measures for epistemic curiosity (EC) and the four measures of self-motives (SM) elicited through questionnaires, and the interaction of each of the six measures with the two information dummies. None of the new measures is significant, and among the twelve interactions only one reaches significance, namely the interaction be-

---

<sup>18</sup>The average marginal effect calculates the average difference in the predicted probability of choosing the performance-based split if all observations would have been in the No Info treatment, compared to if all observations would have been in the control treatment.

Table 1.5: Experiment 1, Additional Probit Regression on Performance-based Choices.

	Model 4	Interactions with No Info Treatment
No Info Treatment	−0.588 (1.7291)	
Female	−0.262 (0.3286)	
No Info X Female	0.227 (0.4726)	
Expected Performance in Septiles	0.536*** (0.0995)	
Risk Attitude	0.019 (0.3769)	−0.054 (0.5074)
Hope of Success	0.047 (0.0666)	0.002 (0.0865)
Fear of Failure	−0.068 (0.0513)	0.122 (0.0757)
Constant	−0.325 (1.3718)	
LogLikelihood	−88.304	
Wald Test	50.598***	
Average Marginal Effects No Info Treatment	0.247*** (0.0517)	

*Notes.* Standard errors in brackets, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

tween self-improvement and the Hidden Information dummy. There is hence, no convincing basis to speculate that the performance curiosity concept might be captured or closely related to existing measures of curiosity or self-motives.

As in Model 3 (Table 1.4), there is no significant difference between the Only Winner and the No Info treatments. Hence, Model 4 shows the robustness of Model 3. To examine the robustness of the main results to the introduction of the new 18 coefficients, we again calculated the average marginal effects to test for the treatment effect. Those are shown at the bottom of Table 1.6. The effect of the No Info treatment is positive and highly signif-

Table 1.6: Experiment 2, Additional Probit Regression: Epistemic Curiosity (EC) and Self-Motives (SM).

	Model 4	Interactions with	
		Hidden Info	Only First
Hidden Info	0.328 (0.4407)		
Only First	0.257 (0.4783)		
Female	−0.155 (0.1901)		
Expected Performance in Septiles	0.628*** (0.0894)		
EC-Interest	−0.012 (0.0683)	0.034 (0.0970)	−0.020 (0.1038)
EC-Deprived	−0.010 (0.0460)	−0.008 (0.0856)	0.051 (0.0941)
SM-Enhance	−0.012 (0.1388)	0.027 (0.2061)	0.005 (0.2038)
SM-Assessment	−0.074 (0.1962)	−0.016 (0.2896)	−0.110 (0.2708)
SM-Verification	−0.029 (0.1397)	0.171 (0.2301)	−0.255 (0.2394)
SM-Improvement	−0.236 (0.1810)	0.574** (0.2601)	−0.232 (0.2357)
Constant	0.204 (0.2889)		
LogLikelihood	−130.392		
Wald Test	72.795***		
		Hidden Info	Only First
Average Marginal Effects		0.235*** (0.0561)	−0.075 (0.0618)

Notes. Standard errors in brackets, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

icant, showing that the probability of choosing the performance-based split is increased by approximately 24% by the basic manipulation. This almost exactly replicates the conclusion from Experiment 1 (recall Table 1.5)



---

## CHAPTER 2

### The Big Robber Game

---

#### 2.1 Introduction

The recent decades have witnessed an astonishing loss of confidence by the general public on financial institutions, large firms, and economic decision makers. The image of riotous protesters at WTO ministerial conferences, G20 summits, and other high-level meetings has become commonplace. Public animosity toward financial institutions, the banking sector, and “Wall Street” reached a 40-year high in 2011 (Owens, 2012). Beyond the circles of academic economists and politicians, confidence in the market has been replaced by the views put forward by former German Chancellor Helmut Schmidt, who coined the term “predatory capitalism” (*Raubtierkapitalismus*; Schmidt, 2003), referring for instance to a large number of scandals where corporate executive officers were viewed as unduly appropriating funds or actively damaging society. As stated by Krugman (2008), “[Americans have] lost confidence in the integrity of our economic institutions.” This has sparked a widespread, cynical view of economic actors as hopelessly-selfish agents capable of inflicting any damage on others for personal gain.<sup>1</sup>

On the other hand, it is clear that self-interest is not the only motivation guiding economic human decisions, and this has long been recognized in economics. As Adam Smith put it, “how selfish soever man may be supposed to be, there are evidently some principles in his nature, which interest him in the fate of others, and render their happiness necessary to him, though he derive nothing from it, except the pleasure of seeing it” (Smith, 1759). The limits of self-interest have been systematically exposed in experiments using

---

<sup>1</sup>The idea that corporate scandals might have even macroeconomic consequences is reflected, for instance, in the model of Myerson (2012), where the possibility of moral-hazard rents extracted by bankers and financial agents can create macroeconomic credit cycles with repeated recessions.

stylized games as the Ultimatum Game (Güth et al., 1982), the Dictator Game (Forsythe et al., 1994), and the Trust Game (Berg et al., 1995), as well as in many distributive-allocation experiments (Engelmann and Strobel, 2004). This evidence has motivated models incorporating various forms of “social preferences,” encompassing fairness concerns, prosociality, and even motivations with an intentional component as reciprocity (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002; Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006; Alger and Weibull, 2013).

The “social-preference revolution” seems to be based on the enticingly simple message that economic agents are not so selfish as usually assumed in economic theory. For example, in a widely-echoed article, Mazar et al. (2008) argue that most people will help others, refrain from actively damaging them, and, if given opportunity, cheat only a little. This message, however, is at odds with the popular (and populist) criticisms of the “economic system,” which showcase the allegedly inordinate levels of selfishness displayed by corporate executives and decision makers in the financial sector. In other words, while social-preference models attempt to defuse the assumption that economic agents are selfish, the popular perception of economic decision makers is that, at least at certain levels, selfishness is rampant.

In view of these opposed trends, it is maybe advisable to make a distinction between social preferences and the experimental observation of prosocial or altruistic behavior. To drive this point home, consider the following example. A decision maker  $i = 1$  can choose among certain allocations for  $n$  participants, including herself,  $x = (x_1, \dots, x_n)$ , with the property that  $\sum_{i=1}^n x_i = C$  for a certain amount  $C > 0$ . Her utility function is given by

$$u(x) = x_1 - 25 \left( \frac{x_1}{C} - \frac{1}{n} \right)^2.$$

This is an example of social preferences as in Bolton and Ockenfels (2000). Suppose this decision maker participates in a Dictator Game where she can freely allocate 10 monetary units between herself and a second participant, that is,  $n = 2$ ,  $C = 10$ ,  $0 \leq x_1 \leq 10$ . A straightforward computation shows that  $u$  is maximized at  $x_1 = 7$ , hence the decision maker will freely

donate 3 out of 10 monetary units to the other player, in line with prosocial behavior as typically observed for the Dictator Game in the lab. That is, this decision maker can be seen as prosocial. Suppose that this very same decision maker is now part of a group of 17 people, each of them endowed with 10 monetary units, and she is given the power to appropriate any amount  $\tau$  from each of the other participants (uniformly). That is,  $n = 17$ ,  $C = 170$ , and the decision maker can freely choose among allocations of the form  $(10 + 16\tau, 10 - \tau, \dots, 10 - \tau)$ , and hence her problem consists of maximizing

$$v(\tau) = 10 + 16\tau - 25 \left( \frac{10 + 16\tau}{170} - \frac{1}{17} \right)^2$$

over  $0 \leq \tau \leq 10$ . A direct computation shows that  $v(\cdot)$  is strictly increasing in its domain, and hence our prosocial decision maker will decide  $\tau = 10$ , appropriating all income from a group of 16 other participants.

In our view, it is uncontested that, first, humans care about how their actions affect others and, second, selfishness is one of the most fundamental and powerful human motivations. The difference between prosocial behavior as observed in the lab and rampant egoism as apparently observed in the field is that the situations triggering the respective responses are fundamentally different. As the example above illustrates, this is readily captured by extant models of social preferences. In this work, we aim to demonstrate empirically that prosocial behavior in the small is fully compatible with morally-outrageous behavior in high-stakes, high-impact decisions where a decision maker can obtain a large personal gain at the expense of significantly damaging a large number of other people. Obtaining such empirical evidence entails three kinds of difficulties, and it is our aim to address them within an experimental paradigm designed to inform the discussion.

First, many of the experimental paradigms on which social-preference models have been based are constrained to bilateral interactions. This is natural, since those games were a first step in research, and keeping the experimental setting as simple as possible facilitates the analysis. Think, however, of the typical situation that critics of financial institutions have in mind, where a single decision maker has the possibility to harm a large

number of people. Clearly, such a situation differs from the experimental paradigms in several dimensions, and hence the latter are not well-suited to predict behavior in the former. For example, in view of corporate scandals, it could be argued that financial institutions select people who are not representative of the general population, and hence prosocial behavior in the large is compatible with selfish behavior at certain levels. Currently-employed experimental paradigms, however, are not appropriate to test for such a selection phenomenon, since one would need an experimental setting where regular laboratory subjects can actually damage a group of other experimental participants (in exchange for significant monetary gain). This is exactly what our design provides.

The second, related difficulty is that games of the bargaining type (Ultimatum, Dictator, and Trust games) are highly stylized and might hence measure a very specific dimension of behavior. Again, this is natural, since they were originally designed to provide demonstrations of the existence of social preferences and motives. But they might be insufficient to explore the full impact of prosocial behavior on economic decisions. In other words, those games might be triggering prosocial behavior in a limited set of situations, which might be compatible with high levels of selfishness being sparked in different circumstances (see also Levitt and List, 2007). Our experiment contributes to the literature supporting this view.

The third difficulty is that, to determine whether prosocial behavior is of sufficient magnitude as to be seen as a major driving force for economic behavior, one needs to address the size of the stakes. This question is hard to examine in standard laboratory experiments, simply because the stakes are too low. There are, however, indications that prosocial behavior can still be found in spite of high stakes. Indeed, experiments in low-income countries have found prosocial behavior with large stakes (Roth et al., 1991; Slonim and Roth, 1998; Andersen et al., 2011).<sup>2</sup> However, an analogous, smoking-gun demonstration for high-income, developed countries is missing, because generating high stakes in such environments would quickly exhaust

---

<sup>2</sup>Andersen et al. (2011) focused on responder behavior in Ultimatum games and showed that rejection rates decreased with high stakes, but were still positive.

even the most generous research budget. Our design makes a methodological contribution by providing a cost-neutral way (in terms of research budget expenditure) to study high-stake settings.

In this contribution, we present a new experimental paradigm, the “Big Robber Game,” which adds to the debate on prosociality vs. selfishness along the three dimensions sketched above. First, it focuses on a situation where a single decision maker can take a significant amount of money *from a large group of other participants*, and is hence closer to phenomena sparking public-opinion concerns in recent years. Second, it provides a qualitatively different paradigm which, together with standard games as the Ultimatum, Dictator, and Trust games, allows for a more complete exploration of the nature and consequences of social preferences. Third, it allows us to test for prosocial behavior by educated subjects in developed countries, with reasonably high stakes (around 100 Dollars/Euros), in a natural, straightforward frame where being selfish inflict serious damage on a group of others.

The basic, pragmatic design idea behind the Big Robber Game is as follows. Consider an average researcher in economics based in the U.S. or Europe, with access to a budget for behavioral experiments. A single session with, say, 30 participants, costs around €/\$ 400 (or more). From the point of view of the average participant, the cost of that individual session is a significant amount of money. Suppose that the experimenter selects a participant randomly from half of the participants in the session and gives him or her the possibility of taking as much as 50% of the earnings of the other half of the players in the session. Then, a player in this game will face a single decision potentially giving him or her the possibility to walk away with €/\$ 100. That decision might seriously damage a relatively large number of people and be considered selfish, even antisocial. The experimental session, however, does not cost more. Our experimenter will spend exactly the same part of his or her research budget as a “regular” experiment would have required.

In our design, we allocated participants to two possible roles, which were framed neutrally (type I and II) but which we will refer to here as “robbers” and “victims.” Robbers were asked the Big Robber question, that is, which

fraction of the earnings of the victims they wished to take away. Crucially, only one robber was selected at the end (and his or her decision actually implemented), but all robbers had to answer the question in advance (as in the strategy method; Selten, 1967; Brandts and Charness, 2011). All victims were informed in advance of the possibility that a fraction of their earnings could be taken away by a robber. There was only one Big Robber decision, with fixed roles, that is, robbers could not be robbed themselves, and there was no possible retaliation. To generate the relevant income, players of both roles provided their decisions for the active and passive roles in a series of intermediate games: Dictator, Ultimatum, and Trust games. This allows us to compare behavior in the Big Robber question to behavior in these standard games, and hence link the different measures of prosociality that they provide. Feedback was not provided until the end of the experiment, and participants did not know the identity or any characteristic of the players they would be matched with in the intermediate games (not even if they were robbers or victims). Therefore, behavior could not be affected by possible group rivalry (see, e.g., Abbink et al., 2010, 2012).

Beyond the primary research question, we were also interested to see whether merely asking the Big Robber question might have an effect on subsequent behavior, since a number of psychological theories (see Section 2.4 below) might predict behavioral effects. Hence, we further divided the set of robbers in two, creating two different treatments. “Ex ante robbers” were asked the Big Robber question at the beginning of the experiment, before they provided their decisions for the intermediate games. “Ex post robbers” were asked the Big Robber question *after* they had made their decisions for the intermediate games (note that, since only victims could be robbed, the Big Robber question did not affect the robbers’ earnings from the intermediate games).

After all decisions had been made, but before they were implemented and payoffs were determined, we asked participants whether they would donate a percentage of their earnings (net of the show-up fee) to a local charity. They were free to specify any fraction. The idea of this “donation question” was

to provide an independent measure of possible guilt, by giving robbers the chance to give a prosocial use to their appropriated earnings.

The theoretical predictions of the selfish *homo oeconomicus* for the Big Robber Game are straightforward: take as much as possible. Received social-preference models predict varying levels of selfish behavior, which we examine in Section 2.6. Those models incorporate parameters of prosociality which can differ across subjects. The intermediate games provide exactly such measures of prosociality. We use the Dictator Game to calibrate the models of Fehr and Schmidt (1999), Bolton and Ockenfels (2000), Charness and Rabin (2002), and Alger and Weibull (2013). As in our example above, the predictions for these models suggest that the agent should in many cases take as much as possible, i.e. 50%; actually, standard social-preference models predict higher levels of selfish behavior than we actually observe.

The results of the experiment show that the paradigm provides new insights into social preferences and prosocial behavior which go beyond the ones that can be gathered with “standard” games as the Dictator, Ultimatum, and Trust games. First, at the stakes level of the Big Robber Game, there is a considerable degree of “selfishness” as predicted, with as much as 56% of robbers deciding to take as much as allowed of the victims’ earnings. However, the behavior of the very same robbers in the Dictator, Ultimatum, and Trust games is within standard ranges. This shows that the Big Robber Game provides indeed a novel view for the debate on prosociality vs. selfishness, and that different situations even within the same experiment can trigger qualitatively different behavior with respect to the perceived prosociality associated with the decisions.

Second, we do find that robbers who decide to take as much as possible behave more selfishly than robbers who do not take the maximum in the Dictator, Ultimatum, and Trust games, and also in the donation question. That is, the Big Robber Game is not orthogonal to other paradigms; on the contrary, behavior in our paradigm does correlate with individual prosociality, and hence is consistent with underlying social preferences as modeled in the literature. Third, we do find some (weak) treatment effects between ex ante and ex post robbers which are compatible with known psychological theories.

However, these effects interact with gender, showing basic differences in the effect that revealing oneself to be selfish has on subsequent behavior. Gender differences were not unexpected, and hence we took care to have a perfectly balanced sample, with exactly half of all robbers within each treatment being female.

Our study is of course related to several branches of the extensive literature on social preferences. Some experimental paradigms in this literature allow players to take money from other players. In the Power-to-Take game (Bosman and van Winden, 2002; Bosman et al., 2005; Reuben and van Winden, 2010), a player can attempt to take the earnings of another player, but the second player can react by destroying his own income and hence the part appropriated by the first player. Players in the “take role” choose considerable take rates in this paradigm. Bosman et al. (2006) consider a version with group decisions and show that behavior is very similar to the individual version. In the Moonlighting Game (Abbink et al., 2000), the first mover can either transfer money to a second mover or steal from him, but the second mover is able to punish. Results showed consistent punishment of stealing behavior although punishment was costly. Bardsley (2008) considered a Dictator Game with the added option to take income away from the recipient (see also List, 2007), and showed that Dictator giving is greatly reduced when the option to take exists. Andreoni (1995) found more selfish behavior when taking from a public account than when giving to a public good.<sup>3</sup> Khadjavi and Lange (2015) replicated the experiment making both the taking and giving options simultaneously available and found that allowing for taking reduces giving. These results are qualitatively in agreement with ours and show that situations where taking is possible generally trigger less prosocial behavior than those framed in terms of bargaining or sharing.

---

<sup>3</sup>There is also a relation to the Common-Pool Resource (CPR) dilemma (Gardner et al., 1990), where multiple players extract resources from a common pool. Among other conditions, Gardner et al. (1990) define CPR dilemmas by Multiple Appropriators, i.e. more than one individual withdraws resources from the common pool, and a common-pool resource. However, the Big Robber Game does not satisfy these two conditions. First, there exists only one Big Robber in the game and second, he or she is not withdrawing from a common pool but only from the individual earnings of each victim who cannot take from others.



The literature has also examined a number of paradigms where antisocial behavior emerges in the laboratory, for instance where participants can destroy the earnings of others or of a group of innocents, due to envy or a concern for relative payoffs (Zizzo and Oswald, 2001; Zizzo, 2004; Abbink and Sadrieh, 2009; Abbink and Herrmann, 2011; Karakostas and Zizzo, 2016). Among those, we single out the experiment of Zizzo (2004), where participants could appropriate other participants' income or transfer it to others. Decisions were made within small (four-player) groups where everybody could be robbed by the other three participants. This creates a situation closer to a social dilemma where a participants' act of stealing could be argued to be self-defense, since other participants will most likely steal from him or her.

Our paradigm goes beyond previous insights by providing a (relatively) high-stakes situation where an individual can take up to half of the earnings of a large group of other participants, while systematically comparing decisions to simultaneous Dictator, Ultimatum, and Trust games. In contrast with the paradigms mentioned above, our focus is on decisions affecting a large group (not bilateral interactions), with high stakes involved, and where the decision to take the earnings of others is non-strategic, that is, can be made with impunity.

Conceptually, our study is aligned with Levitt and List (2007), who argue that behavior is crucially linked to not only the underlying preferences but also the specific characteristics and properties of the situation at hand. The Big Robber Game provides a specific, relevant setting where qualitative behavior becomes predominantly selfish, even though the same participants simultaneously display standard levels of prosocial behavior in Dictator, Ultimatum, and Trust games played during the very same experimental session.

The Chapter is structured as follows. Section 2.2 presents the experimental design in detail. Section 2.3 discusses the results of the Big Robber Game in itself, that is, the decision to take away (or not) the earnings of other players. It also examines the differences in behavior (in the intermediate games) across robbers who behave more or less selfishly. Section 2.4 discusses treatment effects, that is, differences in behavior due to whether

the Big Robber question had already been answered or not when players faced the intermediate games. Section 2.5 examines the answers to the final donation question. Section 2.6 examines the predictions of well-known models of social preferences for the Big Robber Game, performing a simple out-of-sample estimation analysis. Section 2.7 concludes.

## 2.2 Experimental Design

### 2.2.1 Design and Procedures

The objective of the design was to allow for the measurement of a correlate of social preferences where significant personal gain can be obtained at the cost of economically harming a large group of fellow participants. The novel design of the experiment allowed us to accomplish this objective while maintaining an affordable total experiment cost, even though relatively high stakes were involved and the experiment was run in a developed, high-income country (Germany).

The experiment consisted of three parts: the Big Robber question, a sequence of games (Dictator, Ultimatum, and Trust games), which we will refer to as the *intermediate games*, and a final block of questionnaires, including a belief elicitation question and the opportunity to donate a fraction of the individual earnings to a charity organization. For the duration of the experiment participants were assigned two roles: half of the participants were robbers (neutrally labeled “Type I” during the experiment) and the other half were victims (“Type II”). These roles were only relevant for the Big Robber question. All victims knew from the beginning that there is the Big Robber question in which money could be taken away from them by the choice of the robbers. The robbers were allocated in two between-subject treatments, which differed in the timing of the Big Robber question. In the *ex ante treatment*, the Big Robber question was asked before participants played the sequence of games and they knew as much as the victims. In the *ex post treatment*, the order was the reverse one and ex post robbers were not aware of the possibility of robbing until after they provided all their decisions for

the sequence of games. Note that, since robbers could not be robbed themselves this additional piece of information did not affect their payoffs in the sequence of games. This design serves two purposes. First, by comparing across treatments, it becomes feasible to test for possible effects of the Big Robber question on subsequent behavior. Second, it allows us to disentangle robbing behavior from possible confounds due to the order of tasks and whether the involved income has already been generated or not.

We now discuss the three parts in detail. Figure 2.1 presents an overview of the design. At the very beginning, participants were informed that the experiment consisted of multiple parts but no details were provided.

The Big Robber question revealed to the participants that there were two types of players in the experiment. There were 32 participants per session. In each session, 16 participants were randomly assigned to the robber role and the remaining 16 to the victim role. The instructions referred to these roles as “Type I” and “Type II,” but we will continue calling them robbers and victims for concreteness. The robbers were informed that, at the end of the experiment, one of them would be randomly selected and would have the chance to take 50%, 33%, 10%, or nothing from the joint earnings of all victims. It was clear that only victims would be robbed, that is, a robber would never be robbed. Every robber was then asked to provide his or her decision, should he or she be selected (as in the strategy method). In this way, we collected answers to the Big Robber question from all robbers. This question was asked either before or after the sequence of games depending on the treatment. Specifically, in the *ex post* treatment the robbers were initially not provided with any information about the existence of the Big Robber question.

In contrast, *ex ante* robbers and victims in both treatments knew about the Big Robber question from the beginning of the experiment. In particular, victims were informed of the fact that participants of the other type would be asked the Big Robber question, and that one of them would be randomly selected and his or her decision implemented while the *ex ante* robbers answered the Big Robber question.

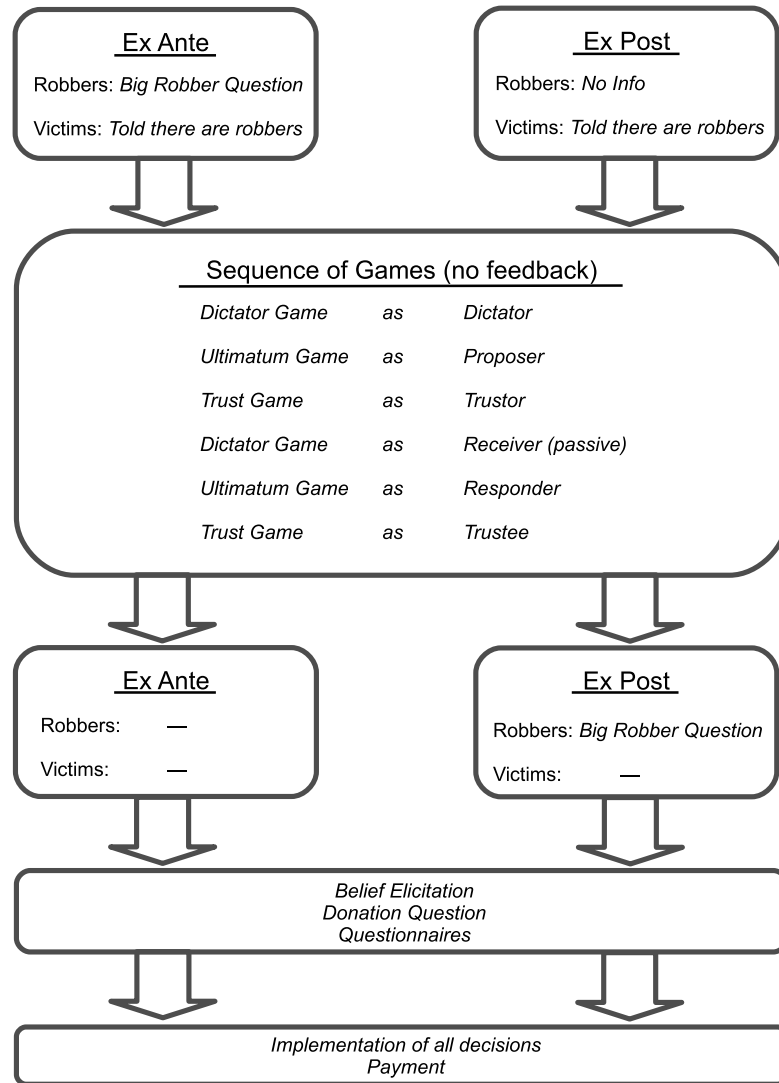


Figure 2.1: Overview of the Experimental Design.

To emphasize the stakes, the Big Robber question also included numerical estimates of the amount of money that would be transferred to the randomly selected robber depending on each possible answer. Those numerical estimates were computed using estimates of behavior in the sequence of games taken from the literature (details are provided below), which delivered an estimated transfer of €100 for a robber deciding to take 50% (respectively €66, €20, and €0 for 33%, 10%, and 0%). We also provided the theoretical maximum transfers derived from the joint profit-maximizing choices. The exact

Table 2.1: Big Robber Choices.

Choice	Proportion in points	Amount in Euros
<input type="checkbox"/>	50% of all points	about €100 (expected, maximum €112)
<input type="checkbox"/>	33% of all points	about €66 (expected, maximum €74.60)
<input type="checkbox"/>	10% of all points	about €20 (expected, maximum €22.40)
<input type="checkbox"/>	0% of all points	€0

format of the Big Robber choices and the provided information is displayed in Table 2.1.

The sequence of intermediate games was played after the Big Robber question in the ex ante treatment and before it in the ex post treatment. The order of those games was fixed and identical for all participants. Each participant played first the Dictator Game (Forsythe et al., 1994) as a dictator, an Ultimatum Game (Güth et al., 1982) as proposer, and a Trust Game (Berg et al., 1995) as trustor. Next, he or she “played” the Dictator Game as receiver, that is, the participant was simply informed that at the end of the game he or she would be matched to some other participant and play the passive receiver role in a Dictator Game. Then, the participant played the Ultimatum Game as responder and, finally, the Trust Game as trustee. Decisions for the last two roles were collected using the strategy method, that is, participants provided contingent answers for each possible proposal of the sender (proposer or trustor, respectively). The order of the games was chosen to minimize the changes between sender and receiver mindsets, hence avoiding an artificial activation of the concept of reciprocity through mere alternation. The framing was in terms of “decisions” and the roles were described neutrally (participants A and B for sender and receiver roles, respectively). Payoffs were realized at the very end of the experiment by randomly matching participants in the corresponding roles in such a way that every participant was matched to a different participant for each of the six games. In particular, no feedback was provided for any of the decisions until the experiment ended. There was no distinction between robbers and victims for the purposes of matching to implement the six intermediate games.

Each game was played with an endowment of 10 points. That is, in each of the games, the sender (dictator, proposer, or trustor) decided how to split 10 points among him- or herself and the other player (with only integer allocations allowed). In the Dictator Game(s), this decision was implemented. In the Ultimatum Game(s), the responder decided whether to accept or reject each possible split (proposal). At the end of the experiment, after matching players and implementing decisions, an acceptance led to the realization of the proposal, and a rejection led to zero payoffs (in that particular game) for both proposer and responder. In the Trust Game, the trustor proposal was implemented, but the part allocated to the trustee was then doubled and the trustee could decide which part of the proceeds (if at all, and constrained to be an integer number of points) to send back to the trustor. The factor of two in the Trust Game(s) was chosen to reduce artificial incentives to cooperate and capture the social preferences in a clean way (Glaeser et al., 2000).

The purpose of the sequence of games was twofold. On the one hand, as most of the literature we use them as (standard) indicators of the social preferences of the participants. On the other hand, the games were a device to generate income for the victims which could then be partially taken away by the robbers.

As commented above, we estimated the income generated by the sequence of games in order to provide numerical estimates of the additional revenue associated to each alternative choice in the Big Robber question. Since each game was played with an initial endowment of 10 points and the amount sent by the trustor in the Trust Game was doubled, the maximum number of points that could be earned for each two players was 80, resulting in a maximum of  $32 \times 40 = 1280$  points for the whole session, of which half (640) corresponds to the victims. With an exchange rate of 35 Eurocents per point, this means that by robbing 50% the selected robber would receive a maximum transfer of €112, and proportionally for the other alternatives, as reported in Table 2.1. To estimate the *expected* transfers, we relied on the literature. We computed the expected earnings relying on the meta-analyses of Oosterbeek et al. (2004) for the Ultimatum Game and Johnson and Mislin (2011) for

the Trust Game.<sup>4</sup> In this way, we obtained proxies for the expected behavior in the sequence of games and were able to compute a numerical estimate of €195.17 for the joint earnings of the victims, leading to estimates (rounding up) of approximately €100, €66, and €20 for the decision to rob 50%, 33%, and 10%, respectively (as reflected in Table 2.1). The actual average joint earnings of victims in our sessions was €179.34, showing that our estimate was reasonably accurate.

The third and last part of the experiment consisted in a block of individual questions. First, participants were asked about their beliefs regarding the Big Robber question. Specifically, they were requested to indicate how many out of 100 people they believed would take 50%, 33%, 10%, or 0%. This belief elicitation question was not incentivized. Second, the participants were given the opportunity to donate a fraction of their total earnings (excluding the show-up fee) to a local charity organization. This was done to examine possible guilt feelings on the part of the robbers. We will discuss this question in Section 2.5 below. Last, the experiment included a questionnaire on demographic data (including gender, age, and field of studies).

After all three parts of the experiment were completed, payoffs were computed. First, the matching algorithm paired the decisions of participants in the six games and computed the generated income. Then, the “Big Robber” was randomly chosen among the 16 robbers and his or her decision was implemented. Then, the donation decision was individually applied to each participant’s earnings. Finally, all payoff-relevant information was presented

---

<sup>4</sup>The joint earnings in a 10-point Dictator Game are obviously constant and equal to 10. According to Oosterbeek et al. (2004), the rejection rate in implementations of the Ultimatum Game played in Germany is about 9.5%, hence the expected joint earnings in a 10-point Ultimatum Game are 9.05 points (10 points with a 90.5% acceptance rate). Johnson and Mislin (2011) contains a Trust Game experiment carried out in Germany with a factor of two, and where the trustor sent on average 58% of the initial endowment (Walkowitz et al., 2009). Hence the expected joint gains (which are independent of the decision of the trustee) for an initial endowment of 10 points are  $4.2 + 2 \times 5.8 = 15.8$ . This leads to a total of  $2 \times (10 + 9.05 + 15.8) = 69.7$  for each set of six two-player games, or  $(1/2) \times 16 \times 69.7 = 557.6$  for all 16 victims in a session. With an exchange rate of 35 Eurocents per point, this translates into €195.17, hence stealing 50% would result in an average transfer of €97.59, which we rounded up to €100.

to the participants. Payment was made anonymously (individually) in a separate room.

### 2.2.2 Power Analysis

Before running the experiment, we conducted a power analysis to determine the sample size. We focused on possible effects in the Dictator Game, and specifically the hypothesis that facing the robbing decision before the Dictator Game might affect the dictator’s decision. As a prior, we used the joint distribution of the meta-analysis of behavior in the Dictator Game by Engel (2011). Starting from this distribution we used three different treatment effect sizes and simulated how many observations would be needed to find a significant difference between the prior distribution and the distribution influenced by the robbing decision. We simulated distributions with deviations of 15, 25, and 30 percentage points from the prior at the focal points represented by transfers of 0 and 5 points (out of 10). Between these values we smoothed the distribution.

Of course, the distribution might be different across groups of participants making different robbing decisions. For example, participants taking 50% are likely to be more selfish than those taking 0%. We expected that robbing decisions would concentrate in two groups, with the remaining categories capturing only a small share.<sup>5</sup> In addition, we expected gender effects, and hence invited enough participants of each gender to ensure an equal amount of male and female robbers. Accounting for these factors, we computed that, with 40 observations for each gender and group of decisions in each treatment, a medium-sized treatment effect ( $\approx 25\%$ ) would still have a power of 80% when restricting to a specific gender. Assuming evenly-split decisions, this means 160 observations per treatment, with 80 men and 80 women each, for a total group size of 320 robbers (and hence 320 victims).

---

<sup>5</sup>An alternative would have been to conduct a pretest to determine base rates for the robbing decision. However, given the nature of this experiment, we decided not to conduct any pretest in order to prevent the possibility that future participants at our lab might have heard about the design.



### 2.2.3 Data

We conducted the experiment at the Cologne Laboratory for Economic Research (CLER). The experiment was programmed in z-Tree (Fischbacher, 2007) and participants were recruited from the student population of the University of Cologne using ORSEE (Greiner, 2015), excluding psychology students. There were 640 participants (334 females). We collected data in 20 sessions scheduled in four consecutive days, 10 according to the ex ante treatment, and 10 according to the ex post treatment.

Subjects belonging to the group of “robbers” were perfectly balanced by gender for each treatment and each session, resulting in a total of 80 male and 80 female ex ante robbers, and 80 male and 80 female ex post robbers. There were 72 male and 88 female victims in the ex ante sessions and 74 male and 86 female victims in the ex post sessions.

The average payoff from the six intermediate games was €11.15 (around \$12 at the time of the experiment), ranging from €4.20 to €19.60. Additionally, participants received a show-up fee of €2.50. Experimental sessions lasted around 50 minutes.

## 2.3 The Big Robber Question

### 2.3.1 Robbers’ Choices: To Rob or Not to Rob

The left-hand side of Figure 2.2 depicts all robbing decisions. A majority of robbers opted for personal gain at the expense of the 16 victims in their session. Out of 320 robbers, 180 (56.25% of the robbers) decided to steal the maximum, i.e. 50% of the victims’ earnings, while 86 further robbers (26.88%) decided to take 33%, only 47 robbers (14.69%) took just 10% (the minimum above zero), and a purely anecdotal 2.19% (just 7 robbers) declined to take anything.

The right-hand side of Figure 2.2 depicts the robbing decisions conditional on gender. There is almost no difference in the robber choices between males and females. Out of 160 male robbers, 94 (58.75%) took 50%, 41 (25.62%) took 33%, 21 (13.12%) took just 10%, and only 4 (2.50%) declined to take




Figure 2.2: Absolute Frequency of Robbing Choices.

*Notes.* The left-hand side shows the robbing decisions pooled for males and females. The right-hand side shows the robbing decisions split by gender.

money from the victims. Out of 160 female robbers, 86 (53.75%) took 50%, 45 (28.12%) took 33%, 26 (16.25%) took just 10%, and only 3 (1.88%) declined to rob. The distributions are not significantly different according to a  $\chi^2$  test ( $\chi^2(3) = 1.216$ ,  $p = 0.7491$ ).<sup>6</sup>

After the Big Robber was selected, there was actual stealing in all 20 sessions. The average robbing earnings (i.e. transfer due to the robbing decision) was €66.83, ranging from €17.85 to €97.65. Counting the earnings from the games, the 20 selected Big Robbers earned an average of €78.23, ranging from €26.25 to €110.25 (not including the show-up fee and the donation decisions). Hence, the experiment truthfully involved high stakes, since it was actually possible to walk away with over €100.<sup>7</sup>

<sup>6</sup>Of the 320 robbers in our experiment, 147 reported having economics-related majors. There were no significant differences between the distributions of robbing choices for them and for the 173 robbers who reported non-economics-related majors.

<sup>7</sup>The remaining 300 robbers who were *not* selected to be Big Robbers earned on average €11.08 (not including the show-up fee and the donation decisions; median €10.85,  $SD = 2.30$ , ranging from €4.20 to €18.20). Since all victims were robbed, their earnings were correspondingly lower (mean €7.19, median €7.00,  $SD = 2.20$ , ranging from €3.15 to €14.35).




Figure 2.3: Robbing Choices by Treatment.

*Notes.* The left-hand side shows relative frequency of robbers-50% by treatment. The right-hand side shows the complete histograms of robbing decisions split by treatment. Bars represent 1 Standard Error of the Mean. \*\*  $p < 0.05$ , test of proportions.

Since the vast majority of robbers took either 50% or 33%, for the subsequent analysis we divided the robbers in two groups. The “more selfish” ones are those who decided to rob 50%, hereafter denoted as *robbers-50%*. The “less selfish ones” are those who took less than 50%, that is, all those who decided to rob 33%, 10%, or nothing at all. We refer to this category as the *robbers-no-50%*.

The left-hand-side of Figure 2.3 depicts the relative frequency of robbers-50% in the ex ante and ex post treatments (recall that there are 160 observations in each treatment). In the ex ante treatment 62.50% of the robbers (100) took 50%, but in the ex post treatment only 50.00% of the robbers (80) took 50%. The difference is significant according to a test of propor-

tions ( $z = 2.254$ ,  $p = 0.0242$ ).<sup>8</sup> The right-hand side of the figure depicts the robbing decisions conditional on treatment. The distributions are significantly different according to a  $\chi^2$  test ( $\chi^2(3) = 9.157$ ,  $p = 0.0273$ ). We can conclude that if the Big Robber question is asked after the games have been played, there is a small shift away from robbing the maximum. However, as we expected the treatment difference is not large and it does not detract from the observation that stealing the maximum is the dominant mode of behavior.

Nevertheless, there are natural avenues of explanation for the treatment difference. A psychologically-motivated explanation might point out that ex post robbers could have developed some empathy towards the victims after playing the 6 games. A more economically-based explanation might point at income effects and attitudes toward risk. First, in comparison to ex ante robbers, ex post robbers were explicitly aware of at least part of the earnings from the games (e.g., they knew that the dictator and the trustor decisions would be implemented as given) and could have an expectation on the total earnings from the games. An increased awareness of those earnings is in practice an income effect which could decrease the motivation to steal. Second, since ex ante robbers had not experienced the six games when making the robbing decision, they could have a less-focused expectation on the earnings accruing from them, i.e. face a more risky prospect regarding those than the ex post robbers. Under risk aversion, this could lead to a larger appropriation decision, in practice compensating for the higher variance of earnings. A related explanation might be that ex ante robbers who took the maximum might be trying to insure themselves against the possibility of having low earnings in the intermediate games, while ex post robbers already had a personal estimate of possible earnings. However, in this case one could argue that robbers who took the maximum in the ex ante treatment should have made more selfish decisions in the Dictator Game than their ex post

---

<sup>8</sup>The difference misses significance if looking only at males: 63.75% (51) of male robbers took 50% in the ex ante treatment, versus 53.75% (43) in the ex post treatment (test of proportions,  $z = 1.285$ ,  $p = 0.1989$ ). Looking at females, the difference is weakly significant: 61.25% (49) took 50% in the ex ante treatment, versus 46.25% (37) in the ex post treatment (test of proportions,  $z = 1.903$ ,  $p = 0.0571$ ).

counterparts, since in this game they have actual control of their earnings, and this hypothesis is not confirmed by the data (see Section 2.4 below).

### 2.3.2 Decision Times

It is a long-standing, well-established observation that decisions where the decision-maker is “closer to indifference” (informally speaking) are harder and result in longer decision times (Dashiell, 1937; Mosteller and Nogee, 1951; Moyer and Landauer, 1967). A recent application to economic data has been provided by Krajbich et al. (2015). The basic observation resonates with the intuition that when alternatives are similarly desirable, the decision maker will be more likely to struggle with the decision and require more time to select a choice.

We relied on this logic to investigate the possible moral struggle faced by robbers. To this end, we recorded the decision times for the Big Robber question. Initially, robbers received an explanation on the robbing decision, detailing the possibility to take part of the victims’ earnings and explaining that a decision on how much to take would have to be made. As they clicked a continuation button, the table with the four possible alternatives was revealed (recall Table 2.1). On this table, robbers had to select a choice and then click a confirmation button. Decision times measure the time elapsed from the appearance of the table with the four alternatives to the clicking of the confirmation button. Hence they include the time needed for reading and understanding the table and making a decision.

Robbers who took the maximum decided significantly faster than other robbers. The mean decision time of robbers-50% ( $N = 180$ ) was 35.24 s, against 42.31 s for robbers-no-50% ( $N = 140$ ). Decision times are clearly significantly different according to a Mann-Whitney-Wilcoxon (MWW) test ( $z = -2.613$ ,  $p = 0.0090$ ).<sup>9</sup> Average decision times are depicted in the left-hand side of Figure 2.4.

The immediate interpretation is that the more selfish robbers faced a less severe moral struggle than the ones who decided not to rob the maxi-

---

<sup>9</sup>The result remains true if we compare robbers-50% only to those who took 33% ( $N = 86$ , mean decision time 43.90 s; MWW test,  $z = -2.568$ ,  $p = 0.0102$ ).




Figure 2.4: Decision Times for the Big Robber Question.

*Notes.* The left-hand side shows the average decision time of robbers-no-50% versus robbers-50%. The right-hand side shows the average decision time of the robbing decision by robbing choice. The bars represent here and subsequently one standard error of the mean. \*\*\*  $p < 0.01$ , MWW test.

mum. The right-hand side of Figure 2.4 depicts the decision times across the different robbing decisions. We observe an inverted U-shape which agrees with this interpretation. More extreme decisions (robbing the maximum or robbing only a little) should reflect a more clear preference, hence shorter decision times, while intermediate decisions might be the result of compromising and balancing tradeoffs (the decision maker is “closer to indifference”), resulting in longer decision times. However, the extreme data point corresponding to purely altruistic behavior (declining to rob) is of course purely anecdotal since there were almost no such observations.

The difference in decision times is also clearly significant when looking at the ex post treatment alone (robbers-50%,  $N = 80$ , average 35.64 s; robbers-no-50%,  $N = 80$ , average 45.53 s; MWW test,  $z = -2.503$ ,  $p = 0.0123$ ). However, although the direction persists, the difference is not significant for the ex ante treatment (robbers-50%,  $N = 100$ , average 34.92 s; robbers-no-50%,  $N = 60$ , average 38.02 s; MWW test,  $z = -0.652$ ,  $p = 0.5144$ ).

Decision times can also be used to study whether the processes underlying a certain decision are more or less intuitive or deliberative in nature, following dual-process theories (see, e.g., Alós-Ferrer and Strack, 2014, for a recent overview) which indicate that intuitive processes are faster.<sup>10</sup> Following this logic, Cappelen et al. (2016) observed that fair decisions in a Dictator Game took on average 38.4 s, whereas selfish decisions took 48.5 s. Relying on this statement, i.e. that fair decisions are faster, they argue that fair decisions might be more intuitive. If we followed the argument of Cappelen et al. (2016), we would have to conclude that the decision to rob as much as possible is more intuitive than the decision to refrain from such behavior. However, decisions with decision times as long as the ones studied here or in Cappelen et al. (2016) clearly always include a large amount of deliberation, and are hence not well-suited for the study of underlying processes, at least in a straightforward way. Inferences of process characteristics in these cases risk a reverse-inference fallacy, i.e. “intuitive” may mean “fast,” but this would not imply that “faster” means “more intuitive” (Myrseth and Wollbrant, 2016). Hence, we favor the simpler interpretation that faster decisions for the robbers who took the maximum indicate a reduced moral struggle in comparison with the robbers who partially overrode their impulse to take the maximum.

### 2.3.3 Beliefs

We elicited the participants’ beliefs by asking how many out of 100 participants they thought would choose to take 0%, 10%, 33%, and 50% of the victims’ earnings. From the resulting distributions, we computed the average percentage that each individual thought would be taken away from the victims. This average is increasing with the type of robber. The few robbers-0% (N=7) believed that on average only 16.73% of the victims’ earnings would be taken away. The averages increased to 26.41% for the robbers-10% (N=47), 33.21% for the robbers-33% (N=86), and 41.33% for the robbers-

---

<sup>10</sup>See Achtziger and Alós-Ferrer (2014), Alós-Ferrer et al. (2016a), and Alós-Ferrer and Ritschel (2018b) for examples of response-times studies of individual decisions in economic settings.

50% (N=180). There is a clear positive correlation between the robbing choice and the beliefs ( $\rho = 0.5928$ ,  $p < 0.0001$ ). The result persists for each gender (males,  $\rho = 0.6233$ ,  $p < 0.0001$ ; females,  $\rho = 0.5658$ ,  $p < 0.0001$ ) and treatment (ex ante,  $\rho = 0.5198$ ,  $p < 0.0001$ ; ex post,  $\rho = 0.6804$ ,  $p < 0.0001$ ). This finding suggests that participants had beliefs consistent with their actual behavior, that is, participants believed that, on average, other participants' behavior would be close to their own.

### 2.3.4 The Selfishness of Big Robbers

It is important to establish that the Big Robber Game measures social preferences and not a completely different construct. To this end, the intermediate games corresponded to the standard ones usually employed to study social preferences, i.e., the Dictator, Ultimatum, and Trust games.

**Dictator Game.** The cleanest measurement of social preferences as studied in the literature is provided by the Dictator Game, since there are no strategic concerns for the dictator decision. For this reason, we placed this decision at the very beginning of the block containing the intermediate games, ensuring that there would be no income effects or other carryover considerations at this point. Robbers who took the maximum were significantly less generous as dictators in the Dictator Game. Robbers-no-50% (N=140) sent on average 3.157 points while robbers-50% (N=180) only sent 1.267 points (see the left-hand extreme of Figure 2.5). The difference is highly significant according to a Mann-Whitney-Wilcoxon test ( $z = 8.210$ ,  $p < 0.0001$ ). The result persists when splitting the data by gender or treatment. Figure 2.6 displays the full distribution of dictators' decisions, separately for robbers-no-50% and robbers-50%. The two distributions are significantly different ( $\chi^2(8) = 79.853$ ,  $p < 0.0001$ ). We conclude that robbers who took the maximum are more selfish, as measured by the Dictator Game, than other robbers.

**Ultimatum Game.** Participants played the Ultimatum Game twice, once as proposer and once as responder. Of course, offers in the Ultimatum Game



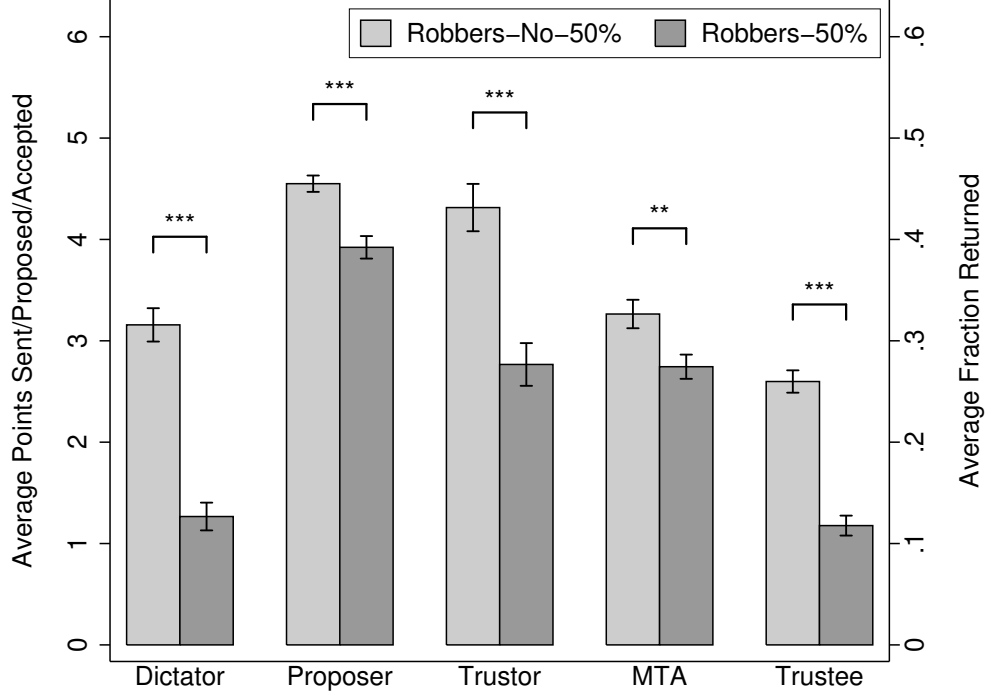


Figure 2.5: Robbers' Decisions in the Intermediate Games.

*Notes.* Comparison of decisions of robbers-50% and robbers-no-50% in the Dictator Game, the Ultimatum Game, and the Trust Game. \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , MWW test.

were larger than in the Dictator Game, due to the strategic aspect of the proposer's decision (avoiding rejection). However, again robbers who took the maximum were revealed to be more selfish than other robbers, this time by their behavior as proposers in the Ultimatum Game. Robbers-50% ( $N = 180$ ) offered 3.922 points on average while robbers-no-50% ( $N = 140$ ) offered 4.550 points (second decision illustrated in Figure 2.5). The difference is highly significant (MWW test,  $z = 4.141$ ,  $p < 0.0001$ ). The effect persists when splitting the data by gender or treatment.

We elicited responder behavior in the Ultimatum Game using the strategy method, yielding 11 decisions per participant (whether to accept or reject proposals of 0, 1, ..., 10 points out of 10). To analyze responder behavior, we computed the smallest accepted offer (or Minimum Threshold of Acceptance, MTA), excluding participants who switched between rejection and

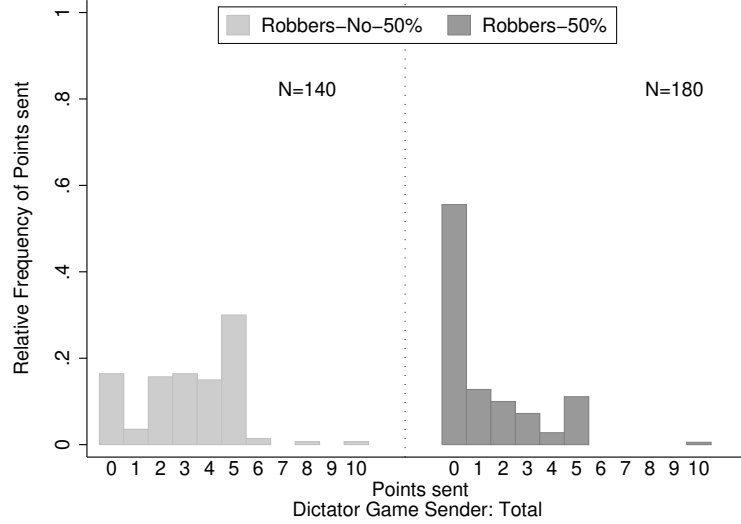


Figure 2.6: Behavior of Robbers in the Dictator Game.

*Notes.* Relative frequency of points sent by the robbers-no-50% and robbers-50% in the Dictator Game.

acceptance multiple times.<sup>11</sup> Robbers-50% ( $N = 178$ ) had an average MTA of 2.730 points while robbers-no-50% ( $N = 135$ ) had a significantly higher MTA of 3.163 points (MWW test,  $z = 2.437$ ,  $p = 0.0148$ ; fourth decision illustrated in Figure 2.5). As ought to be expected, robbers who took the maximum are willing to accept lower offers since a higher level of selfishness goes hand-in-hand with giving priority to purely monetary concerns. The behavior of responders in the Ultimatum Game again indicates that for robbers who took the maximum, social preferences are less marked than for other robbers.

**Trust Game.** Participants played the Trust Game twice, once as trustor and once as trustee. The decision made as trustor was already the third down the line in the block of intermediate games. However, there were again clear differences, with robbers who took the maximum trusting less than other robbers. Robbers-no-50% ( $N = 140$ ) sent on average 4.314 points to the trustee while robbers-50% ( $N = 180$ ) sent on average only 2.767

<sup>11</sup>Out of 640 participants only 19 participants (2.97%) switched more than once.

points. This difference is highly significant (MWW test,  $z = 5.636$ ,  $p < 0.0001$ ; third decision illustrated in Figure 2.5). The effect persists when conditioning on either treatment or gender. Once again, this is evidence that social preferences are less marked for robbers who took the maximum, in the sense that they trust less, presumably because they expect *others* to be more selfish.

We elicited trustee behavior in the Trust Game using again the strategy method, yielding 10 decisions per participant (how much to send back if the trustor sent 1, ..., 10 points out of 10). To analyze trustee behavior, we computed the fraction that was sent back aggregated over all 10 decisions. We again find clear differences. Robbers-no-50% (N=140) returned on average 25.99% of the received points, against only 11.77% for robbers-50%. This difference is highly significant (MWW test,  $z = 8.995$ ,  $p < 0.0001$ ; right-most decision illustrated in Figure 2.5). The effect also persists when conditioning on gender or treatment. Strictly speaking, the decision of the trustee is free of strategic components and is formally equivalent to a Dictator Game. Hence, again we see that robbers who took the maximum behave more selfishly than other robbers.

**Average behavior of robbers.** At the same time, behavior in the Dictator, Ultimatum, and Trust games was well within the standard ranges reported in the literature. The average offer by robbers in the Dictator Game was 20.94%, which is quite close to the grand average of 24.7% reported by Engel (2011) for students. The average offer by robbers in the Ultimatum Game was 41.97%, which was rather close to the grand average of 40.54% (71 studies) reported by Oosterbeek et al. (2004). The average MTA of robbers in the Ultimatum Game was 29.17%, which is within the ranges reported in other Ultimatum Game experiments using the strategy method. For example, Cappelletti et al. (2011) report mean MTAs of 23.00% and 31.42% for an endowment of €15 and €7, respectively. Declerck et al. (2009) found MTAs

of 23.10% and 30.00% when subjects played an Ultimatum Game after and before being matched with another subject, respectively.<sup>12</sup>

For the Trust Game, it is harder to compare our data to the literature because published experiments have a higher design variance than those employing Dictator or Ultimatum games. In our experiment, robbers sent on average 34.44% of the available resources as trustors and returned 17.99% as trustees. These levels seem to be within the standard ranges in the literature as reported, e.g., in the meta-analysis of Johnson and Mislin (2011), although on the lower range of trust and (especially) trustworthiness.<sup>13</sup> The closest design we could find in a published study is Bellemare and Kröger (2007), which employed the strategy method in a laboratory Trust Game played with students in central Europe where transfers were doubled. In that study, trustors sent on average 30.58% of their resources, and trustees send back around 24% of what was available.<sup>14</sup> Our results are also close to those of Eckel and Petrie (2011).

**The Big Robber Game and Social Preferences.** In summary, results are strikingly consistent across all games. Robbers who took the maximum gave less in the Dictator Game, sent less money back in the Trust Game as trustees, made lower offers in the Ultimatum Game, trusted less as trustors in the Trust Game, and accepted lower offers as responders in the Ultimatum Game. In view of this evidence, it can be concluded that the Big Robber Game, as intended, measures a correlate of social preferences *as they have been discussed in the literature until now*, and not a different construct.

---

<sup>12</sup>Armantier (2006) finds slightly higher mean MTAs, ranging from 31.80% to 38.00%. McLeish and Oxoby (2011) find slightly lower MTAs of 26.10% in the baseline condition of their experiment.

<sup>13</sup>In view of the literature, the fact that our results for the Trust Game are in the lower range might be unsurprising. We used a multiplying factor of two, while many studies use a factor of three. Johnson and Mislin (2011) suggest that a lower factor reduces trustors' transfers and trustees' returns. Casari and Cason (2009) argue that the strategy method, which we used, reduces transfers of trustees (although Brandts and Charness, 2011 find no difference). Burks et al. (2003) argue that playing both roles in the Trust Game reduces both overall trust and overall reciprocity; in our setting, players first made a trustor decision and were asked for a trustee decision later on.

<sup>14</sup>We thank Sabine Kröger for providing the aggregate statistics. In that study, however, trustees had an additional endowment.

We remark that the fact that behavior is qualitatively consistent between the Big Robber Game and all other games is of independent interest, since such associations should not be taken for granted at the individual level. Blanco et al. (2011) estimated two parameters of inequality aversion in a within-subject design where participants played several games, including Dictator and Ultimatum games, and found remarkably low correlations within subjects. Galizzi and Navarro-Martínez (2018) found low correlations between behavior in Dictator, Ultimatum, and Trust games in the lab and social behavior in the field. This is in line, e.g., with Stoop et al. (2012), who found clear divergences in cooperation levels between lab and field settings. However, other authors have found positive associations across domains. Dariel and Nikiforakis (2014) found a qualitative correlation for prosocial behavior across games, with participants behaving cooperatively in a public-good game reciprocating higher wages with higher effort levels in a gift-exchange game. Fisman et al. (2007) find a strong positive association between preferences for giving and social preferences referred to distributions among others (which do not affect the own payoffs).

## 2.4 Does the Big Robber Question Affect Behavior?

In this section, we focus on treatment effects. The Big Robber decision is clearly not one that can be taken lightly. It is conceivable that behavior in the Dictator, Ultimatum, and Trust games was affected by having answered this question beforehand. This is precisely the reasoning which brought us to include the ex ante and ex post treatments.

There are indeed a number of natural hypotheses grounded on psychological research. According to the “what-the-hell effect,” an initial loss of self-control can lead to a modal change in which all pretenses are abandoned, for example in the dieting domain (Polivy and Herman, 1985). It has been argued (Achtziger et al., 2015b, 2016, 2018) that prosocial behavior requires self-control in order to override selfish impulses. A natural implication then would be that, if a participant has revealed being selfish in the ex ante Big Robber decision, then he will behave less prosocially afterwards than controls

(ex post robbers). However, the opposite hypothesis could also be justified. According to the “transgression-compliance effect,” people who believe that they have harmed someone show an increased willingness to perform unrelated good deeds later on (Carlsmith and Gross, 1969), as if the latter could compensate the former (see Gneezy et al., 2014, for a recent illustration). This effect might reveal a mechanism to reduce experienced guilt. In our context, it would imply that if a participant has revealed to be selfish in the ex ante Big Robber decision, he or she should behave more prosocially than controls afterwards.

A further possible hypothesis concerns robbers who refrained from taking the maximum. According to the “moral credentialing” effect, humans often act as if an initial good behavior (even an exogenously induced one) provides a license to misbehave later on (Monin and Miller, 2001). For example, people purchasing “green” products are later on less likely to share in a Dictator Game, and more likely to cheat on a task to increase their gains (Mazar and Zhong, 2010; Zhong et al., 2010). According to moral credentialing, robbers who did *not* behave completely selfishly in the ex ante Big Robber decision should behave more selfishly afterwards than controls.

#### 2.4.1 The more selfish robbers

We start with the behavior of robbers who took the maximum possible, i.e. robbers-50%, for which the opposed hypotheses derived from the what-the-hell effect and the transgression-compliance effect apply. In the Dictator Game, ex ante robbers-50% (N=100) gave on average 1.310 points, while ex post robbers-50% (N=80) gave on average 1.213 points. The difference is not statistically significant (MWW test,  $z = 0.695$ ,  $p = 0.4873$ ). However, looking at genders separately allows us to see a different picture. Figure 2.7 shows the relative frequency of points sent by the robbers-50% in the Dictator Game split by gender. In the ex ante treatment, female robbers-50% (N=49) sent on average 1.612 points while in the ex post treatment female robbers-50% (N=37) sent 0.811 points. This difference, which agrees with the possibly guilt-induced transgression-compliance effect, is highly significant according




Figure 2.7: Behavior of Robbers-50% in the Dictator Game.

*Notes.* Relative frequency of points sent by the robbers-50% in the Dictator Game split by treatment and gender (left-hand side: males; right-hand side: females).

to a Mann-Whitney-Wilcoxon test ( $z = 2.442$ ,  $p = 0.0146$ ). In contrast, looking at male robbers only, there is no evidence of the transgression-compliance effect. Ex ante male robbers-50% ( $N=51$ ) sent on average 1.020 points in the Dictator Game, compared to 1.558 points sent by ex post male robbers-50% ( $N=43$ ). The difference goes in the opposite direction, that is, the one predicted by the what-the-hell effect, but does not reach significance (MWW test,  $z = -1.348$ ,  $p = 0.1775$ ).<sup>15</sup>

Moving to the behavior of robbers-50% in the Ultimatum Game, offers in the ex ante treatment ( $N = 100$ , mean 3.760) were significantly lower than the offers of the robbers-50% in the ex post treatment ( $N = 80$ , mean 4.125 points; MWW test,  $z = -1.808$ ,  $p = 0.0705$ ). This is in agreement with the what-the-hell effect. However, this finding is driven by male behavior. Figure 2.8 shows the relative frequency of points proposed by the robbers-50% in the Ultimatum Game split by gender. Male robbers-50% in the ex ante treatment ( $N = 51$ ) offered 3.765 points on average, which was significantly lower than the average 4.209 points offered by the male robbers-50% ( $N = 43$ ) in the ex post treatment (MWW test,  $z = -2.155$ ,  $p = 0.0312$ ). In contrast, there

<sup>15</sup>Out of  $N = 640$  participants, only 3 gave 10 points in the Dictator Game, all of them male: one victim and two robbers. One of the robbers took 50%. Excluding this atypical individual from the test, the average points sent by ex ante male robbers-50% ( $N=50$ ) drop to 0.840 points and the difference between treatments becomes more clear, but still misses significance (MWW,  $z = -1.561$ ,  $p = 0.1185$ ).




Figure 2.8: Behavior of Robbers-50% in the Ultimatum Game.

*Notes.* Relative frequency of points proposed by the robbers-50% in the Ultimatum Game split by treatment and gender (left-hand side: males; right-hand side: females).

were no significant differences for females (ex ante,  $N = 49$ , average 3.755; ex post,  $N = 37$ , average 4.027; MWW test,  $z = -0.319$ ,  $p = 0.7495$ ).

The decision in the Dictator Game and the proposer decision in the Ultimatum Game were the two first decisions in the sequence of games. For those, as commented above we find small effects which depend on gender. The responder decision in the Ultimatum game and the two decisions (as trustor and as trustee) in the Trust Game are the last three decisions in the sequence of games, and behavior might have been affected by the most recent decisions (in the Dictator and Ultimatum games) more than by the previous Big Robber question. Hence, we did not expect strong treatment differences for these decisions. Indeed, we find no significant differences, even when looking at genders separately, in Minimum Thresholds of Acceptance in the Ultimatum Game or in behavior in the Trust Game between ex ante and ex post robbers-50%.

We conclude that the psychological effects caused by the Big Robber decision on more selfish robbers might be present but are, first, generally weak, and, second, dependent on gender. While women seem to be affected by guilt-related considerations as the transgression-compliance effect, men seem to show the opposite tendency as predicted by the what-the-hell effect.



### 2.4.2 The less selfish robbers

The behavior of robbers who declined to take the maximum possible, i.e. robbers-no-50%, might be affected by the moral-credentialing effect, that is, since they did not behave (maximally) selfishly in the Big Robber decision, they might have behaved more selfishly in the subsequent games. However, we find no evidence for such an effect in the Dictator Game. Ex ante robbers-no-50% ( $N = 60$ ) gave an average of 3.200 points, which is not significantly different from the average of 3.125 given by ex post robbers-no-50% ( $N = 80$ ) (MWW test,  $z = -0.172$ ,  $p = 0.8633$ ). The differences remain insignificant when splitting the sample by gender.

In the Ultimatum Game, there are also no differences for the full sample. Ex ante robbers-no-50% ( $N = 60$ ) proposed an average of 4.500 points, which is not significantly different from the average of 4.588 proposed by ex post robbers-no-50% ( $N = 80$ ) (MWW test,  $z = -0.148$ ,  $p = 0.8820$ ). However, splitting the sample by gender reveals (marginally) significant and opposed effects. Ex ante female robbers-no-50% ( $N = 31$ ) offered an average of 4.323 points while ex post female robbers-no-50% ( $N = 43$ ) offered an average of 4.767 points (MWW test,  $z = -1.701$ ,  $p = 0.0889$ ). This difference is in the direction predicted by the moral-credentialing effect. In contrast, the treatment difference for male robbers-no-50% went in the opposite direction, with significantly higher offers in the ex ante treatment ( $N = 29$ , mean 4.690) compared to the ex post treatment ( $N = 37$ , mean 4.378; MWW test,  $z = 1.657$ ,  $p = 0.0974$ ).

As in the case of robbers-50% (and most likely for identical reasons), we do not find any significant treatment differences for robbers-no-50% as responders in the Ultimatum Game or in either role in the Trust Game.

We conclude that, as in the case of more selfish robbers, the psychological effects caused by the Big Robber decision on less selfish robbers are generally weak and depend on gender. While women seem to be affected by the moral-credentialing effect in the Ultimatum Game, men show the opposite tendency.

### 2.4.3 The victims

All victims in the experiment ( $N = 320$ ) learned about the possibility that part of their earnings could be taken away before making their decisions in the sequence of games. It is natural to ask whether this information affected their behavior. For instance, one could speculate that, knowing that part of their earnings might be lost, they might have become more self-centered. One could also speculate with possible negative reciprocity against the robbers, although this is an unlikely motivation because victims did not know whether the partners they would be matched with to play the Dictator, Ultimatum, and Trust games would be victims or robbers (random matching).

The cleanest comparison to find possible effects in victims' behaviors is with ex post robbers ( $N = 160$ ). These robbers were informed about the Big Robber decision after they played the Dictator, Ultimatum, and Trust games, hence their behavior in those games could not be affected by the latter question, and they are in practice a control group. Comparing the behavior of victims to the behavior of ex post robbers, however, reveals no significant differences. In the Dictator Game, victims sent an average of 2.216 points, compared to 2.169 sent by ex post robbers (MWW test,  $z = 0.217$ ,  $p = 0.8282$ ). In the Ultimatum Game, victims proposed an average of 4.359 points, compared to 4.356 proposed by ex post robbers (MWW test,  $z = -0.049$ ,  $p = 0.9609$ ). As responders, the victims' average MTA ( $N = 308$ ) was 2.912, compared to an average MTA of 2.891 for ex post robbers ( $N = 156$ ; MWW test,  $z = -0.042$ ,  $p = 0.9662$ ). In the Trust Game, victims sent an average of 4.013 points, compared to 3.575 sent by ex post robbers (MWW test,  $z = 1.420$ ,  $p = 0.1556$ ). As trustees, victims sent back an average of 20.21%, compared to 18.75% sent back by ex post robbers (MWW test,  $z = 0.797$ ,  $p = 0.4252$ ). Splitting the tests by gender revealed only one significant effect, namely that male victims ( $N = 146$ ) sent an average of 4.480 points while male robbers ( $N = 80$ ) sent an average of 3.625 points, with the difference being significant according to a Mann-Whitney-Wilcoxon test ( $z = 1.793$ ,  $p = 0.0730$ ). However, none of the other nine within-gender comparisons for the five decisions was significant.

We conclude that victims' behavior was unaffected by the knowledge that a part of their earnings might be taken away by robbers.

## 2.5 Donations

Since the Big Robber decision is a high-stakes one, it is natural to ask whether, in some or even most of the cases, the decision to take the maximum might have led to feelings of guilt and regret. There is evidence from social psychology pointing out that guilt can motivate prosocial behavior perceived as reparative, as a way to appease guilty feelings (Malinowski and Smith, 1985; Giner-Sorolla, 2001; Zemack-Rugar et al., 2007). In particular, Lindsey (2005) showed that guilt is associated with increased charity donations. This is also consistent with Andreoni (1989), who argued that people derive direct utility from the act of giving to a charity, with Gneezy et al. (2014), who showed that people were then more likely to donate to charity after making an immoral choice, and with Andreoni et al. (2017), who showed that people give more if exposed to stimuli which activate empathy ("Please give!") but at the same time try to avoid exposure to such stimuli.<sup>16</sup>

We hence decided to include a final question giving participants the opportunity to donate part of their earnings to a charity, a local animal shelter.<sup>17</sup> In the final questionnaire, we also asked for the valuation of the charity organization on a scale from 1 = very bad to 10 = very good. The average evaluation was 6.46 ( $N = 640$ ,  $SD = 2.66$ ), indicating a generally positive view.

The form of the question was as follows. "You can donate a fraction of your total earnings (excluding the show-up fee) to a charity organization (name of animal shelter). How much do you want to donate?" The question was posed after all decisions had been made, but before the actual Big Robber

---

<sup>16</sup>DellaVigna et al. (2012) pointed out that social pressure might be an additional motive for giving which needs to be dissociated from altruism. In our setting social pressure plays no role since all decisions are made anonymously. However the need to appease feelings of guilt might be seen as individual pressure coming from self-perception after facing the Big Robber decision.

<sup>17</sup>The charity was the "Cologne Animal Protection Club of 1868."

of the session was selected and the game decisions were implemented. On average, subjects donated a fraction of 5.31% of their earnings, for a total joint donation of €299.6. Taken as a whole, that is, not differentiating among those who took the maximum and those who did not, the donations of robbers did not differ from those of the victims. Victims ( $N = 320$ ) donated on average 6.10% while robbers (pooled,  $N = 320$ ) donated on average 4.52%. The difference is not significant (MWW test,  $z = -0.122$ ,  $p = 0.9029$ ).

We hypothesized that, if feelings of guilt or regret were associated with the decision to take the maximum, robbers who took 50% would donate more than others. On the contrary, if the decision to take the maximum simply reflects a stronger tendency to rely on pure self-interest, we should observe lower donations. We found that robbers who took the maximum donated significantly less than others. Robbers-50% ( $N = 180$ ) donated an average of only 2.62%, compared to 6.96% donated by robbers-no-50% ( $N = 140$ ). The difference is both substantial and highly significant according to a Mann-Whitney-Wilcoxon test ( $z = -5.934$ ,  $p < 0.0001$ ). The effect remains for each treatment. Ex ante robbers-50% ( $N = 100$ ) donated only 3.12%, while ex ante robbers-no-50% ( $N = 60$ ) donated 6.20% (MWW test,  $z = -3.339$ ,  $p = 0.0008$ ). Ex post robbers-50% ( $N = 80$ ) donated a mere 2.00%, while ex post robbers-no-50% ( $N = 80$ ) donated 7.54% (MWW test,  $z = -4.982$ ,  $p < 0.0001$ ). The left-hand side of Figure 2.9 depicts the average donations of robbers-no-50% and robbers-50% for both treatments. The effect also persists when looking at each gender separately.<sup>18</sup>

A possible criticism of this test is that robbers who took 50% expected higher earnings than those who did not, hence they might have tried to adjust down the donation, expressed as a percentage, while still donating a larger absolute amount. This is not the case. We computed the expected donation under the extremely generous assumption that each individual robber might

---

<sup>18</sup>Pooling all robbers together, there was no significant difference in donations across treatments. There is a small gender difference in donations, with male robbers donating less (3.91%) than female robbers (5.13%). The difference is marginally significant (MWW test,  $z = -1.651$ ,  $p = 0.0988$ ). This gender difference is also found among the victims, with male victims ( $N = 146$ ) donating an average of 3.29%, compared to an average of 8.46% by female victims ( $N = 174$ ). The difference is highly significant (MWW,  $z = -3.691$ ,  $p = 0.0002$ ).

have assumed that he or she would indeed be the selected one (rather than using a probability of  $1/16$ ). Under this assumption, the expected donation of robbers-50% would have been €2.89, compared to an expected donation of robbers-no-50% of €3.80. The difference is again highly significant (MWW test,  $z = -4.114$ ,  $p < 0.0001$ ). That is, no matter how generous the criterion, robbers who took the maximum also donated less in absolute terms.<sup>19</sup> The right-hand side of Figure 2.9 shows the average donations in percentages (left axis) and the expected donations in absolute terms (right axis) for the different robber groups and the victims.

In conclusion, there is no evidence of guilt or regret as captured by the donation decision. One could of course speculate that the donation decision is only moral at an abstract level in the sense that it does not affect the victims. Also, one could invoke a different version of the what-the-hell effect to explain the results. The simplest explanation, however, is that robbers who took the maximum are just more selfish in all dimensions than those who did not.

It is reasonable to ask whether the differences in donations were caused by differences in evaluations, or whether they were rationalized ex post by providing lower evaluations of the charity. In both cases, we would expect robbers-50% to provide worse evaluations. This is, however, not true. Robbers-50% ( $N = 180$ ) rated the charity on average with 6.69, while robbers-no-50% ( $N=60$ ) rated it lower, at 6.28. The difference goes in the opposite direction but is not significant (MWW test,  $z = 1.506$ ,  $p = 0.1320$ ).<sup>20</sup>

---

<sup>19</sup>However, robbers as a whole donated more in absolute terms than victims, which is not surprising since they simply earned substantially more on average. Average expected donations of the victims were €0.63, compared to €3.29 for the robbers. This difference is highly significant (MWW test,  $z = -4.762$ ,  $p < 0.0001$ ).

<sup>20</sup>Interestingly, ex ante robbers-50% ( $N = 100$ ) rated the charity on average with 6.91, while robbers-no-50% ( $N = 60$ ) rated it at 6.13, which is significantly lower (MWW test,  $z = 1.835$ ,  $p = 0.0666$ ). Hence, ex ante robbers who took the maximum donated less in spite of the fact that they valued the charity better. In the ex post treatment the valuations of the robbers-50% ( $N = 80$ , mean 6.43) were not significantly different from the valuations of the robbers-no-50% ( $N = 80$ , mean 6.39; MWW,  $z = -0.100$ ,  $p = 0.9205$ ).

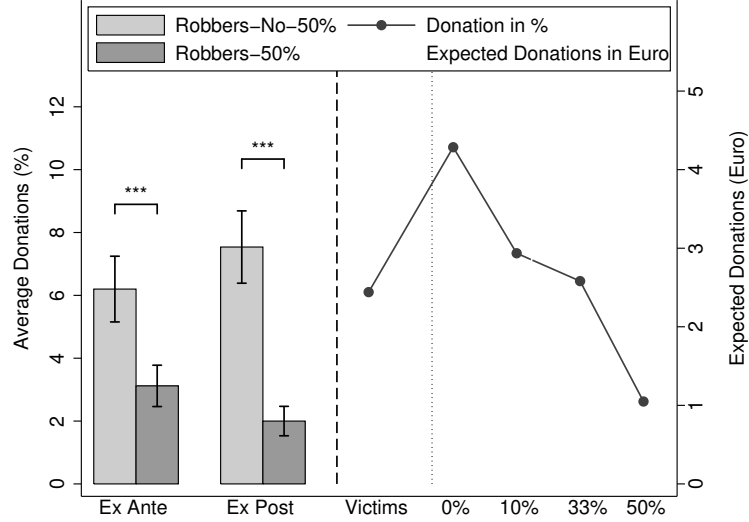


Figure 2.9: Average donations in percentage and expected donations in Euros.

*Notes.* The left-hand side shows the donations of robbers-no-50% and robbers-50% in both treatments. The right-hand side shows the donations in percentage (left axis) and the expected donations in Euros (right axis) for each robbing group and victims. \*\*\*  $p < 0.01$ , MWW test.

## 2.6 Models of Social Preferences

In this section, we examine whether the behavior we observe is compatible with received models of social preferences, focusing on the models of Fehr and Schmidt (1999) (hereafter FS), Bolton and Ockenfels (2000) (hereafter BO), Charness and Rabin (2002) (hereafter CR), and Alger and Weibull (2013) (hereafter AW). The strategy of analysis is a simple out-of-sample exercise as follows. Within the sequence of games, the very first decision of each participant corresponded to that of a dictator in the Dictator Game. From this decision, we deduce the individual parameters of the utility functions proposed in each of the models by FS, BO, CR, and AW. On the basis of these parameters, we derive the predicted behavior in the Big Robber Game for each participant and model. Finally, we compare predicted with actual behavior and examine the “fit” of the models, simply by examining the percentage of decisions in our sample of 320 robbers which are compatible with the model predictions.

Let  $x_r \in [0, 10]$  be the amount sent by the dictator to the receiver. The model of CR reduces to the following parametric family of utility functions for the dictator.

$$U_D^{CR}(x_r) = \begin{cases} 10 - x_r - \rho \cdot (10 - 2x_r) & \text{if } x_r \leq 5 \\ 10 - x_r - \sigma \cdot (10 - 2x_r) & \text{if } x_r > 5 \end{cases}$$

where  $\rho$  and  $\sigma$  are parameters capturing distributional preferences (an additional parameter in the formulation of CR captures reciprocity considerations which play no role here).

This model encompasses the one of FS, with  $\beta = \rho$  being the parameter for advantageous inequality and  $\alpha = -\sigma$  the one for disadvantageous inequality. The model of FS further constrains  $\sigma < 0 < \rho < 1$ . In our setting, the only difference between both models is that with the additional constraints, FS' model cannot explain dictator decisions giving strictly more than 5 to the receiver; as a consequence, strictly speaking 5 out of our 320 observations would remain unclassified in the FS case. Otherwise, the analysis for FS and CR is identical.

In our sample, 123 (38.44%) of the robbers gave 0 in the Dictator Game, which implies  $\rho \leq 1/2$  and  $\sigma \leq 1 - \rho$ ; 130 (40.63%) robbers gave between 1 and 4 in the Dictator Game, which implies  $\rho = 1/2$  and  $\sigma \leq 1/2$ ; 62 (19.38%) robbers gave exactly 5, which implies  $\rho \geq 1/2$  and  $\sigma \leq 1/2$ ; 3 (0.94%) robbers gave between 6 and 9, which implies  $\rho \geq 1/2$  and  $\sigma = 1/2$ ; and 2 (0.63%) robbers gave exactly 10, implying  $\sigma \geq 1/2$  and either  $\rho \geq 1/2$  or  $\rho < 1/2$  together with  $\sigma \geq 1 - \rho$ . The first three columns of Table 2.2 summarize these observations.

Let  $E[\Pi_V]$  denote the expected income of all victims in a session from the intermediate games (from the point of view of the individual robber) and  $E[\Pi_R]$  the own (robber) expected income from those games. Let  $n_V$  be the number of victims, and denote by  $p$  the share the robber chooses to rob. A given robber would expect earnings  $E[\Pi_R] + pE[\Pi_V]$ , and will expect average

earnings of the victims to be  $(1 - p)(1/n_V)E[\Pi_V]$ . For the Big Robber, the CR (or FS) model implies the following utility function.

$$\begin{aligned} U_R^{CR}(p) &= (E[\Pi_R] + pE[\Pi_V]) - \rho \left( E[\Pi_R] + pE[\Pi_V] - (1 - p)\frac{1}{n_V}E[\Pi_V] \right) \\ &= (1 - \rho)E[\Pi_R] + \frac{1}{n_V}E[\Pi_V] (p \cdot (n_V - \rho(n_V + 1)) + \rho). \end{aligned}$$

Note that the parameter  $\sigma$  plays no role because, in expected terms, the robber could not be worse off than the victims.

Denote by  $V$  the ex ante expected earnings of all  $n_V$  victims in a session as communicated in the experiment. For an ex ante robber, since the decision to rob or not was made before the intermediate games were explained and played, all the player could deduce was  $E[\Pi_V] = V$  and  $E[\Pi_R] = V/n_V$ . Hence, for an ex ante robber, the expression above simplifies to

$$U_R^{CR}(p) = \frac{1}{n_V}V [1 + p \cdot (n_V - \rho(n_V + 1))].$$

For an ex post robber, however, the intermediate games had already been played when the Big-Robber decision was made, and expectations could have been adjusted. For instance, those robbers knew that they would be receiving at least  $20 - x_r - x_r^T$ , where  $x_r$  was their decision in the Dictator Game and  $x_r^T$  the decision as a trustor in the Trust Game, and accordingly could adjust  $E[\Pi_V]$  down. However, the optimum of  $U_R^{CR}(p)$  is independent of the exact values of expectations, and in particular there should be no difference between ex ante and ex post robbers according to the CR model. For all  $\rho < \frac{n_V}{n_V + 1} = \frac{16}{17}$  robbers should take 50%, while for  $\rho = \frac{16}{17}$  robbers are indifferent among taking 50%, 33%, 10%, or 0%, and for  $\rho > \frac{16}{17}$  robbers should take 0%. Hence, all robbers who gave strictly less than 5 points in the Dictator Game should take 50%, and we have no prediction for those who gave exactly 5 points or more than 5 points, in the sense that for  $\rho > 1/2$ , their optimum might be to take 50%, 0%, or correspond to full indifference among all possibilities. The second to last column of Table 2.2 summarizes these predictions and the amount of observations compatible with these predictions.



Because of the linearity of their formulation, the models of FS and CR can only explain intermediate values (1 to 4 points) in the Dictator Game through indifference. The model of BO incorporates a nonlinear term to deal with this and related issues, the typical implementation relying on a quadratic functional form. Applied to the Dictator Game, the corresponding utility function is as follows.

$$U_D^{BO}(x_r) = a(10 - x_r) - \frac{b}{2} \left( \frac{10 - x_r}{10} - \frac{1}{2} \right)^2 = a(10 - x_r) - \frac{b}{2} \left( \frac{5 - x_r}{10} \right)^2$$

where again  $x_r$  is the amount the dictator sends to the receiver, and  $a$  and  $b$  are parameters weighting the utility of the own payoff and disutility of the relative payoff, for which BO assume only  $a \geq 0$  and  $b > 0$ . A player's type is fully characterized by the ratio  $a/b$ . The fourth column of Table 2.2 shows the ranges of  $a/b$  implied by the observed behavior of robbers in the Dictator Game.<sup>21</sup>

For the Big Robber, the BO model implies the following utility function.

$$U_R^{BO}(p) = a(E[\Pi_R] + pE[\Pi_V]) - \frac{b}{2} \left( \frac{E[\Pi_R] + pE[\Pi_V]}{E[\Pi_R] + E[\Pi_V]} - \frac{1}{n_V + 1} \right)^2. \quad (2.1)$$

Again, for an ex ante robber  $E[\Pi_V] = V$  and  $E[\Pi_R] = V/n_V$ , hence the expression above simplifies to

$$U_R^{BO}(p) = a \frac{V}{n_V} (1 + p \cdot n_V) - \frac{b}{2} \left( \frac{n_V}{n_V + 1} p \right)^2.$$

Since experimental sessions involved 32 participants, half of which were victims, we have that  $n_V = 16$ . The expected income of all victims was €200, corresponding to 571.43 points (recall that the Dictator Game decisions was how to split 10 points). With these values, a direct computation shows that  $U_R^{BO}(p)$  for ex ante robbers is maximized at  $p = 0.5$  or above for any ratio  $a/b$  larger than 0.00064. In view of the values derived from behavior in

<sup>21</sup>Each possible decision implies a different range, but all decisions in the 1 – 4 range, and analogously for the 6 – 10 range, yield the same predicted behavior in the Big Robber Game, hence we do not differentiate them in the table.

the Dictator Game, all ex ante robbers who gave less than 5 points in the Dictator Game are predicted to take 50%, those who gave 6 points or more are inconsistent with the model, and for those who opted for an equal split, the possible values derived for  $a/b$  do not allow a precise prediction (hence their behavior in the Big Robber Game is always aligned with the model).

Ex post robbers form their expectations about their own payoff from the intermediate games depending on their actions. In our design, the robbers play against a different player in each intermediate game but do not know if the opponent is a robber or victim. If the robbers played only against robbers, the expectations about the victims earnings should be unaffected,  $E[\Pi_V] = V$ . In the opposite extreme, all six interactions of the robber in the intermediate games might have affected victims. It follows that on average, at most one victim (playing in six games) was affected by the robber decisions, implying that  $E[\Pi_V] \in [(n_V - 1)V/n_V, V]$ . From equation (2.1) above, we obtain that the (unconstrained) maximum is reached at

$$p^* = \frac{a}{b} \left( \frac{1}{E[\Pi_V]} (E[\Pi_R] + E[\Pi_V])^2 \right) + \frac{1}{n_V + 1} \left( \frac{E[\Pi_R]}{E[\Pi_V]} + 1 \right) - \frac{E[\Pi_R]}{E[\Pi_V]}$$

and the derivative of this expression with respect to  $E[\Pi_V]$  is

$$\frac{\partial p^*}{\partial E[\Pi_V]} = \frac{a}{b} \left( 1 - \left( \frac{E[\Pi_R]}{E[\Pi_V]} \right)^2 \right) + \frac{n_V}{n_V + 1} \frac{E[\Pi_R]}{E[\Pi_V]^2}.$$

We know that  $E[\Pi_R] \in [20 - x_r - x_r^T, 70 + x_r^T]$ , where the maximum possible payoff of 80 points results if the robber gave the minimum away except for the Trust Game where she gave everything to the trustee and receives the maximum possible in each of the intermediate games. Since  $V = 571.43$  points and  $n_V = 16$ , no matter what the exact expectation adjustment was,  $E[\Pi_R] < (n_V - 1)V/n_V \leq E[\Pi_V]$ . It follows that the derivative of  $p^*$  with respect to  $E[\Pi_V]$  is strictly positive, which implies that, in the BO model, robbers should rob (weakly) more with larger expectations on  $E[\Pi_V]$ .

Analogously, the derivative of the expression of  $p^*$  with respect to  $E[\Pi_R]$  is

$$\frac{\partial p^*}{\partial E[\Pi_R]} = \frac{a}{b} \frac{2}{E[\Pi_V]} (E[\Pi_R] + E[\Pi_V]) - \frac{n_V}{n_V + 1} \frac{1}{E[\Pi_V]}$$

and is larger than zero for all ratios  $a/b$  above  $\frac{n_V}{n_V+1} \frac{1}{2(E[\Pi_R] + E[\Pi_V])}$ . The latter expression cannot be larger than its value for the minimum feasible values of  $E[\Pi_R]$  and  $E[\Pi_V]$ , which are 0 and  $\frac{15}{16}571.43$  points, respectively. With  $n_V = 16$ , this implies that  $p^*$  is increasing in  $E[\Pi_R]$  for all  $a/b$  above (approximately) 0.00088. That is, for such values, in the BO model, robbers should also rob (weakly) more with larger expectations on  $E[\Pi_R]$ . Thus, the optimum  $p^*$  will be larger than the value of  $p^*$  with the smallest possible expectation  $E[\Pi_V] = (n_V - 1)V/n_V$  and the lower bound  $E[\Pi_R] = 0$ , that is,

$$\frac{a}{b} \frac{n_V - 1}{n_V} V + \frac{1}{n_V + 1}$$

but with  $n_V = 16$  and  $V = 571.43$ , this expression is always larger than 0.5 for any  $a/b$  above (approximately) 0.00082. In view of the values derived from behavior in the Dictator Game, all ex post robbers who gave less than 5 points in the Dictator Game are predicted to take 50%, those who gave 6 points or more are inconsistent with the model, and only for those who opted for an equal split, the possible values derived for  $a/b$  do not allow a precise prediction. That is, the model predicts no difference between ex ante and ex post robbers. In particular, the implications are in practice identical to those derived from the CR model.

The model of AW states that an individual is a *homo moralis* if her utility function is the convex combination of selfishness (maximization of own payoff) and Kantian morality (the payoff received when everybody acts as the individual). The weight each subject puts on the moral payoff reveals her degree of morality, i.e.,  $\kappa \in [0, 1]$  with  $\kappa = 0$  being completely selfish (*homo oeconomicus*) and  $\kappa = 1$  being completely moral (or *homo kantianus*).

The AW model considers symmetric games, but it is extended to asymmetric ones as the Dictator Game by recasting it as a role game, that is, the player considers herself as either the dictator or the receiver with probabil-

ity  $\frac{1}{2}$  each (Alger and Weibull, 2013, Section 6.1.1). This yields the utility function

$$U_D^{AW}(x_r) = \frac{1}{2} \left( v \left( \frac{10 - x_r}{10} \right) + \kappa \cdot v \left( \frac{x_r}{10} \right) + (1 - \kappa) \cdot v \left( \frac{y_r}{10} \right) \right)$$

where again  $x_r$  is the amount the dictator sends to the receiver,  $y_r$  is the amount received if the agent is the receiver, and  $v : [0, 1] \mapsto \mathbb{R}$  is a differentiable function with  $v' > 0$  and  $v'' < 0$  representing the well-being from wealth.

Hence, given a function  $v$  we can infer bounds on the degree of morality  $\kappa$  that subjects reveal depending on what they sent in the Dictator Game. For our exercise we define  $v = \sqrt{x}$  which satisfies the required properties for  $v$  stated above.<sup>22</sup> The ranges for the implied  $\kappa$  are displayed in the fifth column in Table 2.2. Note that, as in the case of BO, those who gave 6 points or more are inconsistent with the AW model.

For the Big Robber, the AW model (with a  $\frac{1}{2}$  probability of being a robber or a victim) implies the following utility function.

$$\begin{aligned} U_R^{AW}(p) = \frac{1}{2} \cdot \left( v \left( \frac{E[\Pi_R] + p \cdot E[\Pi_V]}{E[\Pi_R] + E[\Pi_V]} \right) + \kappa \cdot v \left( \frac{(1 - p) \cdot E[\Pi_R]}{E[\Pi_R] + E[\Pi_V]} \right) \right. \\ \left. + (1 - \kappa) \cdot v \left( \frac{(1 - q) \cdot E[\Pi_R]}{E[\Pi_R] + E[\Pi_V]} \right) \right) \end{aligned} \quad (2.2)$$

where  $p$  is as defined above,  $q$  is the share that is taken from the participant in case she is a victim, and  $v$  is defined as in the Dictator Game above. Note that the the second and third terms refer to hypothetical situations where the decision maker becomes a victim. However, the decision maker still has the same expectation  $E[\Pi_R]$  on her earnings from the intermediate games, and the term  $E[\Pi_R] + E[\Pi_V]$  should be interpreted as the sum of this expectation and the expected earnings by all other players in those games.

<sup>22</sup>The model of BO is also cast for general functional forms, but the functional form we use is the one typically employed in the literature. Hence, although using a specific functional form for the model of AW is arbitrary, it keeps the exercise comparable with the previous one.

Once again, for an ex ante robber  $E[\Pi_V] = V$  and  $E[\Pi_R] = V/n_V$ , hence the expression above simplifies to

$$U_R^{AW}(p) = \frac{1}{2} \left( v \left( \frac{1 + p \cdot n_V}{1 + n_V} \right) + \kappa \cdot v \left( \frac{1 - p}{1 + n_V} \right) + (1 - \kappa) \cdot v \left( \frac{1 - q}{1 + n_V} \right) \right)$$

A direct computation shows that, for ex ante robbers,  $U_R^{AW}(p)$  is maximized at  $p = 0.5$  or larger for any level of morality  $\kappa \in [0, 1]$ . Therefore we conclude that all ex ante robbers with  $\kappa \in [0, 1]$  should take 50%.

Regarding ex post robbers, we proceed as in the BO model. The optimal  $p^*$  that maximizes the subject's utility given by (2.2) is

$$p^* = \frac{E[\Pi_V]^2 - \kappa^2 E[\Pi_R]^2}{E[\Pi_V]E[\Pi_R]\kappa^2 + E[\Pi_V]^2}$$

The derivative of this expression with respect to  $E[\Pi_V]$  is

$$\frac{\partial p^*}{\partial E[\Pi_V]} = \frac{E[\Pi_R]\kappa^2 (E[\Pi_V]^2 + E[\Pi_R]^2\kappa^2 + 2E[\Pi_V]E[\Pi_R])}{(E[\Pi_V]E[\Pi_R]\kappa^2 + E[\Pi_V]^2)^2}$$

which is strictly positive for  $\kappa, E[\Pi_V], E[\Pi_R] > 0$ . That is, in the AW model, robbers should rob (weakly) more with larger expectations on  $E[\Pi_V]$ . Note that completely selfish agents, i.e. those  $\kappa = 0$ , will of course rob as much as possible.

Analogously, the derivative of the expression of  $p^*$  with the respect to  $E[\Pi_R]$

$$\frac{\partial p^*}{\partial E[\Pi_R]} = \frac{-E[\Pi_V]\kappa^2 (\kappa^2 E[\Pi_R]^2 + 2E[\Pi_V]E[\Pi_R] + E[\Pi_V]^2)}{(E[\Pi_V]E[\Pi_R]\kappa^2 + E[\Pi_V]^2)^2}$$

is strictly negative for  $\kappa, E[\Pi_V], E[\Pi_R] > 0$ , which implies that robbers should rob (weakly) more with smaller expectations on  $E[\Pi_R]$ .

Thus, the optimum  $p^*$  will be larger than the value of  $p^*$  with the smallest possible expectation  $E[\Pi_V] = (n_V - 1)V/n_V$  and the largest possible  $E[\Pi_R] = 80$ . Straightforward but cumbersome computations show that the corresponding value of  $p^*$  with  $n_V = 16$  and  $V = 571.43$  is always larger than

Table 2.2: Out-of-sample analysis for the models of Charness and Rabin (2002), Bolton and Ockenfels (2000) and Alger and Weibull (2013).

$x_r$	Nr. Obs	Implied $\rho$ (CR)	Implied $a/b$ (BO)	Implied $\kappa$ (AW)	Prediction BR Game	Compatible w/ prediction
0	123	$\leq \frac{1}{2}$	$[0.045, \infty[$	$[0, 0.162]$	Take 50%	100 (81.3%)
1-4	130	$= \frac{1}{2}$	$[0.005, 0.045]$	$[0.162, 0.904]$	Take 50%	59 (45.4%)
5	62	$\geq \frac{1}{2}$	$[0, 0.005]$	$[0.904, 1]$	Undeterm. Take 50%	62 (100%) 20 (32.3%)
6-9	3	$\geq \frac{1}{2}$	N/A	N/A	Undeterm.	3 (100%)
10	2	Undeterm.	N/A	N/A	Undeterm.	2 (100%)
Total	320				CR FS/BO AW	226 (70.6%) 221 (69.1%) 179 (55.9%)

0.5 for any value of  $\kappa \in [0, 1]$ . This implies that ex post robbers with any level of morality should also take 50% in the Big Robber Game.<sup>23</sup> In summary, the AW model predicts the same as the other three models for robbers who gave 4 or less in the Dictator Game, but, while the models of FS, BO, and CR do not provide unique predictions for the 62 robbers who gave 5 in the Dictator Game, the AW model predicts that those robbers should take as much as possible, i.e. 50%. In our sample only 20 participants (32.26%) out of those 62 behaved as predicted by AW.

The last column of Table 2.2 compares the predictions of the models to actual behavior in the Big Robber Game. We obtain that 226 out of our 320 observations (70.63%) can be explained by the model of CR. Further, of the 226 observations, 159 (49.69% of the total, or 70.35% of the explained observations) are such that the robbers took the maximum, while 67 (20.94% of the total, or 29.65% of the explained) are such that the prediction of the model cannot be actually derived (because the participant gave 5 or more

<sup>23</sup>This apparently surprising result has a straightforward intuition. The utility of an AW agent is a convex combination of selfishness and the hypothetical payoff obtained when everybody robs as much as the agent. Selfishness prescribes to rob as much as possible. The other part of the utility combines the large payoff increase obtained when the agent is a robber and robs  $p$  from 16 other agents with the comparatively small payoff decrease obtained when the agent is a victim and  $p$  is robbed from her. Obviously, the payoff increase dominates.

points in the Dictator Game). The 5 observations of robbers who gave strictly more than 5 points are inconsistent with the models of FS, BO, and AW. We obtain that 221 out of 320 observations (69.06%) can be explained by the models of FS and BR. Out of the 221 observations, 159 are such that the robbers took the maximum, while 62 (19.38% of the total, or 28.05% of the explained) are such that the prediction of the model cannot be actually derived (because the participant gave exactly 5 in the Dictator Game). Only the AW model provides a unique prediction for subjects who gave exactly 5 in the Dictator Game. The AW model explains 179 out of 320 observations (55.94%), all of them such that the robbers took the maximum.

In summary, the observations are in line with received models of social preferences as FS, CR, or BO, although a significant percentage of observations are compatible with those models simply because they do not make unique predictions for participants who gave exactly 5 in the dictator game. It should be noted that, as indicated in Table 2.2, the standard models of social preferences we consider predict a larger number of participants taking 50% than we actually observe, and are hence even closer to selfish behavior than the data. The reason is that the very high potential payoff in the Big Robber Game, compared to the Dictator Game, makes it difficult to compare the ranges of the parameters across games. Therefore, each model comes to the conclusion that the robber should take the maximum of 50% whenever she gives less than the equal split in the Dictator Game. This coincides with the majority of our data. The model of AW fares considerably worse, essentially because it predicts that *homo moralis* will always rob as much as possible, which includes even those participants who gave 5 in the Dictator Game.

## 2.7 Discussion

The Big Robber Game is a paradigm which makes high stakes salient, while also emphasizing that the return comes at the cost of actively harming a large group of people. Hence, the paradigm directly captures the idea that a monetary temptation might make people act against society's interests, as in

the many corporate scandals which have sparked public outrage in the recent decades, where “individuals who hold such financial power may be tempted to abuse it for their own personal profit” (Myerson, 2012, p.848). With this paradigm we showed that such corporate scandals are easily reproduced in the lab even with regular university students. An absolute majority of our (large) sample took the maximum possible amount, accepting that their decision would damage a large number of other people. Further, the decision to take the maximum was faster on average than the decision to refrain from it, revealing a weaker moral struggle in the former case, and those who took the maximum also donated less to a charity afterwards, revealing no evidence of guilt. Our results stand in sharp contrast with the hypotheses of Mazar et al. (2008), who argue that most people will cheat only a little if given opportunity.

At the same time, we show that the very same participants who are willing to inflict considerable damage on their peers within the Big Robber Game display standard levels of other-regarding behavior as reflected by Dictator, Ultimatum, and Trust games. That is, we empirically demonstrate the coexistence of prosociality in the small and morally outrageous selfishness in the large, that is, in high-stakes, high-impact decisions affecting large groups.

The behavior we observe, however, is fully compatible with received models of social preferences. First, by calibrating standard models using behavior in the Dictator Game, we show that a large part of our data is consistent with the functional forms commonly-used to capture other-regarding preferences. If anything, those models predict higher levels of selfish behavior in the Big Robber Game than we actually observe. Second, we show that individual behavior in the Big Robber Game stands in a monotonic relation with behavior in the standard games mentioned above. Participants who took the maximum in the Big Robber Game gave less in the Dictator Game, offered less as proposers and accepted more unfair offers in the Ultimatum Game, and transferred less as trustors and returned less as trustees in the Trust Game. That is, behavior in all games, small and big, can be explained within a single account of other-regarding preferences. This is interesting in



itself, in view of previous results on the (in)consistency of social preferences across games (e.g., Blanco et al., 2011).

Our experiment also allowed us to acquire a number of other insights. First, a number of psychological theories predict an effect of previous decisions in the moral domain on subsequent ones. We do find effects as predicted by those theories, but they are small and not systematic (that is, we find them for some games and not for others), and depend on gender. For men, we find evidence that behaving selfishly leads to further selfish behavior down the road, in accordance with the “what-the-hell” effect. For women, behaving selfishly seems to lead to a possibly guilt-induced attempt to behave less selfishly later on, in accordance with the “transgression-compliance” effect. Symmetrically, for women, refraining to behave selfishly earlier might lead to increasingly selfish behavior later, in accordance with the “moral-credentialing” effect. For men, however, refraining to behave selfishly earlier might lead to less selfish behavior later. Hence, while men seem to behave consistently, sticking to being more or less selfish, women act as if decisions added up and one bad canceled one good, and vice versa. We insist, however, that these effects are small and not systematic in our sample.

It is also noteworthy that we find no gender effect whatsoever in the Big Robber decision. Research from psychology suggests that women are more sensitive to social cues and feedback than men (Gilligan, 1982; Roberts and Nolen-Hoeksema, 1989). In line with this, several previous studies have identified gender differences in the degree of inequality aversion (for example, the proportion of egalitarian allocations in the Dictator and Ultimatum games). However, as shown in the review of Croson and Gneezy (2009), the evidence is mixed, and not all experiments find a clear gender difference. In a recent field experiment, Azmat et al. (2016) show that when stakes are high (as in our experiment), women perform worse than men, while the opposite is true whenever the stakes are low. This difference, though, refers to performance in real-effort tasks, while our experiment is about preferential choice.

The most important differences between the Big Robber Game and other games used to capture social preferences are, first, the fact that a single decision can inflict significant damage on a large group of people, and second,

the size of the stakes. There are, of course, by necessity, a number of other differences between the Big Robber Game and other games. In our design, only one robber is selected out of 16 possible ones, and the individual decision is taken before the actual robber is selected. Hence there might be, in a certain sense, a diluted responsibility. However, this interpretation is at odds with evidence from the strategy method (Brandts and Charness, 2011), which shows that behavior is essentially aligned if decisions are made conditional on reaching a certain decision node and if the response is chosen after learning that the node has been reached. Another difference is that, in spite of the neutral framing, the decision in the Big Robber Game is more immediately placed into a moral framework than those in previous games, since it is not in terms of distribution or bargaining but simply in terms of taking away the earnings of others. The Big Robber Game is, simply put, a new paradigm, and not a high-stakes version of previous ones.

Since we aimed for a relatively large experiment in order to be able to test for treatment effects and condition those on gender, we settled on a single basic design rather than complicating the analysis with added variants from the onset. A large number of avenues for future research, involving design variants, are obvious at this point. The upper bound of 50% on what could be taken away from the victims was dictated by practical considerations, in order to avoid damaging the reputation of the lab by having half of the participants walk out of the experiment empty-handed. Future design variants could find ways to relax this constraint. Also, the discretization in only four possibilities reflected the desire to have clear behavioral groups (take nothing, take just a bit, take a significant part but not the maximum, and take the maximum). The baseline result having been thus established, it might be desirable to implement a continuous version. These and many other possibilities are beyond the scope of this work and are left for future research.

In conclusion, the Big Robber Game contributes to the literature in three ways. From the methodological point of view, it provides the possibility to experiment with relatively high stakes without increasing existing research budgets. From the conceptual point of view, it provides a novel empirical demonstration that behavior which is commonly accepted as prosocial in

“small” situations, as captured by standard laboratory games, coexists with behavior which is commonly considered morally questionable in “large” situations (high stakes, high impact). Last, the results echo popular discussions on moral responsibility among economic decision makers by showing that a large part of our sample (containing just regular university students) is willing to inflict significant damage on a relatively large number of people for personal gain, as long as that gain is of sufficient magnitude. Regrettably, our data suggests that, in our Western societies, hundred bucks might do the job.



---

## CHAPTER 3

### The Reinforcement Heuristic in Normal Form Games

---

#### 3.1 Introduction

Reinforcement is one of the most basic processes underlying human learning. Accordingly, it has received widespread attention in psychology, going back to Thorndike’s (1911) “law of effect,” neuroscience (e.g. Holroyd and Coles, 2002; Schönberg et al., 2007), and computer science (Sutton and Barto, 1998). Within microeconomics and game theory, it has been frequently studied as a boundedly-rational behavioral rule (see, e.g. Börgers and Sarin, 1997; Erev and Roth, 1998), as have been other rules, e.g. imitation or myopic best reply (Weibull, 1995; Fudenberg and Levine, 1998). The simplest version of reinforcement learning can be viewed as a heuristic which takes past experiences into account for the choice of upcoming actions and prescribes a shift from actions linked to negative experiences to actions associated with positive rewards: that is, “win-stay, lose-shift.” This heuristic induces a bias towards past-high-reward actions which can conflict with rational behavior (outcome bias; Baron and Hershey, 1988; Dillon and Tinsley, 2008).

Evidence from neuroscience shows that reinforcement-based decisions occur extremely fast in the human brain (Schultz, 1998; Holroyd and Coles, 2002). Indeed, reinforcement is a textbook example of an *automatic process*, as conceived in dual-process theories from psychology (see, e.g., Kahneman, 2003; Strack and Deutsch, 2004; Alós-Ferrer and Strack, 2014). Those theories define automatic (or intuitive) processes as immediate, fast, unconscious, and efficient in the sense of requiring few cognitive resources. For instance, these processes capture impulsive reactions and behavior along the lines of stimulus-response schemes. The dual-process approach postulates that human decisions are mainly influenced by automatic processes and so-called *controlled* (or deliberative) processes. The latter are seen as slow, consum-

ing cognitive resources, not instigated immediately, and reflected upon consciously. Explicit maximization of expected rewards, if conceptualized as a process, would exhibit many if not all of those characteristics.

The relevance of reinforcement for economic decision making was illustrated by Charness and Levin (2005) in a binary-choice, belief-updating task where mistakes (deviations from optimization under correct Bayesian updating) could be traced to a reinforcement heuristic. In essentially the same paradigm, Achtziger and Alós-Ferrer (2014) found evidence of the conflict between reinforcement and rational optimization in the form of response time asymmetries as predicted by an explicit dual-process model. Recent psychophysiological work (Achtziger et al., 2015a) found direct evidence of neural correlates of reinforcement in this paradigm and studied their relation to economic incentives. Further studies relying on this paradigm have examined the interaction of reinforcement and decision inertia (Alós-Ferrer et al., 2016b), the influence of framing on reinforcement decisions (Alós-Ferrer et al., 2017), and the regulation of reinforcement processes through motivational interventions (Hügelschäfer and Achtziger, 2017).

In this work, we take a further step in the study of reinforcement heuristics in economic settings by moving beyond binary-choice tasks and studying the explicit relation between reinforcement and myopic payoff maximization in strategic decisions. Hence, we study reinforcement processes in a more complex setting which results in longer decisions times than, e.g., standard neuropsychological experiments. We concentrate on two-player,  $3 \times 3$  asymmetric normal form games. In this setting, the microeconomics literature has devoted a great deal of attention to myopic best reply, a behavioral rule which maximizes the own payoff assuming the other player will repeat her action, and which can be assumed to have a more deliberative/controlled nature than reinforcement.

Previous work has analyzed paradigms where, by design, reinforcement and more deliberative behavior could either conflict or be aligned (e.g. Achtziger and Alós-Ferrer, 2014). In other settings, however, there might be a great degree of overlap between the prescriptions of reinforcement and those of myopic best reply. In the present work we specifically explore to

what extent reinforcement can act as a shortcut for (apparent) optimization in strategic situations with explicit feedback. Suppose a player's last action delivered the best possible payoff. Reinforcement will then prescribe to repeat the choice (win-stay). By definition, however, that choice is the best reply if the opponent stays put. Likewise, suppose the last action did not deliver the best possible payoff. Reinforcement will prescribe to choose a different action (lose-shift). But, again by definition, the current choice cannot be the myopic optimum, and hence myopic best reply also prescribes to shift. In principle, the "shift" prescribed by reinforcement is arbitrary, but if payoffs are observable (as, e.g., if the payoff table is known), the observable maximum becomes salient and the shift will often be in its direction, leading to an apparent myopic best reply. In view of these observations, we postulate that reinforcement processes might often act as cognitive shortcuts resulting in choices indistinguishable from myopic best reply.<sup>1</sup>

If choices coming from reinforcement and those prescribed by myopic best reply cannot be distinguished, how can this hypothesis be substantiated? There are two possible avenues. The first relies on response times. As explained above, reinforcement processes are automatic and can be expected to lead to shorter response times than alternative processes. In contrast, myopic best reply involves explicit maximization and can be assumed to be deliberative, hence relatively slow. Hence, if the choices favored by both processes are actually due to the involvement of reinforcement processes, one should expect shorter response times (compared to other choices), while if they are due to explicit maximization, response times should be longer.

The second avenue is bygone payoffs. Suppose a player obtains a payoff which is not the maximum possible one given the opponent's strategy. If that maximum payoff is observable, the deviation with respect to it is a cardinal measure of experienced disappointment. Define experienced regret as the difference between the maximum possible payoff (that of the best reply) and the actually obtained payoff. By definition, regret is zero if and only

---

<sup>1</sup>Indeed, the obvious evolutionary reason for the existence of automatic processes is that in certain situations they are adapted and support (near-)optimal decisions while saving cognitive costs.

if the player has chosen a best reply, and strictly positive if not (this will be directly observable in our experiment). Myopic best reply, considered as a behavioral rule, prescribes to change strategy whenever a best reply has not been chosen. In contrast, reinforcement processes are stimulus-response mappings which take the win-loss information as an input. The loss information also carries a measurement of stimulus strength which in turn yields a variation in the responses. Hence, standard formalizations of reinforcement take the probability of a shift to be increasing in the degree of the loss, which is just the experienced regret as defined above. That is, reinforcement should be triggered more often for larger experienced regret. Hence, if observed choices follow from reinforcement processes, one should observe a dependence on experienced regret.

To study these questions, we conducted an experiment ( $N = 144$ ) where participants played  $3 \times 3$  games against other players repeatedly. In order to isolate the decision processes of interest, in our experiment players had full knowledge of their own payoff tables, but were not aware of the payoffs of the opponents. In this way, we aimed to eliminate a number of potential confounds, as e.g. imitation or social preferences. Also, in this simple design the maximum (bygone) payoff associated with a choice is directly observable, and regret is simply the difference between the highest payoff associated with the opponent's choice and the actually received payoff. To make this even simpler, each of the different payoff tables used in the experiment contain only three different numerical payoffs, hence maximum payoffs and regret levels are easily observable.

The experiment was divided in short parts of 13 rounds each, and after each part both the opponent and the payoff table were changed. This was an explicit design decision in order to prevent convergence or long-run effects, and rather concentrate on the decision processes.

Our work is related to several literature strands. Within the reinforcement learning literature, we focus on simple stimulus-response behavioral rules of thumb, capturing the basic idea of a “win-stay, lose-shift” stimulus-response mapping. Of course, the literature on reinforcement also encompasses more involved models. For instance, the game-theoretic cumulative proportional



reinforcement of Laslier et al. (2001) postulates that actions are chosen with probabilities proportional to their cumulative payoffs obtained in the past with that move. Modern reinforcement learning models in computational neuroscience (see Daw, 2012, for an introduction) often include additional factors, e.g. through a perseveration parameter capturing the tendency to repeat or avoid recently chosen actions (e.g., Gershman et al., 2009; Wimmer et al., 2012). Our interest, however, is on Markov behavioral rules conceived as mappings from information at  $t$  (e.g., payoffs and actions of all players) to behavior at  $t + 1$  (probabilities of the feasible actions), as standard in the literature of learning in games (Fudenberg and Levine, 1998). Within this literature, such behavioral rules include imitation (e.g., Vega-Redondo, 1997; Alós-Ferrer and Weidenholzer, 2008) and myopic best reply (e.g. Kandori and Rob, 1995; Alós-Ferrer and Weidenholzer, 2007). This focus is consequential for the analysis. On the assumption that behavior follows Markov rules, one can concentrate on the estimation of probabilities of actions given only the relevant information, e.g. realized and bygone payoffs and actions in the last period. For instance, under such maintained hypotheses, the matching within an experiment is irrelevant, as behavior is assumed to rely only on the information presented to the player (we will, however, also control for matching in our analysis).

Our experiment is also related to the literature on multi-armed bandits (Gittins, 1979, 1989), in the sense that players repeatedly choose an action out of three possibilities. However, contrary to this literature we are not interested in optimal (normative) dynamic strategies, but rather on the actual (one-shot) decision processes employed by human beings in our setting. Since players were explicitly told that the opponent was human, our experiment was not framed in terms of intertemporal optimization or uncovering of optimal actions. Indeed, in our data, we see little evidence for intertemporal optimization or learning effects. On the contrary, we argue that, in our context, even what might be interpreted as one-shot optimization might be actually supported by simpler decision processes.

A more-closely related literature has examined “decisions from experience,” where subjects make repeated individual decisions in stochastic frame-

works but learn the underlying probabilities by making decisions (as opposed to decisions from description, where the priors are induced; Hertwig et al., 2004). Although our setting is interpersonal (subjects play against a human being and not a fixed distribution), insights from this literature are informative. In particular, Erev and Haruvy (2016, Section 1.1.5) remark that decision inertia plays an important role (recall also Alós-Ferrer et al., 2016b), which in our setting would lead to higher rates of win-stay choices compared to lose-shift ones. This is indeed found in our data. Erev and Haruvy (2016) also list “surprise-triggers-change” as one of the main reasons for not repeating the last choice; that is, the probability of inertia decreases when recent outcomes are surprising. In our case, a larger experienced regret plays the role of surprise. The mirror image of this effect is “negative recency,” which is sometimes observed in decisions from experience and implies higher shift rates after surprising positive outcomes (in our case, no regret), even though those reinforce the last choice.

The remainder of the Chapter is structured as follows. Section 3.2 presents the experimental design and procedures. Section 3.3 presents the results, analyzing both choice data and response times. Section 3.4 concludes.

## 3.2 Experimental Design and Procedures

We conducted 6 sessions with 24 participants each for a total of 144 subjects (91 female) at the Cologne Laboratory for Economic Research (CLER). Participants were recruited via ORSEE (Greiner, 2015) from the student population of the University of Cologne excluding students of psychology. The average age was 22.97 years (median 23, range 18–50). The experiment was programmed in z-Tree (Fischbacher, 2007) and sessions lasted on average 60 minutes. The average payoff was 12.28 EUR (SD= 0.94) including a show-up fee of 2.50 EUR.

The experiment involved three different  $3 \times 3$  normal form games (see Table 3.1). The games had a cyclical structure and the three strategies were neutrally labeled:  $\circ$ ,  $\#$ ,  $\div$  for player 1, and  $\sim$ ,  $\bullet$ ,  $\times$  for the opponent. Each table uses only three different payoffs, and each outcome is clearly attainable

for each action the opponent possibly plays. Hence, maximum bygone payoffs and regret levels (given the actual opponent’s strategies) are directly and easily observable.

Each subject played a total of 39 rounds divided into three different parts of 13 rounds each. In each part, a subject faced a fixed game and played against a fixed partner. The subject saw a  $3 \times 3$  payoff table with only her own payoffs. That is, she was not informed about the payoffs of the other player and hence imitation was not feasible, and social preferences were not a concern. The subject chose one of the three actions ( $\circ$ ,  $\#$ ,  $\div$ ) in each round. Rematching was done within blocks of four players after 13 rounds and a new game, i.e. payoff table, was presented. The payoff tables were always reordered and relabeled in such a way that every player saw herself as player 1. The order of the labels for the own strategies was counterbalanced among subjects. Subjects were paid for each decision in all rounds. Alternatively, we could have paid one randomly selected round; Charness et al. (2016) have shown that relying on one or the other method does not significantly affect behavior.<sup>2</sup>

The games were designed to generate different levels of “regret,” defined as the difference between the maximum possible payoff and payoff earned in the previous round. The levels ranged from 0 to 6 (see Table 3.1). By construction, in our games a best reply always reached the maximum payoff available in the payoff table. This was done on purpose to avoid alternative definitions of experienced regret and, hence, potential alternative rules based on reinforcement heuristics.

After each round, the actual play from the previous round was revealed. Subjects saw their own choice, the opponent’s choice, and their own payoff. This information remained on-screen while they made their choice for the next round (obviously, there was no feedback during the first choice of each part). In addition, the column in the payoff table which represented the choice of the other player in the previous round was highlighted.

---

<sup>2</sup>However, Azrieli et al. (2018, Appendix C) conclude that in a dynamic setting with feedback (as our design) paying a randomly selected round is not incentive compatible, because agents have an incentive to experiment.

Table 3.1: The three games used in the experiment.

Game 1				Game 2				Game 3			
	$\sim$	$\bullet$	x		$\sim$	$\bullet$	x		$\sim$	$\bullet$	x
o	6,2	4,4	2,6	o	7,1	4,4	1,7	o	6,1	5,5	1,6
#	2,6	6,2	4,4	#	1,7	7,1	4,4	#	1,6	6,1	5,5
$\div$	4,4	2,6	6,2	$\div$	4,4	1,7	7,1	$\div$	5,5	1,6	6,1
Regret $\in \{0, 2, 4\}$				Regret $\in \{0, 3, 6\}$				Regret $\in \{0, 1, 5\}$			

After completing the 39 rounds of play, participants answered demographic questions, the Maximization and Regret scale (Schwartz et al., 2002), the Faith in Intuition scale (Epstein et al., 1996; Alós-Ferrer and Hügelschäfer, 2012, 2016), and a 15-item Big-Five questionnaire (Lang et al., 2011, p. 560).<sup>3</sup> The results reported below are robust to the inclusion of those measures as controls, but the measures themselves provided no additional insights.

### 3.3 Results

#### 3.3.1 Classification of the Decision Situations

We are interested in the choices that participants made while seeing the previous round's outcomes, and hence we dropped the first round of each part (where no feedback on a previous decision was possible). For all other decisions, if the participant had achieved the highest possible payoff in the previous round, we classified the following decision as a *win* situation, else as a *lose* situation. The complete data set consists of  $144 \times 36 = 5184$  observations, of which 1,794 observations (34.61%) were win situations and 3,390 (65.39%) were lose situations. Regret, defined as the difference between the maximum possible payoff and the realized one, was always 0 in win situations and strictly positive in lose situations. Table 3.2 contains the frequency of

<sup>3</sup>The German versions of the three scales were taken from Greifeneder and Betsch (2006), Keller et al. (2000), and Gerlitz and Schupp (2005), respectively.

Table 3.2: Frequency of the different experienced regret levels in lose situations.

Regret	1	2	3	4	5	6
Number	550	494	552	617	589	588
Frequency (%)	16.22	14.57	16.28	18.20	17.37	17.35

various levels of regret in lose situations, which were not significantly different from a uniform distribution ( $\chi^2$ -test,  $\chi^2_{(5)} = 8.359$ ,  $p = 0.1375$ ).

The “win-stay, lose-shift” version of reinforcement we focus on prescribed staying with the previously chosen option in win situations and shifting away to another option in lose situations. However, in win situations staying with the previous action is also the prescription of myopic best reply in our games, since if a player assumes that the opponent will not change strategy, repeating the strategy which led to a win is optimal. Further, mere inertia (repeating the previous action no matter what; see, e.g., Alós-Ferrer et al., 2016b) also leads to the same prescription. That is, in win situations all three behavioral rules (reinforcement, myopic best reply, and inertia) prescribed to repeat the previous action. There were two possible deviations from this common prescription, i.e. the choice which would have yielded the medium payoff and the one which would have yielded the low payoff (if the opponent stayed put).

In lose situations reinforcement prescribed to shift away from the previous choice to one of the two remaining alternatives. Shifting to the maximum-payoff alternative was aligned with myopic best reply (BR shifts), but shifting to the remaining alternative (which could be a medium- or low-payoff one) was not (non-BR shifts). Since BR shifts just require following the bygone payoff, we focus on a reinforcement rule prescribing BR shifts. By definition, the prescription of inertia was to stay put with the previously-chosen alternative.

The strategy of analysis is as follows. At the first level, we take decisions as a response to the feedback, i.e. we are interested in (Markov) decision rules as studied in the literature on learning in games (e.g., Fudenberg and Levine, 1998). For such an analysis, the origin of the input (feedback) is irrelevant,




Figure 3.1: Average individual choice frequencies.

*Notes.* Frequencies conditional on Win (left-hand side) and Lose (right-hand side) situations. Bars represent 1 Standard Error of the Mean. Significance levels refer to Wilcoxon Signed-Rank tests. \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

since one studies the probabilities of responses conditional on the input. Hence we study subject averages (frequency of choices and response times conditional on choice and situation) with 144 observations (one per player) and conduct non-parametric (within) tests. At the second level, we conduct a robustness analysis and reanalyze the data at the block level (since matching was within blocks of four players), creating 36 independent observations for non-parametric tests. That is, at this level each individual observation averages over all decisions of all four players in a block. At the third level, we analyze the data as a panel through the appropriate regression models, controlling for a number of additional factors.

### 3.3.2 Choice Data

Figure 3.1 depicts the average individual choice frequencies conditional on win and lose situations. In win situations, the decision to “stay” (following reinforcement, myopic best reply, and inertia) was significantly more frequent than any other alternative. Stay decisions, with an average individual frequency of 51.57%, were more frequent than shifts to the medium-payoff (26.35%) or the low-payoff action (22.08%). The differences are significant according to Wilcoxon-Signed Rank tests<sup>4</sup> (stay vs. shift to medium,  $z = 6.130$ ,  $p < 0.0001$ ; stay vs. shift to low,  $z = 6.615$ ,  $p < 0.0001$ ). Shifts to the medium-payoff action were also more frequent than shifts to the low-payoff action, and the difference in distributions is also significant ( $z = 2.599$ ,  $p = 0.0094$ ). Tests at the block level ( $N = 36$ ) yielded the same conclusions (WSR tests, stay (50.73%) vs. shift to medium (26.92%),  $z = 4.550$ ,  $p < 0.0001$ ; stay vs. shift to low (22.36%),  $z = 4.643$ ,  $p < 0.0001$ ; shift to medium vs. shift to low,  $z = 2.353$ ,  $p = 0.0186$ ).

In lose situations, shifts to the myopic best reply were also more frequent than other choices. The average frequency of shifts aligned with myopic best reply was 42.74%, compared to 26.41% of non-BR shifts (WSR test,  $z = 6.295$ ,  $p < 0.0001$ ) and 30.84% of stay choices following inertia (WSR test,  $z = 4.443$ ,  $p < 0.0001$ ). Non-BR shifts were less frequent than inertia decisions (WSR test,  $z = -2.108$ ,  $p = 0.0350$ ). Testing at the block level ( $N = 36$ ) yielded the same conclusions (WSR tests, BR shifts (42.07%) vs. non-BR shifts (26.40%),  $z = 4.454$ ,  $p < 0.0001$ ; BR shifts vs. inertia (31.53%),  $z = 3.849$ ,  $p = 0.0002$ ; non-BR shifts vs. inertia,  $z = -2.090$ ,  $p = 0.0367$ ).

Figure 3.3 (left-hand side) shows that the main results above are robust when choice frequencies are analyzed for each game separately. That is, the decision to stay in win situations, respectively to shift to the myopic best reply in lose situations, was more frequent than other choices in each of the three games.

---

<sup>4</sup>Here and elsewhere, tests are adjusted for multiple comparisons following the Holm-Bonferroni correction.

Table 3.3: Random-effects panel probit regressions for choices.

Reinforcement decision	Model 1	Model 2	Model 3
Win	0.2201*** (0.0618)	0.3737*** (0.0859)	0.3733*** (0.0849)
Lose $\times$ Regret		0.0425** (0.0170)	0.0423** (0.0167)
Normalized Round			-0.1171 (0.0736)
Part 2 Dummy			0.0301 (0.0569)
Part 3 Dummy			-0.0323 (0.0581)
Constant	-0.2075*** (0.0377)	-0.3606*** (0.0738)	-0.2959*** (0.0887)
LogLikelihood	-3429.1594	-3424.0503	-3421.2967
Wald Test	33.0209***	43.1477***	48.5771***

*Notes.* Independent variable takes value 1 if reinforcement was followed. Standard errors (clustered by 36 matching blocks) in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 3.3 presents panel probit regressions with random effects at the individual level and standard errors clustered at the matching block. The independent variable is defined as 1 if the decision followed the prescription of reinforcement / myopic best reply, that is, stay in win situations and shift to the best-payoff action in lose situations. The regressions allow us to control for other variables such as the regret level or learning effects. In all three models, the dummy Win (for win situations) is highly significant and positive, implying that participants were more likely to follow the common prescription of reinforcement and myopic best reply in win situations. That is, win-stay decisions are more likely than lose-shift (to the best reply) even though both follow reinforcement and myopic best reply. This asymmetry is as it should be expected from inertia, since a “stay” decision complies with this additional process (Erev and Haruvy, 2016; Alós-Ferrer et al., 2016b).<sup>5</sup>

<sup>5</sup>This is also confirmed by WSR tests comparing win-stay rates (average 51.57%) and the rate of shifts to the best reply in lose situations (average 42.74%), which are significant both at the individual ( $z = 3.275, p = 0.0011$ ) and the block level ( $z = 3.119, p = 0.0018$ ).



Model 2 adds the interaction  $\text{Lose} \times \text{Regret}$ , which captures the effect of the regret level in lose situations (recall that regret is identically zero for win situations). The coefficient is also significantly positive in Models 2 and 3, indicating that being further away from the maximum payoff resulted in a higher probability of following reinforcement. Model 3 controls for learning effects, adding a coefficient for normalized round within a part (rescaled to range from 0 to 1) and individual dummies for each part. None were significant, indicating no evidence of behavioral change over time. That is, consistent with our Markov-rules approach, at the aggregate level there is no evidence that the reliance on reinforcement changed over time. Additional regressions including the personality questionnaires and demographic factors (Table 3.5, Appendix 3.B) as controls yielded the same qualitative results, while the controls themselves provided no further insights.

Cameron et al. (2008) remarked that a small number of clusters can lead to over-rejection of standard asymptotic tests and recommending bootstrapping the standard errors. According to their simulations group sizes of 30 already showed a rejection rate close to the intended 5% level. Hence, as a further robustness test, we ran a regression bootstrapping the standard errors with 100 repetitions (Table 3.6, Appendix 3.B), which yielded the same conclusions.

In summary, the regressions indicate that “win-stay” in win situations were comparatively more likely than shifting to the myopically best option in lose situations, but that the “lose-shift” choice was more likely with higher regret. Note that the latter result cannot be explained by a behavioral rule based on myopic best reply (which should be impervious to regret), but is consistent with a general reinforcement process for which the strength of a loss does play a role.

### 3.3.3 Cumulative Reinforcement

The behavioral rules we concentrate on rely on information from the most recent period of play only. One can of course ask whether alternative reinforcement rules using information from longer histories can describe the

data better. A prominent example is *cumulative proportional reinforcement* (Laslier et al., 2001), which prescribes to choose the action with the highest cumulative payoff over the whole past. Relying only on the simple rules of reinforcement and inertia, which only take the previous round into account, we can account for 3,406 (65.70%) out of the 5,184 observations in our data set. In contrast, only 2,214 observations (42.71%) agree with cumulative proportional reinforcement, and the majority of those observations are also aligned with simple reinforcement and inertia. The combination of inertia and cumulative proportional reinforcement captures 2,813 observations (54.26%). Only 487 (9.39%) observations are consistent with cumulative proportional reinforcement but not with the other two rules, while 1,080 (20.83%) are compatible with our simple reinforcement rule but not with the other two rules.

We also ran regressions analogous to Table 3.3 for cumulative proportional reinforcement (Table 3.7, Appendix 3.B). That is, the dependent variable was defined as a dummy taking the value 1 if and only if the decision was consistent with this rule. The regression did not yield any significant results for winning or regret levels, but showed a significantly negative time trend, indicating a lower likelihood of following cumulative proportional reinforcement over time. In conclusion, we view these observations as indicators that it is reasonable to assume Markov decisions rules in this context.

### 3.3.4 Response Times

We computed average individual response times conditional on the different situations and choices, dropping the first round of each part. Note that not all participants have response times for each shift or stay, since if for instance a participant never shifted away in win situations, there will be no observations in the corresponding categories. Hence, the tests below will in general have different numbers of observations (tests at the block level, however, always have 36 observations).

Figure 3.2 depicts the averages of response times for shift and stay decisions, conditional on win or lose situations. The fastest decisions are always

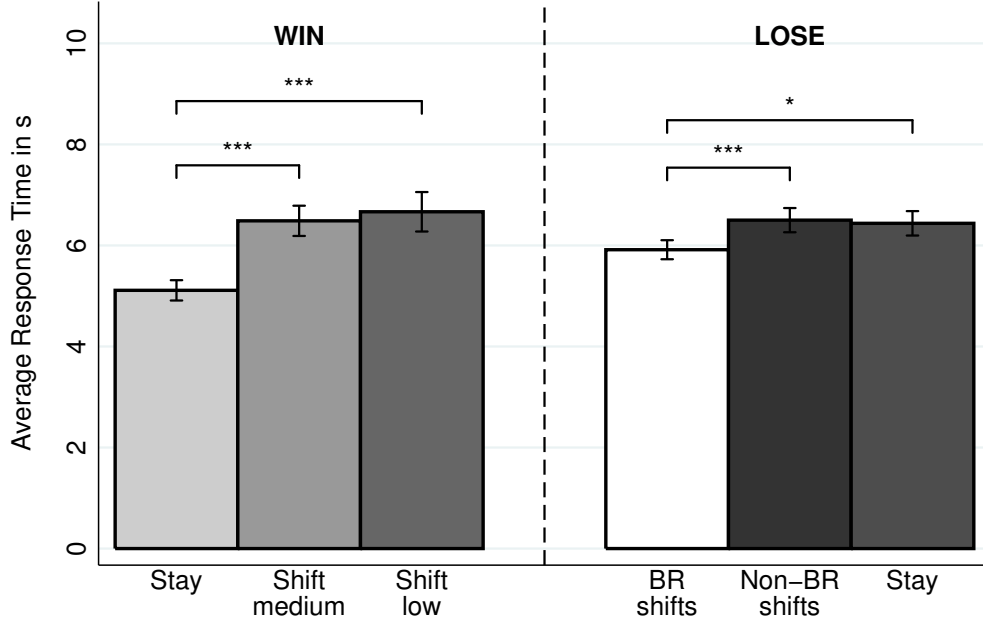


Figure 3.2: Average individual response times.

*Notes.* Response times conditional on Win (left-hand side) and Lose (right-hand side) situations. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , WSR test.

those consistent with the common prescription of myopic best reply and reinforcement. Specifically, in win situations, stay decisions were faster than shifts to the medium-payoff action (average 5.25 s vs. 6.53 s; WSR test,  $N = 126$ ,  $z = -4.342$ ,  $p < 0.0001$ ) and shifts to the low-payoff action (average 5.34 s vs. 6.73 s; WSR test,  $N = 112$ ,  $z = -4.121$ ,  $p < 0.0001$ ). Response times for both kinds of shifts were not significantly different in win situations (shift to medium, 6.34 s; shift to low, 6.53 s; WSR test,  $N = 109$ ,  $z = -0.322$ ,  $p = 0.7475$ ). Testing at the block level ( $N = 36$ ) yields identical conclusions (WSR tests, stay (mean 5.01 s) vs. medium (mean 6.34 s),  $z = -3.629$ ,  $p = 0.0006$ ; stay vs. low (mean 6.48 s),  $z = -3.912$ ,  $p = 0.0003$ ; medium vs. low,  $z = -0.471$ ,  $p = 0.6374$ ).

The same also holds for lose situations. That is, shifts aligned with myopic best reply were faster than non-BR shifts (average 5.91 s vs. 6.50 s; WSR

test,  $N = 142$ ,  $z = -2.988$ ,  $p = 0.0084$ ) and stay decisions (average 5.95 s vs. 6.44 s; WSR test,  $N = 142$ ,  $z = -2.066$ ,  $p = 0.0777$ ). Response times for the latter two were not significantly different in lose situations (non-BR shifts, 6.55 s; stay decisions, 6.45 s; WSR test,  $N = 140$ ,  $z = 0.774$ ,  $p = 0.4391$ ). However, tests at the block level ( $N = 36$ ; adjusted, as always, following Holm-Bonferroni) were not significant in this case (WSR tests, BR shifts (5.98 s) vs. non-BR shifts (6.48 s),  $z = -1.870$ ,  $p = 0.1231$ ; BR shifts vs. stay (6.06 s),  $z = -0.314$ ,  $p = 0.7534$ ; non-BR shifts vs. stay,  $z = 1.901$ ,  $p = 0.1719$ ).

Figure 3.3 (right-hand side) illustrates the response time analysis above for each game separately. The decision to stay in win situations was faster than other decisions for all Games. The decision to shift to the myopic best reply in lose situations was faster than other shifts and stay decisions for Game 3 and faster than stay decisions for Game 1, but the differences were not significant for Game 2.

Table 3.4 presents panel regressions for log-transformed response times, with random effects at the individual level and standard errors clustered at the matching block.<sup>6</sup> The analysis confirms that decisions which agree with the prescriptions of reinforcement (and best reply) are significantly faster than other decisions. To see this for win-stay, we look at the dummy “Stay” which indicates whether the decision was to repeat the previous choice or not, i.e. inertia. Its coefficient is highly significant and negative in all models, indicating that (since the interaction Lose  $\times$  Stay is included in the models) decisions to stay after a win, hence to stay with the best response, were made significantly faster. This is consistent with the interpretation that many such decisions might be the result of relatively automatic processes.<sup>7</sup>

<sup>6</sup>Response times are nonnegative and their distribution is strongly right-skewed, while the distribution of log-transformed response times typically shows a normally-distributed shape (e.g., Fischbacher et al., 2013). Using Shapiro-Wilk W tests at the individual level, the hypothesis of normality was only rejected for 19 of the 144 subjects at the 5% level, and there were only 7 cases with significance between 5% and 10%. In contrast, using regular response times, the hypothesis of normality is rejected in 113 cases at the 5% level, and in 11 further cases at the 10% level.

<sup>7</sup>Some of the non-stay, slower decisions might be the result of more complex decision rules, as e.g. best-responding to a predicted best response.

For lose-shift, we look at the dummy “BR-Shift” which captures shifts to the payoff-maximizing option after a lose situation, as in Figure 3.2 (note that shifting to a best response implies a lose situation). Its coefficient is also significant and negative in all models, again indicating faster decisions.<sup>8</sup>

The regressions also allow us to examine additional questions. This first concerns the difference between win and lose situations. Reinforcement / best-reply decisions were significantly faster in win situations than in lose situations, as revealed by a linear combination test (1) at the bottom of Table 3.4.<sup>9</sup> The dummy for lose situations is not significant, implying that shifts not following the reinforcement prescription did not differ in response times across win and lose situations.

The second additional observation concerns inertia compared to shifts. All models include an interaction term for stay decisions in lose situations ( $\text{Lose} \times \text{Stay}$ ), that is, for following inertia. Concerning inertia and non-BR shifts, the linear combination test (2) in Table 3.4 is only marginally significant, and only in model 3. That is, there is only weak evidence that inertia decisions might be faster than non-BR shifts in lose situations. Concerning inertia and BR shifts, however, the linear combination test (3) comparing BR-shift decisions with stay decisions in lose situations was not significant, in contrast with the non-parametric results.

Models 2 and 3 also add the interaction with the regret level (in lose situations, as in win situations there is no regret), which is not significant, implying that the level of regret experienced did not affect the response time. This is not at odds with our previous finding that higher regret levels induce more reinforcement, because the BR-Shift coefficient already captures the

---

<sup>8</sup>The omitted category are (non-BR) shifts after a win. Since the non-parametric analysis did not find significant differences between shifts to medium and low-payoff actions, we do not distinguish them further. An additional regression including a dummy for shifts to the medium-payoff action in win situations (Table 3.8, Appendix 3.B) yields the same conclusions, with the additional coefficient not being significantly different from 0.

<sup>9</sup>This difference cannot be explained through pure best-reply behavior, since best reply only depends on the opponent’s previous choice and not on whether there was a previous win or loss. Reinforcement yields prescriptions conditional on the situation (win-stay, lose-shift) and could hence account for such difference. Alternatively, one could argue that, for unmodeled reasons, some best-reply decisions (e.g. win-stay) are less cognitively demanding than others.

Table 3.4: Panel regressions for log-transformed response times.

ln(response time)	Model 1	Model 2	Model 3
Stay	−0.2753*** (0.0412)	−0.2752*** (0.0412)	−0.2912*** (0.0411)
Lose	0.0086 (0.0366)	0.0037 (0.0421)	−0.0039 (0.0424)
Lose × Stay	0.2202*** (0.0491)	0.2205*** (0.0487)	0.2233*** (0.0476)
BR-Shift	−0.0633** (0.0321)	−0.0634** (0.0322)	−0.0680** (0.0313)
Lose × Regret		0.0013 (0.0074)	0.0024 (0.0077)
Normalized Round			−0.3446*** (0.0407)
Part 2 Dummy			−0.0790* (0.0409)
Part 3 Dummy			0.0224 (0.0503)
Constant	1.6494*** (0.0440)	1.6494*** (0.0441)	1.8643*** (0.0537)
R <sup>2</sup>	0.0256	0.0256	0.0555
Wald Test	105.5638***	111.5129***	298.8777***
Linear Combination Tests:			
(1) Lose+BR-Shift vs. Stay	0.2205*** (0.0292)	0.2155*** (0.0391)	0.2192*** (0.0414)
(2) Stay+Lose×Stay	−0.0551 (0.0410)	−0.0547 (0.0404)	−0.0678* (0.0384)
(3) BR-Shift vs. Stay+Lose×Stay	−0.0082 (0.0286)	−0.0087 (0.0277)	−0.0002 (0.0270)

Notes. Standard errors (clustered by 36 matching blocks) in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

faster response times in lose situations. Model 3 controls for learning effects as in the regression on choice data. As standard in choice experiments, participants became faster as their familiarity with the interface and the situation increased, as indicated by a significantly negative coefficient for Normalized Round. Participants also became faster after the first part (although the

effect is only weakly significant), but there was no improvement of response times in the third part compared to the first. Additional regressions including questionnaires and demographics (Table 3.9, Appendix 3.B) as controls yielded the same qualitative results. Further, an additional regression bootstrapping the standard errors with 100 repetitions (Table 3.10, Appendix 3.B) yielded the same conclusions.

### 3.4 Discussion

In our experiment, we examined a behavioral rule based on the simplest reinforcement principle, “win-stay, lose-shift,” as a possible cognitive shortcut for myopic best reply. After a previous win, the rule prescribes to repeat the previous choice. After a previous lose situation, the rule prescribes to shift away from it, and available information allows to focus on the specific shift favored by myopic best reply. Reinforcement processes are known to be automatic and associated with short response times, while myopic best reply, which involves explicit maximization, should be a more deliberative, slower process.

In win situations, win-stay was a strong driver of behavior. Such decisions occurred more often and faster than other actions (results confirmed by non-parametric tests both at the individual and block level and by panel regressions). Choice data alone could be justified by either reliance on an automatic, impulsive reinforcement process or a more cognitive, deliberative explicit maximization. However, win-stay decisions were the fastest decisions observed in our experiment, indicating that a more automatic process was at work compared to slower decisions. This leads to the interpretation that stay responses, even though they correspond to myopic bet response, most likely often followed a more automatic process.

In lose situations, shift rates (to the best reply) were sensitive to the magnitude of experienced regret, in agreement with reinforcement learning but in contrast to myopic best reply, which would predict the best-reply rates to be independent from the cardinal difference between experienced and

maximum possible payoffs.<sup>10</sup> Also, shifts to the myopic best reply were more frequent and faster than other choices, confirming the general interpretation of a more automatic underlying process, although response-time evidence was less strong than in the case of win situations.

Regression results on response times suggest that stay decisions (consistent with inertia) might actually be as fast as choices following reinforcement or best reply, in agreement with evidence on inertia as an additional, automatic process (Alós-Ferrer et al., 2016b). This also agrees with the observation that, in lose situations, stay decisions occurred more often than shifts which did not follow the maximum payoff. However, shifts to the myopic best reply were more frequent with higher regret, an effect incompatible with a pure best-reply interpretation, but perfectly aligned with general reinforcement processes in which the strength of the loss does play a role.

Our experiment goes beyond previous paradigms on win-stay, lose-shift rules (which have typically employed binary-choice settings) and shows evidence on the relevance of simple reinforcement processes in strategic settings. The evidence can and should be seen as a general caveat, in the sense that evidence in favor of myopic best reply in games might often be confounded with the workings of reinforcement, the basic building block of human learning (Sutton and Barto, 1998; Daw and Tobler, 2014; Achtziger et al., 2015a). In this and other cases of overlapping behavioral rules, the analysis of response times might provide valuable evidence.

---

<sup>10</sup>In this setting, reinforcement could be seen as a gradual, payoff-dependent version of best reply, giving rise to behavior in line with quantal response models (McKelvey and Palfrey, 1995).



## Appendix 3.A: Frequency and Response Times by Game

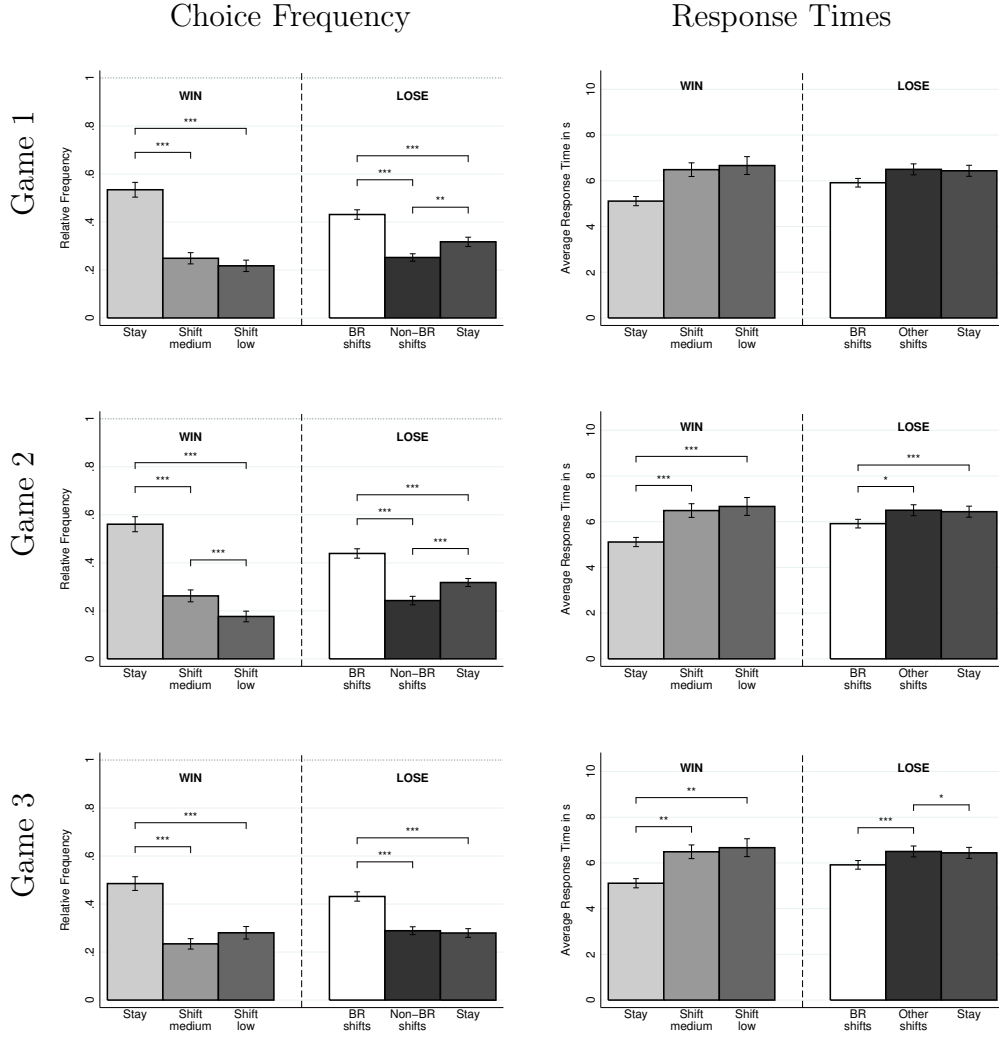


Figure 3.3: Average individual choice frequencies and response times by games.  
*Notes.* Choice frequencies and response times conditional on choice and situation for each game. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , WSR test.

## Appendix 3.B: Additional Regressions

Table 3.5: Random-effects panel probit regressions with additional controls.

Reinforcement decisions	Model 3	Model 4	Model 5
Win	0.3733*** (0.0849)	0.3729*** (0.0850)	0.3737*** (0.0835)
Lose $\times$ Regret	0.0423** (0.0167)	0.0422** (0.0166)	0.0421*** (0.0162)
Normalized Round	-0.1171 (0.0736)	-0.1173 (0.0738)	-0.1175 (0.0739)
Part 2 Dummy	0.0301 (0.0569)	0.0302 (0.0570)	0.0300 (0.0571)
Part 3 Dummy	-0.0323 (0.0581)	-0.0321 (0.0582)	-0.0322 (0.0583)
Female		-0.0908 (0.1003)	-0.0460 (0.1008)
Game Theory Knowledge		0.1279 (0.1157)	0.1193 (0.1193)
Age		-0.0092 (0.0071)	-0.0073 (0.0074)
Faith in Intuition			0.0003 (0.0265)
Openness to Experience			-0.0127 (0.0148)
Conscientiousness			0.0268 (0.0199)
Extraversion			0.0052 (0.0164)
Agreeableness			0.0033 (0.0185)
Neuroticism			-0.0327* (0.0183)
Regret-Scale			0.0245 (0.0198)
Maximization-Scale			-0.0161 (0.0224)
Constant	-0.2959*** (0.0887)	-0.0487 (0.2135)	-0.1677 (0.4408)
Log.Likelihood	-3421.2967	-3419.0832	-3414.9345
WaldTest	48.5771***	53.1115***	61.9844***

*Notes.* Additional controls are demographic factors (Model 4) and personality questionnaires (Model 5). The dummy Game Theory Knowledge takes the value 1 if the subject indicated that she attended a game theory lecture. The coefficients Openness to Experience, Conscientiousness, Extraversion, Agreeableness, and Neuroticism represent the score on each sub-scale of the Big Five questionnaire. Standard errors (clustered by 36 matching blocks) in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 3.6: Regression models while bootstrapping the standard errors.

Reinforcement decisions	Model 1	Model 2	Model 3
Win	0.2201*** (0.0635)	0.3737*** (0.0924)	0.3733*** (0.0912)
Lose $\times$ Regret		0.0425** (0.0181)	0.0423** (0.0178)
Normalized Round			-0.1171 (0.0799)
Part 2 Dummy			0.0301 (0.0655)
Part 3 Dummy			-0.0323 (0.0555)
Constant	-0.2075*** (0.0351)	-0.3606*** (0.0794)	-0.2959*** (0.0912)
Log.Likelihood	-3429.1594	-3424.0503	-3421.2967
WaldTest	11.9980***	17.0684***	22.2866***

Notes. Bootstrapping the standard errors (in parentheses) with 100 repetitions. \*  $p < 0.1$ ,

\*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 3.7: Regressions models for cumulative proportional reinforcement.

Cumulative proportional reinforcement	Model 1	Model 2	Model 3
Win	0.1334*** (0.0399)	0.0624 (0.0485)	0.0615 (0.0502)
Lose $\times$ Regret		-0.0198 (0.0122)	-0.0201 (0.0125)
Normalized Round			-0.1981*** (0.0657)
Part 2 Dummy			-0.0512 (0.0379)
Part 3 Dummy			-0.0748* (0.0382)
Constant	-0.2324*** (0.0216)	-0.1614*** (0.0506)	-0.0115 (0.0609)
Log.Likelihood	-3524.1636	-3522.9853	-3516.1829
WaldTest	12.7985***	15.1394***	28.6499***

Notes. The independent variable takes the value 1 if cumulative proportional reinforcement was followed. Standard errors (clustered by 36 matching blocks) in parentheses. \*  $p < 0.1$ ,

\*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 3.8: Panel regression for log-transformed response times including an additional dummy.

ln(response time)	Model 1	Model 2	Model 3
Stay	−0.2749*** (0.0526)	−0.2748*** (0.0524)	−0.2867*** (0.0522)
Lose	0.0089 (0.0446)	0.0041 (0.0548)	0.0004 (0.0530)
Lose × Stay	0.2198*** (0.0532)	0.2201*** (0.0532)	0.2189*** (0.0519)
BR-Shift	−0.0633** (0.0322)	−0.0634** (0.0323)	−0.0680** (0.0314)
Win and Shift to Medium	0.0006 (0.0546)	0.0007 (0.0544)	0.0079 (0.0530)
Lose × Regret		0.0013 (0.0073)	0.0024 (0.0077)
Normalized Round			−0.3446*** (0.0407)
Part 2 Dummy			−0.0790* (0.0410)
Part 3 Dummy			0.0225 (0.0502)
Constant	1.6491*** (0.0544)	1.6490*** (0.0544)	1.8599*** (0.0564)
R <sup>2</sup>	0.0256	0.0256	0.0554
WaldTest	105.6889***	111.7340***	307.1746***
Linear Combination Tests:			
(1) Lose+BR-Shift vs. Stay	0.2205*** (0.0292)	0.2155*** (0.0390)	0.2192*** (0.0413)
(2) Stay+Lose×Stay	−0.0551 (0.0410)	−0.0547 (0.0405)	−0.0679* (0.0384)
(3) BR-Shift vs. Stay+Lose×Stay	−0.0082 (0.0286)	−0.0087 (0.0277)	−0.0001 (0.0270)

*Notes.* Additional dummy (“Win and Shift to Medium”) takes the value 1 if in the win situations a shift to the medium payoff occurred. Standard errors (clustered by 36 matching blocks) in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 3.9: Panel regression for  $\ln(\text{response time})$  with additional controls.

$\ln(\text{response time})$	Model 3	Model 4	Model 5
Stay	-0.2912*** (0.0411)	-0.2910*** (0.0410)	-0.2909*** (0.0409)
Lose	-0.0039 (0.0424)	-0.0035 (0.0424)	-0.0034 (0.0423)
Lose $\times$ Stay	0.2233*** (0.0476)	0.2230*** (0.0475)	0.2230*** (0.0476)
BR-Shift	-0.0680** (0.0313)	-0.0681** (0.0313)	-0.0681** (0.0314)
Lose $\times$ Regret	0.0024 (0.0077)	0.0023 (0.0077)	0.0023 (0.0077)
Normalized Round	-0.3446*** (0.0407)	-0.3446*** (0.0407)	-0.3446*** (0.0408)
Part 2 Dummy	-0.0790* (0.0409)	-0.0790* (0.0409)	-0.0790* (0.0409)
Part 3 Dummy	0.0224 (0.0503)	0.0224 (0.0503)	0.0224 (0.0503)
Female		0.0194 (0.0649)	0.0441 (0.0661)
Game Theory Knowledge		0.0721 (0.0700)	0.0857 (0.0694)
Age		0.0111* (0.0062)	0.0119* (0.0072)
Faith in Intuition			-0.0119 (0.0186)
Openness to Experience			0.0029 (0.0113)
Conscientiousness			-0.0132 (0.0172)
Extraversion			0.0166 (0.0135)
Agreeableness			0.0160 (0.0104)
Neuroticism			-0.0015 (0.0132)
Regret-Scale			0.0032 (0.0169)
Maximization-Scale			0.0122 (0.0164)
Constant	1.8643*** (0.0537)	1.5855*** (0.1429)	1.4137*** (0.3513)
R <sup>2</sup>	0.0555	0.0626	0.0689
WaldTest	298.8777***	304.1417***	479.6298***
Linear Combination Tests:			
(1) Lose+BR-Shift vs. Stay	0.2192*** (0.0414)	0.2194*** (0.0415)	0.2195*** (0.0414)
(2) Stay+Lose $\times$ Stay	-0.0678* (0.0384)	-0.0681* (0.0384)	-0.0679* (0.0383)
(3) BR-Shift vs. Stay+Lose $\times$ Stay	-0.0002 (0.0270)	-0.0000 (0.0270)	-0.0002 (0.0269)

Notes. Additional control variables as in Table 3.5. Standard errors (clustered by 36 matching blocks) in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 3.10: Panel regression for log-transformed response times while bootstrapping the standard errors.

ln(response time)	Model 1	Model 2	Model 3
Stay	−0.2753*** (0.0382)	−0.2752*** (0.0383)	−0.2912*** (0.0367)
Lose	0.0086 (0.0355)	0.0037 (0.0407)	−0.0039 (0.0400)
Lose × Stay	0.2202*** (0.0449)	0.2205*** (0.0446)	0.2233*** (0.0427)
BR-Shift	−0.0633** (0.0309)	−0.0634** (0.0311)	−0.0680** (0.0290)
Lose × Regret		0.0013 (0.0071)	0.0024 (0.0074)
Normalized Round			−0.3446*** (0.0403)
Part 2 Dummy			−0.0790** (0.0394)
Part 3 Dummy			0.0224 (0.0477)
Constant	1.6494*** (0.0415)	1.6494*** (0.0416)	1.8643*** (0.0476)
R <sup>2</sup>	0.0256	0.0256	0.0555
WaldTest	155.6743***	171.6338***	358.9157***
Linear Combination Tests:			
(1) Lose+BR-Shift vs. Stay	0.2205*** (0.0280)	0.2155*** (0.0379)	0.2192*** (0.0405)
(2) Stay+Lose×Stay	−0.0551 (0.0403)	−0.0547 (0.0402)	−0.0678* (0.0369)
(3) BR-Shift vs. Stay+Lose×Stay	−0.0082 (0.0281)	−0.0087 (0.0274)	−0.0002 (0.0266)

Notes. Bootstrapping the standard errors (in parentheses) with 100 repetitions. \*  $p < 0.1$ ,

\*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

---

## CHAPTER 4

### Multiple Behavioral Rules in Cournot Oligopolies

---

#### 4.1 Introduction

How are economic decisions made? Neoclassical economics avoided addressing this question by focusing on the concept of preferences, whose existence as a formal object organizing choices is equivalent to the consistency of those very same choices (the Weak Axiom of Revealed Preference). That is, choices tautologically reflect preferences, because preferring one option to another just means that the former is chosen over the latter. This is of course a perfectly legitimate approach which identifies choices and the motives behind them and aggregates both into a single object of interest, but it fails to address the question of how choices arise.

This question becomes especially relevant once one moves away from the assumption of full rationality, or the proximate assumption of perfect maximization of stable preferences. This is of course not a new point. The microeconomics literature contains extensive and mature strands analyzing boundedly rational behavior in order to explain certain violations that cannot be accommodated by neoclassical models which are often used as the benchmark of normative behavior. For instance, evolutionary game theory has analyzed different behavioral rules capturing specific deviations from rationality, as e.g. imitation, satisficing behavior, or reinforcement learning. To date, however, these behavioral rules have been treated as “black boxes” and studied in isolation. These evolutionary branches have concentrated on the study of behavioral rules while the neoclassical literature has kept studying fully rational, maximizing agents.

Obviously, economic agents do try to maximize certain objectives. Even more obviously, they frequently fail to do so and follow different impulses instead. In other words, the truth might be in the middle, and it might be

worth to take both views into account when studying human behavior. In particular, it is not clear that an individual either acts fully rational or is always following a behavioral rule over a period of time. Huck et al. (1999) find none of the theoretical learning processes alone can describe observed behavior in their Cournot Oligopoly experiment, however, imitation and myopic best reply play a role for the decisions. Apesteguía et al. (2010) simulate data allowing for a mixture of multiple behavioral rules such as imitation, myopic best reply, fictitious play, and relative payoff maximization. With a relative large weight on imitation they find that the simulation of mixed rules describes their data best. With the idea in mind that the truth about how decisions are made might be in the middle between fully rational agents and heuristic following individuals, we set out in this work to show that multiple behavioral rules govern behavior in a complex economic setting. We will not only look at decisions outcomes itself but also at the characteristics of behavioral rules and present evidence that multiple decision rules codetermine behavior.

In this work, we target a specific behavioral rule, namely *imitation* of successful behavior, which is one of the most prominent boundedly-rational behavioral rules examined in the microeconomics literature. In order to study imitation, we select a standard economic paradigm, namely a Cournot oligopoly. The reason for this decision is that there is both theoretical and empirical literature showing the potential relevance of imitation in this paradigm. Schaffer (1989) showed that the Cournot-Nash equilibrium can be destabilized if firms consider relative payoffs, but the Walrasian equilibrium cannot. Vega-Redondo (1997) showed that if firms follow imitative behavioral rules and make infrequent mistakes, the system converges to the Walrasian equilibrium (in the sense of stochastic stability). Alós-Ferrer and Ania (2005) generalized this result to aggregative games by showing that the outcomes selected by aggregate-taking behavior (a generalization of price-taking behavior) have strong evolutionary stability properties leading to their stochastic stability when agents follow imitative rules. These results have been shown to be empirically relevant in the behavioral laboratory. A number of Cournot-oligopoly experiments (Huck et al., 1999; Offerman et al., 2002; Apesteguía



et al., 2007, 2010) have found partial convergence to Walrasian outcomes, which has then been interpreted as indirect evidence for the presence of imitative behavior.

Our approach is more direct in that we do not aim to examine convergence, but rather imitative decisions in themselves. Following the evidence above, we postulate the existence of (at least) two decision rules. The first, is imitation of observed, successful behavior. The second serves as a proxy for more rational behavior or at least payoff maximization. In the framework of a dynamic Cournot oligopoly, it is natural to identify this process with *myopic best reply*, that is, payoff maximization taking current information on other players' behavior as given. In our context it is reasonable to assume that imitation is a more a heuristic\behavioral rule, where individuals react to a more successful action and respond by imitating this action while myopic best reply, on the contrary, is more rational behavior and reflective process which involves active maximization after considering available information.

To examine the multiple behavioral rules we rely on a formal model, i.e. the dual-process diffusion model (Alós-Ferrer, 2018). We derive testable predictions from the formal drift-diffusion model which incorporates a dual-process theories. *Dual-process theories* postulate that the human mind is mainly influenced by two kinds of processes, called *automatic* and *controlled* (see Kahneman, 2003; Alós-Ferrer and Strack, 2014; see also Evans, 2008 for a detailed review). Automatic processes are defined as immediate, fast, unconscious, and efficient in the sense of requiring few cognitive resources. For instance, these processes capture impulsive reactions and behavior along the lines of stimulus-response schemes. In contrast, controlled processes are slow, consume cognitive resources, are not instigated immediately, and are reflected upon consciously. In other words, they are far closer to the economic idea of rationality. There is, hence, an analogy between how economics has addressed issues of full vs. bounded rationality and how psychology has modeled the origin of decisions. The key difference is that dual-process models assume heterogeneity within the individual, that is, decisions are the result of the interplay between different processes within the individual, while economics has, to date, kept the analysis of maximizing behavior and behavioral

rules in separate strands of the literature. In psychological parlance, imitation as bounded rational behavior can be expected to be *more* automatic than processes behind active preference maximization, i.e. myopic best reply. However, all the processes\behavioral rules we are interested in are skewed towards higher cognitive functions and the paradigms are typically more complex than the ones encountered compared to cognitive or social psychology, and the analysis is correspondingly more involved.

The theoretical dual-process diffusion model (Alós-Ferrer, 2018) delivers testable hypotheses for one of the most basic measures of process data, *response times*, with respect to two different processes, one being more automatic and one being more controlled. Response times are a standard tool in psychology which are slowly being incorporated into the economist’s toolbox (Wilcox, 1993; Moffatt, 2005; Rubinstein, 2007; Piovesan and Wengström, 2009; Achtziger and Alós-Ferrer, 2014; Alós-Ferrer et al., 2016a; Alós-Ferrer and Ritschel, 2018b). One of the basic insights of dual-process theories is that automatic processes are faster than controlled ones, and hence response times have been used as an important source of evidence for the involvement of different decision processes. This does *not* mean that one can simply classify decisions in fast and slow according to some exogenous criterion and conclude that one kind of decisions is more automatic. This is an example of “reverse inference” fallacy (Krajbich et al., 2015) and this is not our approach. In a nutshell, we will derive specific, non-trivial predictions (on response times conditional on specific types of choices) from a dual-process model and will not simply classify more automatic and more controlled processes by faster response times alone.

We conduct two laboratory experiments to test for those non-trivial predictions and analyze differences in the behavioral rules of imitation and myopic best reply in a standard economic setting. Experiment 1 tests the pure predictions of the formal dual-process model while Experiment 2 adds another dimension, i.e. cognitive load, to our setting. In addition, our design will also allow us to consider other likely determinants of behavior. The first is *positive reinforcement*, i.e. the tendency to repeat an action if it was successful (closely linked to the focus on past performance). The second is

*inertia*, which simply means the tendency to repeat the previous action regardless of the previous result. Positive reinforcement is actually a special case of imitation *and* inertia, that is, it corresponds to situations where the player imitates him- or herself. Nevertheless, the main analysis will still focus on the comparison of imitative decisions and myopic best replies.

We confirm in both experiments the following theoretical predictions. Whenever imitation and myopic best reply are in conflict (make different prescriptions), best replies are slower than imitative decisions. In contrast, in situations when imitation and myopic best reply are aligned (make the same prescription), best replies are faster than other responses. We do not find a significant shift in behavior induced by cognitive load, however, in a detailed analysis we show that the cognitive load manipulation affects imitation and myopic best reply in different ways strengthening the evidence for two distinct behavioral rules. Our results suggest that multiple behavioral rules codetermine behavior in a complex, economic setting such as a Cournot oligopoly with imitation being more automatic and myopic best reply being more controlled. Further, our analysis shows that positive reinforcement (imitating yourself) leads to faster decisions than imitating others, but decision inertia seems to play a minor role.

The remainder of the Chapter is structured as follows. Section 4.2 derives our predictions for response times in a Cournot oligopoly. Sections 4.3 and 4.4 present the design and the results of our two experiments. Sections 4.5 and 4.6 discuss the behavioral rules of reinforcement and inertia, respectively. Section 4.7 concludes.

## 4.2 Predictions for Multiple Behavioral Rules

We hypothesize that multiple behavioral rules codetermine behavior. When different behavioral rules are present they can either prescribe different actions or the same action. In the first case, we speak of *conflict* between the rules. For those situations it is easy to identify the different behavioral rules. In the second case, we speak of *alignment*. Alós-Ferrer (2018) proposed the *dual-process diffusion model* (DPDM) which delivers predictions on response

times conditional on situations of conflict or alignment, assuming one of the involved behavioral rule is more automatic, hence faster in expected terms, and the other is more controlled, hence slower in expected terms.

The DPDM assumes that both rules are stochastic in nature, i.e., they carry an amount of noise. For instance, an imitative rule will most of the time (say, with probability larger than  $1/2$ ) select the option which was successful last period. With the remaining probability, however, it might select some alternative action. Likewise, a myopic best reply rule will most of the time indeed select the myopic best reply, but with the remaining probability might select a different action. In the case of binary actions, the processes can be given a micro-foundation as diffusion processes as in the diffusion model of evidence accumulation (Ratcliff, 1978)<sup>1</sup> Alós-Ferrer (2018) shows that, assuming that automatic behavior is swifter than controlled behavior (in the sense of having a stronger trend as diffusion processes, i.e., behaving more in a stimulus-response manner), it follows that its expected response time is shorter. Further, it also follows that its noise level is smaller, that is, the probability that a more automatic behavioral rule deviates from its dominant response is smaller than the probability that a more controlled behavioral rule deviates from its own (possibly different) dominant response. Intuitively, controlled behavior “wanders around” more, while more automatic behavior quickly and efficiently selects a response (which might be a mistake).

For our purpose, the DPDM makes two basic predictions (see Achtziger and Alós-Ferrer, 2014 or Alós-Ferrer, 2018 for further details). For the sake of concreteness, identify the controlled/reflective behavior with more rational behavior and call the action prescribed by this rule “correct” and any other action an “error.” In many cases, these labels actually have a normative meaning, especially if the automatic behavior reflects a heuristic or bias confronted with another rule which is closer to normative behavior. In our case, “correct responses” are just myopic best replies. With this terminology, in case of alignment the more automatic behavioral rule is actually a quick,

---

<sup>1</sup>In this model, evidence accumulation (internal to the decision maker) is modeled as a diffusion process with a trend  $\mu$  and two barriers. Whether the process chooses an option or the other corresponds to whether the upper or the lower barrier is hit first. The response time is the time at which the first barrier is hit.

efficient shortcut to the correct response. In case of conflict, however, it is a quick path to an error (defined as a deviation from the reflective rule).

The first prediction of the DPDM is that, in case of conflict, correct responses are on average slower than errors arising from the automatic behavior. The intuition is simple, since in this case the automatic behavioral rule quickly selects an error, while the controlled rule more slowly leads (mostly) to the correct response. The second, somewhat more surprising prediction of the DPDM is that, in case of alignment, the reverse relation is expected. That is, in case of alignment, correct responses are on average faster than errors. The intuition is that in case of alignment, the automatic behavior is a quick and efficient shortcut to the correct responses, that is, it generates very few errors. The controlled behavior also mostly generates correct responses, but will lead to comparatively more errors than the automatic behavioral rule. Hence, conditional on an error being observed, it is more likely that the response is generated by the slower controlled behavioral rule.

A first test of the DPDM in an economic setting was carried out in Achtziger and Alós-Ferrer (2014), which also contained a simplified version of the model. That study considered a binary-action belief-updating paradigm as in Charness and Levin (2005) where two decision processes influence behavior as compared to our more complex economic setting of Cournot oligopolies. The setting endogenously created situations of conflict and alignment among both behavioral rules, and confirmed the predictions of the DPDM. The results were confirmed to be robust by using different variants of the basic paradigm and adding controls in a regression analysis.

In our experiments, participants take the role of producers in dynamic Cournot oligopolies. Following the evolutionary literature on this setting, we conceive behavioral rules as mappings from last period's information (the actions and profits of all players) to the following period's action (Kandori et al., 1993; Vega-Redondo, 1997; Alós-Ferrer and Ania, 2005). Hence, identifying conflict and alignment is a simple matter. Given the output levels and profits in a given period, myopic best reply identifies the action which would maximize profits given output levels of other players, while imitation simply identifies the action which led to the highest profits last period. If

the two actions are different, the behavioral rules are in conflict. If they are identical, the rules are aligned.

Following the DPDM, we can derive specific hypotheses on response times for the interaction of imitation and myopic best reply. We conceive imitation as an intuitive, heuristic, more automatic behavior and myopic best reply as our proxy for more “rational,” reflexive, controlled behavior. We hence immediately obtain the following two hypotheses.

**H1.** When imitation and myopic best reply are in conflict (prescribe different responses), correct responses (myopic best replies) will be slower than errors following from imitation.

**H2.** When imitation and myopic best reply are aligned, errors (all non-best replies) will be slower than correct responses (myopic best replies and imitation).

A further prediction of dual-process theories reflected in the DPDM is that decisions in case of conflict should take longer than decisions on case of alignment. This is a widely observed phenomenon reflecting the fact that conflict detection and resolution engages the brain’s central executive functions and is time-consuming (see Achtziger and Alós-Ferrer, 2014; Alós-Ferrer, 2018, for further details). Another prediction of the DPDM, in agreement with widespread experimental evidence, is the well-known Stroop effect (Stroop, 1935). In our setting, this effect states that the relation mentioned above holds when one conditions on correct responses, that is, correct responses in case of conflict should be slower than correct responses in case of alignment, even if the additional time needed for conflict detection and resolution is ignored. These additional predictions, show that deviations from myopic best reply is not just random variations, but follow from imitation.

For the sake of concreteness, we will always refer to conflict situations (or being in conflict) as situations when myopic best reply and imitation prescribe different responses. Similarly, we will always refer to alignment situations (or being aligned) as situations when myopic best reply and imitation prescribe the same responses. We will explicitly mention when we deviate from those definitions of conflict and alignment. Further, decisions which

follow myopic best reply (in conflict and alignment) favor the normative correct decisions and will be called correct (or correct decision) throughout this work. Decisions which do not follow myopic best reply and therefore are not correct will be called errors. Those errors can originate from multiple sources and we will label them for clarity. In conflict, errors can either be due to following imitation, i.e. “imitation” errors, or due to other kind of behavior which neither follows imitation nor myopic best reply, i.e. “noise” errors. In alignment, we do not further discriminate the errors according to any behavioral rule and we just label them errors.

### 4.3 Experiment 1

#### 4.3.1 Experimental Design and Procedures

In the main part of the experiment, participants interacted in 4-player Cournot oligopolies (tetrapolies). Each participant participated in three oligopolies (parts), with 17 periods each, which differed in the payoff table and group composition. For each part, we computed a payoff table using a Cournot oligopoly with zero costs and a linear inverse demand function of the form  $P(Q) = a - Q$ , where  $P$  is the price,  $a$  the saturated demand, and  $Q$  the total quantity in the market. However, during the experiment a neutral framing was used and neither firm nor quantities were mentioned. We reduced the action space to four possible actions ( $A$ ,  $B$ ,  $C$ , and  $D$ ). Hence, the whole payoff table has dimensions  $4 \times 20$ , with four rows representing the possible actions and 20 columns labeled  $AAA$  to  $DDD$  representing the possible actions of the opponents. Payoffs were expressed in points, with an exchange rate of 18 Eurocents per 1000 points. The points achieved in all 51 rounds were accumulated and paid at the end of the experiment. The payoff table, containing the profits rounded to the nearest integer, was permanently visible in the upper part of the screen during the corresponding part of the experiment, and hence contained all the information relevant for the (one shot) game.

We chose a setup with discretized actions to make the postulated behavioral rules (myopic best reply and imitation) both feasible and comparable. For instance, a continuous-action setup would have turned myopic best reply in an abstract maximization problem, while imitation would remain a discrete, intuitive rule. Hence, by choosing a discrete setup we act against our hypotheses and reduce the conceptual distance between the two postulated behavioral rules. While many previous experiments with Cournot oligopolies have focused on triopolies to increase the number of independent observations for a given number of participants (Offerman et al., 2002; Apesteguía et al., 2007), we chose to focus on tetrapolies because larger groups make collusion less likely and ensure higher outcome volatility (see Huck et al., 2004).

For all rounds, except the first within each part, participants were given feedback on the actions and profits of every group member in the previous period. The column corresponding to the joint actions of the opponents was highlighted. Myopic best reply corresponds to maximization within that column, that is, the mechanical part of determining a myopic best reply required comparing four numbers only. Hence, the highlighting reduced any time needed to identify the appropriate column of the other players' previous choices.

To make imitation feasible, information about all group members' choices and points earned in the previous round was presented in addition to the payoff table. There were two experimental treatments which differed only in how this additional information was presented. In Treatment *FullInfo*, the choice and points of the other group members were presented in separate boxes, and the box with the highest point amount was highlighted. In Treatment *BestOnly* only one (highlighted) box was shown, displaying the choice which received the highest points in the previous round and the amount of points. The treatments were implemented between subjects, with half the subjects in each treatment in each session and as a robustness test in order to make sure that there were no big differences in response times due to presentational effects (as we will see below, there were none). The *BestOnly* treatment reflects the idea of "imitate the best" as described by Vega-Redondo (1997),



i.e. giving the opportunity to choose the action with the highest profit in the previous round. The FullInfo treatment represents a robustness check with respect to myopic best reply. In the FullInfo treatment the mechanical part of both postulated rules, imitation and myopic best reply, require processing exactly the same number of quantities. Further, this treatment controls for the possibility that the presentation excessively primed participants to imitate.

The reason to have three parts with three different oligopolies was to *avoid* convergence. In contrast to previous experiments with Cournot oligopolies, we are interested in behavioral correlates of each individual action, rather than on eventual convergence. If convergence (or collusion) occurred, we would lose the necessary variance of the data to analyze behavioral rules, and response times would also become meaningless as most of the time participants would just be mechanically repeating a fixed action. To avoid this, we took three measures. First, the payoff tables of each part were different, computed with different demand functions and different quantities. The sequence of the payoff tables was varied across the sessions. Second, the ordering of the quantities from  $A$  to  $D$ , changed with each part, that is, in some parts the assignment of quantities to letters was increasing, in some it was decreasing.<sup>2</sup> The second and third parts of the experiment always had a different payoff table and a reversed ordering of the quantities with respect to the previous part. Third, for each of the three parts the individuals were rematched, with identities reassigned. In each new part, at least two players in the group were different from the previous group. To increase the number of independent observations for the most conservative approach, rematching was done within pre-determined blocks of 8 participants, hence there were a total of 16 fully independent blocks of participants in the experiment. Rematching, working with shorter oligopolies, and the changed payoff tables increase the variance in the behavioral data and diminish the possibilities of collusion.

---

<sup>2</sup>Payoff table 1:  $P(Q) = 150 - Q$ ,  $A = 37.5$ ,  $B = 33.25$ ,  $C = 30$ ,  $D = 18.75$  (or reversed); Payoff table 2:  $P(Q) = 175 - Q$ ,  $A = 43.75$ ,  $B = 38.875$ ,  $C = 35$ ,  $D = 21.875$  (or reversed); Payoff table 3:  $P(Q) = 200 - Q$ ,  $A = 50$ ,  $B = 44.5$ ,  $C = 40$ ,  $D = 25$  (or reversed).

The experiment was conducted at the Cologne Laboratory for Economic Research (CLER) and programmed with z-Tree (Fischbacher, 2007). We recruited 128 participants (82 females; median age 22 years) using ORSEE (Greiner, 2015). During recruitment, we excluded students majoring in economics, psychology, and business as such students might have been taught concepts as Nash equilibrium which might influence their behavior. Students who had already participated in 20 or more experiments were also excluded. A session comprised 32 participants and lasted around 90 minutes. Average earnings were 13.59 EUR, including a show-up fee of 2.50 EUR.

#### 4.3.2 Classification of Decisions

Table 4.1 displays the prescriptions of myopic best reply, imitation, and inertia in the experiment, for the case of decreasing assignment of quantities to letters.<sup>3</sup> Those prescriptions were identical for all three tables. That is, the table shows the prescription of each behavioral rule when a specific combination of one's own choice (row) and the choice of the other players (column) occurred in the previous round. Whenever myopic best reply is in alignment with imitation (that is, both prescribe the same action), the cell of the myopic best reply is shaded in gray. Hence, unshaded entries in the myopic best reply columns indicate conflict with imitation. For reference, we include also the prescriptions of inertia with shaded entries when they coincide with those of imitation. Note that an alignment of imitation and inertia means that the imitative action also follows positive reinforcement (imitating yourself).

With this design and Table 4.1, the identification of behavioral rules for any action is straightforward. For periods 2–17 within each part, we can classify each actual decision of each participant as compatible with myopic best reply, imitation, or inertia. The Venn diagram in Figure 4.1 shows the possible classifications of the decisions. Actions can be classified as imitation, myopic best reply, or inertia, as belonging to any of the intersections (alignments), or as being incompatible with all of them (unclassified).

---

<sup>3</sup>For the analysis of the data, the case of increasing assignment of quantities was simply recorded.

Table 4.1: Overview of prescribed actions.

AAA			AAB			AAC			AAD			ABB			ABC			ABD			ACC			ACD			ADD		
BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta
A	D	A	A	D	A	A	D	A	C	A	A	D	A	A	C	A	A	C	A	A	C	A	A	B	A	A	A	A	B
B	D	A	B	D	A	B	D	A	C	A	B	D	A	B	C	A	B	C	A	B	C	A	B	B	A	B	A	A	B
C	D	A	C	D	A	C	D	A	C	A	C	D	A	C	C	A	C	C	A	C	C	A	C	B	A	C	A	A	C
D	D	A	D	D	A	D	D	A	C	A	D	D	A	D	C	A	D	C	A	D	C	A	D	B	A	D	A	A	D
BBB			BBC			BBD			BCC			BCD			BDD			CCC			CCD			CDD			DDD		
BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta	BR	Imit	Inrta
A	C	A	A	C	A	B	A	A	C	A	A	B	A	A	A	A	A	C	A	A	A	A	A	A	A	A	A	A	A
B	C	B	B	C	B	B	B	B	C	B	B	B	B	B	B	A	B	C	B	B	A	B	B	A	B	A	B	B	B
C	C	B	C	C	B	C	B	C	C	B	C	B	B	C	A	B	C	C	C	C	A	C	C	C	A	C	A	C	C
D	C	B	D	C	B	D	B	D	C	B	D	B	B	D	A	B	D	C	C	D	A	C	D	A	C	D	A	D	D

*Notes.* Overview of prescribed actions for each behavioral rule depending on last period's outcome. Cell entries describe the action prescribed by myopic best reply (BR), imitation (Imit), and inertia (Inrta) when the player previously chose the action given in the row and the opponents chose the actions given in the column. Shaded entries indicate that myopic best reply or inertia are aligned with imitation.

### 4.3.3 Results

The whole data set of choices and response times consists of  $128 \times 48 = 6144$  observations. The first decision within each block is always excluded since for that period there is no feedback concerning previous actions and the behavioral rules considered make no prescriptions.

In order to test our hypotheses, we will initially conduct non-parametric tests. For instance, we can test whether decisions compatible with one behavioral rule are faster than those compatible with another decision rule, conditional, e.g., on conflict among the rules. To do so, we look at all situations where the two rules conflict and build two sets of decisions, those where the prescription of the first rule was followed, and those where the prescription of the second rule was followed. Then we apply the appropriate test (in this case, a Wilcoxon Signed-Rank test).

At this point, we would like to argue that the appropriate unit of observation is the individual participant (but more conservative readers should just wait). The reason is that we are following the logic of evolutionary models (Kandori et al., 1993; Vega-Redondo, 1997; Alós-Ferrer and Ania, 2005), where behavioral rules have a Markovian structure. Under the assumption that we are observing response times generated by behavioral rules, and since those are mere mappings from information (outputs and profits in last period) to actions, it is thoroughly irrelevant how exactly the input of the be-




Figure 4.1: Classification of actions.

havioral rule is generated. The fact that participants were part of tetrapolies which themselves were subgroups of certain blocks plays no role, for we are testing response times which are generated after observation of the input. Following this logic, the appropriate tests consider  $N = 128$  individual sets of observations (participants). For each of those, we compute the average response time when following a given behavioral rule, conditional on conflict or alignment as appropriate, and the average response time when following a different rule, conditional on the same case. We put an additional restrictions on creating those averages by requiring at least two observations of an individual. This might exclude some subjects for the analysis but yields a better interpretation of an individuals average which will not consist of just one observation. We will report in case this additional requirement changes our result.

Econometrically conservative readers will object that the appropriate unit of observation is the group or even the block, because the inputs of decisions depend on the output of decisions of other block members from previous periods. We disagree with this view, but the point is moot. In the sequel, and as a robustness check, we will always report the tests at the group level (that is, where the unit of observation is the block, hence  $N = 16$ ). Since participants were separated into 16 different blocks and they were rematched

only within these blocks, this created completely independent observations. The conclusions always remain unchanged when we test within blocks unless it is stated otherwise. Further, we will later turn to a more detailed regression analysis.

Before we proceed, we remark that the introduction of two informational treatments served merely as a robustness check and there were no significant differences in behavior or response times. Recall that a possible concern was the excessive priming for imitation in the BestOnly treatment. Participants made on average 15.64 imitation decisions in the BestOnly treatment and 15.17 in the FullInfo treatment, which was not significantly different according to a Mann-Whitney-Wilcoxon (MWW) test ( $N = 128$ ,  $z = 0.141$ ,  $p = 0.8880$ ). Hence, there was no priming effect for imitation in the BestOnly treatment. Also, any concerns that the amount of numbers needed to be processed for imitation could cause differences in response times proved to be unfounded. The average response time of imitation decisions was 11.27 s in the BestOnly treatment and 10.89 s in the FullInfo treatment ( $N = 127$ ,  $z = 0.993$ ,  $p = 0.3206$ ). That is, we observe no relevant differences across the treatments and conclude that no presentational effects were present.<sup>4</sup> The BestOnly treatment did not excessively prime the subjects to imitate and did not cause mechanically faster response times for imitation due to different amount of numbers processed. Hence, for the rest of the analysis we will not distinguish them anymore.

Figure 4.2 gives a descriptive overview of the observations and their raw classification according to the behavioral rules of imitation, myopic best reply, and inertia. Table 4.2 shows the classification according to conflict and alignment and reveals that overall the majority of observations happened in conflict situations while only few alignment situations are available for the analysis.

---

<sup>4</sup>We also do not find significant differences between treatments for other behavioral rules. The same is true when testing at the block level except for inertia showing weakly significant more inertia decisions in the BestOnly (mean 168.88) than in the FullInfo treatment (mean 142.88; WSR,  $N = 16$ ,  $z = 1.682$ ,  $p = 0.0927$ ). Further tests conditioning on conflict and alignment find no differences.

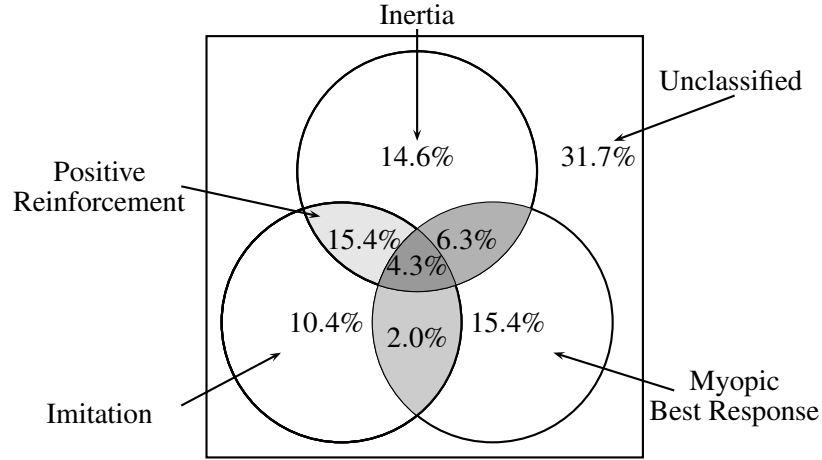


Figure 4.2: Experiment 1, Overview of observations and their classification according to different behavioral rules.

Table 4.2: Experiment 1, Overview of observations in conflict and alignment situations.

Total	6,144	100.00%
Conflict	5,010	81.54%
correct	1,331	21.66%
imitation errors	1,585	25.80%
noise errors	2,094	34.08%
Alignment	1,134	18.46%
correct	387	6.30%
errors	747	12.16%

We now turn to our main hypotheses. H1 states that in case of conflict, correct decisions should be slower than errors following imitation.<sup>5</sup> We hence look at all situations where myopic best reply and imitation conflict, and build two sets of decisions, those where the prescription of myopic best reply was followed, and those where the prescription of imitation was followed. For the purposes of this test, decisions following neither imitation

<sup>5</sup>Note that H1 and H2 yield specific directional predictions which allow us to test the hypothesis using one-sided tests. However, we report two-sided *p-values* for the sake of erring conservatively.

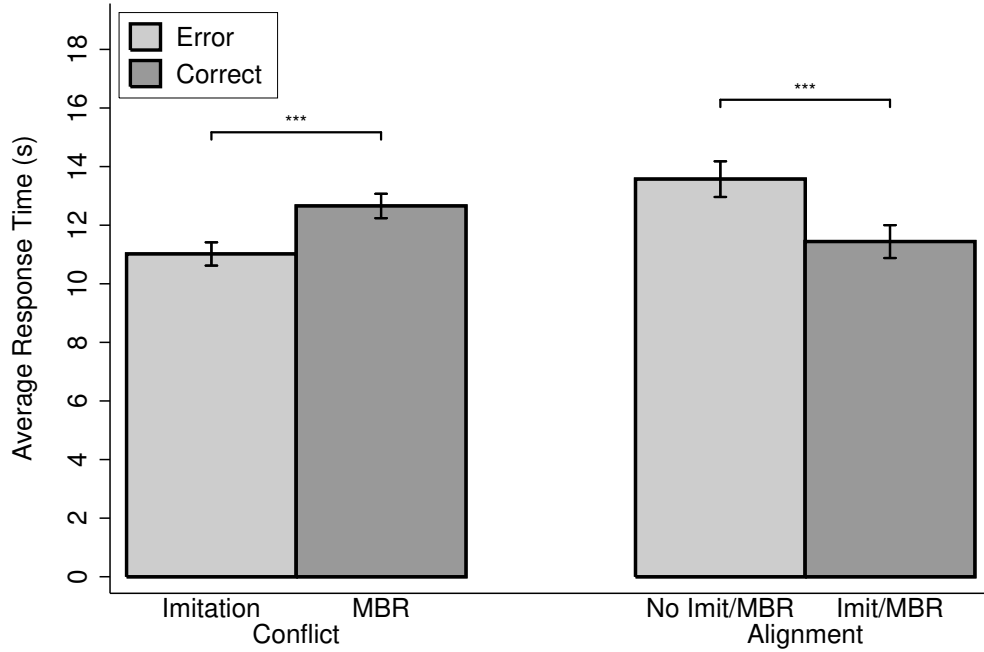


Figure 4.3: Experiment 1, Average response times of correct responses and errors. *Notes.* Left-hand side: Average response times of correct responses (myopic best replies) and errors in line with following imitation in case of conflict. Right-hand side: Average response times of correct responses and errors in case of alignment. Bars represent 1 Standard Error of the Mean. Stars indicate Wilcoxon Signed-Rank tests. \*\*\*  $p < 0.01$ .

nor best reply (noise errors) are discarded. The left-hand side of Figure 4.3 shows the response times averaged over subjects of the correct responses and the imitation errors in case of conflict. As predicted, imitation errors are significantly faster than correct responses in conflict. The response time is 11.02 s for imitation errors and 12.66 s for correct responses. The difference is highly significant according to a Wilcoxon Signed-Rank (WSR) test ( $N = 127$ ,  $z = -5.063$ ,  $p < 0.0001$ ).<sup>6</sup>

Hypothesis H2 states that in case of alignment, correct decisions should be faster than errors (not following myopic best reply and imitation). We hence look at all situations in alignment, and build two sets of decisions,

<sup>6</sup>The final number of observations deviates from 128 because not all subjects had at least two decisions following imitation and myopic best reply.

those where the common prescription of myopic best reply and imitation was followed, and those where something else was decided. The right-hand side of Figure 4.3 shows the response times averaged over subjects of the correct responses and the errors in case of alignment. As predicted, errors are significantly slower than correct responses. The response time is 13.57 s for errors and 11.44 s for correct responses. The difference is again highly significant (WSR,  $N = 84$ ,  $z = 3.233$ ,  $p = 0.0012$ ).

In summary, we confirm the main predictions of the model of response times in Achtziger and Alós-Ferrer (2014) and Alós-Ferrer (2018). Hence, our experimental evidence is compatible with the interpretation that decisions in our experimental Cournot oligopolies were the result of the interaction of two behavioral rules, one more automatic or intuitive corresponding to imitation, and one more controlled or reflexive corresponding to myopic best reply.

We now turn to a more detailed regression analysis. Our data forms a perfectly balanced panel with 48 decisions per participant ( $N = 48 \times 128 = 6,144$  in total). Hence we turn to random effects panel regressions on response times, transformed logarithmically.<sup>7</sup> We further (conservatively) cluster standard errors at the block level.

We consider the conflict and alignment setup between imitation and myopic best reply and the following categories, identified by dummies in the regressions, emerge. The reference group consists of the correct responses in case of alignment. The dummy Conflict indicates that the decision corresponds to a case of conflict between myopic best reply and imitation. In order to directly test for our hypotheses without the recourse to post hoc tests, we “split a dummy” and include the categories Error  $\times$  Conflict and Error  $\times$  Alignment.<sup>8</sup> The category Error  $\times$  Conflict corresponds to all errors in case of conflict, that is, decisions where imitation was in conflict with myopic best reply and the latter was not followed. This category includes both imitation errors and noise errors. The dummy Noise further indicates the non-imitative

<sup>7</sup>The transformation yields distribution which is not as right-skewed and closer to a normal distribution as the regular response times.

<sup>8</sup>That is, we choose an unconventional presentation of regression results for ease of exposition. A more conventional presentation delivers exactly the same conclusions but one needs to recompute the same coefficients through post hoc linear combination tests.



Table 4.3: Experiment 1, Random effects panel regressions for  $\ln(\text{ResponseTimes})$ .

$\ln(\text{ResponseTimes})$	Model 1	Model 2	Model 3
Conflict	0.1154*** (0.0360)	0.1371*** (0.0334)	0.1375*** (0.0333)
Error x Conflict	-0.1809*** (0.0307)	-0.1420*** (0.0224)	-0.1428*** (0.0229)
Error x Aligned	0.1794*** (0.0424)	0.1705*** (0.0363)	0.1700*** (0.0365)
Noise (in Conflict)	0.1794*** (0.0338)	0.1299*** (0.0238)	0.1299*** (0.0242)
Full Info Treatment		-0.0290 (0.0543)	-0.0298 (0.0543)
Normalized Time		-0.3839*** (0.0327)	-0.3839*** (0.0327)
Part 2		-0.2211*** (0.0259)	-0.2299*** (0.0224)
Part 3		-0.3653*** (0.0290)	-0.3640*** (0.0261)
Payoff Table 2			-0.0168 (0.0233)
Payoff Table 3			0.0267 (0.0289)
Collusion		-0.3597*** (0.0283)	-0.3996*** (0.0251)
Constant	2.2313*** (0.0417)	2.6259*** (0.0536)	2.6262*** (0.0554)
R <sup>2</sup>	0.0310	0.1332	0.1343

Notes. Standard errors, clustered by 16 matching blocks, in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

errors in case of conflict, that is, decisions which followed neither imitation nor myopic best reply when both rules were in conflict (remember that these decisions had to be excluded when testing for H1 above). The category Error  $\times$  Alignment corresponds to all errors in case of alignment, that is, decisions which did not agree with imitation and myopic best reply in the cases where those two rules prescribed the same action.

Model 1 in Table 4.3 tests for the basic effects. The coefficient for  $\text{Error} \times \text{Conflict}$  directly represents the difference between imitation errors and correct decisions in case of conflict. The coefficient is negative and highly significant, confirming Hypothesis H1. The coefficient for  $\text{Error} \times \text{Alignment}$  directly represents the difference between errors and correct decisions in case of alignment. The coefficient is positive and highly significant, confirming Hypothesis H2. The coefficient for  $\text{Conflict}$  is positive and highly significant. This shows that correct decisions are slower in case of conflict than in case of alignment, which is essentially the translation of the Stroop effect to our paradigm and is also predicted by the DPDM (Alós-Ferrer, 2018).

The Noise coefficient is also positive and highly significant showing that noise errors are slower than the imitation errors. While we had no strong hypotheses for this category, this could be an indication that the category of noise errors is capturing higher cognitive, more complex decisions or behavioral rules, as e.g. level- $k$  considerations (best-replying to the anticipated best reply of others, etc; see Alós-Ferrer and Buckenmaier, 2018).

Models 2 and 3 add further controls and show that the results just reported are robust. We find learning effects, i.e. the participants become faster the longer they play. This is reflected by the variable Normalized Time (a fraction indicating how many rounds from a specific part had already been completed) and the dummies for parts 2 and 3. Additionally, we control for the possibility of some groups colluding by adding a dummy for groups where collusion was observed. This, however, only affected one group which successfully maintained collusion in the very last part of the experiment. As expected, their decisions were faster, but the inclusion of this dummy does not affect other results. The varying payoff tables also did not affect the response times of the participants.<sup>9</sup>

---

<sup>9</sup>We also ran regressions with demographics and psychological characteristics which yielded the same qualitative results and conclusions for our analysis. After the Cournot oligopoly game was completed, participants were asked to complete a questionnaire including demographics, the Faith in Intuition scale (Alós-Ferrer and Hügelschäfer, 2012, 2016), and a 15-item Big-Five questionnaire (see John et al., 1991; Lang et al., 2011, p. 560).

The regression models confirm our non-parametric analysis while controlling for other important features. Experiment 1 confirms our Hypotheses H1 and H2 suggesting that two interacting behavioral rules with different properties codetermine behavior, one being more automatic and bounded rational (imitation) and the other being more controlled and rational (myopic best reply).

## 4.4 Experiment 2

### 4.4.1 Cognitive Load

The results of Experiment 1 suggest two interacting decision rules codetermine behavior in a complex, dynamic setting such as the Cournot Oligopoly. We found the response times asymmetry as predicted by the formal drift-diffusion model and conclude that imitation, in a dual-process view, is a more automatic, intuitive behavioral rule and myopic best reply a more controlled, deliberative rule. Dual-process theories distinguish automatic and controlled processes along the dimension of consuming cognitive resources. Automatic processes require only few cognitive resources while controlled processes are consciously reflected upon and consume relatively more cognitive resources. In this second experiment we want to further investigate the different nature of imitation and myopic best reply. We, therefore, set out to manipulate the cognitive resources available by implementing *cognitive load* and analyze changes in behavior.

Manipulating cognitive resources is based on the idea that if the working memory is exhausted, e.g. by an additional cognitive load task, the performance decreases. Working memory is a brain system which provides temporary storage and processing of information. Baddeley (1992) proposes a refined model that divides working memory into multiple components. The most prominent system is the central executive which is the attentional-controlling system that, among other features, is responsible for coordinating information. If the central executive is taxed and working memory is

exhausted then processes which require cognitive resources cannot be exerted.

Exhausting the working memory is implemented by adding a secondary task surrounding the main decision task. Cognitive load tasks may consist of memorizing a number sequence (targeting the phonological loop) or visual pattern (targeting the visuo-spatial sketchpad) which has to be recalled after the main decision has been made. These tasks are frequently used in psychology and recently found its way into economics (Carpenter et al., 2013). Carpenter et al. (2013) manipulated the cognitive resources by asking their subjects to memorize a seven-digit number in the cognitive load treatment. Subjects in the load treatment subsequently performed worse in a number strategic games.

With respect to imitation and myopic best reply, the dual-process view leads to the following predictions. Taxing the central executive with cognitive load results in a shift to more automatic behavior and we will observe a higher frequency of imitation decisions in the cognitive load treatment.

**H3.** Under high cognitive load, more imitative decisions will be observed.

In addition, the NoLoad treatment includes a replication of FullInfo treatment of Experiment 1 and the Load treatment itself includes a robustness check of the theoretical predictions tested in Experiment 1.

#### 4.4.2 Experimental Design and Procedures

The cognitive load manipulation was implemented between subjects in two treatments, *NoLoad* and *Load*. The NoLoad treatment was a pure replication of the FullInfo treatment of Experiment 1. The Load treatment had an additional cognitive load task, but was in any other aspect just as the NoLoad treatment (and FullInfo treatment of Experiment 1). Compared to Experiment 1, we increased the exchange rate to 20 Eurocents per 1000 points<sup>10</sup> and the subjects were rematched within blocks of 12 participants after each part. The participants were rematched within a larger pool which

---

<sup>10</sup>The exchange rate was increased because the average payoff in Experiment 1 was slightly below the wage rate demanded by experimental lab.

further decreased the probability of collusion.<sup>11</sup> Besides those points, the most prominent change being the cognitive load task in the Load treatment, there were no other differences compared to the FullInfo treatment of Experiment 1 and the participants played 51 periods in three different tetrapolies, with 17 periods each.

The cognitive load task consisted of memorizing a seven-digit number (as in Carpenter et al., 2013) before each decision in the Cournot Oligopoly. The subjects had 10 s to memorize the number and had to recall the number after the decision in the Cournot Oligopoly. A correct recall of the number was rewarded with additional 750 points. We implemented the treatments in separate sessions each containing one treatment because of the additional task in the Load treatment.

The experiment was conducted at the Cologne Laboratory for Economic Research and programmed with z-Tree. We recruited 144 participants (57 females; median age 23 years) using ORSEE. During recruitment, we excluded students majoring in economics, psychology, and business and students who had already participated in 20 or more experiments. Six sessions in short succession comprised each 24 participants. The sessions in the NoLoad treatment lasted around 1 hour and 25 minutes while the sessions in the Load treatment took longer and lasted around 1 hour and 45 minutes. Average earnings, including the show-up fee of 2.50 EUR, were 13.61 EUR and 20.12 EUR for the NoLoad and Load treatment, respectively. The participants in the Load treatment earned more than the participants in the NoLoad treatment due to the additional earnings of the cognitive load task. After excluding the earnings of the cognitive load task, the average earnings in the Load treatment was 14.06 EUR, including the show-up fee.<sup>12</sup>

---

<sup>11</sup>We are aware that this also reduces the number of the most conservative independent block observations but our focus lies on the subject analysis based on the micro-foundations of the DPDM (Alós-Ferrer, 2018).

<sup>12</sup>There were no significant differences in the earnings from the decision task between the two treatments (MWW test,  $N = 144, z = -1.489, p = 0.1365$ ).

#### 4.4.3 Cognitive Load and Behavior

Before we proceed with the analysis of cognitive load, we establish that the manipulation was working. The participants took the cognitive load task seriously and found it be of medium difficulty. At the end of the experiment we asked how important and how difficult the “additional task,” i.e. the cognitive load task, was. The subjects considered the task to be of importance with a mean answer of 8.73 on a scale ranging from 0 (not at all important) to 10 (very important). The difficulty of the task was evaluated with a mean of 4.58 also on a scale ranging from 0 (not at all difficult) to 10 (very difficult). The medium difficulty of the task was also reflected in the error rates. On average 21% of all recalls of the seven-digit number were incorrect.

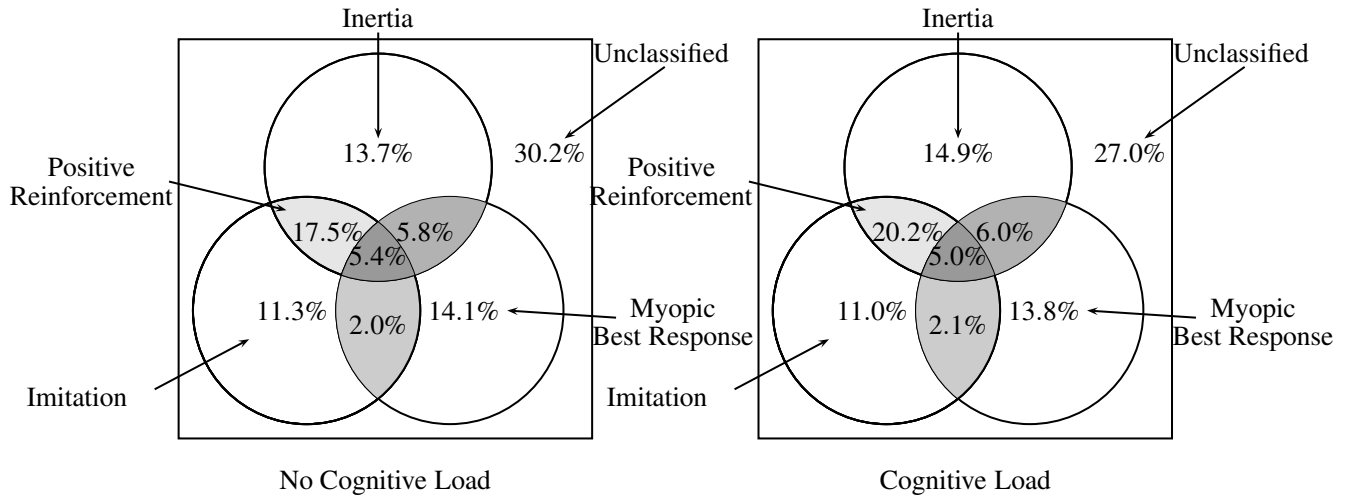


Figure 4.4: Experiment 2, Overview of observations and their classification according to different behavioral rules.

*Notes.* The left-hand side side depicts the no cognitive load treatment (NoLoad) and the right-hand side depicts the cognitive load treatment (Load).

With the design as in Experiment 1 and Table 4.1, the identification of the behavioral rules is straightforward. Figure 4.4 gives a descriptive overview of the observations by treatment and their raw classification according to the behavioral rules of imitation, myopic best reply, and inertia. H3 states that under cognitive load, more imitative decisions will be observed. We hence

Table 4.4: Experiment 2, Overview of observations in conflict and alignment situations.

Total	No Cognitive Load		Cognitive Load	
	3,456	100.00%	3,456	100.00%
Conflict	2,864	82.87%	2,892	83.68%
correct	688	19.91%	684	19.79%
imitation errors	996	28.82%	1,078	31.19%
noise errors	1,180	34.14%	1,130	32.70%
Alignment	592	17.13%	564	16.32%
correct	257	7.44%	247	7.15%
errors	335	9.69%	317	9.17%

look at the number of decisions following imitation in the NoLoad and Load treatment.

Pure imitation, i.e. decisions that follow only imitation and no any other here identified behavioral rule, did not significantly increase under cognitive load. Participants made on average 5.44 pure imitation decisions in the NoLoad treatment and 5.28 in the Load treatment, which was not significantly different according to a Mann-Whitney-Wilcoxon (MWW) test ( $N = 144$ ,  $z = 0.229$ ,  $p = 0.8187$ ). The same conclusion is obtained when looking at overall imitation. Participants made on average 17.40 imitation decisions in the NoLoad treatment and 18.40 in the Load treatment (MWW,  $N = 144$ ,  $z = -0.450$ ,  $p = 0.6527$ ). Hence, there was no significant increase in imitation decisions due to cognitive load.<sup>13</sup>

Table 4.4 gives an overview of all decisions split by conflict and alignment for each treatment and already indicates a strong similarity between both treatments. H3 predicts an increase in imitation under cognitive load which translates to more (imitation) errors in conflict and more correct decisions in alignment situations under cognitive load than under no load. The left-hand side of Figure 4.5 shows the relative frequencies of the errors and the correct responses for both treatments conditional on being in conflict. Imitation

<sup>13</sup>Moreover, further tests also found no differences between treatments for myopic best reply, positive reinforcement, inertia, or unclassified decisions.

errors are not significantly different across treatments. In case of conflict, there were 34.97% imitation errors in the NoLoad and 37.49% imitation errors in the Load treatment (MWW test,  $N = 144$ ,  $z = -0.452$ ,  $p = 0.6516$ ).<sup>14</sup> The right-hand side of Figure 4.5 shows the relative frequencies of the correct responses for both treatments conditional on being in alignment. As in the case of conflict, the correct responses are not significantly more frequent in the NoLoad treatment (43.11%) than in the Load treatment (40.21%; MWW,  $N = 144$ ,  $z = 0.713$ ,  $p = 0.4761$ ).<sup>15</sup>

In summary, we find no evidence supporting Hypothesis H3. Not only do we not find any significant differences in frequencies for myopic best reply, but also no change in behavior for other behavioral rules.

#### 4.4.4 Replication and Robustness

Experiment 2 manipulates cognitive resources but also serves as a replication of the FullInfo treatment of Experiment 1 and a robustness check of Hypotheses H1 and H2 under cognitive load. We follow the same steps as in Experiment 1 and analyze the data for each treatment, NoLoad and Load, separately. We calculated the average response times for each subject when she made a correct response or error when myopic best reply and imitation are in conflict or in alignment. Recall, H1 predicts faster errors in conflict situations and H2 predicts slower errors in alignment situations.

The left-hand side of Figure 4.6 shows the response times for conflict and alignment situations for the NoLoad treatment. As predicted by H1, imitation errors (11.79 s) are significantly faster than correct responses (14.67 s) in conflict situations. The difference is highly significant according to a Wilcoxon Signed-Rank test ( $N = 71$ ,  $z = -4.991$ ,  $p < 0.0001$ ).<sup>16</sup>

---

<sup>14</sup>The correct responses in case of conflict are also not significantly across treatments (NoLoad, 23.81% correct responses; Load 23.87% correct responses; MWW,  $N = 144$ ,  $z = 0.210$ ,  $p = 0.8338$ ).

<sup>15</sup>The analysis conditioning only on responses when the cognitive load task was recalled correctly yields the same conclusion.

<sup>16</sup>The final number of observations deviates from 72 because not all subjects had at least two decisions following imitation and myopic best reply.



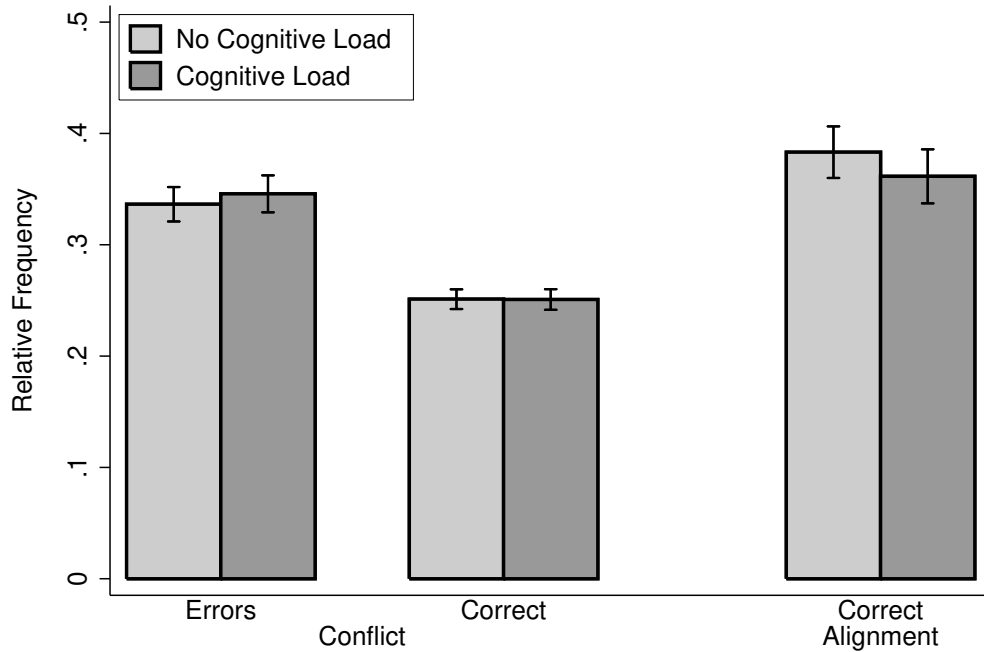


Figure 4.5: Experiment 2, Relative frequency of errors and correct responses for NoLoad and Load treatments.

*Notes.* Left-hand side: Correct responses and imitation errors in case of conflict. Right-hand side: Correct responses in case of alignment.

As predicted by H2, errors (14.79 s) are slower than correct responses (13.08 s) when myopic best reply and imitation are aligned. The difference is significant according to a WSR test despite fewer observations ( $N = 44$ ,  $z = 1.926$ ,  $p = 0.0542$ ).<sup>17</sup> In summary, we replicate the results of Experiment 1 in the NoLoad treatment with fewer numbers of observations.

The right-hand side of Figure 4.6 shows the response times in conflict and alignment situations for the Load treatment. As predicted by H1, imitation errors (8.27 s) are significantly faster than correct responses (10.75 s) in conflict situations. The difference is highly significant according to a Wilcoxon Signed-Rank test ( $N = 70$ ,  $z = -4.878$ ,  $p < 0.0001$ ). As predicted by H2,

<sup>17</sup>If we do not require at least two observation to create an individual's average the resulting test yields only a non-significantly trend (mean errors 14.47 s, mean correct responses 13.17 s; WSR,  $N = 61$ ,  $z = 1.512$ ,  $p = 0.1305$ ).

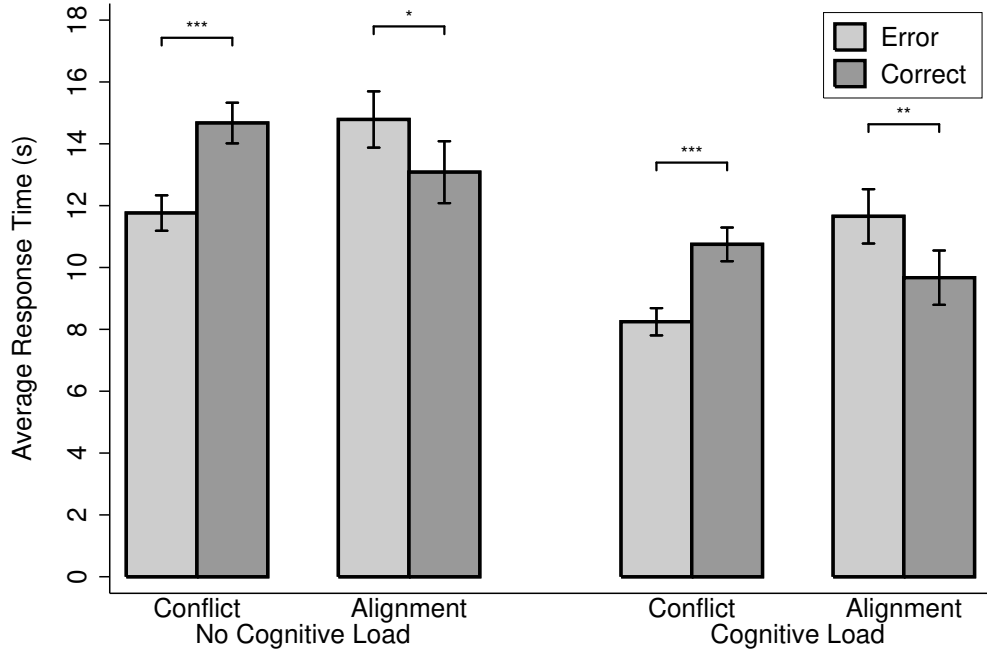


Figure 4.6: Experiment 2, Average response times of correct responses and errors.

*Notes.* Left-hand side: Average response times of correct responses and errors in case of conflict and alignment in the NoLoad treatment. Right-hand side: Average response times of correct responses and errors in case of conflict and alignment for the Load treatment.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , WSR test.

errors (11.65 s) are significantly slower than correct responses (9.67 s) in alignment. The difference is significant (WSR,  $N = 35$ ,  $z = 2.375$ ,  $p = 0.0176$ ). In summary, we also confirm H1 and H2 in the Load treatment serving as a robustness check. The evidence presented above shows that the interaction of two behavioral rules, imitation being more automatic and myopic best reply being more controlled, is robust to cognitive load manipulations and codetermines behavior in a complex, dynamic setting such as Cournot oligopolies.

Before we turn to the regression models as in Experiment 1, we want to point out an interesting finding regarding the response times between the two treatments. Figure 4.7 simply reorders the average response times as shown

in Figure 4.6 and put for each decision the NoLoad and Load treatments next to each other. The picture clearly shows that participants made their decisions significantly faster in the Load than in the NoLoad treatment. A MWW test confirms that response times are significantly faster in the Load than in the NoLoad treatment.<sup>18</sup>




Figure 4.7: Experiment 2, Average response times for the NoLoad and Load treatments.

*Notes.* Left-hand side: Conflict situations for the NoLoad and Load treatments. Right-hand side: Alignment situations for the NoLoad and Load treatments. \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , MWW test.

These response time results in combination with the findings of non-significant differences in (choice) behavior show that the additional cognitive load task made the participants faster while not changing their behavior. This finding has also been explored in a series of cognitive load studies in

<sup>18</sup>Conflict situations: errors ( $p < 0.0001$ ) and correct responses ( $p < 0.0001$ ); Alignment situations: errors ( $p = 0.0291$ ) and correct responses ( $p = 0.0018$ ).

Achtziger et al. (2019), who also offer an explanation and propose it as a test of a successful cognitive load manipulation. The cognitive load, i.e. the additional task of maintaining a number in the working memory, did not slow down the participants but sped up the decision while not changing the choice behavior.

We now turn to a more detailed regression analysis as in Experiment 1 and look at the random effects panel regressions on the log-response times while clustering the standard errors at the block level. We consider the same regression models as in Experiment 1 with the one exception of a different treatment dummy. Starting with regression model 2, the Cognitive Load dummy, i.e. being 1 in the Load treatment and 0 otherwise, is introduced to the regressions.

Model 1 in Table 4.5 tests for the basic effects. The coefficient for  $\text{Error} \times \text{Conflict}$  directly represents the difference between imitation errors and correct decisions in case of conflict. The coefficient is negative and highly significant, confirming Hypothesis H1. The coefficient for  $\text{Error} \times \text{Alignment}$  directly represents the difference between errors and correct decisions in case of alignment. The coefficient is positive and highly significant, confirming Hypothesis H2. The coefficient for Conflict is positive and highly significant. This shows that correct decisions are slower in case of conflict than in case of alignment.

The Noise coefficient is also positive and highly significant showing that noise errors are slower than the imitation errors. As in Experiment 1, we take this as an indication that the category of noise errors is capturing higher cognitive, more complex decision processes or behavioral rules.

Models 2 and 3 add further controls and show that the results just reported are robust. The dummy called Cognitive Load is negative and highly significant showing that the decisions in the Load treatment are significantly faster than the decisions in the NoLoad treatment as seen in the non-parametric analysis. We find a significant time trend, i.e. the participants become faster the longer they play reflected by the variable Normalized Time and dummies for parts 2 and 3. As in Experiment 1, we interpret this as familiarity with the interface or other learning effects, however, our main

Table 4.5: Experiment 2, Random effects panel regressions for  $\ln(\text{ResponseTimes})$ .

$\ln(\text{ResponseTimes})$	Model 1	Model 2	Model 3
Conflict	0.1691*** (0.0304)	0.1621*** (0.0251)	0.1627*** (0.0250)
Error $\times$ Conflict	-0.2548*** (0.0296)	-0.2060*** (0.0112)	-0.2060*** (0.0113)
Error $\times$ Aligned	0.2068*** (0.0347)	0.1717*** (0.0313)	0.1718*** (0.0307)
Noise (in Conflict)	0.2407*** (0.0312)	0.1939*** (0.0158)	0.1946*** (0.0155)
Cognitive Load		-0.3452*** (0.0770)	-0.3446*** (0.0772)
Normalized Time		-0.3116*** (0.0308)	-0.3117*** (0.0308)
Part 2		-0.1701*** (0.0249)	-0.1701*** (0.0209)
Part 3		-0.3054*** (0.0328)	-0.3046*** (0.0313)
Payoff Table 2			0.0245 (0.0278)
Payoff Table 3			0.0431** (0.0174)
Collusion		-0.7449*** (0.0253)	-0.7759*** (0.0344)
Constant	2.1329*** (0.0757)	2.6304*** (0.0675)	2.6067*** (0.0693)
R <sup>2</sup>	0.0467	0.1836	0.1840

Notes. Standard errors, clustered by 6 matching blocks, in parentheses. \*  $p < 0.1$ ,

\*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

results remain robust. The collusion dummy is negative and highly significant, but the previous results and conclusions are not affected. A post-hoc linear combination test shows no significant difference between Payoff Table 2 and Payoff Table 3 (Linear Combination test, Payoff Table 3 - Payoff Table 2, coefficient= 0.0186,  $z = 0.732$ ,  $p = 0.4639$ ).<sup>19</sup>

<sup>19</sup>Additional regressions controlling for demographics and psychological characteristics yielded the same qualitative results and conclusions.

In summary, we find strong evidence in our data suggesting multiple behavioral rules codetermine behavior in Cournot oligopolies. In particular, we find evidence for imitation, a more automatic and intuitive behavioral rule, and myopic best reply, a more controlled and rational behavioral rule.

## 4.5 Reinforcement

In this and the following Section we analyze our data with respect to two other very prominent behavioral rules, i.e. reinforcement and inertia.

Reinforcement is especially important for economics since it captures the focus on past performance. Positive reinforcement is the tendency to repeat whatever has given good results in the past, without paying attention to whether the conditions in which that action was successful have changed (e.g., outcome bias, see Dillon and Tinsley, 2008).

Achtziger and Alós-Ferrer (2014) showed that reinforcement corresponds to a highly automatic process which competes with more controlled processes when feedback has a win-loss valence. Alós-Ferrer and Ritschel (2018b) showed that reinforcement can act as a cognitive shortcut instead of more controlled behavioral rules in case both are aligned. In our paradigm, we can identify decisions following positive reinforcement, namely those where the player obtained the highest profits in the previous period and repeats the decision. That is, those are situations where imitation is aligned with inertia (but not necessarily with myopic best reply), and the player imitates him- or herself. This gives us an opportunity to differentiate between reinforcement and imitation of others. Inertia could also be the driver of behavior which we will address in the next section. In this section we focus on positive reinforcement and imitation of others. Evidence from neuroscience has shown that reinforcement learning is associated with extremely fast and unconscious brain responses (Holroyd and Coles, 2002). Hence, we hypothesize that reinforcement is associated to a more automatic behavioral rule than imitating others, which should be reflected in even shorter response times. Note that the comparison is quite different from the one between imitation and myopic best reply, because positive reinforcement is a subcategory of imitation de-

cisions. That is, whenever positive reinforcement is active, by definition it yields the same prescription as imitation. Hence, when we disentangle decisions following positive reinforcement from other imitative decisions we are comparing two disjoint kinds of imitation (imitating yourself and imitating others) which are neither in conflict nor in alignment.




Figure 4.8: Average response times of reinforcement and imitating-others decisions. *Notes.* Left-hand side: Experiment 1. Right-hand side: Experiment 2 separated by the cognitive load treatments. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , WSR test.

Figure 4.8 displays the response times of decisions where participants imitated themselves vs. those of decisions where they imitated others. The figure shows three subsets of our data, i.e. the pooled data for Experiment 1 (left-hand side) and both treatments (NoLoad and Load) for Experiment 2 (right-hand side). We disentangle the comparison according to whether imitation was in conflict or in alignment with myopic best reply in order to test among comparable decisions, but we expect the same relation in both cases (conflict and alignment refer to imitation and myopic best reply, while

we are testing among two different classes of imitative decisions). Within each subset the left-hand side corresponds to the case where imitation was in conflict with myopic best reply, and hence all imitative decisions were “imitation errors.”

In Experiment 1, reinforcement decisions in this case took an average of 10.28 s and were faster than imitating-others decisions, which averaged 11.91 s. The difference was highly significant (WRS,  $N = 108$ ,  $z = -3.807$ ,  $p = 0.0001$ ). The right-hand side of each subset of data in Figure 4.8 corresponds to the case where myopic best reply and imitation were aligned, hence all imitative decisions were “correct.” Also in this case reinforcement decisions (average 9.86 s) were significantly faster than imitating-others decisions (average 11.38 s; WSR,  $N = 18$ ,  $z = -1.720$ ,  $p = 0.0854$ ).<sup>20</sup>

The analysis for Experiment 2 obtains the same conclusions. In the NoLoad treatment, reinforcement decisions in conflict situations took an average of 10.70 s and were hence faster than imitating-others decisions, which averaged 12.30 s. The difference was highly significant (WRS,  $N = 65$ ,  $z = -3.075$ ,  $p = 0.0021$ ). In case of alignment in the NoLoad treatment, reinforcement decisions (average 8.65 s) were also significantly faster than imitating-others decisions (average 13.55 s;  $N = 8$ ,  $z = -2.521$ ,  $p = 0.0117$ ). In the case of conflict, while under cognitive load, reinforcement decisions took an average of 7.95 s and were hence faster than imitating-others decisions, which averaged 9.54 s. Again, this difference was strongly significant (WRS,  $N = 65$ ,  $z = -2.400$ ,  $p = 0.0164$ ). In the case of alignment, the reinforcement decisions (average 7.58 s) were not significantly different to the imitating-others decisions (average 7.39 s;  $N = 9$ ,  $z = -0.178$ ,  $p = 0.8590$ ).<sup>21</sup>

In summary, we generally confirm that decisions where a participant imitated him- or herself were significantly faster than comparable decisions

---

<sup>20</sup>Notice that requiring at least 2 observations per individual and splitting the imitation decisions into positive reinforcement and imitating others greatly reduces the number of observations in alignment situations. As before we reported a two-sided test, however, note that a one-sided test shows significance at a level of  $p = 0.0427$ .

<sup>21</sup>The number of observations is greatly reduced. Note that a test yields significant faster reinforcement decisions (average 7.21 s) than imitating-others decisions (average 9.88 s) when not requiring at least two observations per individual (WSR,  $N = 29$ ,  $z = -2.562$ ,  $p = 0.0104$ ).



where the participant imitated another participant. This is compatible with the view from psychology and neuroscience that reinforcement processes are extremely fast and more automatic than other processes.<sup>22</sup>

## 4.6 Inertia

In the previous subsection we concluded that imitating oneself is faster than imitating others. Decisions where a player imitates him- or herself are in line with both imitation and inertia, because the previous action is repeated. It is hence legitimate to ask whether the results regarding positive reinforcement are just due to the involvement of a further behavioral rule, namely inertia. Alós-Ferrer et al. (2016b) recently found that decision inertia causes error-rate asymmetries in the belief-updating task of Achtziger and Alós-Ferrer (2014), but the process is considerably weaker than reinforcement and is washed away by the latter. Hence, we will exploratively test for the effects of inertia in the cases where it is *not* aligned with imitation, because the alignment of imitation and inertia corresponds to positive reinforcement. That is, we restrict the tests to decisions not aligned with imitation. In case of conflict between imitation and myopic best reply, this means that we test within correct decisions, i.e. decisions following myopic best replies. In case of alignment between imitation and myopic best reply, it means we test within errors.

Figure 4.9 shows the response times of decisions that are in line with inertia (“stay” decisions) and those that are not (“shift” decisions). The figure displays those response times for three subsets of our data, i.e. the pooled data for Experiment 1 (left-hand side) and both treatments (NoLoad and Load) for Experiment 2 (right-hand side). Within each subset the left-hand side corresponds to response times of correct decisions which are differentiated according to whether the decision followed inertia, resulting in a “stay” best reply, or not, resulting in a “shift” best reply.

---

<sup>22</sup>While it is often useful to think of decision processes in dichotomous terms as “automatic” or “controlled,” this is just a simplification. The automaticity dimension is best conceived of as a continuum, and all that can be concluded when comparing different processes is whether one is *more* automatic than another one or not.




Figure 4.9: Average response times of stay and shift correct decisions and errors. *Notes.* Left-hand side: Experiment 1. Right-hand side: Experiment 2 separated by the cognitive load treatment. Each subset shows stay and shift correct (myopic best reply) decisions in case myopic best reply and imitation were in conflict and stay and shift errors in case myopic best reply and imitation were aligned. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , WSR test.

In Experiment 1, stay best replies averaged 12.10 s compared to 12.24 s for shift best replies, and the difference is not significant (WSR,  $N = 79$ ,  $z = -0.117$ ,  $p = 0.9066$ ). Hence, there is no evidence for an involvement of pure inertia in the correct decisions in Experiment 1. Now consider the case of errors, for which the right-hand side within each subset shows the response times of errors, i.e. non-best replies, in the case where myopic best reply and imitation were aligned. In Experiment 1, stay (inertia) errors were significantly faster (mean 12.41 s) than shift errors (mean 14.62 s; WSR,  $N = 50$ ,  $z = -2.833$ ,  $p = 0.0046$ ). This difference is interesting. In case of alignment between imitation and myopic best reply, errors do not follow from imitation, and shift errors do not follow from inertia either. Although this is

a post hoc interpretation, this result might point out that the shift errors in this case possibly include some decisions following higher-order reasoning, as in the case of the noise errors discussed in the regression analysis of Section 4.3. This suggests even more behavioral rules are at work than we focused on in this work.

The results of the NoLoad treatment in Experiment 2 obtains the same conclusions as in Experiment 1, as expected since it is a replication. In the NoLoad treatment, stay best replies averaged 14.31 s compared to 14.58 s for shift best replies, and the difference is not significant (WSR,  $N = 43$ ,  $z = -0.169$ ,  $p = 0.8658$ ). In the case of alignment, stay (inertia) errors were significantly faster (mean 13.36 s) than shift errors (mean 16.08 s; WSR,  $N = 21$ ,  $z = -1.964$ ,  $p = 0.0496$ ). Hence, there is no strong evidence for an involvement of a pure inertia process for the correct decisions while shift errors might follow higher-order reasoning in the NoLoad treatment.

In the cognitive Load treatment we find the opposite to be true. In the case of conflict, the response times of stay best replies (mean 9.81 s) were significantly faster than shift best replies (mean 11.43 s), in contrast to the findings in Experiment 1 and the NoLoad treatment (WSR,  $N = 40$ ,  $z = -2.554$ ,  $p = 0.0107$ ). Hence, there is evidence for an involvement of pure inertia in the correct decisions. We also find a different result for stay and shift errors in case myopic best reply and imitation were aligned. Stay (inertia) errors were not significantly faster (mean 11.17 s) than shift errors (mean 12.59 s), contrary to the findings in Experiment 1 and the NoLoad treatment (WSR,  $N = 17$ ,  $z = -1.112$ ,  $p = 0.2659$ ).<sup>23</sup>

The influence of inertia on decisions is affected by the cognitive load manipulation. In Experiment 1 and the NoLoad treatment there is no significant difference for stay and shift best replies but there is one in the Load treat-

---

<sup>23</sup>Conservative tests at the block level generally obtain the same results and conclusions. Only the comparison of stay and shift best replies in conflict situations in the Load treatment yields a different conclusion on the block level and is non-significant (stay best replies, mean 9.83 s; shift best replies, mean 11.14 s; WSR,  $N = 6$ ,  $z = -0.943$ ,  $p = 0.3454$ ). Note that the non-significant result for stay and shift errors in the Load treatment are not due to the low number of observations. We also find a non-significant result when not requiring at least two observations per individual (mean stay errors 11.43 s; mean shift errors 11.92 s; WSR,  $N = 50$ ,  $z = -0.729$ ,  $p = 0.4661$ ).

ment. This suggests that under cognitive load the alignment with inertia promoted the use of the shortcut of the very fast heuristic and made decisions significantly faster. For stay and shift errors it is the opposite case, i.e. there is a significant difference for stay and shift errors in Experiment 1 and the NoLoad treatment but there is none in the Load treatment suggesting that cognitive load reduced the complexity of some behavior leading to shift errors. Hence, we find another difference between the behavioral rules of imitation and myopic best reply. The alignment with inertia (whether it is a “stay” or a “shift” decision) affects imitative decisions in a different way than myopic best reply decisions. With our cognitive load manipulation, we do find a difference between the more automatic and intuitive behavioral rule of imitation and the more controlled one of myopic best reply. This finding provides further evidence that these processes are different and codetermine behavior in Cournot oligopolies.

In any case, this is evidence that the interaction with inertia (whether a decision is a “stay” or a “shift”) affects imitative decisions in a different way than myopic best replies. Hence, we aim to conduct a robustness test of our results regarding Hypotheses H1 and H2 controlling for the effects of inertia. Figure 4.10 shows the average response times for Experiment 1 and both treatments of Experiment 2 while excluding all observations which were additionally aligned with inertia. Let us start with H1, which concerns the case of conflict between imitation and myopic best reply. We compare shift best replies to imitating-others errors. The comparison is shown for Experiment 1 in the left-hand side of Figure 4.10. Again, errors (mean 12.04 s) are significantly faster than correct responses (mean 12.61 s; WSR,  $N = 118$ ,  $z = -2.029$ ,  $p = 0.0425$ ), confirming H1 even when all inertia decisions are excluded from the analysis. Hypothesis H2 concerns the case of alignment between imitation and myopic best reply. We are hence comparing shift best replies (which were also imitating-others decisions) to shift errors. The comparison is shown in the right-hand side of the subset for Experiment 1 in Figure 4.10. This time, errors (mean 15.20 s) are not significantly different from correct responses (mean 13.82 s; WSR,  $N = 25$ ,  $z = 1.063$ ,  $p = 0.2879$ ), not confirming H2 when all inertia decisions are excluded from the analy-




Figure 4.10: Average response times of correct decisions and errors excluding inertia decisions.

*Notes.* Left-hand side: Experiment 1. Right-hand side: Experiment 2 separated by the cognitive load treatment. Each subset shows correct decisions and errors in case myopic best reply and imitation were in conflict and in alignment. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , WSR test.

sis.<sup>24</sup> Hence, the predictions of the DPDM (Achtziger and Alós-Ferrer, 2014; Alós-Ferrer, 2018) are confirmed.

Also in the NoLoad treatment of Experiment 2 we find that in conflict situations, errors (mean 12.41 s) are significantly faster than correct responses (mean 14.84 s; WSR,  $N = 68$ ,  $z = -3.605$ ,  $p = 0.0003$ ), confirming H1 even when all inertia decisions are excluded from the analysis and replicating the results of Experiment 1 with fewer observations. Regarding H2, we do

<sup>24</sup>Testing at the block level shows significantly faster errors than correct response in conflict situations and significantly slower errors than correct responses in alignment situations. Also, the test on the individual level not requiring at least two observations per subject yields significantly slower errors (mean 14.51 s) than correct responses (mean 12.88 s; WSR,  $N = 71$ ,  $z = 2.131$ ,  $p = 0.0330$ ).

not obtain significant results and cannot confirm the prediction. Shift errors (mean 14.76 s) are not significantly slower than shift best-replies in alignment (mean 16.24 s; WSR,  $N = 13$ ,  $z = -0.943$ ,  $p = 0.3454$ ).<sup>25</sup>

For cases of conflict in the cognitive load treatment, we find that errors (mean 9.35 s) are significantly faster than correct responses (mean 10.70 s; WSR,  $N = 62$ ,  $z = -3.130$ ,  $p = 0.0017$ ). As in all the other conflict situations we confirm prediction H1. However, as before we do not obtain significant results for H2 in the Load treatment for alignment situations. Errors (mean 9.97 s) are not significantly slower than correct responses (mean 8.51 s; WSR,  $N = 10$ ,  $z = 1.376$ ,  $p = 0.1688$ ), not confirming H2 when all inertia decisions are excluded from the analysis.<sup>26</sup>

One obvious explanation for not obtaining significant results for Hypothesis 2 is the strongly reduced amount of observations after excluding all stay decisions. In case of alignment the number of overall decisions is already relatively low compared to conflict situations and is further split into errors and correct response. When excluding all decisions which are additionally aligned with inertia the overall number of errors and correct response is very low (Experiment 1: 491 errors and 121 correct responses out of 6,144 observations; Experiment 2, NoLoad treatment: 227 errors and 70 correct responses out of 3,456 observations; Experiment 2, Load treatment: 185 errors and 74 correct responses out of 3,456 observations). Not only are there few observations left, there are also only few individuals who have more than one observation for creating the average response times. We cannot confirm the predictions of H2, however, we confirmed the predictions of H1 even when excluding inertia decisions from the analysis in Experiment 1 and both treatments in Experiment 2.

---

<sup>25</sup>The same conclusion obtains testing the block level.

<sup>26</sup>Tests at the block level yield the same conclusion for conflict situations and obtain significant slower errors (mean 11.81 s) than correct responses (mean 9.87 s; WSR,  $N = 6$ ,  $z = 1.782$ ,  $p = 0.0747$ ) in contrast to the results on the individual level.

## 4.7 Conclusion

In two experiments we found evidence suggesting that multiple behavioral rules codetermine behavior. We apply the Dual-Process Diffusion-Model of Alós-Ferrer (2018) which predicts an asymmetry in the response times when different postulated behavioral rules are either in conflict or in alignment. This regularity has been previously found in an abstract belief-updating task involving strictly individual decisions (Achtziger and Alós-Ferrer, 2014). The present work undertakes the analysis for a more complex, interpersonal decision-making task, using Cournot oligopolies as the workhorse. In this economically relevant setting, previous theoretical and experimental work have shown convergence towards the outcomes associated with imitation, even though the game-theoretic setting clearly makes best reply considerations prominent. Hence, we postulated that imitation (being a more intuitive or automatic) and myopic best reply (being a proxy for more rational, reflexive, or controlled behavior) would be the relevant behavioral rules and examined response times as indicators of behavioral involvement.

We found strong evidence for the effects predicted by the dual-process approach. This result contributes to the literature in two directions. First, previous experimental literature on Cournot oligopolies (Huck et al., 1999, 2004; Offerman et al., 2002; Apesteguía et al., 2007, 2010) has concentrated on convergence towards Walrasian outcomes, predicted by theoretical models of imitative behavior (Vega-Redondo, 1997; Alós-Ferrer and Ania, 2005). This is a relevant and legitimate way to test for the actual relevance of the implications of imitative behavior, but it tests on predicted consequences and not on actual behavior. In other words, the actual presence of imitative behavior is indirectly deduced from the fact that one observes empirical convergence to the outcomes which would be predicted if behavior followed (noisy) imitation behavioral rules. In contrast, our analysis targets direct correlates of individual decisions, without relying on (and actually actively precluding) any convergence. In this sense, our response-times data delivers novel evidence which is compatible with the involvement of imitative behavior in economic situations as Cournot oligopolies.

Second, our results find highly significant asymmetric effects in response times as predicted by Alós-Ferrer (2018). This is a subtler and more elaborate prediction than the simple expectation that intuitive decisions should be faster. The asymmetry hinges crucially on the existence of two interacting decision rules with different properties, one being more automatic/intuitive and the other more controlled/reflexive. Hence, our results provide evidence showing that multiple behavioral rules influence and codetermine behavior, and that the dual-process approach arising in psychology delivers interesting insights for economic decision making. Multiple behavioral rules are more than a convenient metaphor, and the thinking and decision making should be considered less as an instantaneous, punctual phenomenon and more as the result of the interaction of different behavioral rules and decision processes in the human brain.

This view, however, does not reduce to a literal import of psychological models into economics. In order to analyze economic decisions, dual-process ideas need to be adapted and refined. Where psychologists speak of low-level processes, we are interested in high-level behavioral rules. While psychologists often concentrate on very short response times, economic decisions are usually slower, and in complex paradigms as the one considered here they are also slower than those typical in the literature on judgment and decision making (in our case, decisions are mostly in the 10-15 seconds window, which is considerably longer than decisions in, e.g., Achtziger and Alós-Ferrer, 2014). Whether dual-process insights remain meaningful for these longer response times is an empirical question, and the answer we have obtained here is clearly positive.

In Experiment 2 we manipulated cognitive load to further investigate differences between imitation and myopic best reply along another dimension. We predicted that under cognitive load a direct increase in more automatic behavior, i.e. more imitative decisions, will be observed. However, we did not find such a direct relation and since there was no change in behavior Experiment 2 includes a robustness check. We are convinced that we had a strong cognitive load manipulation as seen due to the effect on response times and which has successfully induced cognitive load effects in other studies



(Carpenter et al., 2013). The hypothesis H3 implicitly assumes that imitation is the most automatic behavior and all other behavioral rules require much more cognitive resources. However, behavior could shift from more controlled behavior to more automatic behavior, as predicted, but cannot be identified because a shift from myopic best reply or other higher complex reasoning to imitation can be negated by a shift from imitation to even more automatic, intuitive behavior, e.g. random behavior. We did not find such a direct effect of cognitive load, however, we find that imitation and myopic best reply react differently to lack of available cognitive resources. Another, more positive interpretation of our results is that cognitive load makes decisions faster while simultaneously it does not change the observed behavior in terms of decision frequencies.

At a general level, our work adds to the growing body of evidence showing that the dual view of behavior might be better suited to explain economic decision making than either the assumption of fully-rational, optimizing agents or the evolutionary approach replacing agents by boundedly rational behavioral rules. Our results speak in favor of multiple behavioral rules, one of them reflecting evolutionary ideas (in this case, imitative behavioral rules) and the other being closer to rational optimization. The bottom line of our research is that a dual-process model architecture can provide a useful synthesis of previous approaches and help significantly improve our understanding of how economic decisions are made while showing that multiple behavioral rules codetermine behavior in Cournot oligopolies.



---

## CHAPTER 5

### Cognitive Load in Economic Decisions: How To Tell Whether It Works

---

#### 5.1 Introduction

Every non-trivial task, and almost every consequential decision, requires cognitive resources. For easy tasks and decisions, as when one decides on an ice-cream flavor at the parlor, the cognitive demands are low. In other cases, as when one struggles to design an optimal investment plan to finance retirement, cognitive demands are high, and cognitive resources might become a binding constraint.

Economically relevant tasks and decisions might often belong to the second category. In such cases, a decrease in performance is to be expected if cognitive resources become more scarce. Research in psychology has developed a plethora of *cognitive load* manipulations to causally reduce the amount of cognitive resources available for a task, hence illuminating the allocation of resources within the human brain. The most common forms of cognitive load involve taxing cognitive resources through an additional task, as e.g. keeping a number or a graphical configuration in memory or generating random numbers out loud. It has been extensively shown in psychology that such manipulations do impair performance in simple cognitive tasks (Baddeley and Hitch, 1974; Baddeley et al., 1984; Shiv and Fedorikhin, 1999; Hinson et al., 2002; Lavie and de Fockert, 2005; Barrouillet et al., 2007). These manipulations are best understood in the framework of the *working memory* model of Baddeley and Hitch (1974). Working memory is conceived of as a (functional) brain system which processes and stores short-term information, and is assumed to have a limited capacity. Baddeley (1992) divides working memory into multiple subcomponents, which we will briefly describe below,

and cognitive load manipulations differ according to which subcomponent is targeted.

The fact that cognitive load decreases performance is, in itself, not extraordinarily relevant for economics. However, it turns out that cognitive load has the potential to substantially and systematically alter behavior in ways which are of interest for economics. The reason can be explained in a straightforward way following *dual-process theories* (see Kahneman, 2003; Strack and Deutsch, 2004; Alós-Ferrer and Strack, 2014; see also Evans, 2008; Weber and Johnson, 2009 for detailed reviews). These theories, in a simplified version, postulate that the human mind is mainly influenced by two kinds of processes, called *automatic* and *controlled*, with the former capturing intuitive, impulsive, or heavily trained processes of an associative type (e.g., following stimulus-response patterns), and the latter corresponding to more deliberative, reflective processes. In a less simplistic way, the automaticity dimension is often viewed as a continuum (Bargh, 1989), and the actual postulate is that decision processes in the human mind differ in their degree of automaticity, and that often several competing processes influence decisions and performance. A key dimension associated to automaticity is the degree to which processes require the use of cognitive resources as working memory. More automatic processes place little or no requirement on cognitive resources, and are effectively always available. In contrast, more deliberative ones make heavy use of cognitive resources as working memory, and hence become unavailable if the latter are taxed away. Hence, the dual-process view suggests that taxing working memory through cognitive load impairs controlled processes, resulting in a shift to more automatic processes.

In many situations of interest in economics, such a process shift is consequential. In terms of decisions and performance, automatic processes often correspond to cognitive shortcuts or heuristics which might be aligned with deliberation in some or many situations, but might conflict with it, leading to biases, in economically relevant domains as decision making under risk or uncertainty. Thus, tilting the balance toward automatic processes allows to better understand such biases. In terms of preferences and motives, automatic processes reveal intrinsic tendencies (sometimes informally referred to

as a “default mode of behavior”), and hence a shift toward automatic processes might help uncover the roots of many economically relevant human tendencies as altruism or cooperation, to quote just two examples.

Studies relying on cognitive load rely on a triple hypothesis. First, the key procedural assumption is that the particular cognitive load manipulation implemented has taxed cognitive resources to a sufficient extent to induce a shift in processes. While this is often clear in studies in cognitive psychology, due to the sheer number of available studies one can compare to, it is usually far less clear in economic settings, where the number of studies is limited and tasks are far less standardized than in cognitive psychology. Second, the assumption is that the shift to automatic processes will result in an observable change in behavior. This might not always be the case, however. Automatic processes are in themselves not flawed: rather, they have evolved because they economize cognitive resources while delivering a good response in evolutionarily typical situations. Hence, in many situations they will actually prescribe the same response as more deliberative processes. It is only when they are used in an evolutionarily new situation that they will conflict with the latter and prescribe erroneous or suboptimal responses. This gives rise to the key distinction between *alignment* and *conflict* (see Achtziger and Alós-Ferrer, 2014; Alós-Ferrer, 2018): a decrease in performance is not to be expected in a situation of alignment. The third assumption is that the researcher has correctly identified the automatic and deliberative processes playing a role in the specific decision situation, and their respective prescriptions.

The two latter hypotheses are a matter of experimental design, and whether they should be accepted or rejected should be determined on the basis of the collected data within each individual study. Alas, this task becomes impossible in the absence of a clear criterion to decide whether the first hypothesis is correct, that is, whether the particular cognitive load manipulation that the researcher has used has actually successfully induced a process shift. To date, for the class of tasks used in economic experiments, there is no satisfactory test allowing such a conclusion. Actually, there is anecdotal evidence that this problem might currently be a hidden bottleneck

impairing research in experimental economics. A modest number of studies have successfully used cognitive load in explicitly economic tasks (Milinski and Wedekind, 1998; Carpenter et al., 2013; Duffy and Smith, 2014; Schulz et al., 2014; Samson and Kostyszyn, 2015; Buckert et al., 2017; Døssing et al., 2017), but others have found mixed results (Deck and Jahedi, 2015; Allred et al., 2016), and a fair number of studies have reported no effects (Greene et al., 2008; Cappelletti et al., 2011; Cornelissen et al., 2011; Benjamin et al., 2013; Glaser and Walther, 2014; Duffy et al., 2016; Hauge et al., 2016; Dri-choutis and Nayga, 2017). It is reasonable to assume that publication bias might have resulted in an additional number of unsuccessful studies not being circulated.

In this work, we offer a test of the single hypothesis that cognitive load has been successful in inducing a shift toward a higher reliance in automatic processes in the kind of complex tasks typically relevant for economic research. At first glance, this might appear to be an impossible task, because it is fundamental that this test works independently of whether the shift in processes results in decreased performance or an actual shift in decisions (behavior) or not. However, the test is possible if one relies on a different kind of data: response times. The intuition is simple. Independently of whether more automatic processes favor different responses as more deliberative processes or just the same ones, one of the major consequences of automaticity (and, one could easily argue, the reason it provided an evolutionary advantage) is that more automatic processes are typically faster. Hence, elementary dual-process logic dictates that, if a successful shift to more automatic processes is induced, decisions must, on average, become faster. Hence, one obtains the apparently paradoxical conclusion that, if cognitive load is successfully induced, response times for the task of economic interest must become shorter. This simple test of response times becomes the smoking gun showing that cognitive load was successfully induced, which then allows to interpret the behavioral (choice) data. In other words, if response times show that cognitive load did induce a process shift but actual behavior does not change significantly, the researcher is still justified in using the data to draw conclusions about the default mode of behavior, the characteristics of involved

decision processes, or whether there is a process conflict at all. In the absence of the response-time test, such conclusions would be unwarranted.

The test we propose has two clear boundary conditions. First, it will typically only apply to relatively complex tasks where decisions are relatively long. This is because there needs to be a clear difference in response times between the relevant automatic and controlled processes. It is typically assumed that cognitive load causes an increase in response times in parts of the decision process as process selection, task processing, or conflict detection and resolution. Those effects in response times will typically be small (for example, electrophysiological correlates of conflict detection are measured within 200-300 ms of stimuli presentation; Achziger et al., 2014). If the main task is a fast one, however, this effect might dominate, explaining why cognitive load often results in longer response times in psychological paradigms (Gevins et al., 1998; Baddeley et al., 2001; de Fockert et al., 2001) (and, why psychologists have never developed the test we propose). The latter paradigms are set in rather simple settings compared to most economic decisions, and are associated with extremely short response times (often below one second<sup>1</sup>) compared to response times in economic experiments (which, in the cases illustrated in this work, go up to over 20 s).

The second boundary condition concerns the type of manipulation. The kind of cognitive load tasks we consider, which are the most widespread, crucially require no concurrent motor activity or added time due to added demands on attention while the main task is being carried out. The most clear example is keeping a number or a graphical dot configuration in memory, but also internally generating random numbers while carrying out a task fits the bill. Potentially, one could conceive of manipulations where a concurrent task is used, forcing the participant to interrupt the main task to perform the competing one, e.g. answering some unrelated questions as the main

---

<sup>1</sup>Gevins et al. (1998) report response times in an EEG stimulus matching task between 441 and 933 ms. Baddeley et al. (2001) present response times between 35 and 175 s for a block of 40 decisions of adding or subtracting 1 to a single-digit number. de Fockert et al. (2001) report on an fMRI study with a visual selective attention task classifying names with pop stars or politicians for which response times increased from 953 to 1,394 ms under cognitive load.

task develops. This is done in so-called *task-switching paradigms* (Rogers and Monsell, 1995; Monsell, 2003), which however are not usually considered to be cognitive load manipulations. Obviously, such tasks place additional requirements on e.g. attention and motor reactions or cause the main task to be put on hold for a while, which will be reflected in increased response times. Hence, task-switching paradigms of this type are not well-suited for our test.

In this work, we report on four separate experiments in which we employed different types of cognitive load manipulations in different economic experiments: belief-updating tasks, voting, and competitive games (Cournot oligopolies). In all cases, the cognitive load manipulations resulted in shorter response times, even though effects on behavior and performance varied. We also observed that more complex cognitive load tasks led to higher reductions in response times.

The remainder of the Chapter is structured as follows. As a reference for later sections, and for the reader's convenience, Section 5.2 briefly reviews the basics and structure of working memory. This is important because different cognitive load manipulations target different working-memory subcomponents. Sections 5.3 and 5.4 present two separate experiments on Cournot oligopolies where we used two different cognitive load tasks, a light load manipulation and a high load one. Section 5.5 presents a voting experiment with the same cognitive load task as in that of Section 5.4. Section 5.6 presents a belief-updating experiment with two different cognitive load manipulations. Section 5.7 concludes.

## 5.2 Working Memory and Cognitive Load

Understanding cognitive load manipulations requires a discussion of *working memory*, which can be described as the set of functions and resources governing the selection and execution of decision processes. In order to provide the theoretical background for the present studies, we briefly introduce the working memory model of Baddeley (1986, 1996), which is a standard reference in cognitive psychology. This model describes how different working memory



components might be responsible for automatic and controlled processes and their selection. It suggests a supervisory system that controls the switch between automatic and controlled processes. The model distinguishes a *central executive system* from two subordinate (slave) memory systems (components) that are modality-specific. These two components are the *phonological loop* and the *visuospatial sketch pad*. Each of the working memory components has only limited (cognitive) capacity. Accordingly, an overload of these components' resources, for instance through cognitive load manipulations, results in impairments in task performance.

The *phonological loop*, also known as *verbal working memory*, is responsible for the retention of verbally coded material, independently of whether it is presented in written or auditory form. It refreshes stored information through e.g. inner-voice repetition (see e.g. Gathercole and Baddeley (1993)). The storage and maintenance of information carried out by the phonological loop corresponds to cognitive resources which are used more intensely by more controlled processes. Most of the cognitive load manipulations employed in previous economic research target the phonological loop (typically, keeping certain numbers in memory), and accordingly so did several of our manipulations.

The *visuospatial sketch pad* is responsible for the retention of graphically coded material, as e.g. images. Some cognitive load manipulations in psychology avoid the phonological loop and target this subsystem instead by, e.g., asking participants to keep a configuration of dots in memory.

Last, the *central executive* integrates information from various sources and is also seen as the supervisor or controller of the other two working memory components. The functions of the central executive consume much of the (restricted) cognitive resources (Norman and Shallice, 1986). It plays the role of a supervisory system switching between controlled and automatic processes. More generally, it governs the controlled selection or development of strategies in situations which are new in the sense that no specific rules have yet be learned, i.e. when automatic processes are not available. It is also responsible for allocating attention to complex controlled processes and implementing them. Hence, successfully performing complex completing

cognitive tasks (e.g. by inhibiting automatic processes) can be assumed to rely on functions of the central executive. Cognitive load manipulations targeting the central executive are seen as particularly demanding. One of our experiments included a manipulation of this type.

### 5.3 Experiment 1: Cournot Oligopoly Under Light Cognitive Load

#### 5.3.1 Experimental Design and Procedures

For the main decision task, participants interacted each round in 4-player Cournot oligopolies (tetrapolies). Subjects participated in three different oligopolies (parts), with 17 rounds each (total of 51 rounds). For each part, we computed a payoff table using a Cournot oligopoly with zero costs and a linear inverse demand function of the form  $P(Q) = a - Q$ , where  $P$  is the price,  $a$  the saturated demand, and  $Q$  the total quantity in the market. However, during the experiment a neutral framing was used and neither firm nor quantities were mentioned. We reduced the action space to four possible actions, i.e.  $A$ ,  $B$ ,  $C$ , and  $D$  with either increasing or decreasing quantities from  $A$  to  $D$ .<sup>2</sup> Hence, the whole payoff table had dimensions  $4 \times 20$ , with four rows representing the possible actions and 20 columns labeled  $AAA$  to  $DDD$  representing the possible actions of the opponents. Payoffs were expressed in points, with an exchange rate of 20 Eurocents per 1000 points. The points achieved in all 51 rounds were accumulated and paid at the end of the experiment. After the first round the participants were informed about the outcome of the previous round. Participants saw the full payoff table, their own choice and earnings from the previous round, and the previous choice and earnings from the other three group members. While seeing all this information they chose again  $A, B, C$ , or  $D$ . The first round in each part did not provide any information on the previous round and was therefore

---

<sup>2</sup>Payoff table 1:  $P(Q) = 150 - Q$ ,  $A = 37.5$ ,  $B = 33.25$ ,  $C = 30$ ,  $D = 18.75$  (or reversed); Payoff table 2:  $P(Q) = 175 - Q$ ,  $A = 43.75$ ,  $B = 38.875$ ,  $C = 35$ ,  $D = 21.875$  (or reversed); Payoff table 3:  $P(Q) = 200 - Q$ ,  $A = 50$ ,  $B = 44.5$ ,  $C = 40$ ,  $D = 25$  (or reversed).

dropped for the analysis yielding 16 rounds in each part (total of 48 rounds) relevant for the analysis.

In addition to the main decision task we implemented two treatments, i.e. *No Load* and *Load* treatments, within subjects. Each part had 8 rounds of No Load and 8 rounds of Load (excluding the very first round of each part which also did not have any cognitive load). During rounds with cognitive load an additional task was implemented while no additional task was present during no load rounds. The task consisted of two parts and enclosed the main decision task in each round. The first part required the participants to memorize a single-digit number which was displayed for 5 seconds before the Cournot oligopoly screen appeared. During the Cournot oligopoly decision task the participants heard another single-digit number via headphones which was played randomly between 1 and 10 seconds. The participants had to add up the two numbers and enter the sum in another screen after the decision task was done. The cognitive load task was incentivized and correct additions earned additional points.<sup>3</sup>

After the experiment the participants filled out a questionnaire which included questions concerning the importance and difficulty of the cognitive load task, the Faith in Intuition scale (Alós-Ferrer and Hügelschäfer, 2012, 2016) and the Big Five questionnaire.<sup>4</sup>

Alós-Ferrer and Ritschel (2018a) use a similar set-up but focus the analysis on two specific behavioral strategies, i.e. imitation and myopic best reply. They use response times and the theoretical predictions of the Dual-Process Diffusion Model (Alós-Ferrer, 2018) as evidence to show that imitation is more automatic and myopic best reply is more controlled. In this work, we focus on the effects of cognitive load on response times during the Cournot oligopoly. The within-subject design allows us to look at individual differences of response times between the treatments and accounts for heterogene-

---

<sup>3</sup>Each correct addition earned additional 750 points. Participants earned on average 1200 points per round from the Cournot oligopoly decision.

<sup>4</sup>The German version of the Faith in Intuition questionnaire was taken from Keller et al. (2000), who adapted it from Epstein et al. (1996). The German version of the 15-item Big-Five scale was taken from Gerlitz and Schupp (2005) (see Lang et al., 2011, p. 560), who shortened the 44-item version of the Big Five Inventory by John et al. (1991).

ity in response times across subjects. The appropriate test is the Wilcoxon Signed-Rank (WSR) test on the individual level. Since we predict shorter response times in the cognitive load treatment, we will report one-sided  $p$ -values. In the sequel, and as a robustness check, we also report the tests at the group level since those observations are completely independent because the participants are separated into different groups of 4 participants. For reference, we also provide other common measures such as error rates and questionnaire data regarding the cognitive load task.

Our main focus lies on response times. Nevertheless, we will also briefly report the effect of cognitive load on the performance in the decision task. In the Cournot oligopoly task we will use the payoff of the participant as performance measure, i.e. we will compare the payoff for rounds with cognitive load to rounds without cognitive load (excluding the additional payoff from the cognitive load task).

The experiment was conducted at the Cologne Laboratory for Economic Research (CLER), University of Cologne, and programmed in z-Tree (Fischbacher, 2007). Participants were recruited using ORSEE (Greiner, 2015), and were students from the University of Cologne excluding students with a major in Psychology, Economics, and Advanced Business Administration. 64 participants (28 female; age range 18–33 years, mean 24.6 years), participated in exchange for performance-based payment plus a show-up fee. On average the participants earned 18.45 Euro (ranging from 15.00 to 26.50 Euro including the show-up fee of 2.50 Euro) and a session lasted about 1 hour and 45 minutes. The data was collected in two sessions of 32 participants each. Participants played in 16 different and independent groups.

### 5.3.2 Results

The self-reported questionnaire reveals that the cognitive load task was considered to be very important (mean 9.03,  $SD= 1.93$  on a scale from 0 – not at all important to 10 – very important) and that the task was very easy (mean 0.84,  $SD= 1.47$  on a scale from 0 – not at all difficult to 10 – very difficult). This is also reflected in the low error rates of the cognitive load task (mean

error rate 2.41%,  $SD= 3.70$ , ranging from 0% to 20.83%). The questionnaire and error rates show that the subjects took the cognitive load task seriously but also had no difficulties with the task.

We now turn to the response times of the main decision task. The left-hand side of Figure 5.1 depicts the average response time of decisions taken in the No Load and Load treatments. The average response time for decisions in the No Load treatment is 15.22 s while the average response time for decisions in the Load treatment is 14.34 s. As expected, a WSR test confirms our prediction and shows that decisions are significantly faster in the cognitive load treatment ( $N = 64$ ,  $z = 3.110$ ,  $p = 0.0009$ ).<sup>5</sup> The test suggests that the cognitive load task was successfully implemented in the sense that a measurable shift in involved processes occurred. However, the difference in response time, although highly significant, is not very large in value, and hence one can ask whether the process shift was sufficient to lead to a measurable shift in behavior.

We now briefly report on performance using payoff in the Cournot oligopoly rounds as the relevant measure (excluding the show-up fee, the additional payoff from the cognitive load task, and the payoff from first round of each part). The participants earned on average 5.81 Euro during all rounds in the No Load treatment and 5.73 Euro during all rounds in the cognitive Load treatment. This difference is not significant according to a WSR test ( $N = 64$ ,  $z = 0.595$ ,  $p = 0.5517$ ).<sup>6</sup>

---

<sup>5</sup>The conservative test at the group level also reveals significant shorter response times in the Load treatment (WSR test,  $N = 16$ ,  $z = 2.223$ ,  $p = 0.0131$ ).

<sup>6</sup>A conservative test at the group level shows a trend. The group earnings during No Load was 23.23 Euro and during Load 22.92 Euro (WSR,  $N = 16$ ,  $z = 1.500$ ,  $p = 0.1337$ ). Note that a one-sided test with the hypothesis that cognitive load decreases performances yields  $p = 0.0668$ .




Figure 5.1: Experiment 1, Average response times and payoff.

*Notes.* The bars represent one standard error of the mean. Stars indicate a Wilcoxon Signed-Rank test. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## 5.4 Experiment 2: Cournot Oligopoly Under Heavy Cognitive Load

### 5.4.1 Experimental Design and Procedures

In this experiment we used again the Cournot oligopoly setting but employed a different, more demanding cognitive load task. The design of the Cournot oligopoly game was exactly as in Experiment 1. The more demanding cognitive load task consisted of memorizing a seven-digit number, maintaining the number in memory while making the decision in the Cournot oligopoly game, and recalling the number after the decision. There was no auditory load nor did the participants have to add up numbers. Memorizing a number is a common cognitive load task and has been implemented in a variety of experiments (Roch et al., 2000; Hinson et al., 2002; Lavie and de Fockert, 2005; Carpenter et al., 2013; Alós-Ferrer and Ritschel, 2018a). Carpenter et al. (2013) found impaired behavior in a  $p$ -beauty contest while subjects

memorized a seven-digit number. We therefore also required participants to memorize a seven-digit number. The number was displayed for 10 seconds before participants automatically entered the decision phase. After the Cournot decision, the participants had to recall the exact number within 10 seconds. If they succeeded they earned additional 750 points. Rounds with No Load were exactly as in Experiment 1 and consisted only of the Cournot oligopoly decision.

60 Participants (36 female; age range 18–70 years, mean 26.3 years), participated in exchange for performance-based payment plus a show-up fee. The participants earned on average 17.67 Euro (ranging from 12.70 to 21.70 Euro including the show-up fee of 2.50 Euro) and a session lasted about 1 hour and 45 minutes. The data was collected in two sessions with 28 and 32 participants at the CLER. The participants played in 15 different and independent groups.

#### 5.4.2 Results

As in Experiment 1, the self-reported questionnaire reveals that the cognitive load task was considered to be very important (mean 8.88,  $SD= 1.67$  on a scale from 0 to 10). The cognitive load task itself was considered to be of medium difficulty (mean 3.68,  $SD= 2.56$  on a scale from 0 to 10). The increase in difficulty is also reflected in higher error rates for the cognitive load task (mean error rate 13.26%,  $SD= 11.52$ , ranging from 0% to 54.17%).

We now turn to the response times of the main decision task. The left-hand side of Figure 5.2 depicts the average response time of decisions taken in the No Load and Load treatments. The average response time for decisions in the No Load treatment is 14.90 s while the average response time for decisions in the Load treatment is 9.89 s. As expected, a WSR test confirms our prediction and shows that decisions are significantly faster in the cognitive load treatment ( $N = 60$ ,  $z = 6.589$ ,  $p < 0.0001$ ).<sup>7</sup> This cognitive load task had a very strong effect on the response times of the main decision task.

---

<sup>7</sup>The conservative test at the group level obtains the same conclusion (WSR test,  $N = 15$ ,  $z = 3.408$ ,  $p = 0.0003$ ).

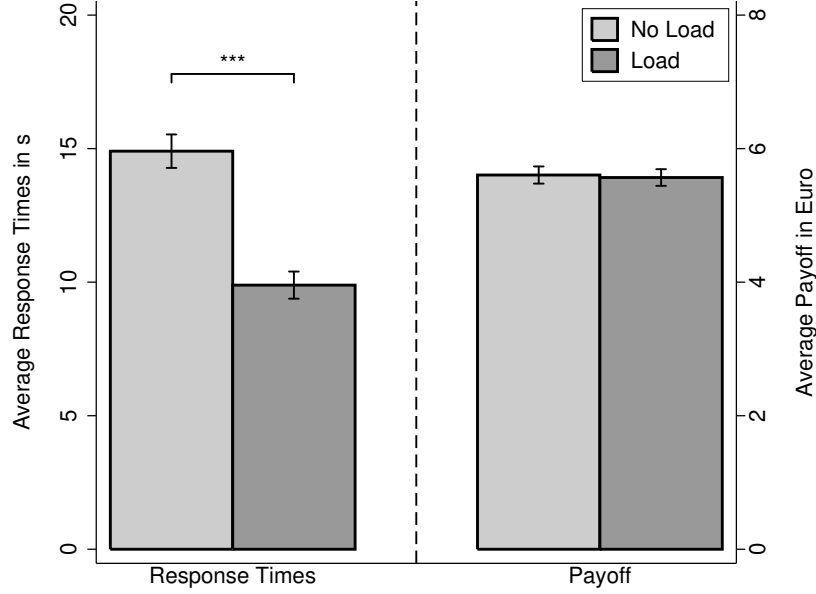


Figure 5.2: Experiment 2, Average response times and payoff.

Notes. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , WSR test.

Experiments 1 and 2 only differ in the cognitive load task which allows us to compare the effect of cognitive load across experiments. As intended, the cognitive Load task was rated significantly more difficult in Experiment 2 (mean 3.68) than the task in Experiment 1 (recall, mean 0.84) according to a Mann-Whitney-Wilcoxon (MWW) test ( $N = 124$ ,  $z = -6.529$ ,  $p < 0.0001$ ). The error rates are also significantly higher in Experiment 2 (13.26%) than error rates in Experiment 1 (recall, Experiment 1 mean 2.41%; MWW,  $N = 124$ ,  $z = -6.698$ ,  $p < 0.0001$ ). The response times in the No Load treatment are not significantly different across experiments (Experiment 1, mean 15.22 s; Experiment 2, mean 14.90 s; MWW,  $N = 124$ ,  $z = 0.460$ ,  $p = 0.6455$ ). We do, however, find significant differences for the Load treatment. The task which was evaluated more difficult also exhibits significantly faster responses in the main decision task (Experiment



1, mean 14.34 s; Experiment 2, mean 9.89 s; MWW,  $N = 124$ ,  $z = 5.775$ ,  $p < 0.0001$ ).<sup>8</sup>

We observe a much stronger response time effect for the main decision task and now turn our attention towards performance. The participants earned on average 5.60 Euro during all rounds of the No Load treatment and 5.57 Euro during all rounds of the cognitive Load treatment. This difference is not significant according to a WSR test ( $N = 60$ ,  $z = -0.272$ ,  $p = 0.7853$ ).<sup>9</sup> Another possible measure of performance could be relative usage of certain strategies, e.g. the research question might ask whether cognitive load shifts behavior towards a specific strategy. Similar to Alós-Ferrer and Ritschel (2018a) we could also test the relative frequency of the simple behavioral rule of imitation under cognitive load. One might expect that participants use such heuristic more often when cognitive resources are not available. In fact, we find significantly more relative imitation in the Load (mean 34.96%) than No Load treatment (mean 31.79%; WSR,  $z = 2.050$ ,  $p = 0.0202$ ).<sup>10</sup>

## 5.5 Experiment 3: Voting

### 5.5.1 Experimental Design and Procedures

For Experiment 3 we considered a complex voting decision. In Cournot oligopolies the participants interacted with each other. In contrast, this study was designed in such a way that the decision task was strictly individual, because no feedback on voting outcomes was provided until the end of the experiment. We followed Forsythe et al. (1993, 1996) (see also Granić, 2017) for the design of our voting experiment. The main decision task was to cast several votes using different voting methods. The participants were split into groups of 6 and voted on four different alternatives in the voting task.

---

<sup>8</sup>A test concerning the difference in response times between treatments across Experiments also is highly significant (Experiment 1, mean difference 0.88 s; Experiment 2, mean difference 5.01 s; MWW,  $N = 124$ ,  $z = 6.355$ ,  $p < 0.0001$ ). Tests at the group level obtain the same qualitative conclusions.

<sup>9</sup>A conservative test at the group level obtains the same conclusion.

<sup>10</sup>This one-sided test is even more significant on the group level ( $p = 0.0006$ ). These tests are not significant in Experiment 1.

Within a group the participants were further divided into three different types of voters with 2 participants each. We provided a table that indicated the payoffs for each type of voter for each of the four different alternatives, i.e. a  $3 \times 4$  payoff table. The participants then filled in their ballot according to the voting method.

The participants voted multiple times in two different voting blocks with a fixed voting method but did not receive any feedback. Therefore, all voting decisions were unaffected by the previous vote because no feedback was revealed and no information of the behavior of other players in the group is revealed, ergo the task can be viewed as an individual decision task. Each voting block consisted of 8 voting decisions each using plurality voting (PV) and approval voting (AV) as voting method. In PV each participant voted for exactly one of the alternatives and the alternative with the most votes won. In AV each participant voted for as many alternatives as she approved of and the alternative with the most approvals won. In case there was a tie among the votes, a random device determined the winner of the vote. In addition there was another “voting” method which always came at the end and served as preference elicitation. For the analysis we only consider the behavior in the PV and AV blocks. At the end of the experiment one voting round was randomly drawn and the winning alternative was determined according to the voting method and the votes of all members of the group.

We implemented the same cognitive load task as in Experiment 2 within subjects. They memorized a seven-digit number before entering the voting stage, voted, and then had to recall the number. In each voting block, 4 different payoff profiles were presented twice, once with cognitive Load and once with No Load. Overall, we collected voting behavior for 16 voting decisions: 4 in the Load and 4 in the No Load treatment for AV and 4 in the Load and 4 in the No Load treatment for PV. At the end of the experiment 1 of the 8 rounds with cognitive load was randomly drawn. If the participants correctly recalled the number they earned additional points.<sup>11</sup>

---

<sup>11</sup>For a correct recall the participants earned 40 points while the earnings from the voting decision ranged between 43 and 93 points.

Since only one randomly drawn round was paid we will not use the payoff as an performance measure in this experiment. We will measure performance as the number of sincere votes and in AV the number of votes cast. A sincere vote is defined as a vote which is in agreement with the subjects preferences and does not include any strategic voting. With respect to cognitive load a possible directional hypothesis could be that cognitive load reduces strategic behavior because less cognitive resources are available. The decrease in strategic behavior could be manifested by an increase in sincere votes for both voting methods. A change in behavior could also be exhibited by a change in the number of votes cast under approval voting.

60 participants (38 female; age range 18–32 years, mean 23.1 years) participated in exchange for performance-based payment plus a show-up fee. Students with a major in Psychology and who already participated in similar voting experiments were excluded from the participants pool. The participants earned on average 18.29 Euro (ranging from 12.00 to 22.20 Euro including the show-up fee of 4.00 Euro) and a session lasted about 1 hour and 15 minutes. The data was collected in two sessions with each 30 participants at the University of Cologne.

### 5.5.2 Results

The cognitive load task was the same as in Experiment 2 except that there were only a total of 8 rounds with cognitive load compared to a total of 24 rounds in Experiment 2. The mean error rate is 9.58% ( $SD= 12.47$ , ranging from 0% to 50.00%) and significantly smaller compared to Experiment 2 (recall, mean 13.26%, MWW,  $z = 2.591$ ,  $p = 0.0096$ ).

The participants voted for only one alternative in plurality voting and may approve of multiple alternatives in approval voting. By design the voting decision in approval voting might potentially take longer and we decided to split the response time analysis by voting method. Figure 5.3 shows the average response time of all decisions in the No Load and Load treatments for plurality voting (left-hand side) and approval voting (right-hand side). In the case of plurality voting, the average response time is 21.77 s in the No

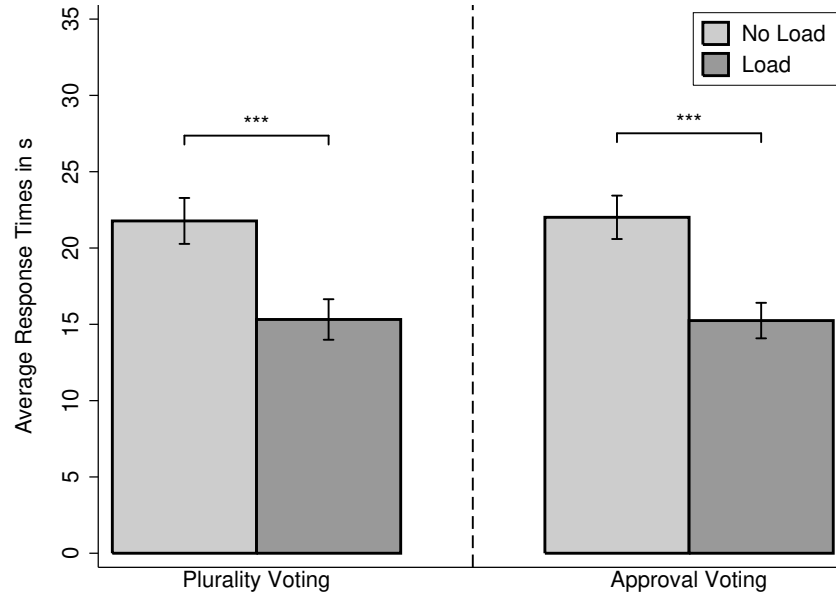


Figure 5.3: Experiment 3, Average response times of voting decisions.

Notes. \*\*\*  $p < 0.01$ , WSR test.

Load treatment and 15.32 s in the Load treatment. As expected, a WSR test confirms our prediction and shows that decisions are significantly faster in the cognitive Load treatment ( $N = 60$ ,  $z = 5.683$ ,  $p < 0.0001$ ). For approval voting, the average response time is 22.01 s in the No Load treatment and 15.25 s in the Load treatment. Again, a WSR test confirms that response times are significantly shorter in the Load treatment compared to the No Load treatment ( $N = 60$ ,  $z = 5.897$ ,  $p < 0.0001$ ).

Figure 5.4 shows the relative frequency of sincere votes for both voting methods and the number of votes under approval voting for the No Load and Load treatments. As performance measure we find significant differences between the amount of sincere votes between No Load and Load treatment for both voting methods. Under plurality voting, the participants voted on average 55.42% sincere in the No Load rounds and 64.17% sincere in the Load rounds. This difference is significant according to a WSR test ( $N = 60$ ,  $z = -2.260$ ,  $p = 0.0119$ ). Under approval voting, the participants voted on average 82.92% sincere in the No Load rounds and 87.50% sincere in the Load

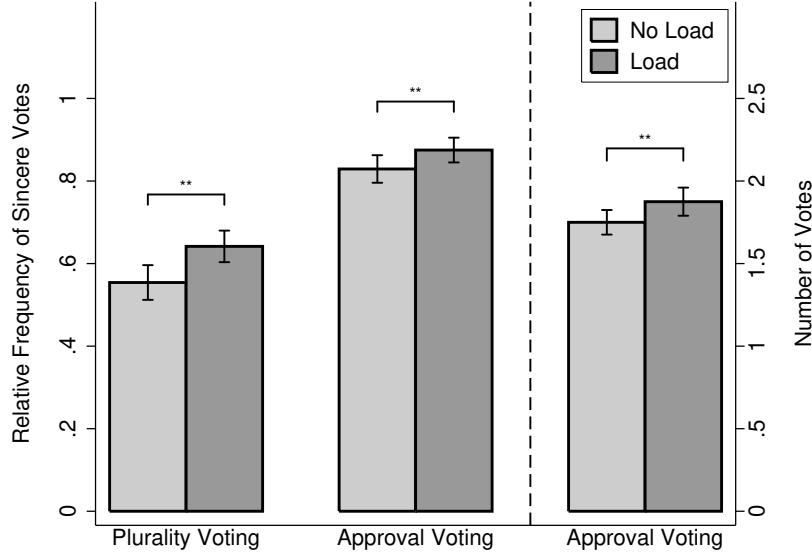


Figure 5.4: Experiment 3, Relative frequency of sincere votes and number of votes.  
*Notes.* \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , WSR test.

rounds. This difference is less pronounced but still significant according to a WSR test ( $N = 60$ ,  $z = -1.683$ ,  $p = 0.0462$ ). Another indicator of change in behavior in approval voting is the number of votes cast which significantly increased between treatments. The average number of votes cast is 1.75 votes in the No Load treatment and 1.875 votes in the Load treatment. According to a two-sided WSR test, this difference is significant ( $N = 60$ ,  $z = -2.522$ ,  $p = 0.0117$ ).

## 5.6 Experiment 4: Bayesian Updating

### 5.6.1 Experimental Design and Procedures

In Experiment 4 we changed again the economic setting and proposed two different cognitive load tasks. The main decision task used the Bayesian Updating Paradigm as implemented in Achtziger and Alós-Ferrer (2014). For this paradigm we expected very short response times between 1,000–5,000 ms (Achtziger and Alós-Ferrer, 2014) and therefore will report response

times in milliseconds. These are very short response times compared to the other experiments presented above. However, the decision task is much more complex in comparison to the other cognitive load studies discussed in the Introduction.

The main decision task was as follows. There were two urns (left and right urn) each with 6 colored balls. The composition of the colored balls in the urns depended on two different state of the worlds. In the “up” state of the world the left urn consisted of 4 black balls and 2 white balls and the right urn of 6 black balls. In the “down” state of the world the left urn consisted of 2 black balls and 4 white balls and the right urn of 6 white balls. The probability of an up and down state was  $1/2$ , known to the participants, and independent across rounds. In each round the participants were asked to choose which urn a single ball should be extracted from (with replacement), and were paid if and only if the ball was of a prespecified color (e.g. black).<sup>12</sup> Then they were asked to choose an urn a second time, a ball was extracted again, and they were paid again if the ball was of the appropriate color. They did this for 60 rounds and the participants earned 18 Eurocents for each ball with the appropriate color drawn from the urns.

The interesting draw was always the second choice of the round. The participant could have used the information of the extracted ball from the first draw to update the beliefs about the current state of the world or followed a simple reinforcement heuristic, i.e. win-stay, lose-shift. The composition of the urns were calibrated in such a way that both behavioral rules were either prescribing the same action, they were in alignment, or prescribed different actions, they were in conflict. Drawing a ball from the right urn perfectly revealed the state of the world and both rules prescribed the same action, i.e. if a winning ball was extracted choose the right urn again, otherwise choose the left urn. Drawing a ball from the left urn did not reveal the state of the world with certainty and updating the beliefs about the state of the world led to a different action than reinforcement. If a winning (black) ball was extracted, Bayesian Updating led to the conclusion that the upper state is more likely

---

<sup>12</sup>The colors of the balls for the composition of the urns and winning balls were counterbalanced. Here we present the case of black being the winning color.

and the best action was to switch to the right urn while if a losing (white) ball was extracted, the down state was more likely and Bayesian Updating prescribed to stay with the left urn. Therefore Bayesian Updating prescribed win-shift, lose-stay, the exact opposite of reinforcement's win-stay, lose-shift. An error in this decision task was defined as not choosing the urn prescribed by Bayesian Updating in both conflict and alignment situations. For the analysis we will condition on conflict and alignment situations and use as a performance measure the error rates in conflict and alignment situations.

In this setting we implemented three treatments, one No Load treatment and two different cognitive load treatments. In the No Load treatment, participants were not put under any kind of load. In the *Phonological Load* and *Central Executive Load* treatments, participants were put under two different forms of load. The Phonological Load treatment required participants to complete the main task while repeating the word “and” every 1.5 seconds, following the rhythm given by a metronome (this device was physically present on the table, i.e. not integrated into the computer program). This manipulation is known to specifically block the phonological loop, which should lead to quick information decay (forgetting; Baddeley (1986), Gathercole and Baddeley (1993)) similar to memorizing a long sequence of digits. The Central Executive Load required participants to name random numbers (from zero to nine) aloud at the rhythm of the metronome. This manipulation is known to seriously impair central executive functions and in addition tax working memory capacity (e.g. attention) to a strong extent.

In all cases, participants received careful instructions on the decision task and in addition were shown how to implement the cognitive load task. Participants were asked to practice the load in the presence of the experimenter and were instructed that successfully conducting this secondary task was a precondition for payment in the main task. They also practiced the load together with the main task (five practice trials). In the experiment itself, participants first started employing the cognitive load right before they started working on the decision task. Accordingly, working memory capacity was already limited by the load manipulation before participants started with the posterior probability task. Participants in the No Load condition started with the

decision task after instructions and practice trials. The articulation of the load task was recorded by means of a computer and provided incentives for actually conducting this task. After the experiment, it was checked whether participants actually carried out the cognitive load task. No participant neglected to perform the secondary task. After the experiment, participants answered a questionnaire and their accumulated earnings were paid.

Participants were 58 university students (20 female, 38 male), randomly allocated to three different treatments, who performed the same Bayesian Updating task.<sup>13</sup> Participants earned on average 11.63 Euro and a session lasted around one hour. The experiment was conducted at the LakeLab (University of Konstanz).

### 5.6.2 Results

Figure 5.5 shows the average response times of the second draw in the No Load, Phonological Load, and Central Executive Load treatments conditional on conflict (left-hand side) and alignment (right-hand side) situations. In conflict situations, we find significantly shorter response times for the two cognitive load treatments compared to the No Load treatment. The average response time is 2,712 ms in the No Load treatment, 1,374 ms in the Phonological Load treatment, and 1,617 ms in the Central Executive Load treatment. The differences between the No Load and cognitive load treatments are significant according to Mann-Whitney-Wilcoxon tests (No Load vs. Phonological Load,  $N = 40$ ,  $z = 2.624$ ,  $p = 0.0087$ ; No Load vs. Central Executive Load,  $N = 38$ ,  $z = 2.017$ ,  $p = 0.0218$ ).<sup>14</sup>

In alignment situations, we find in general shorter response times than in conflict situations. We do not, however, find significantly faster responses in the cognitive load treatments than in the No Load treatment. The average response time is 1,056 ms in the No Load treatment, 1,110 ms in the Phonological Load treatment, and 1,558 ms in the Central Executive Load

---

<sup>13</sup>Two participants of the original 20 in the Central Executive Load treatment had to be excluded from the analysis due to a failure of the recordings necessary to compute their redundancy scores to determine the performance in the load task (we refer to Baddeley (1966) for details).

<sup>14</sup>The one-sided  $p$ -values are Holm-Bonferroni adjusted.






Figure 5.5: Experiment 4, Average response times of the second draw.

*Notes.* Response times conditional on conflict (left-hand side) and alignment situations (right-hand side) . \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , MWW test.

treatment. The differences between the No Load and cognitive load treatments are not significant (MWW: No Load vs. Phonological Load,  $N = 40$ ,  $z = -0.974$ ,  $p > 0.9999$ ; No Load vs. Central Executive Load,  $N = 38$ ,  $z = -1.959$ ,  $p = 0.9749$ ).

Figure 5.6 shows the average error rates according to optimal behavior (following Bayesian Updating) of the second draw in the No Load, Phonological Load, and Central Executive Load treatments conditional on conflict (left-hand side) and alignment (right-hand side) situations. In conflict situations, we do not find strong evidence for the two cognitive load treatments resulting in more errors compared to the No Load treatment. The average error rate is 50.04% in the No Load treatment, 66.94% in the Phonological Load treatment, and 55.66% in the Central Executive Load treatment. The difference between the No Load and Phonological Load treatments just barely misses significance (MWW:  $N = 40$ ,  $z = -1.922$ ,  $p = 0.1093$ ) and we do not

find any significant difference for error rates between the No Load and Central Executive Load treatments (MWW:  $N = 38$ ,  $z = -0.512$ ,  $p = 0.6088$ ).<sup>15</sup>

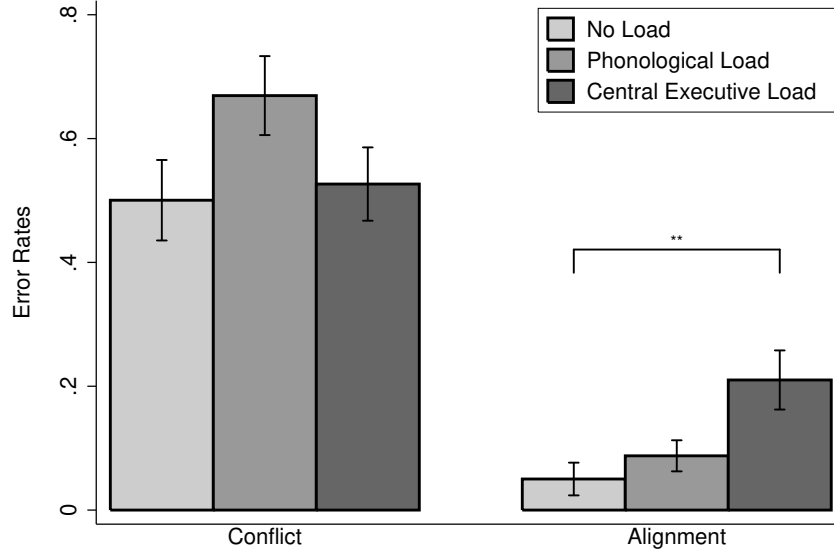


Figure 5.6: Experiment 4, Average error rates of the second draw.

*Notes.* Error rates conditional on conflict (left-hand side) and alignment situations (right-hand side). \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ , MWW test.

In alignment situations, we find in general much lower error rates. The average error rate is 5.02% in the No Load treatment, 8.77% in the Phonological Load treatment, and 21.02% in the Central Executive Load treatment. The difference between the No Load and Phonological Load treatments is not significant (MWW:  $N = 40$ ,  $z = -1.417$ ,  $p = 0.1565$ ), however, we find a significant difference for error rates between the No Load and Central Executive Load treatments (MWW:  $N = 38$ ,  $z = -2.512$ ,  $p = 0.0240$ ).<sup>16</sup>

<sup>15</sup>A directional hypothesis of cognitive load decreasing performance with an one-sided  $p$ -value (and Holm-Bonferroni adjustment) finds significantly more errors under phonological load than under no load ( $p = 0.0547$ ). The conclusion for central executive load does not change for one-sided  $p$ -values.

<sup>16</sup>One-sided  $p$ -values (and Holm-Bonferroni adjustments) find weakly significant more errors under Phonological Load than under No Load ( $p = 0.0783$ ) and significantly more errors under Central Executive Load than under No Load ( $p = 0.0112$ ).

## 5.7 Discussion

In a series of experiments we showed that cognitive load significantly reduces response times in complex, economic decision tasks. This observation is consistent across a variety of very different economic paradigms such as behavior in Cournot oligopoly, in plurality or approval voting, or in belief-based decisions under risk. We also find a consistent reduction in response times for different kinds of cognitive load tasks such as adding a number from memory to a number heard, memorizing a seven-digit number, repeating “and,” and generating a random number every 1.5 seconds. We varied the main decision task and the cognitive load task quite substantially across experiments and always find consistently that under cognitive load the response times are significantly shorter than under no load. Further, some experiments employing economic decisions under risk (lottery choice) have reported response times and can be used to offer further evidence in favor of our test. Whitney et al. (2008) analyze the impact of cognitive load (memorizing a five-letter string and recalling a specific letter) on framing effects while choosing between a gamble and sure outcome, and report that response times decreased significantly from 2,950 ms without load to 2,796 ms with load. Gerhardt et al. (2016) investigate risk attitudes in a lottery-choice experiment with cognitive load. They employed a dot-pattern memorization task as cognitive load manipulation, and report that response times decreased significantly from 3,835 ms without load to 3,449 ms with load. Both papers do find behavioral effects of cognitive load (less gambling and less riskier decisions, respectively). The effect on response times in those papers was unexpected, and the authors simply speculate that participants might have tried to speed up their decisions in order to maintain accuracy in the cognitive load task.

We remark that Experiment 1 and 2 used the same economic paradigm but different cognitive load tasks which differed significantly in difficulty. The load tasks in Experiment 2 was significantly more difficult than the task in Experiment 1 according to self-reported questionnaires and error rates in those load tasks. We reach the same conclusion with our analysis of response times of the main decision task. The response time difference between the No

Load and Load treatments was significantly larger for Experiment 2 (with the more difficult task) than for Experiment 1 (with the less difficult task). This suggests a correlation of response time reduction and reduction of cognitive resources as result of cognitive load task difficulty.

This measure of response times is a very direct measure of the cognitive load effect on the main decision task independent of the change in behavior. It does not rely on a self-reported measure and also not on the performance in the cognitive load task but directly establishes a link between the cognitive load and the main decision of interest. We argue that the reduction in response times under cognitive load is a direct indicator that a cognitive load task has successfully induced a shift in the kind of decision processes used by participants, as a consequence of a successful impairment of available cognitive resources. Then, and only then, the researcher can consider whether and how cognitive load affected the behavior in the main decision task, and the implications of behavioral change or lack thereof. This simple test is readily available in almost all cognitive load experiments.

---

## References

---

- ABBINK, K., J. BRANDTS, B. HERRMANN, AND H. ORZEN (2010): “Inter-group Conflict and Intra-group Punishment in an Experimental Contest Game,” *American Economic Review*, 100, 420–447.
- (2012): “Parochial Altruism in Inter-group Conflicts,” *Economics Letters*, 117, 45–48.
- ABBINK, K. AND B. HERRMANN (2011): “The Moral Costs of Nastiness,” *Economic Inquiry*, 49, 631–633.
- ABBINK, K., B. IRLBUSCH, AND E. RENNER (2000): “The Moonlighting Game: An Experimental Study on Reciprocity and Retribution,” *Journal of Economic Behavior and Organization*, 42, 265–277.
- ABBINK, K. AND A. SADRIEH (2009): “The Pleasure of Being Nasty,” *Economics Letters*, 105, 306–308.
- ACHTZIGER, A. AND C. ALÓS-FERRER (2014): “Fast or Rational? A Response-Times Study of Bayesian Updating,” *Management Science*, 60, 923–938.
- ACHTZIGER, A., C. ALÓS-FERRER, S. HÜGELSCHÄFER, AND M. STEINHAUSER (2014): “The Neural Basis of Belief Updating and Rational Decision Making,” *Social Cognitive and Affective Neuroscience*, 9, 55–62.
- (2015a): “Higher Incentives Can Impair Performance: Neural Evidence on Reinforcement and Rationality,” *Social Cognitive and Affective Neuroscience*, 10, 1477–1483.
- ACHTZIGER, A., C. ALÓS-FERRER, AND A. RITSCHER (2019): “Cognitive Load in Economic Decisions: How To Tell Whether It Works,” Working Paper, University of Zurich.

- ACHTZIGER, A., C. ALÓS-FERRER, AND A. K. WAGNER (2015b): “Money, Depletion, and Prosociality in the Dictator Game,” *Journal of Neuroscience, Psychology, and Economics*, 8, 1–14.
- (2016): “The Impact of Self-Control Depletion on Social Preferences in the Ultimatum Game,” *Journal of Economic Psychology*, 53, 1–16.
- (2018): “Social Preferences and Self-Control,” *Journal of Behavioral and Experimental Economics*, 74, 161–166.
- ALGER, I. AND J. W. WEIBULL (2013): “Homo Moralis–Preference Evolution under Incomplete Information and Assortative Matching,” *Econometrica*, 81, 2269–2302.
- ALLRED, S., S. DUFFY, AND J. SMITH (2016): “Cognitive Load and Strategic Sophistication,” *Journal of Economic Behavior and Organization*, 125, 162–178.
- ALÓS-FERRER, C. (2018): “A Dual-Process Diffusion Model,” *Journal of Behavioral Decision Making*, 31, 203–218.
- ALÓS-FERRER, C. AND A. B. ANIA (2005): “The Evolutionary Stability of Perfectly Competitive Behavior,” *Economic Theory*, 26, 179–197.
- ALÓS-FERRER, C. AND J. BUCKENMAIER (2018): “Cognitive Sophistication and Deliberation Times,” Working Paper, University of Zurich.
- ALÓS-FERRER, C., D.-G. GRANIĆ, J. KERN, AND A. K. WAGNER (2016a): “Preference Reversals: Time and Again,” *Journal of Risk and Uncertainty*, 52, 65–97.
- ALÓS-FERRER, C. AND S. HÜGELSCHÄFER (2012): “Faith in Intuition and Behavioral Biases,” *Journal of Economic Behavior and Organization*, 84, 182–192.
- (2016): “Faith in Intuition and Cognitive Reflection,” *Journal of Behavioral and Experimental Economics*, 64, 61–70.

- ALÓS-FERRER, C., S. HÜGELSCHÄFER, AND J. LI (2016b): “Inertia and Decision Making,” *Frontiers in Psychology*, 7 (169), 1–9.
- (2017): “Framing Effects and the Reinforcement Heuristic,” *Economics Letters*, 156, 32–35.
- ALÓS-FERRER, C. AND A. RITSCHER (2018a): “Multiple Behavioral Rules in Cournot Oligopolies,” Working Paper, University of Zurich.
- (2018b): “The Reinforcement Heuristic in Normal Form Games,” *Journal of Economic Behavior and Organization*, 152, 224–234.
- ALÓS-FERRER, C. AND F. STRACK (2014): “From Dual Processes to Multiple Selves: Implications for Economic Behavior,” *Journal of Economic Psychology*, 41, 1–11.
- ALÓS-FERRER, C. AND S. WEIDENHOLZER (2007): “Partial Bandwagon Effects and Local Interactions,” *Games and Economic Behavior*, 61, 1–19.
- (2008): “Contagion and Efficiency,” *Journal of Economic Theory*, 143, 251–274.
- ANDERSEN, S., S. ERTAÇ, U. GNEEZY, M. HOFFMAN, AND J. A. LIST (2011): “Stakes Matter in Ultimatum Games,” *American Economic Review*, 101, 3427–3439.
- ANDREONI, J. (1989): “Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence,” *Journal of Political Economy*, 97, 1447–1458.
- (1995): “Warm-Glow Versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments,” *Quarterly Journal of Economics*, 110, 1–21.
- ANDREONI, J., J. M. RAO, AND H. TRACHTMAN (2017): “Avoiding the Ask: A Field Experiment on Altruism, Empathy, and Charitable Giving,” *Journal of Political Economy*, 125, 625–653.

- APESTEGUÍA, J., G. AZMAT, AND N. IRIBERRI (2012): “The Impact of Gender Composition on Team Performance and Decision Making: Evidence from the Field,” *Management Science*, 58, 78–93.
- APESTEGUÍA, J., S. HUCK, AND J. OECHSSLER (2007): “Imitation—Theory and Experimental Evidence,” *Journal of Economic Theory*, 136, 217–235.
- APESTEGUÍA, J., S. HUCK, J. OECHSSLER, AND S. WEIDENHOLZER (2010): “Imitation and The Evolution of Walrasian Behavior: Theoretically Fragile But Behaviorally Robust,” *Journal of Economic Theory*, 145, 1603–1617.
- ARMANTIER, O. (2006): “Do Wealth Differences Affect Fairness Considerations?” *International Economic Review*, 47, 391–429.
- AZMAT, G., C. CALSAMIGLIA, AND N. IRIBERRI (2016): “Gender Differences in Response to Big Stakes,” *Journal of the European Economic Association*, 14, 1372–1400.
- AZMAT, G. AND N. IRIBERRI (2010): “The Importance of Relative Performance Feedback Information: Evidence from a Natural Experiment using High School Students,” *Journal of Public Economics*, 94, 435–452.
- (2016): “The Provision of Relative Performance Feedback: An Analysis of Performance and Satisfaction,” *Journal of Economics & Management Strategy*, 25, 77–110.
- AZRIELI, Y., C. P. CHAMBERS, AND P. J. HEALY (2018): “Incentives in Experiments: A Theoretical Analysis,” *Journal of Political Economy*, 126, 1472–1503.
- BADDELEY, A. (1986): *Working Memory*, Oxford: Clarendon Press.
- (1992): “Working Memory,” *Science*, 255, 556–559.
- (1996): “Exploring the Central Executive,” *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 49, 5–28.



- BADDELEY, A., D. CHINCOTTA, AND A. ADLAM (2001): “Working Memory and the Control of Action: Evidence From Task Switching,” *Journal of Experimental Psychology: General*, 130, 641–657.
- BADDELEY, A. AND G. HITCH (1974): “Working Memory,” *The Psychology of Learning and Motivation*, 8, 47–89.
- BADDELEY, A., V. LEWIS, M. ELDRIDGE, AND N. THOMSON (1984): “Attention and Retrieval from Long-Term Memory,” *Journal of Experimental Psychology: General*, 113, 518–540.
- BADDELEY, A. D. (1966): “The Capacity for Generating Information by Randomization,” *Quarterly Journal of Experimental Psychology*, 18, 119–129.
- BANDIERA, O., I. BARANKAY, AND I. RASUL (2005): “Social Preferences and the Response to Incentives: Evidence from Personnel Data,” *Quarterly Journal of Economics*, 120, 917–962.
- BARDSLEY, N. (2008): “Dictator Game Giving: Altruism or Artifact?” *Experimental Economics*, 11, 122–133.
- BARGH, J. A. (1989): “Conditional Automaticity: Varieties of Automatic Influences in Social Perception and Cognition,” in *Unintended Thought*, ed. by J. S. Uleman and J. A. Bargh, Guilford, 3–51.
- BARON, J. AND J. C. HERSHEY (1988): “Outcome Bias in Decision Evaluation,” *Journal of Personality and Social Psychology*, 54, 569–579.
- BARROUILLET, P., S. BERNARDIN, S. PORTRAT, E. VERGAUWE, AND V. CAMOS (2007): “Time and Cognitive Load in Working Memory,” *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 570–585.
- BELLEMARE, C. AND S. KRÖGER (2007): “On Representative Social Capital,” *European Economic Review*, 51, 183–202.

- BÉNABOU, R. AND J. TIROLE (2006): “Incentives and Prosocial Behavior,” *American Economic Review*, 96, 1652–1678.
- (2011): “Identity, Morals, and Taboos: Beliefs as Assets,” *Quarterly Journal of Economics*, 126, 805–855.
- BENJAMIN, D. J., S. A. BROWN, AND J. M. SHAPIRO (2013): “Who is ‘Behavioral’? Cognitive Ability and Anomalous Preferences,” *Journal of European Economic Association*, 11, 1231–1255.
- BERG, J., J. DICKHAUT, AND K. MCCABE (1995): “Trust, Reciprocity, and Social History,” *Games and Economic Behavior*, 10, 122–142.
- BERGER, J., C. HARBRING, AND D. SLIWKA (2013): “Performance Appraisals and the Impact of Forced Distribution—An Experimental Investigation,” *Management Science*, 59, 54–68.
- BLANCO, M., D. ENGELMANN, AND H. T. NORMANN (2011): “A Within-Subject Analysis of Other-regarding Preferences,” *Games and Economic Behavior*, 72, 321–338.
- BLANES I VIDAL, J. AND M. NOSSOL (2011): “Tournaments without Prizes: Evidence from Personnel Records,” *Management Science*, 57, 1721–1736.
- BOLTON, G. E. AND A. OCKENFELS (2000): “ERC: A Theory of Equity, Reciprocity, and Competition,” *American Economic Review*, 90, 166–193.
- BÖRGERS, T. AND R. SARIN (1997): “Learning Through Reinforcement and Replicator Dynamics,” *Journal of Economic Theory*, 77, 1–14.
- BOSMAN, R., H. HENNIG-SCHMIDT, AND F. VAN WINDEN (2006): “Exploring Group Decision Making in a Power-to-Take Experiment,” *Experimental Economics*, 9, 35–51.
- BOSMAN, R., M. SUTTER, AND F. VAN WINDEN (2005): “The Impact of Real Effort and Emotions in the Power-To-Take Game,” *Journal of Economic Psychology*, 26, 407–429.

- BOSMAN, R. AND F. VAN WINDEN (2002): “Emotional Hazard in a Power-to-Take Experiment,” *Economic Journal*, 476, 147–169.
- BRANDTS, J. AND G. CHARNES (2011): “The Strategy versus the Direct-response Method: A First Survey of Experimental Comparisons,” *Experimental Economics*, 14, 375–398.
- BRUNSTEIN, J. C. AND H. HECKHAUSEN (2008): “Achievement Motivation,” in *Motivation and Action*, ed. by J. Heckhausen and H. Heckhausen, New York: Cambridge University Press, 137–183.
- BUCKERT, M., J. OECHSSLER, AND C. SCHWIEREN (2017): “Imitation Under Stress,” *Journal of Economic Behavior & Organization*, 139, 252–266.
- BURKS, S. V., J. P. CARPENTER, AND E. VERHOOGEN (2003): “Playing Both Roles in the Trust Game,” *Journal of Economic Behavior and Organization*, 51, 195–216.
- CAMERON, A. C., J. B. GELBACH, AND D. L. MILLER (2008): “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 90, 414–427.
- CAPPELEN, A. W., K. O. MOENE, E. Ø. SØRENSEN, AND B. TUNGODDEN (2013): “Needs versus Entitlements—An International Fairness Experiment,” *Journal of the European Economic Association*, 11, 574–598.
- CAPPELEN, A. W., U. H. NIELSEN, B. TUNGODDEN, J.-R. TYRAN, AND E. WENGSTRÖM (2016): “Fairness is Intuitive,” *Experimental Economics*, 19, 727–740.
- CAPPELLETTI, D., W. GÜTH, AND M. PLONER (2011): “Being of Two Minds: an Ultimatum Experiment Investigating Affective Processes,” *Journal of Economic Psychology*, 32, 940–950.
- CARLSMITH, J. M. AND A. E. GROSS (1969): “Some Effects of Guilt on Compliance,” *Journal of Personality and Social Psychology*, 11, 232–239.

- CARPENTER, J., M. GRAHAM, AND J. WOLF (2013): “Cognitive Ability and Strategic Sophistication,” *Games and Economic Behavior*, 80, 115–130.
- CASARI, M. AND T. N. CASON (2009): “The Strategy Method Lowers Measured Trustworthy Behavior,” *Economics Letters*, 103, 157–159.
- CHARNESS, G. AND U. GNEEZY (2008): “What’s in a Name? Anonymity and Social Distance in Dictator and Ultimatum Games,” *Journal of Economic Behavior and Organization*, 68, 29–35.
- CHARNESS, G., U. GNEEZY, AND B. HALLADAY (2016): “Experimental Methods: Pay One or Pay All,” *Journal of Economic Behavior and Organization*, 131, 141–150.
- CHARNESS, G. AND D. LEVIN (2005): “When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect,” *American Economic Review*, 95, 1300–1309.
- CHARNESS, G. AND M. RABIN (2002): “Understanding Social Preferences with Simple Tests,” *Quarterly Journal of Economics*, 117, 817–869.
- CORNELISSEN, G., S. DEWITTE, AND L. WARLOP (2011): “Are Social Value Orientations Expressed Automatically? Decision Making in the Dictator Game,” *Personality and Social Psychology Bulletin*, 37, 1080–1090.
- CROSON, R. AND U. GNEEZY (2009): “Gender Differences in Preferences,” *Journal of Economic Literature*, 47, 448–474.
- DAHME, G., D. JUNGnickel, AND H. RATHJE (1993): “Psychometric Properties of a German Translation of the Achievement Motives Scale (AMS): Comparison of Results from Norwegian and German Samples,” *Diagnostica*, 39, 257–270.
- DARIEL, A. AND N. NIKIFORAKIS (2014): “Cooperators and Reciprocators: A Within-subject Analysis of Pro-social Behavior,” *Economics Letters*, 122, 163–166.

- DASHIELL, J. F. (1937): “Affective Value-Distances as a Determinant of Aesthetic Judgment-Times,” *American Journal of Psychology*, 50, 57–67.
- DAW, N. D. (2012): “Model-Based Reinforcement Learning as Cognitive Search: Neurocomputational Theories,” in *Cognitive Search: Evolution, Algorithms and the Brain*, ed. by P. M. Todd, T. T. Hills, and T. W. Robbins, Cambridge, MA: MIT Press, 195–208.
- DAW, N. D. AND P. TOBLER (2014): “Value Learning through Reinforcement: The Basics of Dopamine and Reinforcement Learning,” in *Neuroeconomics: Decision Making and the Brain*, ed. by P. W. Glimcher and E. Fehr, London: Academic Press, 283–298, 2nd ed.
- DE FOCKERT, J. W., G. REES, C. D. FRITH, AND N. LAVIE (2001): “The Role of Working Memory in Visual Selective Attention,” *Science*, 291, 1803–1806.
- DECK, C. AND S. JAHEDI (2015): “The Effect of Cognitive Load on Economic Decision Making: A Survey and New Experiments,” *Personality and Social Psychology Bulletin*, 78, 97–119.
- DECLERCK, C. H., T. KIYONARI, AND C. BOONE (2009): “Why Do Responders Reject Unequal Offers in the Ultimatum Game? An Experimental Study on the Role of Perceiving Interdependence,” *Journal of Economic Psychology*, 30, 335–343.
- DELLAVIGNA, S., J. LIST, AND U. MALMENDIER (2012): “Testing for Altruism and Social Pressure in Charitable Giving,” *Quarterly Journal of Economics*, 127, 1–56.
- DELQUIÉ, P. (1993): “Inconsistent Trade-Offs Between Attributes: New Evidence in Preference Assessment Biases,” *Management Science*, 39, 1382–1395.
- DILLON, R. L. AND C. H. TINSLEY (2008): “How Near-Misses Influence Decision Making Under Risk: A Missed Opportunity for Learning,” *Management Science*, 54, 1425–1440.

- DØSSING, F., M. PIOVESAN, AND E. WENGSTRÖM (2017): “Cognitive Load and Cooperation,” *Review of Behavioral Economics*, 4, 69–81.
- DRICHOUTIS, A. C. AND R. NAYGA (2017): “Economic Rationality Under Cognitive Load,” MPRA Working Paper 81111. Available at: <https://ideas.repec.org/p/pra/mprapa/81111.html>.
- DUFFY, S., J. J. NADDEO, D. OWENS, AND J. SMITH (2016): “Cognitive Load and Mixed Strategies: On Brains and Minimax,” MPRA Working Paper 71878. Available at: [https://mpraub.uni-muenchen.de/71878/1/MPRA\\_paper\\_71878.pdf](https://mpraub.uni-muenchen.de/71878/1/MPRA_paper_71878.pdf).
- DUFFY, S. AND J. SMITH (2014): “Cognitive Load in the Multi-Player Prisoner’s Dilemma Game: Are There Brains in Games?” *Journal of Behavioral and Experimental Economics*, 51, 47–56.
- DUFWENBERG, M. AND G. KIRCHSTEIGER (2004): “A Theory of Sequential Reciprocity,” *Games and Economic Behavior*, 47, 268–298.
- ECKEL, C. C. AND R. PETRIE (2011): “Face Value,” *American Economic Review*, 101, 1497–1513.
- ENGEL, C. (2011): “Dictator Games: A Meta Study,” *Experimental Economics*, 14, 583–610.
- ENGELMANN, D. AND M. STROBEL (2004): “Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments,” *American Economic Review*, 94, 857–869.
- EPSTEIN, S., R. PACINI, V. DENES-RAJ, AND H. HEIER (1996): “Individual Differences in Intuitive-Experiential And Analytical-Rational Thinking Styles,” *Journal of Personality and Social Psychology*, 71, 390–405.
- EREV, I. AND E. HARUVY (2016): “Learning and the Economics of Small Decisions,” in *Handbook of Experimental Economics, Volume 2*, ed. by J. Kagel and A. E. Roth, Princeton University Press, 638–716.

- EREV, I. AND A. E. ROTH (1998): “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *American Economic Review*, 88, 848–881.
- EVANS, J. S. (2008): “Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition,” *Annual Review of Psychology*, 59, 255–278.
- FALK, A. AND U. FISCHBACHER (2006): “A Theory of Reciprocity,” *Games and Economic Behavior*, 54, 293–315.
- FATAS, E., A. MORALES, AND A. JARAMILLO-GUTIÉRREZ (2015): “Social Competition in Low-Information Cournot Markets,” CBESS Discussion Paper 15-15, School of Economics, University of East Anglia.
- FEHR, E., M. NAEF, AND K. M. SCHMIDT (2006): “Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments: Comment,” *American Economic Review*, 96, 1912–1917.
- FEHR, E. AND K. M. SCHMIDT (1999): “A Theory of Fairness, Competition, and Cooperation,” *Quarterly Journal of Economics*, 114, 817–868.
- FERSHTMAN, C., U. GNEEZY, AND J. A. LIST (2012): “Equity Aversion: Social Norms and the Desire to Be Ahead,” *American Economic Journal: Microeconomics*, 4, 131–144.
- FESTINGER, L. (1954): “A Theory of Social Comparison Processes,” *Human Relations*, 7, 117–140.
- FILIPPIN, A. AND P. CROSETTO (2016): “A Reconsideration of Gender Differences in Risk Attitudes,” *Management Science*, 62, 3138–3160.
- FISCHBACHER, U. (2007): “z-Tree: Zurich Toolbox for Ready-Made Economic Experiments,” *Experimental Economics*, 10, 171–178.
- FISCHBACHER, U., R. HERTWIG, AND A. BRUHIN (2013): “How to Model Heterogeneity in Costly Punishment: Insights from Responders’ Response Times,” *Journal of Behavioral Decision Making*, 26, 462–476.

- FISMAN, R., S. KARIV, AND D. MARKOVITS (2007): “Individual Preferences for Giving,” *American Economic Review*, 97, 1858–1876.
- FORSYTHE, R., J. L. HOROWITZ, N. E. SAVIN, AND M. SEFTON (1994): “Fairness in Simple Bargaining Experiments,” *Games and Economic Behavior*, 6, 347–369.
- FORSYTHE, R., R. B. MYERSON, T. A. RIETZ, AND R. J. WEBER (1993): “An Experiment on Coordination in Multi-Candidate Elections: The Importance of Polls and Election Histories,” *Social Choice and Welfare*, 10, 223–247.
- FORSYTHE, R., T. A. RIETZ, R. B. MYERSON, AND R. J. WEBER (1996): “An Experimental Study of Voting Rules and Polls in Three-Candidate Elections,” *International Journal of Game Theory*, 25, 355–383.
- FRANZEN, A. AND S. POINTNER (2012): “Anonymity in the Dictator Game Revisited,” *Journal of Economic Behavior and Organization*, 81, 74–81.
- FUDENBERG, D. AND D. K. LEVINE (1998): *The Theory of Learning in Games*, Cambridge, Massachusetts: The MIT Press.
- GALIZZI, M. M. AND D. NAVARRO-MARTÍNEZ (2018): “On the External Validity of Social-Preference Games: A Systematic Lab-Field Study,” *Management Science*, forthcoming.
- GARDNER, R., E. OSTROM, AND J. M. WALKER (1990): “The Nature of Common-Pool Resource Problems,” *Rationality and Society*, 2, 335–358.
- GATHERCOLE, S. E. AND A. D. BADDELEY (1993): *Working Memory and Language*, Hillsdale, NJ: Lawrence Erlbaum Associates.
- GERHARDT, H., G. P. BIELE, H. R. HEEKEREN, AND H. UHLIG (2016): “Cognitive Load Increases Risk Aversion,” SFB 649 Discussion Paper No. 2016-011. Available at: <http://sfb649.wiwi.hu-berlin.de/papers/pdf/SFB649DP2016-011.pdf>.



- GERLITZ, J.-Y. AND J. SCHUPP (2005): “Zur Erhebung der Big-Five-basierten Persönlichkeitsmerkmale im SOEP,” *German Institute of Economic Research (DIW) Research Notes 2005-04*.
- GERSHMAN, S. J., B. PESARAN, AND N. D. DAW (2009): “Human Reinforcement Learning Subdivides Structured Action Spaces by Learning Effector-Specific Values,” *The Journal of Neuroscience*, 29, 13524–13531.
- GEVINS, A., M. E. SMITH, H. LEONG, L. MCEVOY, S. WHITFIELD, R. DU, AND G. RUSH (1998): “Monitoring Working Memory Load During Computer-Based Tasks with EEG Pattern Recognition Methods,” *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 40, 79–91.
- GILLIGAN, C. (1982): *In a Different Voice*, Harvard University Press.
- GINER-SOROLLA, R. (2001): “Guilty Pleasures and Grim Necessities: Affective Attitudes in Dilemmas of Self-Control,” *Journal of Personality and Social Psychology*, 80, 206–221.
- GITTINS, J. C. (1979): “Bandit Processes and Dynamic Allocation Indices,” *Journal of the Royal Statistical Society, Series B*, 41, 148–177.
- (1989): *Multi-Armed Bandit Allocation Indices*, New York: Wiley.
- GLAESER, E. L., D. I. LAIBSON, J. A. SCHEINKMAN, AND C. L. SOUTTER (2000): “Measuring Trust,” *Quarterly Journal of Economics*, 115, 811–846.
- GLASER, M. AND T. WALTHER (2014): “Run, Walk, or Buy? Financial Literacy, Dual-Process Theory, and Investment Behavior,” Working Paper. Available at: <https://ssrn.com/abstract=2167270>.
- GNEEZY, U., A. IMAS, AND K. MADARÁSZ (2014): “Conscience Accounting: Emotion Dynamics and Social Behavior,” *Management Science*, 60, 2645–2658.

- GNEEZY, U., K. L. LEONARD, AND J. A. LIST (2009): “Gender Differences in Competition: Evidence from a Matrilineal and a Patriarchal Society,” *Econometrica*, 77, 1637–1664.
- GNEEZY, U., M. NIEDERLE, AND A. RUSTICHINI (2003): “Performance in Competitive Environments: Gender Differences,” *Quarterly Journal of Economics*, 118(3), 1049–1074.
- GRANIĆ, G. (2017): “The Problem of the Divided Majority: Preference Aggregation Under Uncertainty,” *Journal of Economic Behavior and Organization*, 133, 21–38.
- GREENE, J. D., S. A. MORELLI, K. LOWENBERG, L. E. NYSTROM, AND J. D. COHEN (2008): “Cognitive Load Selectively Interferes with Utilitarian Moral Judgment,” *Cognition*, 107, 1144–1154.
- GREGG, A. P., E. G. HEPPER, AND C. SEDIKIDES (2011): “Quantifying Self-Motives: Functional Links Between Dispositional Desires,” *European Journal of Social Psychology*, 41, 840–852.
- GREIFENEDER, R. AND C. BETSCH (2006): “Lieber die Taube auf dem Dach!” *Zeitschrift für Sozialpsychologie*, 37, 233–243.
- GREINER, B. (2015): “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE,” *Journal of the Economic Science Association*, 1, 114–125.
- GÜTH, W., R. SCHMITTBERGER, AND B. SCHWARZE (1982): “An Experimental Analysis of Ultimatum Bargaining,” *Journal of Economic Behavior and Organization*, 3, 367–388.
- HANNAN, R. L., R. KRISHNAN, AND A. H. NEWMAN (2008): “The Effects of Disseminating Relative Performance Feedback in Tournament and Individual Performance Compensation Plans,” *The Accounting Review*, 83, 893–913.

- HAUGE, K. E., K. A. BREKKE, L.-O. J. O. JOHANSSON-STENMAN, AND H. SVEDSÄTER (2016): “Keeping Others in our Mind or in our Heart? Distribution Games Under Cognitive Load,” *Experimental Economics*, 19, 562–576.
- HERTWIG, R., G. BARRON, E. U. WEBER, AND I. EREV (2004): “Decisions from Experience and the Effect of Rare Events in Risky Choice,” *Psychological Science*, 15, 534–539.
- HINSON, J. M., T. L. JAMESON, AND P. WHITNEY (2002): “Somatic Markers, Working Memory, and Decision Making,” *Cognitive, Affective, & Behavioral Neuroscience*, 2, 341–353.
- HOFFMAN, E., K. MCCABE, K. SHACHAT, AND V. SMITH (1994): “Preferences, Property Rights, and Anonymity in Bargaining Games,” *Games and Economic Behavior*, 7, 346–380.
- HOLROYD, C. B. AND M. G. COLES (2002): “The Neural Basis of Human Error Processing: Reinforcement Learning, Dopamine, and the Error-Related Negativity,” *Psychological Review*, 109, 679–709.
- HOLT, C. A. AND S. K. LAURY (2002): “Risk Aversion and Incentive Effects,” *American Economic Review*, 92, 1644–1655.
- HUCK, S., H.-T. NORMANN, AND J. OECHSSLER (1999): “Learning in Cournot Oligopoly - An Experiment,” *Economic Journal*, 109, C80–C95.
- (2004): “Through Trial and Error to Collusion,” *International Economic Review*, 45, 205–224.
- HÜGELSCHÄFER, S. AND A. ACHTZIGER (2017): “Reinforcement, Rationality, and Intentions: How Robust is Automatic Reinforcement Learning in Economic Decision Making?” *Journal of Behavioral Decision Making*, 30, 913–932.
- HYDE, J. S., E. FENNEMA, AND S. J. LAMON (1990): “Gender Differences in Mathematics Performance: A Meta-Analysis,” *Psychological Bulletin*, 107, 139–155.

- JOHN, O. P., E. M. DONAHUE, AND R. L. KENTLE (1991): “The Big Five Inventory—Versions 4a and 54,” Berkeley: University of California, Berkeley, Institute of Personality and Social Research.
- JOHNSON, N. D. AND A. A. MISLIN (2011): “Trust Games: A Meta-Analysis,” *Journal of Economic Psychology*, 32, 865–889.
- KAHNEMAN, D. (2003): “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *American Economic Review*, 93, 1449–1475.
- KANDORI, M., G. J. MAILATH, AND R. ROB (1993): “Learning, Mutation, and Long Run Equilibria in Games,” *Econometrica*, 61, 29–56.
- KANDORI, M. AND R. ROB (1995): “Evolution of Equilibria in the Long Run: A General Theory and Applications,” *Journal of Economic Theory*, 65, 383–414.
- KARAKOSTAS, A. AND D. J. ZIZZO (2016): “Compliance and the Power of Authority,” *Journal of Economic Behavior and Organization*, 124, 67–80.
- KELLER, J., G. BOHNER, AND H.-P. ERB (2000): “Intuitive und heuristische Urteilsbildung — verschiedene Prozesse? Präsentation einer deutschen Fassung des ‘Rational–Experiential Inventory’ sowie neuer Selbstberichtsskalen zur Heuristiknutzung,” *Zeitschrift für Sozialpsychologie*, 31, 87–101.
- KHADJAVI, M. AND A. LANGE (2015): “Doing Good or Doing Harm: Experimental Evidence On Giving And Taking In Public Good Games,” *Experimental Economics*, 18, 432–441.
- KONOW, J. (2003): “Which is the Fairest One of All? A Positive Analysis of Justice Theories,” *Journal of Economic Literature*, 41, 1188–1239.
- KÖSZEGI, B. (2006): “Ego Utility, Overconfidence, and Task Choice,” *Journal of the European Economic Association*, 4, 673–707.
- KRAJBICH, I., B. BARTLING, T. HARE, AND E. FEHR (2015): “Rethinking Fast and Slow Based on a Critique of Reaction-Time Reverse Inference,” *Nature Communications*, 6:7455.

- KRUGLANSKI, A. W. (1989): *Lay Epistemics and Human Knowledge: Cognitive and Motivational Bases*, New York: Plenum.
- KRUGMAN, P. (2008): “Crisis of Confidence,” *The New York Times*, April 14th.
- KUHNEN, C. M. AND A. TYMULA (2012): “Feedback, Self-Esteem, and Performance in Organizations,” *Management Science*, 58, 94–113.
- KUZIEMKO, I., R. W. BUELL, T. REICH, AND M. I. NORTON (2014): ““Last-Place Aversion”: Evidence and Redistributive Implications,” *Quarterly Journal of Economics*, 129, 105–149.
- LANG, F. R., D. JOHN, O. LÜDTKE, J. SCHUPP, AND G. G. WAGNER (2011): “Short Assessment of the Big Five: Robust Across Survey Methods Except Telephone Interviewing,” *Behavior Research Methods*, 43, 548–567.
- LANG, J. W. AND S. FRIES (2006): “A Revised 10-Item Version of the Achievement Motives Scale: Psychometric Properties in German-Speaking Samples,” *European Journal of Psychological Assessment*, 22, 216–224.
- LASLIER, J.-F., R. TOPOL, AND B. WALLISER (2001): “A Behavioral Learning Process in Games,” *Games and Economic Behavior*, 37, 340–366.
- LAVIE, N. AND J. DE FOCKERT (2005): “The Role of Working Memory in Attentional Capture,” *Psychonomic Bulletin & Review*, 12, 669–674.
- LEVITT, S. D. AND J. A. LIST (2007): “What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World?” *Journal of Economic Perspectives*, 153–174.
- LINDSEY, L. L. M. (2005): “Anticipated Guilt as Behavioral Motivation: An Examination of Appeals to Help Unknown Others Through Bone Marrow Donation,” *Human Communication Research*, 31, 453–481.
- LIST, J. A. (2007): “On the Interpretation of Giving in Dictator Games,” *Journal of Political Economy*, 115, 482–493.

- LITMAN, J., T. HUTCHINS, AND R. RUSSON (2005): “Epistemic Curiosity, Feeling-of-Knowing, and Exploratory Behaviour,” *Cognition & Emotion*, 19, 559–582.
- LITMAN, J. A. (2008): “Interest and Deprivation Factors of Epistemic Curiosity,” *Personality and Individual Differences*, 44, 1585–1595.
- LITMAN, J. A. AND P. MUSSEL (2013): “Validity of the Interest-and Deprivation-Type Epistemic Curiosity Model in Germany,” *Journal of Individual Differences*, 34, 59–68.
- LOEWENSTEIN, G. (1994): “The Psychology of Curiosity: A Review and Reinterpretation,” *Psychological Bulletin*, 116, 75–98.
- MALINOWSKI, C. I. AND C. P. SMITH (1985): “Moral Reasoning and Moral Conduct: An Investigation Prompted by Kohlberg’s Theory,” *Journal of Personality and Social Psychology*, 49, 1016–1027.
- MAZAR, N., O. AMIR, AND D. ARIELY (2008): “The Dishonesty of Honest People: A Theory of Self-Concept Maintenance,” *Journal of Marketing Research*, 45, 633–644.
- MAZAR, N. AND C.-B. ZHONG (2010): “Do Green Products Make Us Better People?” *Psychological Science*, 21, 494–498.
- MCCLELLAND, D. C., R. KOESTNER, AND J. WEINBERGER (1989): “How Do Self-Attributed and Implicit Motives Differ?” *Psychological Review*, 96, 690–702.
- MCKELVEY, R. D. AND T. R. PALFREY (1995): “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, 10, 6–38.
- MCLEISH, K. N. AND R. J. OXOBY (2011): “Social Interactions and the Salience of Social Identity,” *Journal of Economic Psychology*, 32, 172–178.
- MILINSKI, M. AND C. WEDEKIND (1998): “Working Memory Constrains Human Cooperation in the Prisoner’s Dilemma,” *Proceedings of the National Academy of Sciences*, 95, 13755–13758.

- MOFFATT, P. G. (2005): “Stochastic Choice and the Allocation of Cognitive Effort,” *Experimental Economics*, 8, 369–388.
- MONIN, B. AND D. T. MILLER (2001): “Moral Credentials and the Expression of Prejudice,” *Journal of Personality and Social Psychology*, 81, 33–43.
- MONSELL, S. (2003): “Task Switching,” *Trends in Cognitive Sciences*, 7, 134–140.
- MOORE, D. A. AND W. M. KLEIN (2008): “Use of Absolute and Comparative Performance Feedback in Absolute and Comparative Judgments and Decisions,” *Organizational Behavior and Human Decision Processes*, 107, 60–74.
- MOSTELLER, F. AND P. NOGEE (1951): “An Experimental Measurement of Utility,” *Journal of Political Economy*, 59, 371–404.
- MOYER, R. S. AND T. K. LANDAUER (1967): “Time Required for Judgments of Numerical Inequality,” *Nature*, 215, 1519–1520.
- MYERSON, R. B. (2012): “A Model of Moral-Hazard Credit Cycles,” *Journal of Political Economy*, 120, 847–878.
- MYRSETH, K. O. R. AND C. E. WOLLBRANT (2016): “Commentary: Fairness is Intuitive,” *Frontiers in Psychology*, 7, 654.
- NIEDERLE, M. AND L. VESTERLUND (2007): “Do Women Shy Away from Competition? Do Men Compete Too Much?” *Quarterly Journal of Economics*, 122, 1067–1101.
- NORMAN, D. A. AND T. SHALLICE (1986): “Attention to Action: Willed and Automatic Control of Behavior,” in *Consciousness and Self-Regulation; Advances in Research and Theory*, Vol. 4, ed. by R. J. Davidson, G. E. Schwartz, and D. Shapiro, NY: Plenum Press, 1–18.

- OFFERMAN, T., J. POTTERS, AND J. SONNEMANS (2002): “Imitation and Belief Learning in an Oligopoly Experiment,” *Review of Economic Studies*, 69, 973–997.
- OOSTERBEEK, H., R. SLOOF, AND G. VAN DE KUILEN (2004): “Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis,” *Experimental Economics*, 7, 171–188.
- OWENS, L. A. (2012): “Confidence in Banks, Financial Institutions and Wall Street, 1971–2011,” *Public Opinion Quarterly*, 76, 142–162.
- PIOVESAN, M. AND E. WENGSTRÖM (2009): “Fast or Fair? A Study of Response Times,” *Economics Letters*, 105, 193–196.
- POLIVY, J. AND C. P. HERMAN (1985): “Dieting and Binging: A Causal Analysis,” *American Psychologist*, 40, 193–201.
- RATCLIFF, R. (1978): “A Theory of Memory Retrieval,” *Psychological Review*, 85, 59–108.
- REUBEN, E. AND F. VAN WINDEN (2010): “Fairness Perceptions and Prosocial Emotions in the Power to Take,” *Journal of Economic Psychology*, 31, 908–922.
- ROBERTS, T.-A. AND S. NOLEN-HOEKSEMA (1989): “Sex Differences in Reactions to Evaluative Feedback,” *Sex Roles*, 21, 725–747.
- ROCH, S. G., J. A. LANE, C. D. SAMUELSON, S. T. ALLISON, AND J. L. DENT (2000): “Cognitive Load and the Equality Heuristic: A Two-stage Model of Resource Overconsumption in Small Groups,” *Organizational Behavior and Human Decision Processes*, 83, 185–212.
- ROGERS, R. D. AND S. MONSELL (1995): “The Costs of a Predictable Switch Between Simple Cognitive Tasks,” *Journal of Experimental Psychology: General*, 124, 207–231.



- ROTH, A. E., V. PRASNIKAR, M. OKUNO-FUJIWARA, AND S. ZAMIR (1991): “Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study,” *American Economic Review*, 81, 1068–1095.
- RUBINSTEIN, A. (2007): “Instinctive and Cognitive Reasoning: A Study of Response Times,” *Economic Journal*, 117, 1243–1259.
- SAMSON, K. AND P. KOSTYSZYN (2015): “Effects of Cognitive Load on Trusting Behavior—An Experiment Using the Trust Game,” *PloS one*, 10, e0127680.
- SCHAFFER, M. E. (1989): “Are Profit-Maximisers the Best Survivors?” *Journal of Economic Behavior and Organization*, 12, 29–45.
- SCHMIDT, H. (2003): “Das Gesetz des Dschungels [The Law of the Jungle],” *Die Zeit*, December 4th.
- SCHOENBERG, E. J. AND E. HARUVY (2012): “Relative Performance Information in Asset Markets: An Experimental Approach,” *Journal of Economic Psychology*, 33, 1143–1155.
- SCHÖNBERG, T., N. D. DAW, D. JOEL, AND J. P. O’DOHERTY (2007): “Reinforcement Learning Signals in the Human Striatum Distinguish Learners from Nonlearners During Reward-Based Decision Making,” *The Journal of Neuroscience*, 27, 12860–12867.
- SCHULTZ, W. (1998): “Predictive Reward Signal of Dopamine Neurons,” *Journal of Neurophysiology*, 80, 1–27.
- SCHULZ, J. F., U. FISCHBACHER, C. THÖNI, AND V. UTIKAL (2014): “Affect and Fairness: Dictator Games under Cognitive Load,” *Journal of Economic Psychology*, 41, 77–87.
- SCHWARTZ, B., A. WARD, J. MONTEROSSO, S. LYUBOMIRSKY, K. WHITE, AND D. R. LEHMAN (2002): “Maximizing Versus Satisficing: Happiness is a Matter of Choice,” *Journal of Personality and Social Psychology*, 83, 1178–1197.

- SEDIKIDES, C. AND M. J. STRUBE (1995): “The Multiply Motivated Self,” *Personality and Social Psychology Bulletin*, 21, 1330–1335.
- (1997): “Self-Evaluation: To Thine Own Self Be Good, To Thine Own Self Be Sure, To Thine Own Self Be True, and To Thine Own Self Be Better,” *Advances in Experimental Social Psychology*, 29, 209–269.
- SELTEN, R. (1967): “Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments,” in *Beiträge zur experimentellen Wirtschaftsforschung*, ed. by H. Sauermann, J. C. B. Mohr (Paul Siebeck), Tübingen, 136–168.
- SHIV, B. AND A. FEDORIKHIN (1999): “Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making,” *Journal of Consumer Research*, 26, 278–292.
- SLONIM, R. AND A. E. ROTH (1998): “Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic,” *Econometrica*, 63, 569–596.
- SMITH, A. (1759): *The Theory of Moral Sentiments*, Indianapolis: Reprinted 1981 by Liberty Fund (D. D. Raphael and A. L. Macfie, editors).
- SPENCER, S. J., C. M. STEELE, AND D. M. QUINN (1999): “Stereotype Threat and Women’s Math Performance,” *Journal of Experimental Social Psychology*, 35, 4–28.
- STOOP, J., C. N. NOUSSAIR, AND D. VAN SOEST (2012): “From the Lab to the Field: Cooperation among Fishermen,” *Journal of Political Economy*, 120, 1027–1056.
- STRACK, F. AND R. DEUTSCH (2004): “Reflective and Impulsive Determinants of Social Behavior,” *Personality and Social Psychology Review*, 8, 220–247.
- STROOP, J. R. (1935): “Studies of Interference in Serial Verbal Reactions,” *Journal of Experimental Psychology*, 35, 643–662.

- SUTTON, R. S. AND A. G. BARTO (1998): *Reinforcement Learning: An Introduction*, Cambridge, MA: MIT Press.
- THORNDIKE, E. L. (1911): *Animal Intelligence: Experimental Studies*, NY: MacMillan (see also Hafner Publishing Co., 1970).
- TROPE, Y. (1980): “Self-Assessment, Self-Enhancement, and Task Preference,” *Journal of Experimental Social Psychology*, 16, 116–129.
- VAN DE VEN, N., M. ZEELENBERG, AND E. VAN DIJK (2005): “Buying and Selling Exchange Goods: Outcome Information, Curiosity and the Endowment Effect,” *Journal of Economic Psychology*, 26, 459–468.
- VEGA-REDONDO, F. (1997): “The Evolution of Walrasian Behavior,” *Econometrica*, 65, 375–384.
- WALKOWITZ, G., H. HENNIG-SCHMIDT, AND C. OBERHAMMER (2009): “Experimenting over a Long Distance—A Method to Facilitate Intercultural Experiments and its Application to a Trust Game,” Working Paper, Department of Economics, University of Bonn.
- WEBER, E. U. AND E. J. JOHNSON (2009): “Mindful Judgment and Decision Making,” *Annual Review of Psychology*, 60, 53–85.
- WEIBULL, J. (1995): *Evolutionary Game Theory*, Cambridge, Massachusetts: The MIT Press.
- WHITNEY, P., C. A. RINEHART, AND J. M. HINSON (2008): “Framing Effects Under Cognitive Load: The Role of Working Memory in Risky Decisions,” *Psychonomic Bulletin & Review*, 15, 1179–1184.
- WILCOX, N. T. (1993): “Lottery Choice: Incentives, Complexity, and Decision Time,” *Economic Journal*, 103, 1397–1417.
- WIMMER, G. E., N. D. DAW, AND D. SHOHAMY (2012): “Generalization of Value in Reinforcement Learning by Humans,” *European Journal of Neuroscience*, 35, 1092–1104.

- ZEMACK-RUGAR, Y., J. R. BETTMAN, AND G. J. FITZSIMONS (2007):  
“The Effects of Nonconsciously Priming Emotion Concepts on Behavior,”  
*Journal of Personality and Social Psychology*, 93, 927–939.
- ZHONG, C.-B., G. KU, R. B. LOUNT, AND J. K. MURNIGHAN (2010):  
“Compensatory Ethics,” *Journal of Business Ethics*, 92, 323–339.
- ZIZZO, D. J. (2004): “Inequality and Procedural Fairness in a Money Burning and Stealing Experiment,” *Research on Economic Inequality*, 11, 215–247.
- ZIZZO, D. J. AND A. J. OSWALD (2001): “Are People Willing to Pay to Reduce Others’ Incomes?” *Annales d’Economie et de Statistique*, 39–65.