UNIVERSITY OF ESSEX

DOCTORAL THESIS

# Essays in Behavioural and Experimental Economics

*Author:*
Ludovica ORLANDI

*Supervisor:*
Professor Friederike MENGEL

*A thesis submitted in fulfillment of the requirements for the degree of Doctor of Philosophy*

*in the*

Department of Economics
University of Essex

February 2021

# Declaration of Authorship

I, Ludovica ORLANDI, declare that this thesis titled, "Essays in Behavioural and Experimental Economics" and the work presented in it are my own. I confirm that:

- No part of this thesis has been submitted for another degree.
- Chapter 2 is co-authored with Professor Friederike Mengel and Professor Simon Weidenholzer.
- Chapter 1 and 3 are exclusively mine.

UNIVERSITY OF ESSEX

# *Abstract*

Department of Economics

Doctor of Philosophy

**Essays in Behavioural and Experimental Economics**

by Ludovica ORLANDI

This thesis contains three essays in Behavioural and Experimental Economics.

In Chapter 1, I investigate what effect growing up as an only-child has on social and risk preferences, trust, and behaviour in strategic interaction. To do that, I use data from a laboratory experiment in which university students participated in three different games (a public goods game, a response game, and an investment game) and in a task for eliciting risk preferences. I find that only children are less likely to punish for misbehaviour and more likely to choose social welfare maximising actions in response games. Only-children are, however, significantly less likely to cooperate in one-shot public goods games and are less trusting than subjects with siblings.

In Chapter 2, Professor Friederike Mengel, Professor Simon Weidenholzer and I study decisions in infinitely repeated games. Using data from previous prisoner's dilemma games and data from our own experiment, we assess subjects' cooperation rates and show that realized match length of early matches has a significant impact on the cooperation rates of subsequent matches. We furthermore demonstrate that the effect is in line with a class of learning models displaying the "power law of practice".

In Chapter 3, I use data from a dictator game where information about receiver's deservingness is endogenously determined by the receiver's decision to disclose information. I therefore explore whether receivers voluntarily provide information about their deservingness and the effect of endogenous information on dictators' decision. I find that receivers having deservingness levels far from the median level of deservingness in a session are significantly less likely to disclose information about their deservingness. Furthermore, the choice to disclose information is related to the receivers' beliefs about dictators' actions. I also show that dictators give more when information is endogenous than in the case in which it is exogenously provided.

# *Acknowledgements*

I would first like to express my sincere gratitude to my supervisor, Professor Friederike Mengel, for her invaluable advice, patience and motivation. Her guidance, continuous support and encouragement has been invaluable throughout my Ph.D study. Without her the research contained in this thesis would not have been possible.

I would also like to express my deep gratitude to Professor Simon Weidenholzer for his collaboration in Chapter 2, his support, advice and recommendations on my research. It has been a great experience for me to work with him. I am also very grateful to the chair of my supervisory panel, Dr. Lingqing Jiang. She has been very supportive all throughout the last years of my Ph.D study and has given me useful comments on my work.

I have been fortunate to share my Ph.D experience with exceptional friends. I would like to thank Ludovica Giua, Claudio Deiana, Elisabetta Leni, Enes Duysak, Annalivia Polselli and Andrea Albertazzi for all the stimulating discussions and for making my life in Colchester memorable. A very special thanks goes to Patryk Bronka for his continuous support in this stressful final year.

I would also like to thank the EssexLab for the support to my experiments and all participants that took part in my studies and enabled this research to be possible.

Most of all I would like to thank my parents for their unconditional support, wise counsel and sympathetic ear. They have always been there for me.

Lastly, I am grateful to the Economics Department of the University of Essex and the Economic and Social Research Council (ESRC) for the financial support during my Ph.D study.

The research here included has been funded by the British Academy Leverhulme Small Research Grant, ESSEXLab Seedcorn funds and the Economic and Social Research Council.

# Contents

**Declaration of Authorship** iii

**Abstract** v

**Acknowledgements** vii

**1 The Effect of Growing up as an Only-Child on Preferences Structure and Strategic Behaviours** 1
1.1 Introduction . . . 4
1.2 Subjects pool and procedure . . . 10
1.3 Experiment Design . . . 14
1.3.1 Public Goods Game . . . 14
1.3.2 Response Game . . . 17
1.3.3 Investment Game . . . 18
1.3.4 Risk Task . . . 20
1.4 Results . . . 21
1.4.1 Public Goods Game . . . 21
1.4.2 Response Game . . . 27
1.4.3 Investment Game . . . 32
1.4.4 Risk Task . . . 35
1.5 Discussion and Conclusions . . . 37

**2 History matters: Match length realizations and cooperation in indefinitely repeated games** 41
2.1 Introduction . . . 38
2.2 The Effect of Match Length Realizations . . . 41
2.2.1 The prisoner's dilemma . . . 41
2.2.2 Empirical Analysis . . . 44
2.2.3 Learning . . . 48
2.3 Case Studies . . . 52
2.3.1 Cooperation in Continuous Time . . . 53
2.3.2 Finite and Indefinitely Repeated Linear Public Good Games . . . 57
2.3.3 Strategy Revision Opportunities . . . 59
2.4 Discussion and Conclusions . . . 61

**3 Do subjects disclose outcomes strategically? An experiment on endogenous information in a dictator game.** 67
3.1 Introduction . . . 70
3.2 Experiment Design . . . 78
3.3 Subjects pool and procedure . . . 85
3.4 Results . . . 87
3.4.1 Distribution of outcomes . . . 87
3.4.2 Participants' decisions . . . 89

Receivers' disclosure decision . . . 89
Dictators' transfer decision . . . 96
3.4.3 Beliefs . . . 99
Beliefs - Manipulation check . . . 99
Receiver's beliefs . . . 101
Dictator's beliefs . . . 103
Receivers' beliefs and dictators' decision . . . 108
3.4.4 The effect of endogenous information on dictators' transfer . . . 110
3.5 Conclusion . . . 112

**A Chapter 1 Appendix A: Experimental Instructions** **115**

**B Chapter 1 Appendix B: Additional tables** **127**

**C Chapter 2 Appendix C: Additional details learning models** **131**
C.1 Comparative Statics of Choice Probabilities . . . 131
C.2 Further details on simulations . . . 133

**D Chapter 2 Appendix D: Additional Information Experiments** **135**

**E Chapter 2 Appendix E: Additional Tables** **137**
E.1 Additional Tables for Section 2.2.2 . . . 137
E.2 Additional Tables for Section 2.2.3 . . . 145
E.3 Additional Tables for Section 2.3.1 . . . 148
E.4 Additional Tables for Section 2.3.2 . . . 149
E.5 Additional Tables for Section 2.3.3 . . . 150
E.6 Additional Tables for Section 2.4 . . . 151

**F Chapter 2 Appendix F: Additional Figures** **153**

**G Chapter 3 Appendix G** **157**
G.1 Experimental Instructions . . . 157

**H Chapter 3 Appendix H** **169**
H.1 Hypothetical dictator game (Stage 5) . . . 169

**I Chapter 3 Appendix I** **175**

**Bibliography** **179**

# List of Figures

1.1 Mean beliefs, contributions and predicted contribution over time for subjects with siblings and only-children . . . 24
1.2 Average tokens returned by tokens received and only-child status . . . 35
1.3 Frequency of gamble choices by only-child status . . . 37

2.1 Left: Prisoner's dilemma (PD) game with payoff parameters $\mathtt{T} > \mathtt{R} > \mathtt{P} > \mathtt{S}$ and $\mathtt{T} + \mathtt{S} < 2\mathtt{R}$. Middle: Normalized game where joint defection payoff $\mathtt{P}$ is subtracted from each cell and all payoffs are divided by $\mathtt{R} - \mathtt{P}$ (difference between mutual cooperation and defection payoffs). Right: Payoffs in the game induced by Grim-Trigger and Always Defect. GT played against GT yields a payoff of 1 in all $\mathbb{E}[T]$ stages. AD (GT) played against GT (AD) yields once a payoff of $1 + g$ ($-l$) and zero in the remaining $\mathbb{E}[T] - 1$ stages. . . . 42
2.2 Distribution of the difference between theoretical median match length and realized match length overall (left panel) and separately for $\delta = 0.5, 0.75$ and $\delta = 0.9$. 44
2.3 Left Panel: Effect Size for $\Delta^{1st}_{above}$ in simulated data for the three models (bars) compared to effect size in Dal Bó and Fréchette (2018) data (dashed line) in two samples: below median values of SizeBAD (left three bars and thick dashed line) and above median values of SizeBAD (right three bars and thick dashed line). Right Panel: Relative Effect size of $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$ in simulated data for the three models (bars) in two samples: below median values of SizeBAD (left three bars) and above median values of SizeBAD (right three bars). The dashed line indicates $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$. . . . 51
2.4 Kernel density estimates of the difference between theoretical median match length and realized match length overall (left panel) and for the 1st third of matches. . . . 53
2.5 Different effect sizes obtained for the treatment differences between deterministic and stochastic treatments in average cooperation rates (left panel) and initial cooperation rates (right panel) depending on match length realizations. Bigoni et al: original effect size in Bigoni et al. (2015); Replication: replication treatment; Inverse: inverse treatment; MatchStoch: treatment with match length realization drawn at match level; MinMatch: only individuals from MatchStoch who had the smallest number of matches 1-7 above median (specifically 1) in this treatment; MaxMatch: only individuals from MatchStoch who had the largest number of matches 1-7 above median (specifically 6) in this treatment; . . . 56

3.1 Distribution of the number of words encoded by participants' role . . . 89
3.2 Relative rank distribution by choice to reveal the task outcome . . . 91
3.3 Average stated beliefs by task outcome and participant's role . . . 100
3.4 Average stated beliefs by task outcome and choice to reveal . . . 101
3.5 Dictators and receivers' average stated beliefs by task outcome . . . 104
3.6 Dictators' guess of receivers' disclosure choice across different outcome levels vs. actual disclosure rates . . . 107

A.1 Screen shot of the Public Goods Game Conditional contribution decision . . . . . 124
A.2 Screen shot of a dictator game from the Response Game . . . . . 125
A.3 Screen shot of a response game game from the Response Game . . . . . 125
A.4 Screen shot of Player A's choice from the Investment Game . . . . . 126
A.5 Screen shot of the Risk Task . . . . . 126

C.1 Match Length Realization in our simulations. . . . . 134

F.1 Meta Data used in Section 2.2.2: Kernel density estimates of the difference between theoretical median match length and realized match length overall (left panel) and separately for $\delta = 0.5, 0.75$ and $\delta = 0.9$ (right panel). . . . . 153
F.2 Section 2.2.2 The Effect of Experience (Number of Matches Played in a Session) on Estimated Coefficient $\Delta^{1st}_{above}$. Black diamonds show the estimated coefficient $\Delta^{1st}_{above}$ on cooperation in the last third of matches depending on the total number of matches played in the session. The black line shows fitted values using a square polynomial. Squares and light gray dashed line show the effect of the first half of matches on the second half and triangles and dark gray dashed line the effect of the first quarter on the last quarter of matches. . . . . 154
F.3 Section 2.3.1: Illustration of how sequence of "inverse" match lengths is generated for the two sessions. . . . . 154
F.4 Section 2.3.3: Illustration of treatment effect for different permutations of two matching groups. Histogram and estimated kernel density. Left Panel: treatments with revision opportunities. Right panel: treatments without revision opportunities. . . . . 155

G.1 Screen shot of the Encryption Real-effort task . . . . . 165
G.2 Screen shot of the receiver's decision to reveal the task outcome . . . . . 165
G.3 Screen shot of the dictator's decision to give . . . . . 166
G.4 Screen shot of the receiver's guess about the number of tokens received in the four alternative task-outcome scenarios . . . . . 166
G.5 Screen shot of the dictator's guess about the number of tokens expected by the receiver in the four alternative task-outcome scenarios . . . . . 167
G.6 Screen shot of the guess of revealing rates across outcome categories . . . . . 167

I.1 Rank index distribution by choice to reveal the task outcome . . . . . 175

# List of Tables


1.1 Basic characteristics of the sample . . . 11
1.2 Gambles characteristics . . . 20
1.3 Consistency of behaviours with different contribution preferences . . . 22
1.4 OLS estimates of the effect of only-child on unconditional contribution . . . 23
1.5 Beliefs about other group members' contribution and Contribution determinants 26
1.6 Proportion of decisions consistent with different models . . . 28
1.7 Proportion of Social Welfare maximizing decisions and the effect of being Only-child . . . 29
1.8 Logit regression determinant of reciprocity . . . 31
1.9 OLS estimates of the effect of only-child on amount sent to player B . . . 33
1.10 Gamble choice by only-child . . . 36

2.1 Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level. Observations stem from 141 sessions spread across 15 papers. . . . 45
2.2 Summary Statistics of the different treatments conducted to replicate Bigoni et al. (2015) . Number of Sessions (S) and observations (N) in the different conditions. Average Match Length (Avg ML), average match length in the 1st third of the experiment (Avg ML 1st third), median match length and median match length in the 1st third, average cooperation rate (avg coop) and average initial cooperation rate (avg coop initial). Stars indicate statistical significance (*** 1%, ** 5%, * 10 %) of the difference to the deterministic case in random effects OLS regression with standard errors clustered at session level (see Appendix Tables E.12 and E.13). . . . 54
2.3 Average cooperation rates across all rounds in the finite sessions (column (1)) and across all sessions with a probabilistic ending (column (2)) as in Table 3 in Lugovskyy et al. (2017). We further split the sessions with probabilistic ending in those with initially short matches (S1) and those with initially long matches (S23). Below each cooperation rate we show how the finite setting compares to the rate in question (as in Lugovskyy et al. (2017)). The last column shows the p-value when comparing initially short and long sequences. Following Lugovskyy et al. (2017) standard errors are clustered at the participant level in all regressions. 58
2.4 Efficiency measure from Embrey, Mengel, and Peeters (2019). Treatment Effect is the difference between substitutes and complements. The two rightmost columns split out the groups with short initial matches from the rest. As in Embrey, Mengel, and Peeters (2019)'s main analysis matches 7-10 are considered. Stars are from t-tests with standard errors clustered at group level.$^{***}p < 0.01$, $^{**}, p < 0.05$, $^{*}p < 0.1$. . . . 60

3.1 Share of subjects who decide to reveal their outcome from the task . . . . . 89
3.2 Logit regression determinants of choice to reveal task outcome . . . . . 94
3.3 Dictators' decision to give by whether receiver's task outcome is revealed . . . . 96
3.4 OLS estimates of the effect of information on giving decisions . . . . . 97
3.5 Average stated beliefs by choice to reveal . . . . . 102
3.6 Average receivers' stated beliefs and dictators' transfer decisions by information scenario and dictators' type . . . . . 110
3.7 Transfer decisions by information scenario and origin of information . . . . . 111

B.1 Response Game - Games description . . . . . 128
B.2 Logit regression determinant of reciprocity . . . . . 130

C.1 Payoff parameters used for simulations. . . . . 134

D.1 Basic Sample Characteristics of Lab Experimental Sample used in Section 2.3.1. Mean Age, fraction of participants born in UK, fraction female, mean risk attitude (0, most risk averse, 10 least risk averse), fraction displaying high trust, fraction answering Logic 1 question correctly and fraction answering Logic 2 question correctly. . . . . 135

E.1 Initial (first stage) cooperation rate in remaining matches depending on whether at least $\frac{2}{3}$ of first X-th of matches where above theoretical median length. . . . . 138
E.2 Initial (first stage) cooperation rate in remaining matches depending on whether at least $\frac{2}{3}$ of first X-th, second X-th, third X-th,... of matches where above theoretical median length. . . . . 139
E.3 Placebo test. Initial (first stage) cooperation rate in the *1st* third of matches explained by dummy $\Delta^{3rd}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 3rd third of the experiment were longer than the theoretical median match length. . 139
E.4 Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level. Observations stem from 141 sessions spread across 15 papers. With paper fixed effects. . . . . 140
E.5 Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by share of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by share of matches longer than theoretical median length in 1st and 2nd third of matches. Standard errors clustered at session level. Observations stem from 141 sessions spread across 15 papers. . . . . 141
E.6 Reduced Threshold. Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{1}{2}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. . . . . 142
E.7 Reduced Threshold. Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{3}{4}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. . . . . 143

E.8 Focus on Short Matches. Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{below}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were shorter than the theoretical median match length. . . . 144
E.9 Simulated Reinforcement Learning Model. Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level. . . . 145
E.10 Simulated Reinforcement Learning Model with Counterfactuals. Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level. 146
E.11 Simulated Fictitious Play Learning Model. Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level. . . . 147
E.12 Random Effects OLS regression of average cooperation rate on treatment dummies and covariates (age, gender, birthplace, nationality). Standard errors clustered at session level. Columns (3) and (4) only use data from match stochastic and deterministic sessions. . . . 148
E.13 Random Effects OLS regression of initial cooperation rate on treatment dummies and covariates (age, gender, birthplace, nationality). Standard errors clustered at session level. Columns (3) and (4) only use data from match stochastic and deterministic sessions. . . . 148
E.14 Initial cooperation rates across all rounds in the finite sessions (column (1)) and across all sessions with a probabilistic ending (column (2)) as in Table 3 in Lugovskyy et al. (2017). We further split the sessions with probabilistic ending in those with initially short matches (S1) and those with initially long matches (S23). Below each cooperation rate we show how the finite setting compares to the rate in question (as in Lugovskyy et al. (2017)). The last column shows the p-value when comparing initially short and long sequences. Following Lugovskyy et al. (2017) standard errors are clustered at the participant level in all regressions. . . . 149
E.15 Efficiency measure from Table 2 (all matches) in Embrey, Mengel, and Peeters (2019). Treatment Effect is the difference between substitutes and complements. The two rightmost columns split out the groups with short initial matches from the rest. . . . 150
E.16 First round cooperation in the 2nd and 3rd third of matches regressed on mean and median match length in the 1st third. We selected the two discounted factors with most papers in the meta-study (8 papers for $\delta = 0.75$ and 4 papers for $\delta = 0.5$). 151

G.1 Table 1 . . . 159

H.1 Kerschbamer (2015)'s hypothetical dictator game . . . 172
H.2 Determination of (X, Y)- Score (Kerschbamer, 2015) . . . 172
H.3 Sample Xscore and Yscore distribution . . . 173

I.1 OLS estimates of the effect of information on giving decisions - no inconsistent subjects . . . . . 176
I.2 OLS estimates of the effect of information on giving decisions - alternative thresholds . . . . . 177
I.3 OLS estimates of treatment effect on transfer decisions . . . . . 177

# Chapter 1

# The Effect of Growing up as an Only-Child on Preferences Structure and Strategic Behaviours

# *Abstract*

In the psychological and economic literature there is considerable evidence supporting the fact that being an only-child has an impact on social preferences. Results are however, fragmented and sometimes contrasting. Our research comprehensively studies the effect that growing up as an only-child has on social preferences, risk preferences, trust and behaviour in strategic interaction. We set up an experiment where 276 participants sequentially participate in four different simple games: a public goods game, a response games, an investment game and a task for eliciting risk preferences. We find evidence of consistent difference in social preferences of only-children and non only-children, in particular with negative reciprocity and social welfare maximizing actions. Only-children are less likely to punish for misbehaviour and more likely to choose social welfare maximizing actions in response games. Furthermore, only-children cooperate significantly less in one shot public goods games. Finally, in line with earlier results we find that only-children are less trusting than subjects with siblings. [1]

[1]I particularly would like to thank my supervisor Professor Friederike Mengel, for having provided invaluable guidance, support and help. I also thank Professor Matthias Sutter and his Research Group for helpful comments and suggestions. I am grateful to seminars participants at the University of Essex, the Max Planck Institute for Research on Collective Goods, FAIR at NHH Norwegian School of Economics, as well as participants at the 2018 SABE/ IAREP conference at Middlesex University and the 2018 SABE/ IAREP Elsevier Best PhD Student Paper Competition committee. All errors are mine. I gratefully acknowledge financial support from the British Academy Leverhulme Small Research Grant(grant number: SG162637) and the ESSEXLab Seedcorn Grant.

## 1.1 Introduction

In psychological literature there is considerable evidence concerning the impact of being an only-child on preferences structure. Economists, however, have devoted comparatively little attention to differences in preferences and have mostly concentrated on differences in economic outcomes between only-children and subjects with siblings.

Understanding some of the determinants of preferences and of their process of evolution could provide a relevant contribution to the economic literature. Preferences structure has in fact been demonstrated to be able to affect social and economic outcomes, and to play an important role in the formation of social capital. Works by Carpenter and Seki (2011), and Barr and Serneels (2009) show the existence of a correlation between social preferences and subject productivity at work and wages. Similarly, time preferences have been linked with diverse economic outcomes. Higher discount rates seem to negatively affect school performance, health, labour supply and lifetime income (Golsteyn, Grönqvist, and Lindahl, 2014; Sutter et al., 2013). Finally, the risk component of preferences has been found to be a good predictor of investment decisions, health and occupational choices. Subjects who are more risk seeking are more likely to be self-employed, to smoke and to invest in stocks (Dohmen et al., 2011). Therefore, understanding determinants of preferences structure can help to interpret and predict society's economic outcomes and to develop more efficient policies.

This research contributes to the understanding of the sources of preferences structure by experimentally documenting a relation between the structure of subjects' preferences and the condition of growing up as an only-child.

Previous research in Economics has already provided evidences of differences in economic outcomes between subjects with siblings and only-children. For example, Björklund and Jäntti (1994) use representative data from Finland-Sweden-US to study the relationship between the number of siblings, and both earning and school attainment. They find that having one or two siblings brings about better outcomes both

in earnings and education. Similar results are presented by Hauser and Kuo (1998) and Black, Devereux, and Salvanes (2005). The latter use a dataset representative of the total Norwegian population and report that being an only-child is associated with an average 0.25 year less schooling than subjects having up to three siblings. Conversely, Juhn, Rubinstein, and Zuppann (2015) exploit data from the US Survey of Youth to study how family size affects long term outcomes and obtain results suggesting that having siblings decreases both educational attainment and labour market participation. Similar results, in terms of school progress and college enrolment, are also provided by Conley and Glauber (2006) and Rosenzweig and Zhang (2009). Other examples of differences in outcomes are provided by the study of Knoef and Kooreman (2011). They analyse datasets from 12 European countries to investigate intra-family links and find that only-children are significantly more involved in parental informal care and family activities. Finally, Becker (1981) reports a positive correlation between the number of a person's siblings and their future number of children and a paper of Lampi and Nordblom (2010) shows a much higher concern of single-children about their relative earnings and job position. Even though results are mixed, they seem to suggest that a relevant difference between economic outcomes of only-children and subjects with siblings exists.

By contrast, the literature on the relationship between the only-child status and preference structure is scarce. Economic literature has mainly focused on the more general understanding of the impact of family and social environment on preferences development. Several studies have demonstrated the important role of early environmental conditions in shaping personality and preferences structure. Cunha and Heckman (2007) and Heckman and Mosso (2014), for example, provide evidence of the strong influence of initial family conditions on an individual's trust, cooperativeness, altruism and risk-aversion. Dohmen et al. (2012) study the mechanism of attitude transmission and observe that only-children are generally less risks-seeking than children with siblings, probably as a consequence of parental influence or smaller family networks. Cappelen et al. (2019) provide evidence of early childhood education having an impact on children's social preferences and consider siblings as a potential

moderator of this effect. They show that the more egalitarian effect of the Preschool program is significantly higher across only-children and they argue that this could be due to only-children's lack of experience in conflict resolution at home. Deckers et al. (2015) come closer to studying the effect of siblings' number on preference structure. They use parent's years of education and household income as measures of socioeconomic status, and show that children from families with higher socioeconomic status are more patient, less risk seeking and more altruistic. They furthermore document how the number of siblings represents one of the dimensions that systematically differ by socioeconomic status of a family and constitute a possible pathway through which a child's preference structure is affected. To our knowledge, however, only the study of Cameron et al. (2013) has specifically addressed the question of the effect of growing up as an only-child on preferences development. Their study assesses the behavioural effects of growing up as an only-child during the One-Child Policy by the analysis of behaviours in small experiments run on a Chinese population sample. They exploit the One-Child Policy as an exogenous variation to identify the causal impact of being an only-child and find that single children of the new Chinese society are significantly less trusting, less trustworthy and more risk-averse. Cameron et al. (2013)'s results are, though, based on a very specific national sample and therefore may not be representative of a more general population. Differences in results already observed when studying the effect of siblings' number on economic outcomes from different national samples seems in fact to suggest some likelihood of obtaining different estimations from different samples (Juhn, Rubinstein, and Zuppann, 2015; Hauser and Kuo, 1998; Black, Devereux, and Salvanes, 2005). One way our research contributes, hence, is by providing a different sample consisting of students from a UK University [2].

For the purpose of our study, it is worth underling that the above cited literature

[2] Additional evidence about different preference structures between only-children and subjects with siblings is also provided by studies not directly aimed at understanding the determinants of individuals' preferences. An example of this is the study of Glaeser et al. (2000). This study was developed to test various kinds of measure of trust, but it turns out also to provide significant evidence of a lower level of trustworthiness amongst only-child subjects. Finally, Angerer et al. (2015) and Fehr, Bernhard, and Rockenbach (2008) developed respectively a study on the relation between altruism, risk attitudes and time preferences of children and on the process of evolution of other regarding preferences with age, and found a common positive effect of being an only child on altruistic behaviours.

concerning the relationship between socioeconomic background and preference development is mainly focused on subjects in their childhood. The comparability of different results requires therefore the assumption that the observed disparities also persist in adulthood. This assumption however does not appear to be too strong as indeed the literature on preference development seems to suggest a consistency or augment of differences with age (Cunha and Heckman, 2007). The predominant characteristic of previous studies of being focused on the study of differences amongst subjects in their early years, brings us to underline another contribution of our research. We develop an analysis of the impact of growing up as an only-child amongst individuals in adulthood which not only operate as a test for previous findings in childhood, but also allows us to collect insights on the persistence of some characteristics of preferences in society. Furthermore, by collecting evidence of difference between single-children and subjects with siblings' preferences in adulthood we are able to better understand what can be the actual impact of this characteristic on individual choices in real life. Economics and political decisions, which are more likely to have social consequences, are typically made in adulthood.

Psychological literature on the other hand has provided an extensive analysis of the mechanism concerning the development of varying preferences across diverse individuals and populations. Most of the studies have stressed the importance of the learning process and socialization during childhood as a fundamental factor in the future manifestation of prosocial behaviours (Eisenberg and Fabes, 1998; Solomon et al., 1988; Staub, 1979). Piaget (1973) suggests that an individual's morality evolves during life and is shaped by experiences of cooperation and partnership with others. Similarly, according to Brody (1998), siblings represent the greatest source of repeated interaction for the development of social-skills during childhood. Through the process of interacting with siblings, children can experience and practice different roles and thus gain interpersonal skills. Even conflicts between siblings are significant in the development of cognitive and non-cognitive skills, and in the understanding of others' emotions and perspectives. As, in fact, is pointed out by Kitzmann, Cohen, and Lockwood (2002) "siblings help learning to manage conflict and obtain from them the

most efficient outcome". Conversely, only-children are viewed as lacking in this critical social learning background, and thus, are more likely to develop less cooperative and more selfish attitudes (Falbo and Polit, 1986). Psychology literature has also documented the existence of behavioural differences between only-children and children with siblings. Amongst these, only-children have been associated with lower interpersonal skills, more aggressive attitudes, being less trustworthy and averse to risk-taking (Downey and Condron, 2004). However, results are often conflicting. Many scholars have, in fact, developed discussions on the effect of siblings in prosocial-development, relying on the reverse model of resource dilution. According to this model, parents are endowed with limited resources such as money, energy and time. The more children in the family, the fewer parental resources can be allocated to each of the siblings. The educational process and social development can therefore be negatively affected by a large sibship size (Blake, 1989; Downey, 1995). Nevertheless, these mixed findings could potentially be explained by the different account of the number, order and spacing of siblings taken by various studies. It should also be noted that most studies in psychology are not incentivized. Overall, there have been different interpretations of the effect of siblings on the development of preferences; however, the general literature seems to agree that having a few children in the household has a positive impact on social skills. Under the supervision of parents, siblings are more likely to learn and generalize the social-skills acquired from positive and effective interaction, as well as how to avoid incurring unproductive conflicts. Assuming that overseen interaction between siblings is easier to realize when the number of children is small, we could expect that the advantages of this learning process are more evident when comparing only-children with children having a few siblings.

Our research contributes to the existing literature both in terms of methodology and research question. We study comprehensively the effect that growing up as an only-child has on social preferences, risk preferences, trust and behaviours in strategic

interaction within a coherent framework[3].

For the analysis, we start from the psychological idea of siblings as a source of interaction from where social-skills can be acquired and developed, and link this to the economic theory of preferences development as a consequence of early environmental conditions. We therefore try to understand whether differences exist between single-children and non single-children's preferences and whether the only-child status can be considered a predictor of less cooperative, altruistic, trusting and trustworthy preferences and of a different process of decision making and beliefs updating over repeated interactions as a consequence of lack of previous experience in social interaction. We also collect detailed questionnaire measures on family environment and socioeconomic characteristics of participants in the experiment and of their parents. Past research has in fact demonstrated that being an only-child can be endogenous to family background. Parents who have only one child may differ from those who have many children. Behavioural differences observed between single children and subjects with siblings may not therefore represent the actual causal impact of being an only-child, but they may be an expression of a broader set of different family background characteristics.

We elicit social preferences, risk preferences, trust and behaviours in interaction by running four different experiments with the same sample of subjects. More precisely, we ask participants to sequentially play a public goods game, a combination of dictator games and response games, an investment game and a task for eliciting risk preferences. After the experiment, we issue a questionnaire to disclose socioeconomic background information and assess the only-child status of each participant.

The public goods game allows us to address the question of heterogeneous preferences between only-children and non only-children and collect insights on eventual differences in the process of making decisions and updating beliefs in an interaction

[3]Social preferences models assume that an individual's utility is not only affected by her own material payoff, but it also depends on others' behaviours and material payoff. Social preferences include: altruism (Eckel and Grossman, 1996), difference-aversion, social-welfare and negative and positive reciprocity (Charness and Rabin, 2002; Falk and Fischbacher, 2006). Trust represents a subject's beliefs about others' trustworthiness together with preferences to take risk(Fehr, 2009).

environment. We then investigate more deeply the role that different social motivations (such as difference-aversion, social welfare and reciprocity) can have in determining behaviours of players in the two groups. With the investment game and the risk-preference game we finally test for differences in the level of trust, trustworthiness and risk aversion. Results document a consistent difference in social preferences of only children and non only-children, with negative reciprocity and social welfare maximizing actions in particular. Only-children are less likely to punish for misbehaviour and appear to be significantly more likely to choose social welfare maximizing actions in response games. In line with previous research, we further find that only-children are less trusting than subjects with siblings (Cameron et al., 2013; Glaeser et al., 2000). Finally, the detailed information about siblings collected through the questionnaire allows us to extend the analysis of differences in preferences between only-children and subjects with siblings to a more specific sibship structure and disentangle across different effects. We find that sibship structure differently affects subjects' preferences in different games. In particular sisters seem to play a role in determining different outcomes between single-children and subjects with siblings.

The reminder of the paper is organised as follow: Section 1.2 describes the compositions of our participants' sample and the experiment procedure. Section 1.3 introduces the experimental design. Section 1.4 presents and discusses main results and Section 1.5 concludes.

## 1.2 Subjects pool and procedure

The experiment was programmed in z-Tree (Fischbacher, 2007) and conducted at the EssexLab of the University of Essex between April and December 2017. Ethical approval was obtained by the Social Sciences Faculty Ethics Sub-Committee of the University of Essex. All 276 participants were recruited amongst undergraduate and master students, using the EssexLab Online Recruitment System. Table 1.1 summarises the

main characteristics of the participating students[4].

TABLE 1.1: Basic characteristics of the sample

| Variable | Observations | Mean | Std. Dev. |
|---|---|---|---|
| Female | 276 | 0.66304 | |
| Average participants age | 276 | 23.0217 | 6.3147 |
| Only_child | 276 | 0.14856 | |
| Average number of siblings | 276 | 1.934783 | 1.702219 |
| First-born | 232 | 0.375 | |
| Middle-born | 232 | 0.3707 | |
| Last-born | 232 | 0.2543 | |
| Age spacing | 100 | 4.53 | 3.608926 |
| Average number of sisters | 100 | 1.09331 | 1.064534 |
| Average number of brothers | 100 | 0.96342 | 1.067187 |
| Mother's age when born | 272 | 28.64338 | 5.188836 |
| Average mother's age | 272 | 51.48897 | 7.089821 |
| Average father's age | 274 | 53.85036 | 10.91138 |
| Nursery | 276 | 0.8007246 | |
| Time lived at home | 276 | 16.57971 | 5.235976 |
| Parents are married | 276 | 0.71385 | |
| Mother has higher edu. | 276 | 0.5145 | |
| Father has higher edu. | 276 | 0.4819 | |

We collect information relative to those family characteristics that have been shown to be relevant for subjects' only-child status in order to be able to control for them in our experiment results and thus obtain more reliable estimates of the causal impact of being an only-child on preferences structure.Our data on family environment include the years a subject has lived at home with parents, the marital status of parents, number, age and gender of older/younger/twin siblings, number, age and gender of older/younger/twin step-brothers or sisters and whether a subject attended nursery school. While socioeconomic information on parents comprises: age of parents, highest qualification achieved by parents, religion, ethnicity and country of origin of parents. Furthermore, as already outlined in the discussion above, a contribution of our research is to provide an estimate of the effect of being an only-child on preferences structure for a general population in its adulthood. We therefore collect information about socioeconomic, demographic and study backgrounds characteristics of subjects

[4]The number of observations for "Age spacing, Average number of sisters and Average number of brothers" is lower than that for the other variables as the questionnaire has been modified in the last 8 sessions by adding more detailed questions about the gender of siblings and their age difference.

in order to be able to test whether our students sample provides a good representation of the general populations. The share of only-children taking part in the experiment (approximately 15%) is relatively small compared to that of subjects with siblings. The low representation of only-children in our sample is due to the choice of a specific recruitment procedure. In order not to cause self-selection or priming effect into the experiment we decided not to ask participants specific questions about the number of their siblings before the end of the experiment. This choice, however, did not allow us to control for a balanced representation of single-children and subject with siblings in our participants' pool. Nevertheless, the only-children sample size we observed, is in line with the literature. Herrmann, Thöni, and Gächter (2008), report that the average share of only-children taking part into an experiment run in universities of fifteen different countries, is less than 14%. Furthermore, the share of only-children they report to participate into the experiment in the UK University (Nottingham University), is even lower than that of participants in our sample.

We ran a total of fifteen sessions with 12 to 28 subjects in each session. Sessions lasted on average 100 minutes and the average earnings were £19. In each session, we ran four games and a questionnaire with the same sample of participants. All participants played the four games in sequential order. In the two games involving two roles of decision making (the investment game and the response game), players were asked to make sequential decisions in each one of the roles and therefore play the same game twice. Instructions for each of the four games were presented separately. Participants received instructions for each game before the start of the specific game and were given some time to read them and ask questions[5]. The game was started only after everyone has finished reading the instructions and was certain to have understood them. Games were played simultaneously by all participants in the experiment session. The subsequent game with its instructions was therefore presented only after everyone has made decisions in the previous game. In all games except the repeated public goods game task, participants were not communicated any results about their decision or

[5]When first reading the instructions ,participants were informed that the experiment consisted of four games, but they were only provided with instructions for Game 1. Instructions can be found in the Appendix A.

other participants' decisions. The final earnings from the experiment were communicated only after the completion of the questionnaire. At the end of the experiment two out of the four games were selected for payments. Players were then told which games would have been relevant for their earning and were paid according to the decisions they made in the two selected games. This payment procedure was explained to participants at the beginning of the experiment. We chose to implement this payment procedure in order to maintain incentives from each game while minimizing confusion between games and avoiding averaging payoffs across games (a similar procedure has been implemented also by Charness and Rabin (2002) and Blanco, Engelmann, and Normann (2011)).

The games we chose for our experiment comprise a two tasks public goods game (Fischbacher and Gächter, 2010), a response game (Charness and Rabin, 2002), an investment game (Berg, Dickhaut, and McCabe, 1995) and a task for eliciting risk preferences (Eckel and Grossman, 2002). We also added a task for eliciting risk preferences since risk is often related to trust (Ben-Ner and Putterman, 2001; Cook and Cooper, 2003). Collecting direct evidence on participants' risk attitudes, could therefore help to better interpret behaviours observed from other games and is interesting in itself. Finally, we exploited the structure of the ten periods repeated public goods game task to assess beliefs about other' contributions, a measure of trust comparable with that obtained from other games, and to assess the process of beliefs updating and decision making of only-children and subjects with siblings in strategic interactions.

We ran eight sessions with the sequence (public goods game- response game- investment game- risk game) and seven with the sequence (response game- investment game- risk game- public goods game). We decided to run these specific two sequences as they allow us to vary the collocation of the public goods game from the first position to the last. The repeated task of the public goods game is, in fact, the only game where players receive feedback about their decisions and those of others. By moving the public goods game from the first to the last position, we can therefore test for sequence effect and whether receiving feedback affects participants' decisions in the following games. We did not find difference between these two sequences, and therefore pooled

the data and present general results in the following sections.

## 1.3 Experiment Design

We now describe in detail each of the different tasks.

### 1.3.1 Public Goods Game

We ask participants to solve two tasks: the first task consists of a one-shot public goods game, while the second task is a ten period repeated public goods game. All participants play both tasks in the same order.

The first task of the game provides a first measures of participants' cooperation preferences and requires participants to make an unconditional and a conditional contribution decision. In the unconditional contribution decision, each player is asked to state an unconditional choice of contribution for the public account. The conditional contribution decision, instead, is developed on the basis of the "Strategy method"[6] and is aimed at eliciting players' contribution preferences. It consists of a one-shot public goods game where players, rather than interacting with other players, are asked to indicate the amount of endowment that they would contribute, given each possible rounded average contribution choice of other players in the group. The game is played once and is implemented by showing the participant a "contribution table" reporting a list of the others' 21 possible average contributions (0,...,20)[7]. Players are then required to state their choice of contribution for any possible listed alternative. This first task is structured in an incentive compatible way, so that players are motivated to reveal their actual preferences.[8]

In the second task of the game we, instead, implement a variant of the public goods game experiment of Fischbacher and Gächter (2010).

[6] See Selten (1967).

[7] A screen shot of the decision screen is included in Appendix A, Figure A.1.

[8] At the end of the first task, one subject is randomly selected in each group. Final payoffs are then calculated taking into account the "Unconditional contribution decision" of three participants in the group and the corresponding "Conditional contribution decision" of the selected subject.

Our public goods game looks at contribution decisions and elicits beliefs about others' contributions in a sequence of ten periods repeated interactions with fixed groups. Beliefs about others' contributions, obtained from the repeated game, represent our first measure of trust and can furthermore help to disentangle different motives underlying the classification of social preferences. The study of contribution choices in the 10 periods game, instead, enables us to test the hypothesis of differences in the process of beliefs updating and contribution decision making between only-children and subjects with siblings.Both economic and psychological literature underline the importance of early environmental conditions for the development of cognitive/non-cognitive skills. In particular, in psychology it is often argued that only-children are more likely to lack the interaction experiences from which the ability to understand others' perspective and incentives from the environment of interaction can be acquired (Falbo and Polit, 1986; Kitzmann, Cohen, and Lockwood, 2002; Brody, 1998). We could therefore expect only-children to use different heuristics in updating beliefs, have diverse determinants of contribution choices, and their choices to develop according to different trends in an environment of repeated interaction.

We address all these questions by studying behaviour across the ten periods public goods game. Information on beliefs and heterogeneity of cooperation preferences allows us to assess the effect of different components on the contribution choices of subjects in the two groups. Meanwhile, the choice of an invariant structure for the groups' composition across periods enables us to compare the actual preferences of players (obtained from the first task schedule of contribution decisions) with choices revealed in an environment that allows for strategic play and therefore test for strategic plays[9].

The basic decision framework is a standard linear public goods game. Participants are randomly allocated to groups of four people and not informed about the identity of their partners. Each player is endowed with 20 tokens and is asked to choose how many of these to contribute for a public account. We assume a linear payoff function,

[9]Fischbacher and Gächter (2010) find a high degree of consistency between expressed preferences and actual contribution behaviours. Deviations from the expressed preferences could therefore provide insights on strategic play.

such that the payoff of a subject is given by:

$$\pi_i(g) = y - g_i + \alpha G$$

where

$$G = \Sigma_{j=1}^{N} g_j$$

is the value of the public account equal to the sum of all the contributions of players within the same group. $\alpha$ is the marginal per capita return on $G$ and is assumed to be (= 0.4).

This framework does not provide any incentive for the player to contribute to the public account and according to the standard theory, players should have a unique dominant strategy that is not to contribute anything $(g_j = 0, \forall j)$. However, the social return from a single contribution (=1.6) is higher than the private cost of contributing and thus, there exists a Pareto-dominant outcome $(g_j = 20, \forall j)$ which differs from the theoretical one. Furthermore, while the first equilibrium $(g_j = 0)$ is social-inefficient, the second one could provide a social welfare maximization.

The composition of groups is fixed across the ten periods of the game and players are made aware of this. Other than choosing a certain amount of endowment to contribute, in each period, players also have to report their beliefs about the average contribution of the other three players in that period. We elicit beliefs in an incentive compatible way, so that subjects are motivated to guess the exact amount contributed by others.[10] At the end of each period participants are provided feedback about the actual average contribution of the other group members and about their own earnings in that period. The payoff from the second task is given by the total sum of payoffs obtained in each one of the ten periods of iteration of the game.

[10]Participants are communicated in the instructions that additional earnings are possible based on the accuracy of their guesses: they receive three tokens if they exactly estimate others' contribution, two tokens if their estimate is higher or lower than others' actual average contribution by one unit, one token if the estimate differs from others' actual average contribution by two units; no additional earnings are given if the estimated average differs by more than two units from the actual others' contribution. At the end of the game one of the ten periods is selected and earnings obtained from the correct guess in that period are added to the total earnings from the public goods game.

Finally, the total payoff from the public goods game is calculated as the sum of earnings from the two tasks of the game.

### 1.3.2 Response Game

We use response games in order to understand the extent to which reciprocity, difference-aversion and social-welfare concerns drive the two groups behaviours.

In this task participants are asked to make decision in a sequence of two-person sequential move games. Games we implement are taken from a design of Charness and Rabin (2002). From the original experiment, however, we select a subset of games, already used in the experiment of Chen and Li (2009) and which allows us to account for the sensitivity of players' choices to the cost of self-benefit in the analysis.[11]

The response game comprises a mix of five dictator games and sixteen response games.[12] The dictator games allow us to test for various distributional motivations without including reciprocity concerns. The response games, instead, directly assess the role of reciprocity in decisions, by looking at how the second mover decision varies relative to the action that first mover has forgone[13]. The game also provides a rich set of alternative options that enables us to account for the effect that cost of reward, cost of punishment or the absence of private-costs have on the reciprocal choice of a subject.

Each subject is asked to make choices in a subset of twelve different games[14]. Each game is played twice by each subject so that everyone plays both roles. The order with which the two different roles are played is randomized across participants to enable us to control for order effect. Every time a participant makes a new decision, he/she is paired with a different subject in the experiment session. In order to avoid strategic play, participants are made aware of this procedure of random matching occurring at

[11] Payoffs from the original setting were changed in order to make them comparable to payoffs from other games.

[12] The Appendix shows a description of the set of games. See Table B.1 in Appendix B.

[13] Screen shots of dictator games and response games are shown in Figure A.2 and A.3 of the Appendix.

[14] In each session of the experiment half of the participants play GameSet 1 while other play GameSet2. GameSet1 includes: Dict1, Dict3, Resp1, Resp3, Resp6, Resp9, Resp10, Resp13, Resp16-19. GameSet2 includes: Dict2, Dict4, Dict5, Resp2, Resp4, Resp5, Resp7, Resp8, Resp11, Resp12, Resp14 and Resp15. See Table B.1 in Appendix.

each step of decision. No feedback about decision outcomes is given to participants at any time. Decisions made by the second mover in the response games are therefore based on the assumption that the first mover has forgone the outside option and left the second mover to make the choice determining final outcomes.

At the end of the game two out of the twenty-four decisions made are randomly selected for each player and the sum of payoff collected in both gives the final payoff from the response game.

### 1.3.3 Investment Game

We use the investment game in order to collect further information on the extent to which trust, reciprocity, difference-aversion and social welfare concerns differently affect the preferences structure of single-children and subjects with siblings.

The game is characterized by two distinct roles, role A and role B. Each participant in the experiment is asked to sequentially play both roles, however, the order of the roles is varied across participants to enable us to control for an eventual order effect in decision making. Players are explained that pairings are changed for each role, in the instructions. Each subject playing in the A role, receives an endowment of 10 tokens and is given the possibility to transfer all, part or nothing of her endowment to an anonymous paired player B. Precisely, player A can transfer to player B any integer number between 0 and 10. Player A is also informed of the fact that any transferred amount will be multiplied by 3 by the experimenter before being delivered to player B. Participants playing in the B role do not receive any endowment, but are instead given the opportunity to return all, part or none of the tripled received amount. Player B can return any integer number between zero and three time the amount the anonymous player A has transferred to her. In order to collect more observations, and better estimate the extent to which player B conditions her decision on the received amount, we elicit player B's decisions by the use of the strategy method. We therefore do not communicate to player B the actual amount the paired player A has transferred to her, but instead we ask her to state the amount she wants to return, for each of the possible

amounts player A has transferred. Specifically, player B is shown a table reporting the 11 possible values she has received from player A (from 0 to 30) and is asked to state the corresponding amount she would return for each of the 11 possibilities[15]. The self-regarding model would predict that player B keeps all received amounts, while player A, knowing that no amount will be returned by player B, does not transfer any positive amount. Any deviations from this assumption, provides us with some inference about trust and social-preferences amongst players. While the choice to transfer a positive amount to the second player can be rationalized by trust or altruistic other-regarding preferences, the choice to return a positive sum to the first players can be explained by reciprocity or unconditional other-regarding preferences.

At the end of the game, one role is randomly selected for each player and the player's payoff from the investment game is defined according to her decision and the decision of her partner in that role. Specifically, for those players for which role B is selected to be payoff relevant, earnings are calculated by considering the conditional choice they made, given the actual transfer made by the matched player A. Player A's payoff is calculated as the received amount (10 tokens), minus the amount transferred, plus the tokens player B decides to return. Player B's payoff is, instead, given by the amount transferred from player A minus the returned amount.

The investment game provides an example of decision environment comparable to that of the public goods game. It has, in fact, been broadly argued that most of the subjects who contribute for the public good, do so as they are expecting others to contribute as well: a possible explanation for contributions could therefore be trust in the others (Sugden, 1984; Keser and Winden, 2000; Fischbacher, Gächter, and Fehr, 2001; Herrmann and Thöni, 2003). However, the investment game is characterized by a strategically different setting. It therefore allows us to obtain results comparable with those from the public goods game, but it also helps in gathering further evidence for their interpretation.

[15] A screen shot of the table shown to Player A is included in Appendix, Figure A.4.

### 1.3.4 Risk Task

We decided to elicit risk preferences as it could prove to be a useful tool for interpreting the results from other games. Different studies have underlined the important correlation between risk and trust. Ben-Ner and Putterman (2001) suggest that trust in investment games, where no previous information about the other player is given, can be mainly explained by previous experience and players' own risk attitudes. Consequently, subjects who are more averse to risk are also less trusting. The link between trust and risk extends also to the interpretation of cooperation choices and therefore to the results from the public goods game. Cooperation could, in fact, be interpreted as a decision to trust other subjects by taking a risk on them (Cook and Cooper, 2003).

The risk task we implement is taken from Dave et al. (2010) and represents a variant of the method developed by Eckel and Grossman (2002). Participants are shown a table listing six gambles and are asked to choose the one that they want to play (see Table 1.2)[16].

TABLE 1.2: Gambles characteristics

| Gamble Choice (50/50) | Low Payoff | High Payoff | Expected return | Standard deviation |
|---|---|---|---|---|
| Gamble_1 | 28 | 28 | 28 | 0 |
| Gamble_2 | 24 | 36 | 30 | 6 |
| Gamble_3 | 20 | 44 | 32 | 12 |
| Gamble_4 | 16 | 52 | 34 | 18 |
| Gamble_5 | 12 | 60 | 36 | 24 |
| Gamble_6 | 2 | 70 | 36 | 34 |

Each gamble may have either a high payoff or a low payoff, with equal probability of occurring (50%). The first gamble gives a certain payoff (28 tokens), while the expected payoff increases with risk in the subsequent four gambles (from gamble 2 to gamble 5). Gamble 6 is instead characterized by the same expected payoff as gamble 5, but has a higher standard deviation. The increase in risk is represented by the increase in the standard deviation across payoffs of each gamble. The design of these gambles allows us to distinguish risk-averse players, who should choose gambles with the lower standard deviation (from 1 to 4), from those who are risk-neutral and should

[16] See also Figure A.5 in Appendix.

choose the gamble with the higher expected return and those who are risk-seeking and should choose gamble 6.

At the end of the game, the computer randomly selects one payoff to realize (either the high payoff or the low payoff) for each player. The final payoff from the risk game is then determined by applying the computer selection to the payoffs for the chosen gamble.

## 1.4 Results

We will first analyse behaviours of only children and subjects with siblings in each game separately, then we will provide a synthesis and compare results across different games.

### 1.4.1 Public Goods Game

The conditional contribution decisions made in the first task of the public goods game provide us with a contribution schedule which describes how much each participant is willing to contribute as a function of others' contribution. We use this contribution schedule and classify players in different categories of preferences types, according to the slope (of a linear regression) of the schedule and the mean contribution in the schedule. Free-rider types are those subjects who contributes 0 regardless of what others contribute (mean=0, slope=0). Perfect conditional co-operators are subjects who contribute exactly the same amount others contribute (mean=10, slope=1). Unconditional co-operators, on the other hand, are subjects who contribute their whole endowment for every possible average contribution of the others (mean=20, slope=0). Table 1.3 shows the proportion of only-children and subjects with siblings who made contribution decisions consistent with different preferences types. From the table we can observe that both groups are characterized by heterogeneity across different preferences types. Most of the subjects, both for only-children and subjects with siblings, have a positive mean of contribution and a positive slope. They therefore contribute

an amount which is increasing with others' contribution, but which does not match the others' contribution. The other predominant types of preferences in both groups are given by free-riders, perfect conditional co-operators and subjects contributing an amount decreasing in others' contribution. There are additionally some participants who contribute an unconditional positive amount (positive mean and slope=0). We perform a test of proportions for the null-hypothesis that the proportion of only-children and non only-children is the same across all different preferences types and find we are unable to reject the null-hypothesis for all preferences types. Therefore, first results seem to suggests that there is no difference in cooperative preferences between only-children and subjects with siblings.

TABLE 1.3: Consistency of behaviours with different contribution preferences

| . | Total # observations | Free Riders | Perfect Conditional Cooperators | Unconditional Positive Amount | Positive mean and slope | Negative slope Positive mean |
|---|---|---|---|---|---|---|
| Only children | 41 | 9.76% (4) | 17.07% (7) | 7.32% (3) | 56.10% (23) | 9.76 (4) |
| Subj. with siblings | 235 | 8.09% (19) | 10.64% (25) | 5.96% (14) | 60.85% (143) | 14.04% (33) |
| Overall | 276 | 9.21% (23) | 11.59% (32) | 6.16% (17) | 60.14% (166) | 13.41% (37) |

Having collected this first evidence, we continue the study of behaviours in the public goods game and we move on by analysing the contribution decisions of the two groups in the unconditional contribution decision task. We use an OLS regression, with robust standard errors clustered on subject level and estimate the effect that being an only-child has on unconditional contributions. Table 1.4 shows regression results of the effect of being an only-child on unconditional contribution decision. The only-child dummy variable represents the fundamental variable of interest for our analysis and is equal to one for all subjects who do not have any sibling or step-brother/sister. We find that only-children contribute a significantly lower amount (p-value $< 0.01$) and this negative effect persists also after controlling for demographic and family-background characteristics (p-value $< 0.01$).

As we collected detailed information on sibship structure in the questionnaire, we

TABLE 1.4: OLS estimates of the effect of only-child on unconditional contribution

| Variables | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| Only_child | -2.355*** | -2.247*** | -2.542*** | | |
| | (0.756) | (0.859) | (0.845) | | |
| #siblings in HH | | | | 0.360* | 0.352* |
| | | | | (0.200) | (0.201) |
| age | | 0.103** | 0.066 | 0.058 | 0.062 |
| | | (0.050) | (0.066) | (0.067) | (0.067) |
| Female | | 0.310 | 0.327 | 0.330 | 0.211 |
| | | (0.679) | (0.725) | (0.718) | (0.738) |
| Constant | 7.574*** | 7.042* | 7.452* | 6.007 | 5.928 |
| | (0.340) | (3.617) | (4.309) | (4.389) | (4.483) |
| Observations | 276 | 276 | 272 | 272 | 269 |
| R-squared | 0.027 | 0.176 | 0.228 | 0.211 | 0.214 |
| Controls | No | | | | |
| Birth order | | No | No | No | Yes |
| Ethnicity | | Yes | Yes | Yes | Yes |
| Nationality | | Yes | Yes | Yes | Yes |
| Religion | | Yes | Yes | Yes | Yes |
| Nursery | | No | Yes | Yes | Yes |
| Time_home | | No | Yes | Yes | Yes |
| Par_maritalStat | | No | Yes | Yes | Yes |
| AgeMoth_birth | | No | Yes | Yes | Yes |
| ParentsEdu | | No | Yes | Yes | Yes |

Robust standard error clustered on subject level.
Standard errors in parentheses
*** p<0.01, **p<0.05, * p<0.1

try to disentangle the observed effect by introducing in the analysis specific information about siblings. Specifically, we add as control variables, the total number of siblings in the household and the birth order of a subject (model 5).[17] [18] As shown in model 4 and 5, we find that the overall number of siblings in the household is driving the result. Therefore, the more siblings a subject has, the more she is likely to contribute a positive amount in the unconditional contribution task.

In order to better interpret the observed difference in contribution decisions between the two groups, we extend the analysis to subjects' beliefs. A possible candidate to explain lower contribution decisions in an environment where no feedback about others' contribution is given and where subjects are motivated by social preferences, are, in fact, the beliefs about others' behaviours. We use the beliefs stated in the first

[17] Birth order refers to the order a subject is born in her family. Birth order in our sample goes from first born up to eight born.

[18] We also tried to disentangle the effect of being an only-child by including other detailed information about siblings, such as gender composition and age gap, but we didn't find any of these characteristics to be significant in driving the result and therefore do not report the results.

contribution period of the ten periods repeated public goods game as a measure of the "home-grown" beliefs of a subject. As no previous interaction has occurred, we can consider these beliefs representative of an individuals' intuitive beliefs about others' contribution. The red dashed-dotted lines in Figure 1.1 show the average period beliefs for subjects with siblings and only-children.

FIGURE 1.1: Mean beliefs, contributions and predicted contribution over time for subjects with siblings and only-children

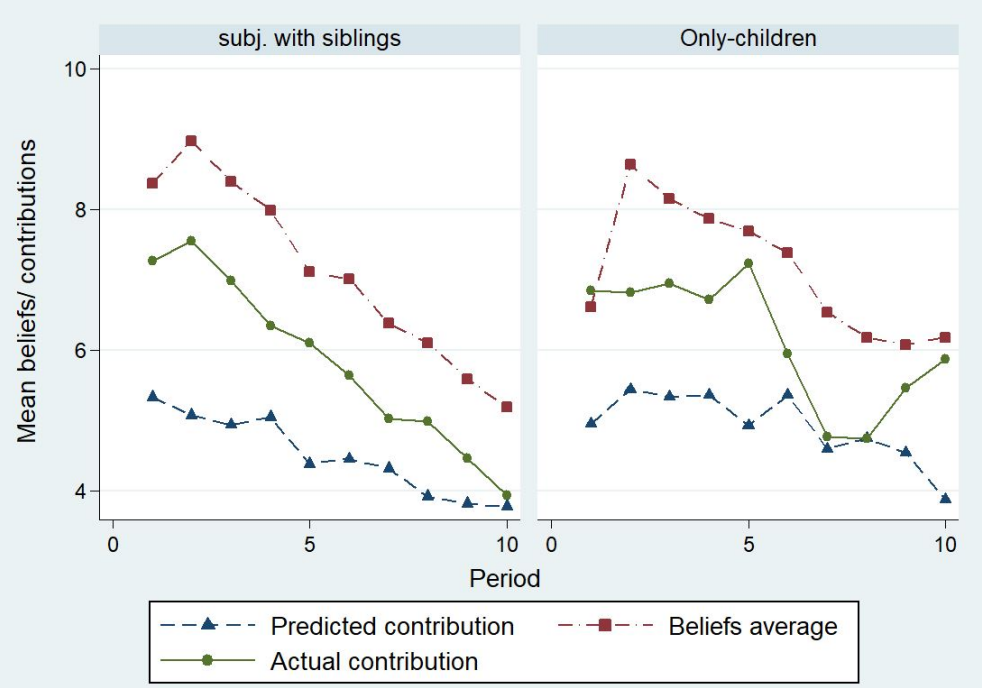


Only-children's average first period beliefs (6.61 tokens) are significantly lower than those of subjects with siblings (8.37 tokens) (p-value $< 0.05$ Mann Whitney rank-sum test). This result is consistent with the above classification of preferences types and the lower contribution choice of single-children. The pessimistic beliefs held by a consistent share of conditional co-operators can, in fact, reasonably result in lower contribution decisions. Moreover, that only-children hold significantly lower beliefs confirms the frequent finding that only-children are less trusting than subjects with siblings (Glaeser et al., 2000; Cameron et al., 2013).

We then extend our analysis to the study of behaviours of single-children and subjects with siblings in the ten periods repeated game in order to understand whether observed differences in contribution and beliefs persist in an environment of repeated

interaction. Furthermore, we exploit the structure of the repeated public goods game to investigate the effect of being an only-child on the process of beliefs updating and contribution decision making. Figure 1.1 shows mean beliefs and contributions over time for only-children and subjects with siblings. Only-children and subjects with siblings' contributions and beliefs across the ten periods are not significantly different from each other(Kruskal-Wallis test with group averages over all 10 periods as independent observations, p = 0.361 for contributions and p = 0.410 for beliefs)and both decrease over the ten periods following a statistically significant decreasing trend ($p < 0.01$, Cuzick's Wilcoxon-type test for trend). It seems therefore that when the game is repeated and strategic play introduced the difference in behaviours between the two groups tends to disappear.

We finally conclude the analysis of behaviours in the ten period public goods game by investigating the process of beliefs updating and development of contribution decisions amongst the two groups of interest. We firstly study the formation of beliefs by the use of an OLS regression with robust standard errors clustered at the subject level (Table 1.5). Model 1, which includes just the dummy variable for only-children, simply confirms the insignificant impact of being an only-child on average beliefs about others' contribution. Model 2 instead assesses the effect of time, subjects' previous period beliefs and average contribution of others in the previous period, on beliefs. Our results are in line with findings from Fischbacher and Gächter (2010) . Both beliefs at (t-1) and others' contribution in (t-1) have a significant positive effect on beliefs about other group members. However, while Fischbacher and Gächter (2010) observe an insignificant negative effect of period, we find that period does have a significant negative effect. Nevertheless, this difference could be explained by the fact that we use an invariant structure of groups, while their experiment is based on a random matching in each period. In model 3 we add demographic and family background controls and observe that "only-child" remains statistically insignificant. Finally, in model 4, we test whether the effect of period, beliefs at (t-1) and others' contribution in (t-1) have a different effect on beliefs updating of single-children and subjects with siblings. The insignificant coefficient for "only-child" interacted with period and others' contribution

in (t-1), suggests that these elements have the same effect on beliefs updating process both of single-children and of subjects with siblings. However, we find a small, but significant effect of the coefficient for "only-child" interacted with beliefs at (t-1), indicating that only-children assign a slightly higher weight to beliefs held in previous period while updating their beliefs for the following period.

TABLE 1.5: Beliefs about other group members' contribution and Contribution determinants

| | Beliefs about others' contribution | | | | Contribution determinants | | | |
|---|---|---|---|---|---|---|---|---|
| Variables | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Only_child | 0.018 | | 0.246 | -0.230 | 0.307 | | -0.001 | -1.254 |
| | (0.498) | | (0.265) | (0.755) | (0.666) | | (0.518) | (0.895) |
| Period | | -0.114*** | -0.124*** | -0.129*** | | -0.173*** | -0.177*** | -0.199*** |
| | | (0.024) | (0.025) | (0.028) | | (0.029) | (0.030) | (0.032) |
| Others'contribution (t-1) | | 0.377*** | 0.380*** | 0.395*** | | | | |
| | | (0.032) | (0.030) | (0.032) | | | | |
| Average beliefs (t-1) | | 0.506*** | 0.480*** | 0.461*** | | | | |
| | | (0.030) | (0.030) | (0.034) | | | | |
| Beliefs about others'contribution | | | | | | 0.372*** | 0.378*** | 0.349*** |
| | | | | | | (0.044) | (0.045) | (0.045) |
| Predicted Contribution | | | | | | 0.331*** | 0.308*** | 0.327*** |
| | | | | | | (0.051) | (0.053) | (0.048) |
| Female | | | 0.252 | 0.248 | | | -0.505 | -0.522 |
| | | | (0.182) | (0.180) | | | (0.367) | (0.371) |
| Age | | | 0.018 | 0.016 | | | 0.025 | 0.029 |
| | | | (0.016) | (0.015) | | | (0.034) | (0.034) |
| Only_child x Period | | | | 0.024 | | | | 0.117 |
| | | | | (0.065) | | | | (0.078) |
| Only_child x Others'contribution (t-1) | | | | -0.078 | | | | |
| | | | | (0.079) | | | | |
| Only_child x Average beliefs(t-1) | | | | 0.112* | | | | |
| | | | | (0.062) | | | | |
| Only_child x Average beliefs about oth. | | | | | | | | 0.191 |
| | | | | | | | | (0.166) |
| Only_child x Predicted Contribution | | | | | | | | -0.157 |
| | | | | | | | | (0.218) |
| Constant | 7.115*** | 1.706*** | 3.071*** | 3.225*** | 5.829*** | 2.671*** | 6.304*** | 6.404*** |
| | (0.186) | (0.272) | (0.930) | (0.978) | (0.253) | (0.335) | (1.813) | (1.836) |
| Observations | 2,680 | 2,412 | 2,376 | 2,376 | 2,680 | 2,680 | 2,640 | 2,640 |
| Number of periods | 9 | 9 | 9 | 9 | 10 | 10 | 10 | 10 |
| Controls | No | No | Yes | Yes | No | No | Yes | Yes |
| Ethnicity | | | Yes | Yes | | | Yes | Yes |
| Nationality | | | Yes | Yes | | | Yes | Yes |
| Religion | | | Yes | Yes | | | Yes | Yes |
| Nursery | | | Yes | Yes | | | Yes | Yes |
| Time_home | | | Yes | Yes | | | Yes | Yes |
| Par_maritalStat | | | Yes | Yes | | | Yes | Yes |
| AgeMoth_birth | | | Yes | Yes | | | Yes | Yes |
| ParentsEdu | | | Yes | Yes | | | Yes | Yes |
| R^2 | 0.01 | 0.52 | 0.53 | 0.53 | 0.01 | 0.39 | 0.4 | 0.4 |

Robust standard error clustered on subject level.
*** p<0.01, ** p<0.05,* p<0.1

Having documented that there exists no significant difference between single-children and subjects with siblings' beliefs across the ten periods of the game, we move on in the analysis and try to study whether instead a difference between the two groups occurs when looking at the ten period contribution decisions. The left panel of table 5 reports results of an OLS regression with robust standard errors clustered on the subject level, for the effect of being an only-child on contribution decision. Regression results show

that being an only-child has a negative effect on contribution decision, but this effect is not significant when we take into account the panel data structure of the ten periods game. Model 2 confirms the finding of Fischbacher and Gächter (2010), that both beliefs, predicted contribution and period affect the contribution decision of subjects. In model 4, we finally test for a different effect of beliefs, predicted contribution[19] and period on contribution choices of only-children and subjects with siblings. However, the insignificant effect of the only-child dummy variable interacted with beliefs, predicted contribution and period, suggests that the three explanatory variables have the same effect on only-children and non only-children's decisions.

### 1.4.2 Response Game

We exploit data from the response game in order to directly differentiate among various social preferences (difference-aversion, social-welfare and reciprocity) and try to better interpret the differences observed between preferences and behaviours of only-children and subjects with siblings in the public goods game. First, we investigate the effect of being an only-child on distributional preferences (competitive, difference-averse and social-welfare) by computing the proportion of single-children and subjects with siblings, who make decisions consistent with different distributional models.[20] Table 1.6 shows the proportion of decisions consistent with different models for only-children and for subjects with siblings.

Further, it reports proportions distinguishing two categories: dictator games and all games. This distinction enables us to assess subjects' behaviours both in the case where reciprocity is not involved (dictator games) and where, instead, it may play a role in players' decisions (response games). Our analysis focuses only on player B's behaviour, to simplify the comparability of results across different games. From Table 1.6 it is possible to observe that difference-aversion and social-welfare preferences are the

[19] The variable "predicted contribution" represents subjects' preferences and it corresponds to the contribution decision expressed in the Conditional contribution task for the level of "others' contribution" revealed by the stated beliefs in that period.

[20] The above statistics of consistency of behaviours with distributional models rely on the classification proposed by Charness and Rabin (2002).

TABLE 1.6: Proportion of decisions consistent with different models

| . | Game | Tot # observations | Narrow Self-interest | Competitive | Difference Aversion | Social Welfare |
|---|---|---|---|---|---|---|
| Only-Children | Dictator | 99 | 78.79% (78) | 66.67% (66) | 84.85% (84) | 87.88% (87) |
| | Response | 321 | 81.62% (262) | 65.73% (211) | 82.24% (264) | 89.72% (288) |
| | All games | 420 | 80.95% (340) | 65.95% (277) | 82.86% (348) | 89.29% (375) |
| Subj. with Siblings | Dictator | 591 | 78.34% (463) | 69.54% (411) | 86.97% (514) | 89.51% (529) |
| | Response | 1887 | 78.70% (1485) | 65.92% (1244) | 80.92% (1527) | 85.96% (1622) |
| | All games | 2478 | 78.61% (1948) | 66.79% (1655) | 82.36% (2041) | 86.80% (2115) |
| Overall | Dictator | 690 | 78.41% (541) | 69.13% (477) | 86.67% (598) | 89.28% (616) |
| | Response | 2208 | 79.12% (1747) | 65.90% (1455) | 81.11% (1791) | 86.50% (1910) |
| | All games | 2898 | 78.95% (2288) | 66.67% (1932) | 82.44% (2389) | 87.16% (2526) |
| | | | Hypotesses and p-values Only-children = Subj.with Siblings | | | |
| | Dictator | | 0.9205 | 0.5664 | 0.5653 | 0.6275 |
| | Response | | 0.2335 | 0.9463 | 0.5762 | 0.068 |
| | All games | | 0.2766 | 0.737 | 0.8063 | 0.1597 |

most effective in explaining subjects' behaviours both for only-children and subjects with siblings and both for the case where reciprocity is excluded and the case where it is considered. These results are largely consistent with those from Charness and Rabin (2002). The last two rows of Table 1.6 report p-values for the test of proportions for the null hypothesis that the proportion of decisions consistent with each different distributional model is the same for only-children and subjects with siblings. P-values (p-value $< 0.10$) suggest that being an only-child has a significant positive effect on the likelihood of making choices consistent with social welfare maximizing preferences in response games. In order to collect more information and to better interpret this difference, we extend the analysis and focus solely on the likelihood of only-children and subjects with siblings of taking actions consistent with social welfare models. To this scope, we compute the proportion of subjects with siblings and only-children who make social welfare maximizing decisions and we try to extend previous evidence by including this time in the analysis also decisions made when playing in the role of

player A. Further, we exclude from the study of proportion of social welfare preferences, four games that are characterized by same aggregate payoffs for different outcomes [21]. Table 1.7 shows the proportion of social welfare maximizing decisions made by only-children and subjects with siblings for both role A and B.

TABLE 1.7: Proportion of Social Welfare maximizing decisions and the effect of being Only-child

| . | Only-Children | Subj. with Siblings | Overall | Hypotheses and p-values Only-children = Subj.with Siblings |
|---|---|---|---|---|
| **Player A** | 58.84% (203) | 53.82% (1077) | 54.56% (1280) | 0.0839 |
| **Player B** | 59.51% (244) | 55.32% (1300) | 55.94% (1544) | 0.1146 |
| **Overall** | 59.21% (1477) | 54.63% (2377) | 55.31% (2820) | 0.0196 |

The last column presents p-values for the test of proportion for the null hypothesis that only-children and subject with siblings are equally likely to make decisions consistent with social welfare model. Results confirm previous observations. Only-children are more likely to take social welfare maximizing actions both when playing in the role of player A and when we consider all behaviours and roles at the aggregate level. The difference, however, vanishes again when comparing behaviours of the two groups deciding as player B in both dictator and response games. Results seem therefore to suggest that the two groups' attitude towards social welfare maximizing actions differs to some extent and it is sensitive to the specific game structure. This finding is also supported by behaviours revealed in the two tasks of the public goods game. Observed stronger attitudes toward social welfare maximizing actions would in fact lead to expect higher choices of contribution across only children when analysing data from the public goods games. Results from the one shot public goods game instead show that only-children are significantly less cooperative than subject with siblings. However, the difference between only-children and subjects with siblings disappears when we assess subjects' decisions in the different setting of the repeated public goods game, confirming that differences in behaviours between the two groups could actually depend on the specific game structure.

[21] The games excluded from the analysis are: Dict 5, Resp 9 and Resp 10.

We finally complete our analysis on social-preferences, by directly studying the effect of being an only-child on reciprocity. We study reciprocity concerns amongst the two groups by looking at behaviours of only-children and subjects with siblings in games of positive and negative reciprocity. Games of positive reciprocity are those where the entry of player A into the game is associated with good intentions whereas game of negative reciprocity are those games where the entry of player A is associated with bad intentions. Table 1.8 shows the logit model we use to estimate the probability that player B punishes player A in games of negative reciprocity.

Models 1-2-3 and 4 present the results of logit regressions where the dependent variable is the probability B punishes A. Model 1 looks at the effect that factors such as the damage to A if B punishes (which is equal to the loss of A's payoff if B punishes), the difference between B and A's payoff if B punishes and the cost of punishment (which is calculated as the difference in B's payoff when B chooses to punish over the alternative action), have on the probability of choosing the action to punish. All these factors have a statistically significant effect on B's likelihood to punish (p-value $< 0.01$)[22]. We therefore include them in the following model as additional explanatory variables. Model 2 presents the results of the logit regression where we add as independent variable a dummy for only-child and we control for demographic and family background characteristics. Results of model 2 suggest that being an only-child significantly decreases the likelihood to punish (p-value $< 0.1$) for misbehaviour. In model 3 we further interact each of the three covariates determinant of punishment introduced in model 1 with the only-child dummy, so that to try to assess whether they have a different effect in determining punishment decisions of only-children and subjects with siblings. Results suggest that there exists a small, but significantly different effect of two covariates on only-children and subjects with siblings' choice of punishment. "Damage to A if B punishes" interacted with the only-child dummy exhibits a marginal effect of 0.026 (p-value $< 0.1$), which indicates that the size of the damage that can be inflicted to A has a stronger effect in determining punishment decisions across only-children. However, the interaction between the only-child dummy and the variable "B's payoff ahead of A

[22] Results are consistent both in terms of magnitude and sign with those from Chen and Li (2009).

TABLE 1.8: Logit regression determinant of reciprocity

| | Probability B punishes,A | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Only_child | | -0.077* | -0.082 | |
| | | (0.042) | (0.053) | |
| younger-brother | | | | -0.057 |
| | | | | (0.042) |
| younger-sister | | | | 0.104** |
| | | | | (0.052) |
| older-brother | | | | 0.076 |
| | | | | (0.063) |
| older-sister | | | | -0.040 |
| | | | | (0.043) |
| Age spacing | | | | 0.005 |
| | | | | (0.005) |
| damage to A if B punishes | 0.047*** | 0.049*** | 0.046*** | 0.043*** |
| | (0.004) | (0.004) | (0.004) | (0.006) |
| B's payoff ahead of A if punishes | -0.107*** | -0.113*** | -0.099*** | -0.112*** |
| | (0.014) | (0.014) | (0.015) | (0.025) |
| cost to punish | -0.312*** | -0.313*** | -0.301*** | -0.252*** |
| | (0.016) | (0.015) | (0.018) | (0.027) |
| damageA x Only_child | | | 0.026* | |
| | | | (0.016) | |
| Bahead x Only_child | | | -0.143* | |
| | | | (0.073) | |
| cost x Only_child | | | -0.078 | |
| | | | (0.069) | |
| Female | | 0.091*** | 0.092*** | 0.043 |
| | | (0.033) | (0.033) | (0.042) |
| age | | -0.002 | -0.002 | 0.001 |
| | | (0.002) | (0.002) | (0.004) |
| Observations | 1,518 | 1,499 | 1,499 | 529 |
| Controls | No | | | |
| Ethnicity | | Yes | Yes | Yes |
| Nationality | | Yes | Yes | Yes |
| Religion | | Yes | Yes | Yes |
| Time_home | | Yes | Yes | Yes |
| Par_maritalStat | | Yes | Yes | Yes |
| AgeMoth_birth | | Yes | Yes | Yes |
| ParentsEdu | | Yes | Yes | Yes |
| Log-likelihood | (-728.80) | (-667.27) | ( -663.52) | (-219.36) |
| Pseudo R | 0.172 | 0.232 | 0.236 | 0.256 |

Tables show marginal effects. Robust standard error clustered on subject level.
*** p<0.01, **p<0.05, * p<0.1

if punishes" has a marginal effect of -0.143 (p-value $< 0.1$) suggesting that the concern for payoff gap instead has a stronger effect in reducing only-children's likelihood of punishment. The latter result is interesting also because it provides evidence that corroborates the connection between distributional and reciprocal preferences amongst only-children, that was observed while assessing distributional preferences between groups. In the first part of analysis of behaviours in response games we observed, in fact, that only-children are significantly more likely to choose social welfare maximizing actions in games where reciprocity is involved (response games). Similarly, when we instead directly study reciprocity concerns between groups, we find that the probability of negative reciprocity is increased by the option of obtaining more equal payoffs across players. Finally, in model 4, we try to exploit questionnaire information about family structure in order to disentangle the effect of being an only-child on negative reciprocity. We therefore include in the model four dummy variables reporting information on siblings' age and gender. Each dummy assumes value one when the subject has at least one younger brother and/or a younger sister and/or an older brother and/or an older sister. Results in model 4 show that the effect in this case is driven by having at least one younger sister. Having at least one younger sister increases, in fact, the probability to punish by 10.4 % (p-value$<$0.05)[23]. We also run a similar analysis for games of positive reciprocity, but as we cannot find any significant effect of being an only-child on positive reciprocity, we do not report the results in the main text[24].

### 1.4.3 Investment Game

We analyse data from the investment game to obtain a measure of trust comparable with that elicited by studying beliefs in the public goods game and collect more information about different social preferences between only-children and subjects with

[23] To test whether the lower choice of punishment of only-children is determined by a specific sibship structure, we also run regressions where we control for family size, birth order and overall number of younger and older brothers and sisters. The results, however, do not reveal any significant effect for these components and we therefore do not report them in the main text.

[24] Table B.2 in Appendix shows the logit model we use to estimate the probability that player B rewards player A in games of positive reciprocity.

siblings.

We first examine behaviours of single-children and subjects with siblings when playing role A[25]. The average amount sent to player B is 3.84 tokens for subjects with siblings and 3 tokens for only-children. The average amount sent by the two groups is statistically significantly different (p=0.09, Mann Whitney rank-sum test). We run an OLS regression in order to test for the effect of being an only-child on the amount sent and control for demographic/ family background characteristics. Table 1.9 reports results of the OLS regression, with robust standard errors clustered on subject level, for the effect of being an only-child on the amount sent to player B.

TABLE 1.9: OLS estimates of the effect of only-child on amount sent to player B

| Variables | (1) | (2) | (3) |
|---|---|---|---|
| Only_child | -0.874* | -1.120** | |
| | (0.462) | (0.481) | |
| tot. younger_sist | | | -0.044 |
| | | | (0.542) |
| tot. younger_broth | | | -0.228 |
| | | | (0.569) |
| tot. older_sist | | | 0.758** |
| | | | (0.374) |
| tot. older_broth | | | 0.504 |
| | | | (0.327) |
| age | -0.007 | -0.045 | -0.060 |
| | (0.016) | (0.028) | (0.038) |
| Female | -0.214 | -0.095 | 0.652 |
| | (0.385) | (0.408) | (0.639) |
| Constant | 3.193*** | 6.520*** | 2.180 |
| | (0.679) | (2.930) | (1.988) |
| Observations | 264 | 260 | 121 |
| R-squared | 0.025 | 0.089 | 0.179 |
| Controls | | | |
| Ethnicity | Yes | Yes | Yes |
| Nationality | Yes | Yes | Yes |
| Religion | Yes | Yes | Yes |
| Nursery | No | Yes | Yes |
| Time_home | No | Yes | Yes |
| Par_maritalStat | No | Yes | Yes |
| AgeMoth_birth | No | Yes | Yes |
| ParentsEdu | No | Yes | Yes |

Robust standard error clustered on subject level.

*** p<0.01, **p<0.05, * p<0.1

[25]The analysis of subjects' behaviours in the investment game is based on a smaller sample (264 observations rather than 276) as we lost some observations due to a crash in the software during the experiment.

Results confirm the observed difference in amount sent by the two groups by showing that only-children pass a statistically significant lower amount of their endowment (0.87 tokens less, p-value$<$0.1) and that the significant difference persists also after controlling for demographic and family-background characteristics (model 1 and 2). As the sent amount represents a measure of trust, these results corroborate our previous findings of only-children's significantly lower beliefs as observed in the public goods game. Furthermore, our negative only-children's trust results are consistent with that from Cameron et al. (2013), who elicit trust measures in a similar experiment setting and find that only-children are significantly less trusting. We then conclude the study of trust amongst only-children and subjects with siblings by directly trying to disentangle the sibship components determining differences in trust. We investigate the siblings' characteristics driving the results in model 3. In this model we add as control variables for the determinants of trust the number of older brothers and sisters and the number of younger brothers and sisters that a subject has [26]. Results indicate that the difference in trust between only-children and subject with siblings is mainly determined by the number of older sister of a subject. The higher is the number of older sister of a subject the higher is the amount that the subject is willing to transfer to the second player. This result is aligned with Harbaugh et al. (2002) and Hardin (2001), who show that younger siblings are more trusting and suggest as an explanation for these results, that the higher trust originates from the greater opportunity to experience positive reciprocity from older siblings. Our result is consistent with these findings and furthermore provide a detailed analysis for the effects of older sisters and brothers.

Having studied behaviours of single-children and subjects with siblings when playing role A, we move on and analyse choices made in role B.

The average amount returned by only-children is 4.12 tokens while it is 3.95 tokens for subjects with siblings. Figure 1.2 shows the average amount returned by amount received and only-child status.

[26] As we have done in previous analysis, we control also in here for different sibship structure specifications, but report results only for the siblings' characteristics that have a significant effect.

FIGURE 1.2: Average tokens returned by tokens received and only-child status

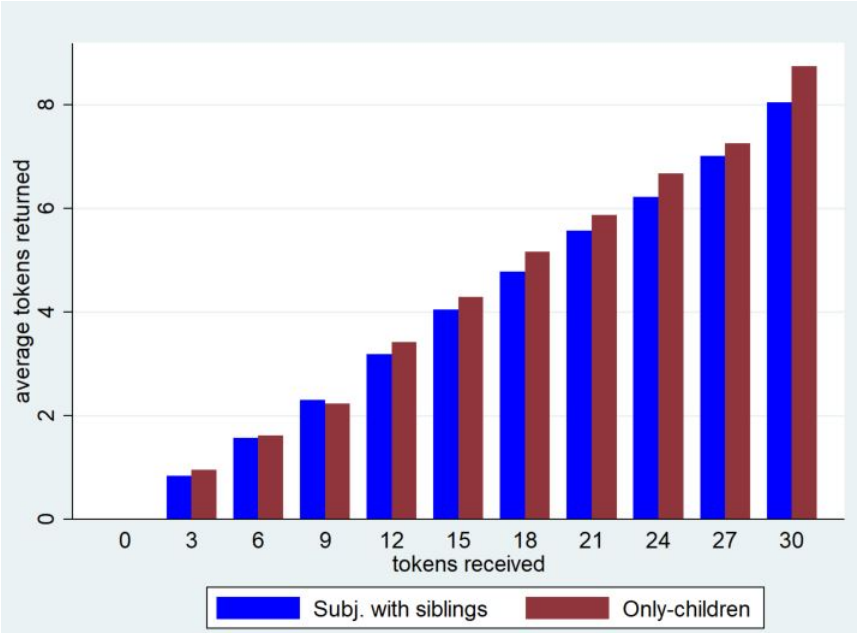


The average amount returned by single-children and subjects with siblings across different levels of received amount is not statistically different (p-value $> 0.1$ across all levels of received amount, Mann Whitney rank-sum test). However, these results are not in line with previous findings from Glaeser et al. (2000) and Cameron et al. (2013). A possible explanation could be given by the fact that we use the strategy method for eliciting player B's choices, whereas other studies are based on sequential games, where player B only has to state a choice given the actual amount she receives.

### 1.4.4 Risk Task

We use the risk game to investigate whether risk attitudes differ between single-children and subjects with siblings and whether they can help explain the difference in behaviours in other games.

As argued in the design description of the game, gambles numbers are increasing in expected return and standard deviation associated with each gamble. The choice of a specific gamble number is therefore a good parametric summary index of risk preferences. Table 1.10 presents the number of subjects who chose each gamble by

only-child status. Figure 1.3, instead, shows the frequency of gamble choice by only-children and subjects with sibling. The mean gamble choice made by all subjects is 2.67 and the median chosen gamble is gamble 3. Only-children's mean gamble choice is 2.63, while subjects with siblings mean gamble choice is 2.67. A mean test fails to reject the null hypothesis of no differences in mean gamble choice by only-child status (p=0.18, Mann Whitney rank-sum test). Furthermore, a test of proportions fails to reject the null hypothesis that only-children and subjects with siblings are equally likely to choose each one of the alternative gambles (p-values in Table 1.10, test of proportions across all alternative gambles). Results from the risk game suggest that no significant difference exists between only-children's and subjects with siblings' risk preferences. Different levels of risk-aversion seems therefore not to represent a possible explanation for the observed different behaviours between single-children and subjects with siblings. Our results are not in line with the findings in the work by Cameron et al. (2013) that only-children are significantly more risk-averse. Differences in results might be explained by the different experiment procedure. Cameron et al. (2013), in fact, elicit risk attitudes by asking participants to make an investment into an account that can yield, with equal probability, three times the invested amount or zero, and measure risk aversion by looking at the average amount invested. We instead present participants a broader sets of alternative outcomes that enable us to better disentangle across different level of risk.

TABLE 1.10: Gamble choice by only-child

| Gamble choice | Only-children(%) | | Subj. with siblings(%) | | All subjects(%) | | p-values Only-child=Subj. with siblings |
|---|---|---|---|---|---|---|---|
| 1 | 11 | (26.83) | 59 | (25.11) | 70 | (25.36) | 0.820 |
| 2 | 10 | (24.39) | 54 | (22.98) | 64 | (23.19) | 0.843 |
| 3 | 11 | (26.83) | 65 | (27.66) | 76 | (27.54) | 0.912 |
| 4 | 5 | (12.20) | 38 | (16.17) | 43 | (15.58) | 0.517 |
| 6 | 4 | (9.76) | 19 | (8.09) | 23 | (8.33) | 0.720 |
| Total | 41 | | 235 | | 276 | | |

FIGURE 1.3: Frequency of gamble choices by only-child status

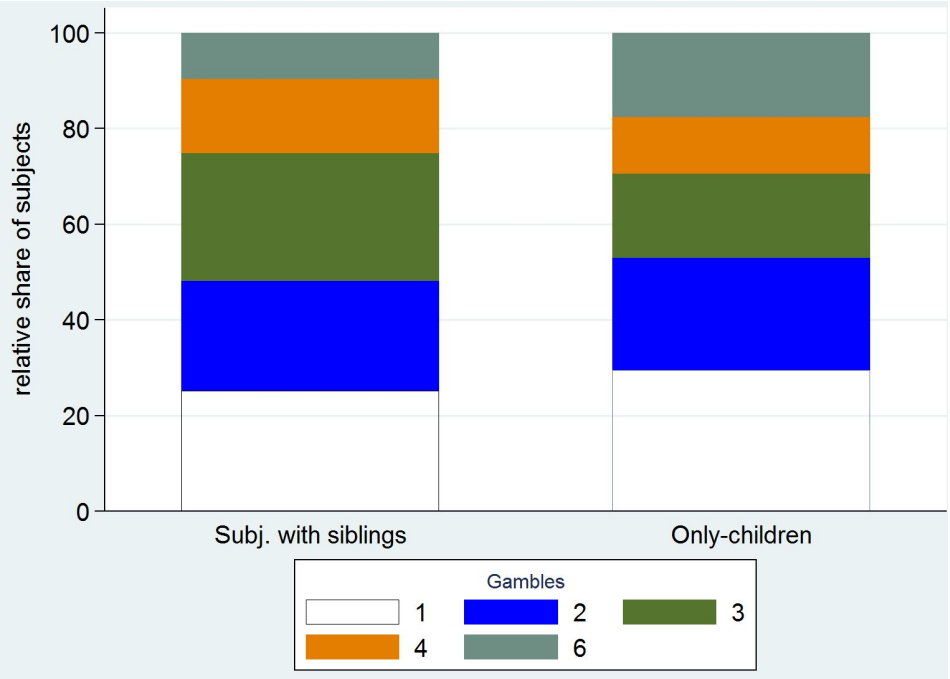


## 1.5 Discussion and Conclusions

The existence of differences in social-preferences, trust and risk preferences between only-children and subjects with siblings has been assessed and documented in economic and psychological studies. However, most of these studies have separately assessed the impact of being an only-child on different aspects of preferences and are mainly based on children. There is instead a lack of studies investigating only-children and subjects with siblings' preferences within a consistent setting and amogst subjects in their adulthood. In this paper we try to address this question and study the effect of being an only-child on social-preferences, trust, risk preferences and strategic behaviours in a coherent framework. Running four different games on the same sample of 276 students, we find evidence of differences in social preferences, trust and cooperation decisions between only-children and subjects with siblings.

Specifically, we observe that only-children and subjects with siblings differ both to the extent of distributional preferences and negative reciprocity. Results from the response game show that only-children are significantly less likely to punish actions that reflect negative intentions while they appear significantly more likely to choose

social welfare maximizing actions. The analysis of behaviours between the two groups in the response game reveals also an interesting link between reciprocity and distributional preferences. Differences in behaviours between the two groups, in fact, become statistically significant when games outcomes involve both distributional and reciprocity concerns. We collect evidence of a first connection between these two elements when, analysing distributional preferences, we observe that only-children are significantly more likely to choose social welfare maximizing actions in games where reciprocity is involved. This evidence is then corroborated by the finding of a significantly stronger effect of equal payoffs in decreasing punishment decisions across only-children. Overall, the evidence from our analysis of behaviours in the response game suggests therefore that only-children's actions are more strongly affected by distributional concerns when reciprocity is involved in the decision environment. However, results are not unambiguous. The difference in social welfare maximizing attitudes between the two groups and the link between reciprocity and distributional concerns disappears when different game structures are considered. No difference in social welfare maximizing actions is observed when looking at behaviours in dictator games, while contrasting results derive from the analysis of decisions made in public goods games. Observed stronger attitudes toward social welfare maximizing actions would suggest higher choices of contribution across only-children when analysing data from the public goods game. Results from the one shot public goods game instead show that only-children are significantly less cooperative than subjects with siblings. However, the difference between only-children and subjects with siblings disappears again when we analyse subjects' behaviours in the different setting of a repeated public goods game, where reciprocity is involved, confirming that differences in behaviours between the two groups could also depend on the specific game structure. Our results corroborate therefore Cameron et al. (2013)'s findings and show that the difference in altruism observed in the literature (Angerer et al., 2015; Fehr, Bernhard, and Rockenbach, 2008) when studying children's behaviour in dictator games disappears when only-children and subjects with siblings' behaviour is instead assessed across adults. We furthermore document a link between only-children's distributional preferences

and reciprocity which might help explaining some contrasting results discussed in the literature. Many existing findings on the effect of growing up as an only-child rely, in fact, on games where very different behavioural motives might play a role. Our results suggest that taking into account correlations amongst different motives might help explaining and understanding these different outcomes.

We also document evidence of difference in trust between single-children and subjects with siblings. We use beliefs stated in the first contribution period of the ten periods repeated public goods game as a measure of trust across the two groups, and show that single-children's beliefs are significantly lower than those stated by subjects with siblings. This finding is also supported by the significantly lower contribution decisions made by only-children in the unconditional contribution task of the public goods game. Pessimistic beliefs of a consistent share of subjects motivated by social preferences can, in fact, reasonably explain lower contribution choices. The lower trust we document by the analysis of decisions made in the public goods game is furthermore in line with results of Cameron et al. (2013) and Glaeser et al. (2000) and it is confirmed by the results we observe in the investment game, which is designed to provide a more direct test for differences in trust across only-children and subjects with siblings. We therefore provide evidence of trust representing a trait that consistently differs between only-children and subjects with siblings' preferences both across games and age samples. Through the results of the public goods game we furthermore show how the observed different trust levels can help to interpret differences in behaviour between the two groups and amongst different game settings.

In addition to these main findings, we use the detailed information about siblings collected through the questionnaire to extend the analysis of differences in preferences between only-children and subjects with siblings to a more specific sibship structure and disentangle across different effects. Lampi and Nordblom (2010), Saroglou and Fiasse (2003) and Beck, Burnet, and Vosper (2006) suggest, in fact, that different number, order and spacing of siblings can have a different impact on moral and social development of an individual. We find that sibship structure differently affects subjects' preferences in different games. In public goods games, for example, the overall

number of siblings increases the average contribution, in the response games, instead, subjects with siblings' higher likelihood to punish is determined by having at least one younger sister and in the trust game the higher choice of amount to transfer is driven by the number of older sisters. Overall, the effect of sibship components is therefore not decisive, but our data seem to suggest that sisters may play a stronger role in determining different outcomes between single-children and subjects with siblings.

# Chapter 2

# History matters: Match length realizations and cooperation in indefinitely repeated games

# History matters: Match length realizations and cooperation in indefinitely repeated games[1]

FRIEDERIKE MENGEL [2] LUDOVICA ORLANDI[3]

University of Essex University of Essex

and Lund University

SIMON WEIDENHOLZER[4]

University of Essex

[1]We thank Maria Bigoni, Anna Dreber, Matt Embrey, Guillaume Frechette, Dan Friedman, Drew Fudenberg, Ed Hopkins, Volodymyr Lugovskyy, Ryan Oprea, Ronald Peeters, Giancarlo Spagnolo and Robert Sugden for helpful comments and for providing data access. We also thank audiences at the LEG 2019 workshop at Bar Ilan, the University of East Anglia, the University of Heidelberg, the SAET 2019 conference, and at the Vienna University of Economics and Business Adminstration for their helpful comments. Financial support by the British Academy (SG162637) is gratefully acknowledged.

[2]Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ *and* Department of Economics, Lund University, Tycho Brahes vag 1, Lund (SE). *e-mail*: fr.mengel@gmail.com.

[3]Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ. *e-mail*: lorlan@essex.ac.uk.

[4]Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ. *e-mail*: sweide@essex.ac.uk.

# *Abstract*

Experimental studies of infinitely repeated games typically consist of several indefinitely repeated games ("matches") played in sequence with different partners each time, whereby the length, i.e. the number of stages, of each game is randomly determined. Using a large meta data set on indefinitely repeated prisoner's dilemma games (Dal Bó and Fréchette, 2018) we demonstrate that the realized match length of early matches has a substantial impact on cooperation rates in subsequent matches. We show that the effect is in line with a class of learning models displaying the "power law of practice". We then study three cases from the literature where realized match length has a strong impact on treatment comparisons, both in terms of the size and the direction of the treatment effect. These results have important implications for our understanding of how people learn in infinitely repeated games as well as for experimental design.

## 2.1 Introduction

Infinitely repeated games are of enormous importance in many areas of Economics, but also in Politics, Sociology, Biology and many other subjects. The theory of infinitely repeated games has delivered key insights on how repeated interaction changes incentives and can sometimes enable a wide array of outcomes including, for example, cooperation in social dilemmas. However, it does not always offer sharp predictions. For instance, in the prisoner dilemma both cooperation and defection are equilibrium actions provided that players are sufficiently patient. The multiplicity of outcomes gives an important role to lab experimental research on infinitely repeated games to narrow down what we can expect empirically in these games. This research has provided key insights on e.g. the determinants of cooperation in social dilemmas, the role of monitoring, or the differences between discrete and continuous time.[1]

Infinitely repeated games are implemented in the lab as "indefinitely repeated games" using a random continuation probability, as originally proposed by Roth and Murnighan (1978). After every round of play, there is a fixed known probability $\delta$ that the game continues for an additional round and a probability $1-\delta$ with which the match ends. A match refers to a supergame (a "repeated game"), and a round (or stage) is one play of the stage game. To allow for learning, experiments typically feature repetitions. Participants play several matches and are rematched in between. The length of each match (in terms of number of rounds) is determined randomly, giving rise to a sequence of match length realizations. For a given termination probability, $\delta$, participants may thus experience different realizations of match length and consequently different sequences of match lengths. This design is unproblematic if either the realized sequences of match length realizations "correctly" represent the infinitely repeated game or if

[1] See Dal Bó and Fréchette (2018) for a review of this extensive literature.

match length realization is irrelevant for behaviour.[2] Achieving the former can be difficult as there are considerable practical difficulties involved in getting a large enough sample of different sequences of match length realizations.[3] If the number of match length realizations is "small", it is difficult to know whether the realized distribution is "representative enough" for the infinitely repeated game under consideration, as this will depend on how people learn and which moments of the distribution are the most critical for their learning. For this reason it is important to understand whether match length realization influences behaviour.

In this paper we first demonstrate that the sequence of match length realizations has a substantial, robust and highly statistically significant effect on behaviour. Using a data set on the infinitely repeated prisoner dilemma compiled by Dal Bó and Fréchette (2018) we show that when subjects initially experienced relatively long matches subsequent cooperation rates are substantially higher. Specifically, when most matches in the first part of an experiment were "long" (above theoretical median length), then cooperation rates are 44% higher in subsequent matches. Intuitively, participants who experience longer matches become more optimistic about the relative benefits of conditionally cooperative strategies and cooperate more. Moreover, by comparing the impact of long matches in the first third to the impact of long matches in the middle third of an experiment on cooperation in the final third, we demonstrate that the effect of early matches is as least as important as the one of recent ones. This observation is consistent with the "power law of practice" which describes the phenomenon of initially steep- and then flattening out- learning curves. We show that fairly simple reinforcement learning models in the spirit of Erev and Roth (1998) are consistent with these findings.

Our results hence show that the environment in which early interactions take place matters for subsequent interactions as people learn from match length realizations. As such our results speak to our understanding of what is often termed "cultural differences". They can help explain why people growing up in different social backgrounds

[2] In standard theory only expected match length (captured by $\delta$) should matter for behavior. Hence, according to standard theory match length realization should indeed be irrelevant for behaviour.

[3] We discuss some of these difficulties in detail in Section 2.4.

(characterized by more or less stable interactions) or people coming from different work environments (characterized by more or less turnover) might show different behaviour in the exact same situation. Our results also show that we would not expect these groups to behave differently in *all* situation. If the environment is very unfavorable to cooperation we would not expect a behaviour difference between these groups based on our results in this paper. The results also have implications for evidence based policy making. If a policy (e.g. designed to increase cooperative behavior) is evaluated over a certain fixed period, it is possible that the results of the evaluation are affected by early match length realizations even if they are exogenous to the policy evaluated.

Our findings do not only yield insights into how people learn in indefinitely repeated games, they also have important methodological implications for experimental design. The length of each match is typically drawn at the session level, meaning that all subjects in a given session experience the same sequences of match length draws. In fact, all papers in the Dal Bó and Fréchette (2018) meta study use this or a very similar design. The number of different sequences of match length realizations for a given treatment ranges between 1 and 10 across the different papers contained in the meta-study. Given our results discussed above we would expect that - with such small numbers of match length realizations - treatment comparisons can be affected. We provide three case studies of papers from the existing literature, which were not part of the Dal Bó and Fréchette (2018) meta study: a continuous time prisoner's dilemma (Bigoni et al., 2015), a public good game (Lugovskyy et al., 2017) and oligopoly games (Embrey, Mengel, and Peeters, 2019). We show that - for each of them - treatment effects differ depending on match length realization. We also run our own experiments and show that in some cases the conclusions drawn from the research might have been different for different match length realizations.

Our paper contributes to a substantial and active literature on indefinitely repeated games, much of it summarized by Dal Bó and Fréchette (2018). Several researchers have documented a positive effect of the length of the immediately preceding match on cooperation (see e.g. Camera and Casari 2009, Dal Bó and Fréchette 2011; 2018,

Fréchette and Yuksel 2017, Bernard, Fanning, and Yuksel 2018).[4] In the context of infinitely repeated trust games, Engle-Warnick and Slonim (2006) find some evidence that there is more trust and trustworthiness in sessions that initially featured long matches as compared to sessions starting out with short ones. As they observe, this gap could have been due to individual subject or session effects since there was already more trust and trustworthiness in the beginning of the initially long sessions. We add to this literature by providing the first comprehensive analysis of the long lasting effects of (the entire sequence of) match length realization on cooperation in infinitely repeated social dilemmas. To the best of our knowledge our paper is also the first to discuss in detail the potential implications of these findings for measuring cooperation levels and for experimental design.

The paper is organized as follows. In Section 2.2 we demonstrate the main empirical finding of an effect of match length realization on cooperation. We also introduce the learning models and show that it can be explained by learning. Section 2.3 contains our discussion of the case-studies and Section 2.4 concludes. Additional theory, tables, figures and information on our own experiments can be found in an Appendix.

## 2.2 The Effect of Match Length Realizations

### 2.2.1 The prisoner's dilemma

We consider agents who play a $2 \times 2$ indefinitely repeated prisoner's dilemma like the one illustrated in the left panel of Figure 2.1. Payoffs satisfy $\mathtt{T} > \mathtt{R} > \mathtt{P} > \mathtt{S}$ and $\mathtt{T} + \mathtt{S} < 2\mathtt{R}$ such that mutual defection is the only Nash equilibrium of the stage game but mutual cooperation maximizes joint payoffs. Following Dal Bó and Fréchette (2018) we can normalize payoffs so that we only have two parameters, see middle panel of figure 2.1. The continuation probability $\delta$ indicates the probability with which the game continues for one more round. The number of stages in the indefinitely repeated game is

[4] A similar positive effect is document for the behaviour of the previous opponent, in the sense that subjects are more likely to cooperate when they have been previously matched with somebody starting out with cooperation.

hence a random variable $T$. It is common in modern experiments to play several such indefinitely repeated games. Usually participants are rematched at the end of one repeated game and play a new game with a new partner. Each such repeated game is often referred to as a "match". Typical experiments differ in the number $M$ of such matches implemented, the expected length of a match (given by $\mathbb{E}[T] = \frac{1}{1-\delta}$) as well as the realized match length. We index the round of play within a match by $t$ and the match by $m$. $T^m$ is the realized match length of match $m$, i.e. the number of stages in match $m$.

| | C | D |
|---|---|---|
| C | R | S |
| D | T | P |

PD game.

| | C | D |
|---|---|---|
| C | $\frac{R-P}{R-P} = 1$ | $\frac{S-P}{R-P} = -\ell$ |
| D | $\frac{T-P}{R-P} = 1+g$ | $\frac{P-P}{R-P} = 0$ |

Normalized game.

| | GT | AD |
|---|---|---|
| GT | $\mathbb{E}[T]$ | $-\ell$ |
| AD | $1+g$ | 0 |

GT vs. AD game.

FIGURE 2.1: Left: Prisoner's dilemma (PD) game with payoff parameters $T > R > P > S$ and $T + S < 2R$. Middle: Normalized game where joint defection payoff $P$ is subtracted from each cell and all payoffs are divided by $R - P$ (difference between mutual cooperation and defection payoffs). Right: Payoffs in the game induced by Grim-Trigger and Always Defect. GT played against GT yields a payoff of 1 in all $\mathbb{E}[T]$ stages. AD (GT) played against GT (AD) yields once a payoff of $1 + g$ ($-l$) and zero in the remaining $\mathbb{E}[T] - 1$ stages.

A substantial experimental literature has studied how payoff parameters affect cooperation in the prisoner's dilemma.[5] One particularly successful approach, proposed by Blonski and Spagnolo (2015), analyzes a setting where agents can only choose among the repeated game strategies "Grim-Trigger" (GT) and "Always Defect" (AD) and payoffs are given by the expected sum of payoffs of the induced indefinitely repeated game shown in the right panel of Figure 2.1 (see also Blonski, Ockenfels, and Spagnolo (2011) and Dal Bó and Fréchette (2011)). Provided "Grim-Trigger" can sustain cooperation in a subgame perfect Nash equilibrium ($\mathbb{E}[T] \geq 1 + g$), the resulting

[5] See e.g. Embrey, Fréchette, and Yuksel (2017) and Mengel (2018) for contributions analyzing finitely repeated PD games and Dal Bó and Fréchette (2018) for a survey of the literature on the indefinitely repeated version.

game constitutes a coordination game. The size of the basin of attraction of AD, denoted by SizeBAD, is defined as the threshold probability of choosing GT that has to be exceeded to make GT a best response.[6] Formally,

$$\texttt{SizeBAD} = \begin{cases} 1 & \text{if } \mathbb{E}[T] < 1 + g \\ \frac{\ell}{\mathbb{E}[T]+\ell-g-1} & \text{otherwise} \end{cases} \tag{2.1}$$

Note that SizeBAD is decreasing in $\mathbb{E}[T]$ (respectively $\delta$), conveying the intuitive idea that cooperation is easier to sustain under longer expected match durations. Dal Bó and Fréchette (2018) show that SizeBAD indeed predicts cooperation rates very well in a meta-study of indefinitely repeated prisoner's dilemma experiments. They also show that the length of the immediately preceding match has an effect on cooperation rates in the subsequent match. They suggest that this is either due to a minority of participants who may not understand how match lengths are determined or due to how participants update their overall evaluation of the value of cooperation through experience. They write "there is an interesting - as yet unexplored - question regarding the way that humans learn in infinitely repeated games. Is the impact of the realized length constant throughout or is the impact more important early on?"

The remainder of this section is dedicated to answering this question. Doing so will not only yield valuable insights into how people learn, but it will also make an important methodological point regarding how indefinitely repeated games should be implemented in the lab. As we will see below not only does the length of the immediately preceding match matter, but the entire sequence of match length realizations is important. Further, addressing the question posed by Dal Bó and Fréchette (2018), the impact of realized match length is not constant throughout. Early matches matter at least as much as later matches and sometimes more. We will now demonstrate these patterns empirically (Section 2.2.2) and then show that they are in line with several simple learning models (Section 2.2.3).

[6] This corresponds to the probability of *AD* in the mixed strategy equilibrium of the game in the right panel of Figure 2.1.

### 2.2.2 Empirical Analysis

To study empirically whether there exists a persistent effect of match length in early matches we use the data collected by Dal Bó and Fréchette (2018). They use data from 141 different sessions of indefinitely repeated prisoner's dilemma experiments with 2415 participants (see Table 3 in Dal Bó and Fréchette (2018)). In all papers contained in their data set the sequences of match length realizations is drawn at session level, i.e. all subjects in a given session faced the same sequence of match length realizations. Figure 2.2 shows the distribution of the difference between theoretical median match length and realized match length in the meta-study.[7] The left panel aggregates games with different discount factors. It can be seen that match lengths are, as expected, concentrated around the median with a good amount of variation on both sides. The right panel shows separate graphs for the three most common discount factors $\delta = 0.5$, 0.75 and $\delta = 0.9$. The figure shows that for the longer games with $\delta = 0.75$ and $\delta = 0.9$ most matches are somewhat shorter than what we would expect. However, in all cases, there is a good amount of variation.

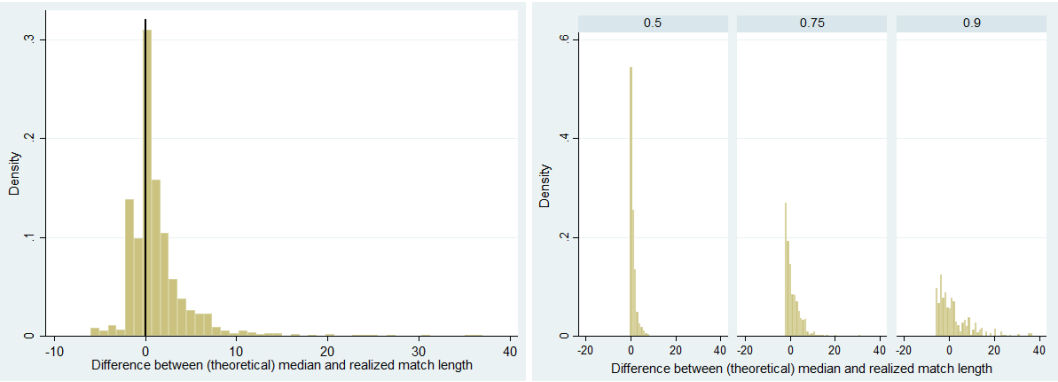


FIGURE 2.2: Distribution of the difference between theoretical median match length and realized match length overall (left panel) and separately for $\delta = 0.5, 0.75$ and $\delta = 0.9$.

We use this variation to study how match length realization in early matches affects subsequent cooperation. We define early matches as the 1st third of matches in a session and create a dummy variable $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ out of these

[7] Appendix Figure F.1 shows kernel density estimates.

early matches were (weakly) longer than the theoretical median length.[8] The dummy takes the value 1 in 44% of sessions. Analogously, we can also define dummies $\Delta^{2nd}_{above}$ and $\Delta^{3rd}_{above}$ which take the value 1 in 42% and 47% of sessions, respectively.

| | *Effect of Match Length Realization on subsequent cooperation* | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\Delta^{1st}_{above}$ | 0.142*** | 0.101** | 0.226*** | 0.125** | 0.101** | 0.207*** |
| | (0.053) | (0.039) | (0.056) | (0.058) | (0.041) | (0.062) |
| $\Delta^{2nd}_{above}$ | | | | 0.085 | 0.069* | 0.032 |
| | | | | (0.060) | (0.039) | (0.068) |
| `SizeBAD` | | -0.765*** | -0.539*** | | -0.809*** | -0.659*** |
| | | (0.069) | (0.094) | | (0.068) | (0.149) |
| `SizeBAD` $\times$ $\Delta^{1st}_{above}$ | | | -0.296*** | | | -0.241** |
| | | | (0.094) | | | (0.104) |
| `SizeBAD` $\times$ $\Delta^{2nd}_{above}$ | | | | | | 0.048 |
| | | | | | | (0.144) |
| Constant | 0.321*** | 0.974*** | 0.747*** | 0.294*** | 0.994*** | 0.844*** |
| | (0.028) | (0.074) | (0.098) | (0.032) | (0.074) | (0.153) |
| $\delta$ f.e. | NO | YES | YES | NO | YES | YES |
| Test $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$ | - | - | - | 0.6903 | 0.6063 | 0.0989 |
| Observations | 34,319 | 34,319 | 34,319 | 18,536 | 18,536 | 18,536 |
| R-squared | 0.021 | 0.219 | 0.223 | 0.034 | 0.251 | 0.255 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE 2.1: Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level. Observations stem from 141 sessions spread across 15 papers.

Table 2.1 shows the results of regressing cooperation in the final third of matches on $\Delta^{1st}_{above}$ as well as `SizeBAD` and an interaction. Early match length has a substantial impact on cooperation in later matches. If at least $\frac{2}{3}$ of these early matches are "long", then

[8] The reasoning behind these choices is the following. We split matches in three groups (early, middle and late) rather than e.g. two is that it allows us to compare the effect of early (1st third) and middle (2nd third) matches on cooperation in late (3rd third) matches. This allows to address the question whether early experience or recent experience is more important for cooperation. Appendix Table E.1 shows results for alternative splits. The reason we use a dummy is that (i) theoretical medians differ with $\delta$, which means that we cannot just use match length directly, and that (ii) it makes regression results more easily interpretable. Appendix Table E.5 shows results when we use the share of matches above median instead. Last, the reason that we use $\frac{2}{3}$ as a cutoff for the share of long matches is that it produces relatively balanced groups, though some other cutoffs would have produced that too. Appendix Tables E.6-E.7 show the results with alternative cutoffs.

cooperation rates are higher for the remainder of the experiment as shown by the positive coefficient on the dummy $\Delta^{1st}_{above}$ in column (1). The effect size is substantial, with cooperation rates being 44% higher when initial matches were long as compared to when they were short. As expected, the table also shows a negative impact of `SizeBAD` on cooperation rates (columns (2)-(3)). Interestingly, there is also an interaction effect between `SizeBAD` and $\Delta^{1st}_{above}$. If early matches are long, then the detrimental effect of `SizeBAD` is more pronounced. This is intuitive as longer early matches could allow participants to better learn the incentives coming from the game parameters. Conversely, the interaction term also shows that the positive effect of early match length realization is stronger the more favorable the climate is for cooperation. In fact according to Table 2.1, the effect is positive if and only if `sizeBAD` is smaller than 0.77.

**Robustness** Appendix E.1 contains tables showing that these results are qualitatively robust to the inclusion of paper fixed effects (Table E.4), to considering different thresholds (Tables E.6-E.8) or to using the share of matches above median instead of a dummy variable (Table E.5). A placebo test shown in Appendix Table E.3 where we regress cooperation in the 1st third of matches on $\Delta^{3rd}_{above}$ shows that the results in Table 2.1 are fundamental and not e.g. driven by correlations of match lengths within sessions or observed or unobserved heterogeneity across papers or treatments, e.g. caused by different ways researchers implement match length draws.[9]

**Early vs Recent Matches** Next we ask what is more important for cooperation in the final third of the experiment, early experience, i.e. match length in the 1st third, or recent experience, i.e. match length in the 2nd third of matches? Columns (4)-(6) in Table 2.1 show the results of regressing cooperation in the final third of matches on both dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. The table shows very clearly that early experience in the 1st third of matches is very important. In all specifications the coefficient on $\Delta^{1st}_{above}$ is at least as large as that for $\Delta^{2nd}_{above}$ and exhibits a higher level of statistical significance. The

[9] We would not expect realized match length of final matches, which have not yet been played, to affect cooperation in the beginning of the experiment. Hence we would expect zero coefficients on $\Delta^{3rd}_{above}$ and the corresponding interaction term. We do indeed find that these coefficients are close to zero and statistically not significant.

interaction effect with SizeBAD is also more important for these matches. Early matches seem at least as important as recent matches and potentially, more important.[10]

**Experience** Does the effect vanish with more experience, i.e. if enough matches are played in the experiment? To answer this question we rerun specification (1) of Table 2.1 restricting the sample to sessions with (i) at most 12 matches in total, (ii) 12-24 matches, (iii) 24-36 matches etc.[11] Appendix Figure F.2 shows that a positive effect size can be found even in sessions that feature at least 72 matches in a session. The figure also shows a possible downward trend in coefficient sizes as more matches are played, but if at all the trend is slow and suggests that at the very least 80 matches would have to be played in a session for the coefficient to vanish. This can quickly become infeasible especially if the discount factor $\delta$ is high. Note also that there is a compositional effect in this analysis as sessions with more matches tend to have smaller $\delta$ in the meta-study (t-test, $p < 0.0001$). As with the geometric distribution we would expect more extreme outliers in match length realizations when $\delta$ is higher, the compositional effect should artificially exacerbate the effect of experience, i.e. make it seem that with more matches there is less of an effect of match length realization. That we see very little in terms of a downward trend despite this suggests that adding more matches will not easily eliminate the impact of early match length realizations.

To sum up, the results in this section have shown that there can be substantial and non-trivial effects of realized early match length on cooperation rates in the rest of the experiment. Hence, which match length realizations are drawn can potentially affect research results. This is particularly likely if few draws are made (e.g. only one draw per session or treatment). In Section 2.3 we will study three case studies highlighting this point.

[10] Appendix Table E.2 compares the importance of early and recent matches for more different splits. Specifically the table compares the impact of match length realization in the first X-th, second X-th, third X-th...of matches on cooperation in the last X-th of matches, where X ranges from 2,...,10. For all X=2,...,9 the coefficient on the first X-th of matches is larger than that of the (X-1)th Xth of matches.

[11] We choose multiples of 12 to cut the sample as (i) they are close to the 25th, 50th and 75th percentile of match numbers in the overall sample (25th percentile is 11, 50th is 23 and 75th is 34) and (ii) 12 divides by 2, 3 and 4 without remainder allowing us to split the total number of matches in halfs, thirds and quarters as in Appendix Figure F.2.

We have also seen that early matches matter at least as much as recent matches. This is in line with a substantial body of evidence on both human and animal learning which shows that learning curves tend to be steeper initially and then flatter. This observation is known as "power law of practice" and according to Erev and Roth (1998) dates back to at least Blackburn (1936). In the next subsection we will outline a class of learning models that implies this property and show that the results above can be explained by such a learning model.

### 2.2.3 Learning

To study how match lengths can affect learning we discuss several simple learning models which display the "power law of practice". These are straightforward adaptations of previously studied models to an environment where payoffs depend on stochastic realizations of match length. To this end, we consider a set of agents which are recurrently matched to play a series of indefinitely repeated PD games. Following much of the literature (Dal Bó and Fréchette, 2018; Embrey, Fréchette, and Yuksel, 2017), we restrict attention to the $GT$ and $AD$ strategies.[12] The payoffs of the game induced by these strategies are given in the right panel of figure 2.1, where the expected match length $\mathbb{E}[T]$ is now replaced by the actual match length realization $T^m$. In the learning models the choices of agents are determined by propensities which are updated after each match. Propensities can be interpreted as beliefs (as in fictitious play, see e.g. Mookherjee and Sopher 1997) but can also incorporate a much wider set of feelings, such as e.g. familiarity or habituation (as in reinforcement learning, see e.g. Erev and Roth 1998, Boergers and Sarin 1997). Each agent $i$ is endowed with an initial propensity for each strategy, denoted by $\psi_s^{i,0}$ for strategies $s \in \{GT, AD\}$, which may capture pre-game experience, initial prepositions as well as initial beliefs. The models we study differ only in how these propensities are updated after each match.

Under *reinforcement learning without counterfactuals* (see e.g. Erev and Roth (1998) and Roth and Erev (1995)) agents increase the propensity of the strategy chosen by the

[12] A theoretical justification for why it is sensible to restrict to these strategies is provided in Blonski and Spagnolo (2015) and Blonski, Ockenfels, and Spagnolo (2011).

payoff received, i.e.

$$\psi_s^{i,m+1} = \psi_s^{i,m} + \mathbb{1}(s^{i,m} = s)\pi(s, s^{-i,m}, T^m),$$

where $s^{-i,m}$ denotes the strategy of $i$'s opponent in match $m$, $\mathbb{1}(s^{i,m} = s)$ indicates whether agent $i$ uses strategy $s$ in match $m$ or not, and $\pi(s, s^{-i,m}, T^m)$ gives the payoff earned with strategy $s$ in this case. Note that the propensity for strategies not chosen does not change. This can be seen as a pretty simplistic way of learning which is solely driven by ones' own experience. Match length realizations influence learning through their role in determining payoffs in the game. Note that this form of reinforcement learning features what Erev and Roth (1998) call "force of habit" where frequently chosen actions are reinforced more frequently. This is not the case under *reinforcement learning with counterfactuals* (see e.g. Vriend 1997 and Rustichini 1999 and the special cases in Erev and Roth 1998 and Camerer and Ho 1999) where also strategies that have not been played are reinforced and propensities for all strategies evolve according to

$$\psi_s^{i,m+1} = \psi_s^{i,m} + \pi(s, s^{-i,m}, T^m).$$

In environments were payoffs are stable, in the sense that they do not feature an exogenous stochastic element, reinforcement learning with counterfactuals is very closely related to (smooth) fictitious play (see e.g. Fudenberg and Kreps 1993 and Fudenberg and Levine 1998) where agents play a (smooth) best response to the belief that future play will follow the past empirical distribution (see e.g. Cheung and Friedman 1997 and Camerer and Ho 1999). The equivalence holds because looking back to previous earnings of strategies is equivalent to forming beliefs based on past behaviour and then computing expected payoffs based on these beliefs. In order to specify (smooth) *fictitious play* in the present environment we need to specify beliefs about play of the others as well as beliefs about match length realizations. Here we show a version where agent $i$ simply uses the average previous match length realization, given by $\bar{T}^m = \frac{1}{m}\sum_{k=1}^{m} T^k$, and the share of her opponents previously choosing grim trigger,

given by $\sigma^{-i,m} = \frac{1}{m}\sum_{k=1}^{m} \mathbb{1}(s^{-i,k} = GT)$, where $\mathbb{1}(s^{-i,k} = GT)$ equals one if agent $i$'s opponent had chosen $GT$ in period $k$ and is zero otherwise. We adopt a propensity based formulation of fictitious play as in Camerer and Ho (1999) and Hopkins (2002) where propensities are simply given by the expected payoffs $\hat{\pi}$ given these beliefs, i.e.

$$\psi_s^{i,m+1} = \hat{\pi}(s, \sigma^{-i,m}, \bar{T}^m).$$

It remains to specify how these propensities translate into choices. For all three models we use Luce's linear probability choice rule where choice probabilities are linearly proportional to propensities.[13] More formally agent $i$'s probability to choose GT in match $m$ is given by

$$p_{GT}^{i,m} = \frac{\psi_{GT}^{i,m}}{\psi_{GT}^{i,m} + \psi_{AD}^{i,m}} \tag{2.2}$$

and AD is chosen with the remaining probability, i.e. $p_{AD}^{i,m} = 1 - p_{GT}^{i,m}$.

To quantify the effect of match length realizations on cooperation rates in these models, we have run several simulations for each of these three learning models, using key game parameters from the meta dataset of Dal Bó and Fréchette (2018)'s data. The payoff parameters of the underlying PD games were chosen to correspond to the 25th, 50th and 75th percentile of the distribution of `SizeBAD`. In each simulation run, 16 agents were matched against each other in a round robin tournament and match length realizations were drawn at the session level using the discount factor $\delta = 0.75$, thus targeting the median group size and the most frequent discount factor in Dal Bó and Fréchette (2018). Our exercise contains 4000 simulated experiments played over 15 matches for each of the three payoff configuration leading to 960,000 indefinitely repeated PD games for each of the three learning models. More details can be found in Appendix C.2.

For each of these three datasets we run regressions identical to the one shown in column (1) of Table 2.1, thus measuring the impact of long match length realizations

[13] This is a special case of the power probability form and is the most frequently used in the reinforcement learning literature. Under fictitious play the exponential or logit choice rule is more commonly used. Its main advantage lies in being able to be directly deduced from the maximization of an underlying perturbed payoff function (see e.g. Fudenberg and Levine (1998)).

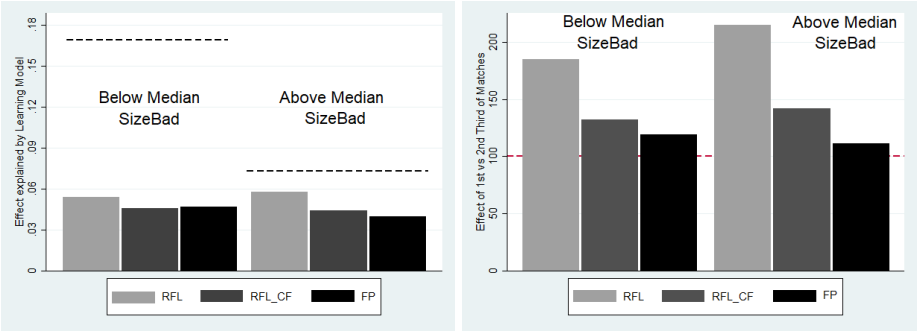


FIGURE 2.3: Left Panel: Effect Size for $\Delta^{1st}_{above}$ in simulated data for the three models (bars) compared to effect size in Dal Bó and Fréchette (2018) data (dashed line) in two samples: below median values of SizeBAD (left three bars and thick dashed line) and above median values of SizeBAD (right three bars and thick dashed line). Right Panel: Relative Effect size of $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$ in simulated data for the three models (bars) in two samples: below median values of SizeBAD (left three bars) and above median values of SizeBAD (right three bars). The dashed line indicates $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$.

in the first third of matches on initial cooperation rates in the remainder of matches. Standard errors are clustered at the run level. We then compare the effect size obtained in these regressions to the empirically observed one. As we have seen that there is an interaction effect between SizeBAD and the effect of early match length realization (Table 2.1) we also split the sample in below and above median SizeBAD.

The left panel of Figure 2.3 shows the results. Reinforcement learning (without counterfactuals) shows a slightly larger effect compared to the other models.[14] For all models the effect size does not differ across the two subsamples of low and high SizeBAD ($p > 0.1$). This differs from the human data, where for small values of SizeBAD early match length realization has a much larger effect size, though the difference is just outside of conventional levels of statistical significance ($p = 0.1753$). In this sample the effect size obtained purely through the learning models is about a third of the overall effect size. In the sample where defection is relatively attractive (high SizeBAD), the simulated effect size is about 80% of the empirical effect size with human players.

The right panel of Figure 2.3 shows the ratio between $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$ for our three learning models. It can be seen that this ratio is always above 100%, i.e. $\Delta^{1st}_{above} > \Delta^{2nd}_{above}$. This illustrates the consequences of the "power law of practice", which implies that

[14] Given the substantial sample size in the simulations all the effects obtained are highly statistically significant ($p < 0.0001$).

early realizations of match lengths may be more important for cooperation rates than later experience.[15] To understand why this is the case,[16] note that in all models the probability of choosing $GT$ and $AD$ are determined by the ratio of their propensities. In the case of the reinforcement learning models these propensities are generally increasing over time and consequently higher in later matches. Thus, the impact of a long match and of the associated high payoffs on the choice probabilities will be larger early on when propensities are small than later when they are large (keeping the ratio of propensities and the play of the opponent constant). Thus, whatever happens early on shifts the pattern of play in a more pronounced way than later experiences.[17]

We have seen that all three learning models can explain the direction of the empirically observed effect of match length realization. However, we have also seen that effect sizes are larger empirically than what we would expect from the models. This is particularly the case when sizeBAD is small, i.e. when cooperation is relatively attractive. One possible explanation why effect sizes are larger with human players is that they may stop learning after some rounds, while the simulated learning models keep learning. Hence for the simulated learners the effect of initially long matches is (partially) corrected when later matches are shorter.

## 2.3 Case Studies

We will discuss three applications to illustrate how match length realizations can affect treatment comparisons when indefinitely repeated games are compared with finitely repeated games (subsections 2.3.1 and 2.3.2) or when indefinitely repeated games are compared with other indefinitely repeated games (subsection 2.3.3). The three cases highlighted are not part of the Dal Bó and Fréchette (2018) meta-study data used in

[15] This figure does not show effect sizes from human players as they are much larger than the scale of the figure if sizeBAD is small (below median). In this case the estimated effect size of $\Delta^{1st}_{above}$ is $0.169^{***}$ and for $\Delta^{2nd}_{above}$ it is 0.001. When sizeBAD is above median these effect sizes are much closer.

[16] We explore these arguments more formally in the appendix.

[17] Appendix Tables E.9 and E.10 contain the corresponding regression tables and show that indeed early matches are more important than recent matches in these models. There is no statistically significant difference, though, in the importance of early and more recent matches in the fictitious play model (Table E.11).

Section 2.2.2 and feature a continuous time prisoner's dilemma (subsection 2.3.1), a public good game (2.3.2) and oligopoly games (2.3.3).[18]

### 2.3.1 Cooperation in Continuous Time

Our first case study is the paper "Time Horizon and Cooperation in Continuous Time" by Bigoni et al. (2015) published in *Econometrica*. Bigoni et al. (2015) compare cooperation rates in a prisoner's dilemma played in deterministic and stochastic continuous time.[19] They consider games of short (20 seconds) and long (60 seconds) expected length, where here we focus on the short games (which is where they find a treatment effect). The deterministic short game lasts 20 seconds. The stochastic short game has a continuation probability of $\delta = \frac{992}{1000}$ and every 0.16 seconds it ends with probability $1 - \delta$. This means that the expected match length in the continuous game is 20 seconds just as in the deterministic game. The expected median length is 13.86 seconds. Bigoni et al. (2015) focus on average cooperation rates in a match. They find that in short games cooperation is higher under deterministic than under the stochastic horizon.

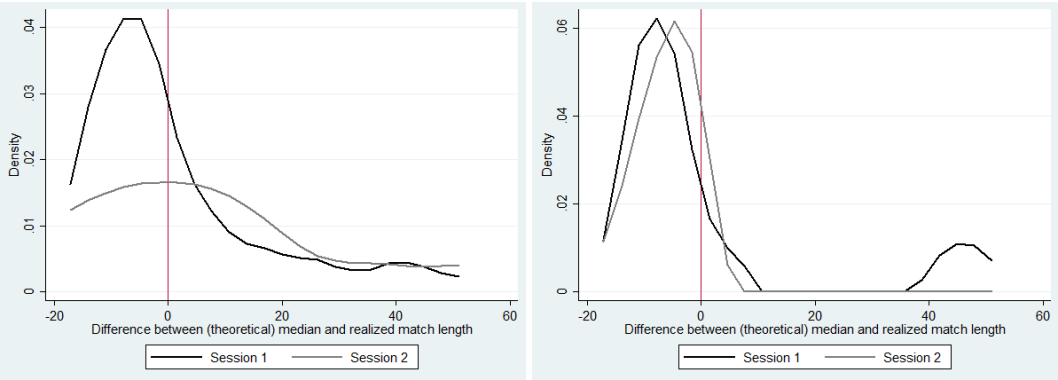


FIGURE 2.4: Kernel density estimates of the difference between theoretical median match length and realized match length overall (left panel) and for the 1st third of matches.

[18] Our selection of case studies followed four criteria: (i) the paper should *not* be already included in the meta-study used in Section 2.2; (ii) it should be on an indefinitely repeated social dilemma; (iii) it has to feature different match length realizations across sessions and (iv) data are publicly available or were made available to us.

[19] This important research programm combines elements from Dal Bó (2005) studying the role of deterministic vs. stochastic horizon in discrete time and Friedman and Oprea's (2012) study of discrete vs. continuous time under a deterministic horizon.

We now study how this result might be affected by match length realizations. Bigoni et al. (2015) conduct two sessions for each treatment condition. In each session there are 24 participants who play 23 matches. Match length is drawn at the session level, i.e. all participants face the same sequence of match lengths. Figure 2.4 shows kernel density estimates of the difference between theoretical median match length and realized match length for the two sessions. The left panel shows the entire session and the right panel only the 1st third of the experiment, specifically the first 8 matches (out of 23). It can be seen that in both sessions the vast majority of matches at the beginning of the experiment (right panel) were shorter than theoretical median length.

| | *S* | N | *Avg ML* | *Avg ML 1st third* | *Median ML* | *Median ML 1st third* | *Avg Coop* | *Avg Coop Initial* |
|---|---|---|---|---|---|---|---|---|
| Deterministic | 4 | 2208 | 20 | 20 | 20 | 20 | 54.04 | 73.95 |
| Replication | 2 | 1104 | 22.94 | 17.6 | 11.04 | 8.48 | 39.58*** | 51.90*** |
| Inverse | 2 | 1104 | 17.86 | 24.0 | 17.56 | 21.66 | 50.90 | 73.36 |
| Match Stoch | 4 | 2208 | 19.97 | 19.57 | 13.44 | 13.52 | 47.40 | 61.18** |

TABLE 2.2: Summary Statistics of the different treatments conducted to replicate Bigoni et al. (2015) . Number of Sessions (S) and observations (N) in the different conditions. Average Match Length (Avg ML), average match length in the 1st third of the experiment (Avg ML 1st third), median match length and median match length in the 1st third, average cooperation rate (avg coop) and average initial cooperation rate (avg coop initial). Stars indicate statistical significance (*** 1%, ** 5%, * 10 %) of the difference to the deterministic case in random effects OLS regression with standard errors clustered at session level (see Appendix Tables E.12 and E.13).

To study whether this realization of match lengths could have affected the treatment effect we first replicated Bigoni et al. (2015)'s experiment. We conducted four sessions of the deterministic condition and then two sessions with the same match-length realizations as Bigoni et al. (2015) ("Replication"). Those sessions were conducted as exact replications of their study. See Appendix D for further details. We further conducted two sessions with inverse match length realizations ("Inverse"). For the inverse sessions we determined a sequence of match lengths $(T^m)_{m=1}^{23}$ as follows. For each realized match length $T^m$ in the Replication we compute $\Pr(x \leq T^m)$ and then replace the $m$-th entry in the sequence by the value $T'$ that satisfies $\Pr(x \leq T') = 1 - \Pr(x \leq T^m)$.

Appendix Figure F.3 illustrates how the "inverse" match length sequences are constructed. Last, we conducted 4 sessions where we randomize the sequence of match lengths at the match level ("Match Stoch"). Hence in this treatment we have 96 different realized match length sequences as opposed to just two.

Table 2.2 gives an overview of the different treatments we conducted as well as the average and median match lengths. The table shows that - compared to the deterministic case - both average and median match length are short in the replication treatments, particularly in the 1st third of the experiment. There the median match length is only 8.48 seconds, much shorter than the 20 seconds in the deterministic case or than the theoretical median of 13.86 seconds. In the inverse condition these match length realizations are naturally longer with the median match length in the 1st third being 21.66, just above the deterministic condition. Last, as expected, when match lengths are drawn at the match level, then, by the law of large numbers, both average and median lengths are close to the theoretical averages and medians.

How does match length realization affect average cooperation rates and the treatment comparison? First, it should be noted that we manage to replicate Bigoni et al. (2015)'s result quite closely. Between the deterministic and replication treatment there is a 14.46 percentage point difference in average cooperation (Table 2.2) compared to Bigoni et al. (2015)'s 10.9 percentage point difference (Table II in Bigoni et al. (2015)). We fail to replicate the result, though, when we use inverse match lengths. Here the difference in average cooperation rates to the deterministic case is only 3.14 percentage points and not statistically different from zero. With match level draws (Match Stoch) we find a difference to the deterministic case of 6.64 percentage points which is less than half of the effect size than in the replication, but more than twice the effect size of the inverse condition. The difference between the match stochastic condition and the deterministic case is not statistically significant at the 10% level.[20] Having a closer look at the data, we do find, however, that average *initial* cooperation rates (in the first stage of each match) do differ significantly between the Match Stoch and deterministic

[20] In this treatment we use twice the number of observations as in Bigoni et al. (2015). It is still possible that this effect becomes statistically significant with a larger sample size.

environments with an effect size of about half of that found in the replication.

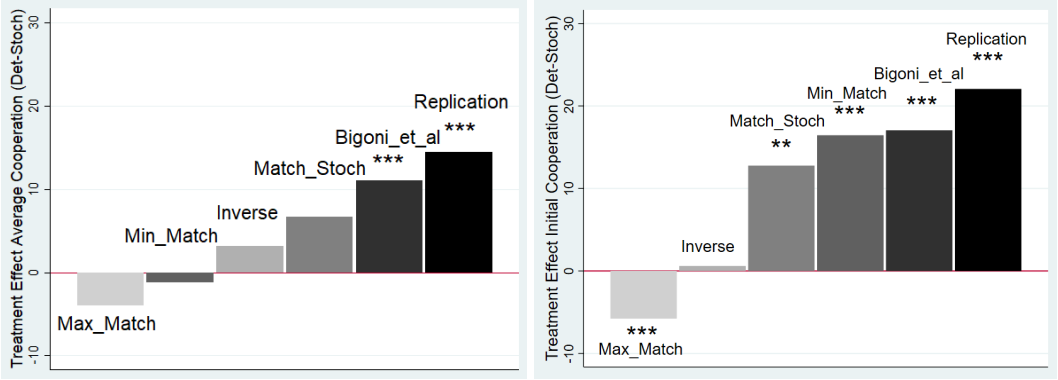


FIGURE 2.5: Different effect sizes obtained for the treatment differences between deterministic and stochastic treatments in average cooperation rates (left panel) and initial cooperation rates (right panel) depending on match length realizations. Bigoni et al: original effect size in Bigoni et al. (2015); Replication: replication treatment; Inverse: inverse treatment; MatchStoch: treatment with match length realization drawn at match level; MinMatch: only individuals from MatchStoch who had the smallest number of matches 1-7 above median (specifically 1) in this treatment; MaxMatch: only individuals from MatchStoch who had the largest number of matches 1-7 above median (specifically 6) in this treatment;

Figure 2.5 illustrates the different effect sizes that can be obtained for the comparison between the deterministic and stochastic game depending on match length realizations. The largest effect size is obtained in our replication of Bigoni et al. (2015)'s original study where we use the same match length realizations as them. This is true for both average cooperation rates (left panel) and initial cooperation rates (right panel). It can also be seen, though, that the treatment difference for average cooperation rates is not statistically significant for any other match length realization. For initial cooperation rates the treatment difference is smaller in the treatment with match-level randomization and statistically not different from zero for the inverse treatment. We also analyzed two sub-groups from the match-stochastic treatments: those with the smallest and those with the largest share of early matches with above median length. For the latter (MaxMatch) we even find a statistically significant negative treatment effect, specifically higher initial cooperation rates in the stochastic game. This exercise illustrates how treatment comparisons can yield entirely different conclusions depending on match length realization. We should also note that - despite the fact that treatment

effects can be strongly impacted by match length realizations - we do not consider this an overall unsuccessful replication. Our exact replication was very successful in the treatment with "many match length realizations" (MatchStoch) the direction of the effect always goes in the same direction as in the original study.

We will get back to the question of how to measure the "correct" treatment effect in Section 2.4. Before we do so, we study two more applications showing how treatment comparisons can be affected by match length realizations.

### 2.3.2 Finite and Indefinitely Repeated Linear Public Good Games

Our second case study is the paper "An experimental study of finitely and infinitely repeated linear public goods games" by Lugovskyy et al. (2017) published in *Games and Economic Behavior*. The finitely repeated games they study all have a match length of 5 rounds, while for the indefinitely repeated public good games they draw three sequences of match lengths (using discount factor $\delta = 0.8$). Average match length in the 1st third of sequence 1 is below the mean of 5 used in the finite sessions, specifically 4.4 rounds. By contrast, in sequences 2 and 3 it is above, specifically 6 and 6.6, respectively. Hence initial matches are substantially shorter in sequence 1 compared to the other sequences. Overall, however, the three sequences are very similar with average match length across all 15 matches equalling 5.3, 5.4 and 5.7, respectively. In both the finite and indefinitely repeated sessions participants play 15 matches.

The first hypothesis Lugovskyy et al. (2017) test is that "contributions in repeated games with sequences that have probabilistic end rounds will be greater than or equal to those in repeated games with sequences that have known end rounds". They evaluate this hypotheses by comparing behavior in finite and probabilistic settings for four different pairs of treatments which differ in group size, MPCR and whether participants make a binary contribution choice or not.

Table 2.3 shows the results of this analysis. The first two columns ("Finite" and "Prob All") reproduce the analysis in Table 3 in Lugovskyy et al. (2017). The analysis shows that in two of the four treatments cooperation is higher in the finitely repeated

| | *All Rounds Cooperation Rate* | | | | |
|---|---|---|---|---|---|
| Decision Setting | Finite | Prob All | Prob S1 | Prob S23 | S1 vs S23 |
| $N = 4$, MPCR= 0.3 | 15.0 | 22.4<br>$<^{**}$ | 10.36<br>$>$ | 28.41<br>$<^{***}$ | $p = 0.000$ |
| $N = 4$, MPCR= 0.6 | 39.4 | 44.3<br>$<$ | 33.52<br>$>$ | 48.39<br>$<$ | $p = 0.088$ |
| $N = 2$, MPCR= 0.6 | 41.1 | 38.3<br>$>$ | 31.71<br>$>$ | 42.64<br>$<$ | $p = 0.084$ |
| $N = 2$, MPCR= 0.6, Binary | 54.5 | 41.2<br>$>^{***}$ | 36.84<br>$>^{**}$ | 42.65<br>$>^{**}$ | $p = 0.398$ |

TABLE 2.3: Average cooperation rates across all rounds in the finite sessions (column (1)) and across all sessions with a probabilistic ending (column (2)) as in Table 3 in Lugovskyy et al. (2017). We further split the sessions with probabilistic ending in those with initially short matches (S1) and those with initially long matches (S23). Below each cooperation rate we show how the finite setting compares to the rate in question (as in Lugovskyy et al. (2017)). The last column shows the p-value when comparing initially short and long sequences. Following Lugovskyy et al. (2017) standard errors are clustered at the participant level in all regressions.

game and in the other two it is higher in the indefinitely repeated (probabilistic) game. One each of these comparisons is statistically significant. These and other analysis lead Lugovskyy et al. (2017) to conclude "We do not, however, find consistent evidence that overall cooperation rates are affected by whether the number of decision rounds is finite or determined probabilistically."

When we split out the sessions in those with initially short and those with initially long matches, though, we might have reached a different conclusion. The column "Prob S1" shows cooperation rates as well as comparisons in the session with initially shorter matches. In this case all four comparisons point into the same direction: more cooperation in the finitely repeated game. Only one of the comparisons is statistically significant. It should be noted, however, that the first ($N = 4$, MPCR= 0.3) and third comparisons ($N = 2$, MPCR= 0.6) are both just outside 10 percent statistical significance ($p = 0.150$, $p = 0.102$) in a comparison that is somewhat underpowered.

In the sessions with initially long matches the picture is very different. In this case three out of four comparisons point towards less cooperation in the finitely repeated game. Out of the statistically significant comparisons one each is pointing towards more and one towards less cooperation in the finitely repeated game. Hence while the

sessions with initially long matches show more of a similar picture than the overall sample, the sessions with initially short matches behave quite differently and would lead to a different conclusion. It should also be noted that, except for the last comparison ($N = 2$, MPCR$= 0.6$, Binary), the differences in average cooperation rates across the initially short and long sessions are always statistically significant.[21] Appendix Table E.14 shows that similar conclusions hold when we consider first round cooperation rates only.

While on balance it seems to us that Lugovskyy et al. (2017)'s overall conclusion is likely to be robust once "many" match length realizations are considered, the case study shows again how easily different conclusions could have been reached with different match length realizations.

### 2.3.3 Strategy Revision Opportunities

Our third case study is the paper "Strategy Revision Opportunities and Collusion" by Embrey, Mengel, and Peeters (2019) published in *Experimental Economics*. Embrey, Mengel, and Peeters (2019) explore how the possibility of being able to change a repeated game strategy during the course of play (i.e. to use "behaviour" strategies) affects cooperative behaviour in stylized oligopoly experiments. Their main treatment variations compare games of strategic substitutes and strategic complements with and without revision opportunities (RO). They find that without RO (when strategies have to be encoded upfront) there is more cooperation in games of substitutes than in games of complements. With RO there is more cooperation in games of strategic complements than with substitutes, but the latter difference is not statistically significant. The column "All groups" in Table 2.4 shows their main treatment effects in terms of efficiency, i.e. in terms of the percentage of the difference between joint profit maximizing payoff and Nash equilibrium payoff realized in the stage game.

[21] This difference is not driven by the shorter matches themselves. If we restrict attention to cooperation rates in the last third of matches only we find a difference (S1-S23) of $-18.03$, $p = 0.000$ for treatment $N = 4$, MPCR$= 0.3$, of $-17.55$, $p = 0.082$ for treatment $N = 4$, MPCR$= 0.6$, of $-7.66$, $p = 0.260$ for treatment $N = 2$, MPCR$= 0.6$ and of $-16,42$, $p = 0.133$ in treatment $N = 2$, MPCR$= 0.6$, Binary.

| *with RO* | All groups | Short First Matches | Other Groups | Short vs Others |
|---|---|---|---|---|
| Substitutes | 21.0 | 7.7 | 28.2 | |
| Complements | 26.4 | 44.2 | 16.9 | |
| Treatment Effect | - 5.4 | -36.4*** | 11.3** | $p = 0.0030$ |
| *without RO* | All groups | Short First Matches | Other Groups | Short vs Others |
| Substitutes | 22.7 | 30.9 | 18.2 | |
| Complements | 9.4 | 13.8 | 7.0 | |
| Treatment Effect | 13.3** | 17.1 | 11.2 | $p = 0.6888$ |

TABLE 2.4: Efficiency measure from Embrey, Mengel, and Peeters (2019). Treatment Effect is the difference between substitutes and complements. The two rightmost columns split out the groups with short initial matches from the rest. As in Embrey, Mengel, and Peeters (2019)'s main analysis matches 7-10 are considered. Stars are from t-tests with standard errors clustered at group level.$^{***}p < 0.01$, $^{**}, p < 0.05$, $^{*}p < 0.1$.

Embrey, Mengel, and Peeters (2019) use a discount factor of 0.875 implying a median match length of 7 stages. There are six different matching groups with different match length realizations. We will group those into two categories: (i) those where first matches are short, specifically where the first three (out of ten) matches all have a below median match length and (ii) all other matches. Table 1 in Embrey, Mengel, and Peeters (2019) shows that two matching groups (groups 2 and 6) fall into category (i). It should be noted that these matching groups do not have fewer stages overall than others. In fact group 6 has the most stages overall of all groups.[22]

Table 2.4 shows the average efficiency, defined as the percentage of the difference between joint profit maximizing payoff and Nash equilibrium payoff realized in the stage game. In all cases we follow Embrey, Mengel, and Peeters (2019) and focus on the average efficiency across matches 7-10.[23] We then compare average efficiency in games of strategic substitutes and complements separately for games with and without revision opportunities (RO).

The table shows that treatment effects depend on match length realizations. The overall negative, but statistically insignificant treatment effect (difference between substitutes and complements) with revision opportunities seems driven by the groups

[22] Alternatively we could split the sample into two equal sized categories by focusing on the length of the first two matches only, as there are three groups were both the first and second match are below median length. In this case qualitatively the same conclusions do hold.

[23] Appendix Table E.15 shows all matches

with short initial matches, where the effect is almost seven times larger and statistically highly significant. In the other groups the treatment effect reverses sign and is also statistically significant. The difference between groups with short initial matches and other groups is statistically significant at the 1% level. In this case, hence, diametrically opposite conclusions could be reached when matches with short or long initial realizations are studied. This is also illustrated in the left panel Appendix Figure F.4 which shows the treatment effect for all possible selections of two match length realizations.

Without revision opportunities we see a slightly different pattern. There is an overall positive and statistically significant treatment effect. The effect is somewhat larger with short first matches and somewhat smaller in the other groups, but both are positive. The effect is not significant in either of the subgroups, presumably due to lower sample size. Hence in this case the treatment comparison seems robust to match length realizations as is also illustrated in the right panel of Appendix Figure F.4.

Note that when revision opportunities are ruled out then all updating has to take place across (as opposed to within) matches. The analysis in this Section and the fact that match length realization seems to have a stronger impact with revision opportunities could suggest that within match learning might also play an important role for the effect of match length realizations. This could be an additional reason why the effect sizes observed with human players in the Dal Bó and Fréchette (2018) meta study are larger than with computer simulated learners who learn only across matches (see Section 2.2.3). In the next section we will discuss implications for learning and experimental design in more detail.

## 2.4 Discussion and Conclusions

We have seen that the realized match length of early matches in indefinitely repeated games has a substantial impact on cooperation rates in subsequent matches. We also studied three cases from the literature where realized match length has a strong impact on treatment comparisons, both in terms of the size and the direction of the treatment

effect. As discussed in the Introduction our results have implications for policy evaluation and can shed some light on what is often termed "cultural differences". They can help explain why people growing up in different social backgrounds (characterized by more or less stable interactions) or people coming from different work environments (characterized by more or less turnover) might show different behaviour in the exact same situation. Our results also have important implications for our understanding of how people learn in infinitely repeated games as well as for experimental design. We now discuss these in turn.

Our results have shown that people do learn and what they learn is affected by match length realization. Models displaying the "power law of practice" are able to explain the key patterns in our results quite well and can explain part of the empirical effect of match length realization. One interesting question is which moment of the distribution of match length realizations is most important for learning. Indefinitely repeated games are implemented using a mean expected match length that derives from the discount factor in the infinitely repeated game considered. Appendix Table E.16 shows, however, that the median match length realization seems a more important determinant for participants' behaviour than the mean.[24] This raises the question of which sequence of indefinitely repeated games "correctly" represents the infinitely repeated game one ultimately has in mind. This question has been answered theoretically under standard game theoretic assumptions. But, given how people seem to learn in these games, it might be necessary to rethink this question. For example, in our first case study (subsection 2.3.1), it is the match stochastic treatment, which as expected, closely matches the mean match length of the deterministic case. The median match length of the deterministic case is better matched by the inverse treatment, though. Which of those is the more relevant comparison depends on which of these moments is more important for how people learn. If it indeed turns out that median

[24] There is research in other contexts suggesting that the median experience might be relevant and that people understand information based on median/rank better especially when there is a lot of skewness in the distribution, which is also the case with match length realizations (Aldrovandi, Brown, and Wood, 2015; Wood, Brown, and Maltby, 2012). One example is consumption of alcohol, where people seem to have a good sense of how their consumption compares to the median, but not to the mean (Wood, Brown, and Maltby, 2012).

match length is the key statistic determining learning, then future research in both theory and experiments is needed to build and test new models of learning which can accommodate this fact.

We now discuss the consequences for experimental design when treatment comparisons can be affected by match length realizations. First note that drawing the same sequence of match length realizations for all treatments does not solve the issue. Both Lugovskyy et al. (2017) and Embrey, Mengel, and Peeters (2019) do exactly that, i.e. draw one set of sequences and use it for all treatments. As both of these case studies have illustrated this does not solve the issue, as there can be interactions of the treatment effect with match length realization.

The most obvious solution to the problem is to simply use as many match length realizations as possible, e.g. by randomizing match length at the match level, as we did in our "MatchStoch treatment". The main downside of this approach is that it can induce waiting times as all participants in a session or matching silo have to wait for the longest match to end before being rematched. This also restricts the total number of matches that can be played overall. As the overall effect on length is increasing in the number of matched subjects, the effect on expected length can be potentially mitigated by matching silos at either the cost of shorter matches or possibly allowing for rematching of the same individuals. Instead of using many match length realization one could try and use many matches and hope that the effect of early match length realization washes out over time. As our analysis in Section 2.2.3 has shown this does not seem to be the case for match numbers that can reasonably fit in a two-hour experimental session.

What could be other possible solutions? Fréchette and Yuksel (2017) studied three alternatives to the standard random termination method used to examine infinitely repeated games in the laboratory. In a method used e.g. by Sabater-Grande and Georgantzis (2002), Cabral, Ozbay, and Schotter (2014) or Vespa (2019), a fixed (known) number of rounds are played with certainty, and payoffs in these rounds are discounted at a known rate $\delta$. After the rounds with certainty are played, there is a fixed known probability $\delta$ that the match continues for an additional round, and payoffs in

these rounds are no longer discounted. Andersson and Wengstroem (2012) and Cooper and Kuehn (2014) use a similar method that also starts with a fixed number of rounds with payoff discounting, but is then followed by the coordination game induced by considering only two particular strategies in the infinitely repeated game - namely, the Grim trigger strategy and the strategy of always defecting. Fréchette and Yuksel (2017) also propose their own method - block random termination - in which subjects receive feedback about termination in blocks of rounds. Specifically, the play as in the standard method, but in blocks of a pre-announced fixed number of rounds. Within a block, subjects receive no feedback about whether or not the match has continued until that round, but they make choices that will be payoff-relevant, contingent on the match actually having reached that point. Once the end of a block is reached, subjects are told whether the match ended within that block and, if so, in what round; otherwise, they are told that the match has not ended yet, and they start a new block. Subjects are paid for rounds only up to the end of a match, and all decisions for subsequent rounds within that block are void.

Can these methods solve the problem of match length realizations? The first method could possibly mitigate the impact of match length realizations, as at least very short matches are avoided. On the other hand by systematically eliminating short matches, this could distort the results and in a social dilemma e.g. lead towards higher cooperation. In addition, as overall match length is still random, this method does not eliminate the problem. The second method does not have an uncertain match length, but it has the downside that the number of repeated game strategies allowed needs to be restricted ex ante. This can be undesirable in some contexts. Further, Fréchette and Yuksel (2017) found that neither of these methods induces behaviour that is consistent with the presence of dynamic incentives. Last, with block random termination the length of blocks is crucial. If it is very short, then effect of match length realization will at best be mitigated.[25] With long blocks, there is the downside that the experiment lasts "unnecessarily" long and as a consequence only fewer matches can be played. An open question is whether the fact that many stages are played affects behaviour even

[25] If blocks are of length one the method is the same as the standard method

if participants learn afterwards that some of them were payoff-irrelevant. Similar it could be that participants' behaviour reacts only to information on how many stages were payoff relevant as opposed to how many stages were played.[26] One possibility to address the latter concern might be to use the block random method and give feedback about match length realization only at the very end of the experiment. This, however, would make learning much harder for participants.

How should one then design experiments to get the "correct" treatment effect? By running sessions with "many" match length realizations - as in our match stochastic treatment - we should be able to adequately capture the average effect for the stochastic game. Having many match length realizations also allows to study how sensitive the specific results obtained are to match length realization. In many cases results may be robust to match length realization. If it is not possible to run "many" different sequences, the solution we recommend at the moment is to use "very different" match length realizations (across sessions or matching silos) and attempt to bound the treatment effect by comparing the (initially) shortest with the (initially) longest realizations. This still, of course, requires a good enough sample size for each match length realization. For future research, studying the other possibilities mentioned and finding out which method works best seems a worthwhile programme to us.

[26]Some indication that this might be the case is given by Fréchette and Yuksel (2017) themselves. They find that, even when the discount rate is the same, if interactions are expected to be longer defection increases and the use of the Grim strategy decreases.

## Chapter 3

# Do subjects disclose outcomes strategically? An experiment on endogenous information in a dictator game.

# *Abstract*

In the economic literature there is convincing evidence that dictators are more generous towards more deserving receivers. While previous studies have provided evidence that individuals use information about receivers' deservingness to decide how much to give in contexts of exogenous information, little is known about decisions made in a context of endogenous information. To start closing this gap, we study distributional decisions in a context where information about receiver's deservingness is endogenously determined by the receiver's decision to disclose information. Using a one-shot dictator game and varying deservingness by exploiting differences in outcomes from a real-effort task, we explore whether receivers voluntarily provide information about their deservingness-type and the effect of information on dictators' decision. We find that receivers disclose information about outcomes close to the median, but withhold information about outcomes significantly higher or lower than the median outcome. Furthermore, the choice to disclose information is strongly related to the receivers' stated beliefs about dictators' actions. Results also show that dictators give more in the case information is endogenous than in the case in which it is exogenously provided.[1]

[1]I particularly would like to thank my supervisor Professor Friederike Mengel, for having provided invaluable guidance and help. I also thank Professor Bertil Tungodden and his Research Group for helpful comments and suggestions. I am grateful to seminars participants at the University of Essex and FAIR at NHH Norwegian School of Economics. All errors are mine. I gratefully acknowledge financial support from the ESSEXLab Seedcorn Grant.

## 3.1 Introduction

Many experimental studies involving bargaining games such as the ultimatum game and the dictator game, have provided evidence of individuals being systematically more benevolent than the self-interested homo economicus introduced by the neoclassical economic theory (Thaler, 1988; Güth and Tietz, 1990). The existing literature has shown, in fact, that a substantial fraction of individuals is concerned about others' payoffs and intentions and is willing to sacrifice its pecuniary gains in order to help others (Camerer, 2003). Empirical studies and laboratory experiments assessing this type of behaviour have furthermore documented how different contextual cues and conditions can affect individuals' willingness to give (Engel, 2011; Tversky and Simonson, 1993) and have pointed at the context about the deservingness of the receiver as one of the most relevant factors amongst these. The evidence that people condition their behaviour on others' characteristics that account for deservingness is well-established across disciplines. Subjects regularly show more generosity towards receivers who are hard working rather than lazy (Alesina and Angeletos, 2005) and make higher transfers to those subjects who have been affected by bad luck, but can't be held accountable for their misfortune (Konow, 2000; Cappelen et al., 2013; Mollerstrom, Reme, and Sørensen, 2015; Lefgren, Sims, and Stoddard, 2016; Tinghög, Andersson, and Västfjäll, 2017).

While previous studies have provided substantial evidence that individuals use information about receivers' deservingness to decide how much to give in a context where the information is exogenously imposed, very little attention has been devoted to the study of decisions made in a context of endogenous information. Understanding the effect of endogenous information about the deservingness of the receiver on distributional decisions could, however, provide a relevant contribution to the economic literature. Introducing endogenous information into the decision environment

has been demonstrated to affect individuals' behaviour in many different decision situations. For instance, endogeneity of information can generate uncertainty or information asymmetries into the decision environment, but also alter normative considerations that are embedded in it. As shown in the literature, these factors can induce subjects to behave more self-interestedly either by allowing them to capitalize on others' lack of information to avoid sanctions, or by changing the extent to which social norms impact on fairness consideration (Roth and Murnighan, 1982; Mitzkewitz and Nagel, 1993; Dana, Cain, and Dawes, 2006; Cialdini, Reno, and Kallgren, 1990; Bicchieri, 2005; Shang and Croson, 2005; Krupka and Weber, 2013). A representative example of the effect of endogeneity of information in the context of dictator games is provided by the work of Dana, Weber, and Kuang (2007). The authors study subjects' behaviour in the case information about the receiver's payoff depends on the dictator's decision to disclose this information. They show that subjects are significantly less generous in the manipulation, relative to the game where information is provided, and often justify giving themselves the highest payoff by choosing to avoid information about the effect of their choice on the receiver' s payoff. These results lead the authors to conclude that "transparency plays an important role in the choices people make and kinds of outcomes that result" in studies measuring fairness. Studying individuals' behaviour in the case information about receiver's deservingness is not directly provided to the decision-maker could therefore help to capture behavioural motives that affect giving in the real world, but are not observable in a context of full information. This paper contributes to the understanding of distributional choices in the real-world by experimentally documenting an effect of endogeneity of information about the deservingness of the receiver on transfers made in a dictator game.

To our knowledge only two studies have investigated giving decisions in the case information regarding the deservingness of the receiver is endogenous to the decision environment. The first is a study by Thunström et al. (2016). Here the authors use a one-shot dictator game to explore the extent to which subjects actively seek or avoid information about the deservingness of the receiver they are matched with. They find that the majority of subjects choose to become aware of the deservingness type of the

receiver and adjust their transfers according to the information received – more is given to deserving receivers and less to receivers represented as less deserving. No difference in dictator' decisions is, however, documented between endogenous and exogenous treatment. In the second study, Fong and Oberholzer-Gee (2011), study, instead, dictators' transfer decisions in the case information about the deservingness of the receiver depends on the choice of the dictator to buy information. In line with Thunström et al. (2016), they show that a significant number of dictators is willing to pay in order to know the type of the receiver. They, however, find that when information is endogenous, aggregate transfers fall significantly. The authors argue that this result can be explained by the combination of two factors: first, subjects who buy information use it to give less to the less preferred receivers' type and, second, the cost of information reduces the overall endowment to split.

Even though results about the effect of endogenous information regarding the deservingness of the receiver presented by these two papers are mixed, they seem to corroborate the idea that even small changes to the decision context (in this case the cost of information) can produce significant variation in individuals' behaviour. Likewise, the finding that dictators actively seek information about receivers' deservingness and act on this frame in both decision situations, confirms the extensive evidence about the importance of this kind of information on distributive choices that we discussed at the beginning of the introduction.

It is, however, noteworthy that both of the studies mentioned above have introduced endogeneity in the decision environment by focusing solely on dictators and by giving them the possibility to seek or avoid information about the deservingness of the matched receiver.

Yet, in many real-life situations giving may occur in a context where information about receivers' characteristics that account for deservingness is not exogenously provided or determined by the dictator' s decision to reveal or avoid it. It may, in fact, be that the availability of information is determined by the choice of the receiver itself to disclose the information.

Assessing giving decisions in a richer decision-making environment, in which also

behaviour and beliefs of receivers are accounted for, can therefore contribute to better understand the underlying motivations determining moral behaviour in the real-world and help developing a better model of human decision-making.

This paper adds to the research on distributional choices and information both in terms of methodology and research question. Our study extends the analysis of the effect of information about deservingness on distributional decisions to a context where information about the receiver's deservingness is endogenously determined by the receiver's decision to disclose information. We therefore not only focus on the effect of information about deservingness on dictators' decision, but also explore how receivers' deservingness affects a subject's decision to disclose information about her deservingness. We believe this is the first study that tries to identify the effect of information about the receiver's deservingness on both dictators and receivers' decisions and explore the relation between these two decisions within a coherent framework. Investigating the interplay between these two decisions can be particularly relevant for a better understanding of the process of production and consumption of information in the real-world. This parallel design enables us to investigate how dictators interpret lack of information actively determined by the decision of the receiver and whether there exists a relationship between giving decisions and beliefs. On the other hand, it also allows us to assess how receivers value information about their deservingness and whether they perceive a benefit to hiding or revealing it to others.

We address this research question by developing a novel experiment which connects two strands of the economic literature: the deservingness literature and the disclosure theory. The deservingness literature, as described above, provides us with robust evidence of the extent to which information about receivers' deservingness affects individuals' giving decisions in the case information is available- subjects are more generous towards more deserving receivers. On the other hand, the disclosure theory offers us a theoretical background through which we can introduce endogeneity of information in the decision environment and model predictions regarding the mechanism of disclosure of information.

The central idea of the disclosure theory is that market forces can bring information senders (for example firms) to voluntarily disclose complete information about the quality of their type (product), provided the information disclosed is verifiable and the costs of disclosure are negligible (Viscusi, 1978; Grossman and Hart, 1980; Grossman, 1981; Milgrom, 1981; Jin, Luca, and Martin, 2015). The dynamic that explains this prediction relies on the restrictive assumptions that information receivers (consumers) hold correct beliefs about not-disclosed information and have the ability to make correct inference about it. The process that leads to a full disclosure of information is iteratively generated by a combination of information receivers' scepticism and information senders' ability to reveal information. In particular, when no information is available, best information senders have an incentive in revealing information about the quality of their type in order to distinguish themselves from other information senders. This incentive applies, then, iteratively to all subsequent quality-types of information senders, leading to an equilibrium where all not disclosed information is viewed with scepticism and associated with the worst quality. Therefore, in equilibrium all information senders reveal information on the quality of their type, with the exception of information senders producing the worst quality type, who are indifferent between revealing and not-revealing.

Building on these findings and elaborating on the evidence from the deservingness literature in dictator games we therefore try to understand whether the prediction of the disclosure theory applies to the information about deservingness of the receiver. We test whether subjects voluntarily disclose information about their deservingness and whether non-disclosed information is viewed with scepticism and interpreted as a signal of low deservingness.

We investigate the unravelling (full disclosure) predictions developing an experiment that is sufficiently complex to capture the main strategic tensions of the theory, but at the same time maintains the game structure close to the standard paradigm used in the literature for the study of the effect of information regarding deservingness on distributional choices. The choice of this specific design allows us to introduce the

changes required to address our research question in the standard environment studied in the literature without affecting the comparability of results.

In our experiment, subjects make decisions in the context of a modified one-shot dictator game, where there exists information asymmetry about receiver's deservingness and the receiver can decide whether to reveal her deservingness-type before the dictator makes her choice. To vary deservingness across receivers we exploit differences in the outcome that receivers obtain by participating in a real-effort task. Each session outcome distribution is common knowledge and receivers are shown where their outcome places in the distribution and whether this is above or below the median session outcome. Following the receiver's decision to disclose information about her outcome from the task, the dictator is asked to choose how to split a fixed amount which is given as payment for the work that both receiver and dictator have done on the task. In the case the receiver has disclosed her outcome, the dictator knows it and can make her distributive choice based on certain information. In the case where instead the receiver has chosen to withhold information, the dictator must guess the real outcome from the receiver's choice to hide it and from the distribution of outcomes, which is common knowledge.

We are interested in the study of the specific case where the recipient can only choose whether to reveal or not information about her deservingness to the dictator. We believe that this is a form of decision that individuals face daily when in possession of information that they think could be of interest to the other parties or affect their decision. For example, one might choose to not report a low mark on her cv, but instead highlight a good one, if she thinks that it could have an impact on her chances of getting a job. Similarly, one might omit to tell a friend that she went out partying in the weekend, but report instead that she had an accident, if she thinks this could affect her friend's choice to help her with an assignment with which she has felt behind. Focusing on the scenario where the receiver cannot misrepresent information, furthermore, allows us to address our research question in two additional ways. First, it enables comparability of results with the existing literature by recreating a decision scenario similar to the one previously assessed in the literature on endogenous information

about deservingness of the receiver. Indeed, both in Fong and Oberholzer-Gee (2011) and Thunström et al. (2016)'s study, the only information scenarios that are considered are the one where the dictator is informed (because she acquired information) and the one where instead no information is available (information is avoided). Second, it allows us to reproduce the assumptions underlying the disclosure theory and therefore make our data more easily interpretable in light of the theory. Prior research has also provided evidence of how dishonest information about receivers' deservingness can affect fairness considerations (Cappelen, Cappelen, and Tungodden, 2018; Bortolotti et al., 2017). The possibility of lying could therefore have a different effect on the choice to report information of receivers and the corresponding transfers made by dictators. We, however, leave the question of how endogenous information might affect subjects' behaviour in the case where dishonest information can be disclosed to future research.

We also collect measures of subjects' social preferences by using a non-parametric approach introduced by Kerschbamer (2015). Past research has in fact demonstrated that differences in individuals' social preferences can help explaining different distributive choices [2].

By eliciting receivers' first order beliefs and dictators' second order beliefs (beliefs about receiver' beliefs) we collect direct insights on the dynamics driving subjects' choices in our experiment. Specifically, the study of beliefs allows us to assess the extent to which subjects hold correct beliefs about others' actions or beliefs and how they respond to these beliefs.

We finally directly investigate the effect of endogenous information about recipients' deservingness on distributional choices by comparing dictators' behaviour in our endogenous treatment with behaviour observed in two additional treatments, where information is exogenously provided.

For the analysis, we start from the evidence from the deservingness literature that

[2]Self-interest, fairness, and efficiency have been identified in the literature as the three primary motives through which social preferences can explain distributive choices (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Konow, 2000; Andreoni and Miller, 2002; Charness and Rabin, 2002; Cappelen et al., 2007; Almås et al., 2010; Almås, Cappelen, and Tungodden, 2020).

dictators care about the receiver' s deservingness-type and give more to more deserving receivers and we link this with the idea that receivers and dictators have opposite interests. In our setting we assume therefore that dictators want to guess the correct receiver's outcome and give accordingly, while receivers want the dictator to guess a high outcome as this can bring to higher endowment transfers. We conceptualise dictators' and receivers' interest in this way as we believe that conflict of interests between agents mirrors an important feature of many real-life situations where an agent has to make a distributive choice in a situation of asymmetric information. Recalling the example above, a student might be happy to help a friend that has felt behind with an assignment if the reason for which the friend needs help is that she has been involved in an accident, but not if the reason is instead that she went out partying in the weekend. Yet, the student might not have direct access to this information. Conversely, the friend who has been out partying might only be interested in being helped and finishing the assignment. We therefore try to understand whether the arguments of the disclosure theory apply to this decision scenario. We test whether receivers voluntarily disclose their outcome from the task and whether only those receivers whose outcome corresponds to the lowest outcome in the session are indifferent between revealing or not. We further examine dictators' guesses about non-disclosed outcomes. Hence, we test whether they interpret the lack of information as a signal of low outcome and give less to receivers who hide their outcome, and make instead transfers to other recipients according to the disclosed outcome.

Results show that the unravelling prediction of information about the deservingness of the receiver does not fully apply to the behaviour of subjects observed in our experiment. We find that receivers disclose information about outcomes close to the median, but withhold information about outcomes that are both significantly higher and lower than the median outcome in a session. The distance of a subject's outcome from the median outcome of the session, and not from the worst achieved outcome, seems therefore to drive the choice to withhold information in our setting. Receivers' choice to disclose information is, however, strongly related to stated beliefs about dictators' actions. Receivers choose in fact the disclosure option associated with the highest

expected endowment transfer from the dictator. Although when we look at the aggregate sample of dictators we do not observe differences in giving decisions between dictators who are informed and those who are not, we find that some effect of being informed about receiver's outcome exists when we assess the sample split according to preference types. In particular, more self-interested subjects make significantly lower transfers when information about the receiver's outcome is not revealed. Transfers of more other-regarding dictators are instead independent from the receiver's choice to disclose information. Results also show that dictators give more in the case information is endogenous than in the case in which it is exogenously provided.

The reminder of the paper is organised as follow: Section 3.2 introduces the experimental design. Section 3.3 describes the compositions of the participants' sample and the experiment procedure. Section 3.4 presents and discusses main results and Section 3.5 concludes.

## 3.2 Experiment Design

The experiment is a modified version of the standard one-shot dictator game. In our setting dictators have to decide on the split of an endowment which is not "manna from heaven", as in most of the literature on dictator games, but it is given as a payment for the work that participants completed participating in a real-effort task. Dictators do not have ex-ante information about the outcome of the work of the receiver in the task, but the receiver can decide to make this information available before the dictator decides how to distribute the joint earnings.

The experiment consists of five stages.

**First Stage** : In the first stage participants play an encryption task for 15 minutes. During this task participants are provided with a table in which each letter of the alphabet is associated with a different random number. They are therefore asked to encrypt a sequence of words by replacing each letter of a word with the corresponding number

shown in the table[3]. The words to be encrypted are presented one at time and participants can start working on a new word only if the previous one has been encrypted correctly. Each word encoded correctly counts as one. Participants are communicated that the goal of the task is to encrypt as many words as possible in the given time, however, they are not provided with information about next stages of the experiment. Previous evidence from the literature (Erkal, Gangadharan, and Nikiforakis, 2011) has shown that in experiments where subjects earn their endowment, information about a subsequent redistributive phase can cause a selection effect based on other-regarding preferences. Subjects might in fact self-select according to a specific preferences type in different endowment categories. However, the specific preferences type might, in turn, lead to observe particular distributive behaviour across subjects in a given category. In order to avoid our results to be affected by selection based on other-regarding preferences, we therefore wait to introduce the distributive decision until the task is completed and limit the information about the aim of the task to the fact that its result might affect decisions in the following stages.

We introduce deservingness considerations in the decision environment by presenting the endowment received in next stage as the payment for the work completed in this first stage of the experiment. We choose to implement this specific design as it has been used in many other experimental studies on distributive choices and it has revealed to be an effective way to make deservingness consideration a relevant part of the distributive problem (Cappelen et al., 2013; Cappelen, Sørensen, and Tungodden, 2010; Cherry and Shogren, 2008; Ruffle, 1998). This specific design, furthermore, allows us to directly link a subject's productivity with her deservingness and obtain a simple measure of deservingness to use throughout our analysis. A participant's deservingness in our experiment is therefore simply given by the correct number of words encoded during the encryption task.

[3] A screenshot of the encryption task is presented in Appendix G, Figure G.1

**Second Stage** : In the second stage of the experiment participants are matched in groups of two and are assigned either to the role of dictator or receiver[4]. In this stage participants are also provided with information both about their outcome in the encryption task and the outcome of other participants playing in the same session. Specifically, participants are shown the number of words that they encoded in the task and also the maximum, minimum and median number of words encoded in the session. This stage allows us to introduce common knowledge about the distribution of outcomes in the session and therefore test the extent to which subjects use the information about the outcome distribution to make inference about not disclosed information or to decide whether to disclose their outcome. Showing participants where their outcome approximately places in the task-outcome distribution and whether their position falls above or below the median outcome of the session, helps furthermore to make more salient deservingness concerns in the decision environment. Both absolute outcome and relative outcome with respect to the median have been used in the literature to vary deservingness across participants and have proved to be effective procedures to affect subjects' perception of deservingness (Cherry and Shogren, 2008; Ruffle, 1998).

At the end of stage two, receivers are also assigned to an outcome type. Participants are explained that the receivers' outcome type depends partially on a subject's outcome from the task, and partially on a random shock and therefore it does not necessarily bring information about the receivers' work in the task. Before types are assigned, a coin is tossed virtually for each receiver. If the coin shows head, the outcome considered for the type assignment corresponds to the same one achieved in the task, while if the coin shows tails the outcome that is considered corresponds to half of the one from the task. Based on an outcome distribution that accounts for the eventual variation in receivers' outcome, receivers are then assigned to a "High-type" if their post-shock outcome is above or equal to the median outcome of the new distribution, and to a "Low-type" if their post-shock outcome is instead below the median outcome[5]. The

[4]To avoid framing effects, in the experiment the receiver is referred to as *Player A* and dictator as *Player B*.

[5]Also in this case, in order to avoid framing effects, we vary the name of types that are presented to participants in the experiment. The receiver is shown a "H-type" if her outcome is equal or above the median and "L-type" if instead her outcome is below the median.

allocation of receivers to an outcome type introduces an element of luck in the decision environment, allowing us to address the issue of different distributional behaviour observed in the literature when luck is a determinant of receivers' outcome and therefore add robustness to our findings. Evidence from the literature shows that giving decision might change when an individual's outcome is affected by luck and that subjects are more likely to give to a receiver who has been affected by bad luck[6]. Following this literature, we could therefore expect that the decision of the receiver to disclose her outcome from the task could also vary depending on whether a luck component is a determinant of her type. By assigning receivers to an outcome type that is the result of a combination of performance in the real-effort task and luck, we can assess whether there is an interaction effect between these two components on the receivers' choice to disclose information. Note that the choice to allocate receivers to a type that depends on the new post-shock distribution of outcomes enables us to address the question of the effect of both good and bad luck. While participants with an outcome above the median, if affected by the shock, might be assigned to a "Low-type", participants with an outcome below the median session might benefit from the shift in the post-shock distribution and be assigned to a "High-type".

**Third Stage** : In the third stage of the experiment participants are introduced to the distributive decision. Participants are told that together with their matched player, who completed the same task, they earned 100 tokens. These tokens, however, are given only to participants assigned to the role of dictators in each group. Dictators can hence decide how to split the 100 tokens with their matched receiver by choosing to transfer any integer number between 0 and 100. Before making their choice, dictators are always provided with information about the outcome-type assigned to the matched receiver, however, they are provided with information about the receiver's outcome from the task only if the receiver chooses to make this information available. Endogeneity is therefore directly introduced in the decision environment by asking the

[6]See Fong (2001), Fong (2007), Cappelen, Sørensen, and Tungodden (2010), and Cappelen et al. (2007) for examples of the effect of luck on distributive decisions.

receiver to make the active choice to disclose or hide her task outcome (and therefore her place in the outcomes distribution) to the dictator, before she decides on the split of the common endowment[7]. The study of receivers' choice to disclose information across different outcome categories (deservingness categories) enables us to assess the extent to which deservingness considerations affect the choice to disclose information and collect insights on eventual differences amongst deservingness levels. It therefore helps us to test whether the hypothesis of unravelling of information applies to the deservingness context and all receivers, but those achieving the lowest outcomes, choose to disclose information.

After receivers make their decision, dictators are informed about the disclosure choice of the matched receivers. Dictators are aware that the only two possible actions of the receiver are: *reveal, not reveal*. If the receiver chooses *reveal*, the dictator is provided with both information about the receiver's outcome from the encryption task and her outcome-type. If the receiver chooses instead *not reveal* , the dictator receives information about the receiver's outcome type and a message saying "task outcome: not revealed"[8]. If information is not revealed, hence, the dictator is required to choose how to distribute the endowment in a situation of asymmetric information. Following the evidence from the deservingness literature, we could expect subjects to care about information regarding receivers' outcome and be willing to transfer more to those receivers who achieved a higher outcome. If we apply the argument of the disclosure theory to our design, we could furthermore expect dictators to make negative inference about not disclosed information and in line with their preferences give less. We address all these questions by comparing distributive decisions across different situations of information disclosure and levels of receivers' outcome.

**Fourth Stage** : In the fourth stage we collect information about beliefs of dictators (second-order beliefs) and receivers (first-order beliefs). The literature provides considerable evidence of the existence of a correlation between individuals' beliefs and

[7] Appendix G, Figure G.2 shows a screen shot of the decision screen presented to receivers.
[8] A screen shot for the dictators' decision screen can be found in Appendix G, Figure G.3.

behaviour[9]. Studying participants' beliefs can therefore help to improve our understanding of the actual underlying motives driving subjects' behaviour in other stages. In particular, assessing subjects' beliefs enables us to address the questions of whether subjects hold correct beliefs about others' choices and whether they best respond to their beliefs.

Receivers are asked to state their beliefs about the amount given by the dictator, while dictators are asked to guess the amount that the receiver expects to receive. In order to collect more observations, and better estimate the extent to which deservingness (outcomes) affects expectations and behaviour, we elicit dictators and receivers' beliefs in four different scenarios of outcome disclosure. We ask participants to state their beliefs relative to the scenarios where the receiver chooses to disclose her outcome and the outcome disclosed corresponds to three different levels of outcomes: actual outcome of the receiver, actual outcome of the receiver increased by 50% or reduced by 50%[10]. Beliefs are then collected also relative to the scenario where the receiver chooses to not disclose her outcome[11]. Participants are made aware that with 10% probability the actual disclosure decision of the receiver is changed by the computer and therefore beliefs stated in any information scenario can become relevant for additional earnings. We therefore elicit beliefs in an incentive compatible way so that subjects are motivated to correctly guess the amount received or expected in all scenarios[12]. Only in the case

[9]See Fischbacher and Gächter (2010) and Gächter and Renner (2010) for evidence of correlation between beliefs and choices.

[10]In case the value of the increased or reduced outcome exceeds one of the extremis of the outcomes distribution, we replace the corresponding value with the maximum or minimum outcome of the session. Alternatively, if a participant's outcome corresponds to one of the values at the extremis of the distribution, or it is very close to them, we replace the increased or reduced outcome with an outcome that is equal to the (rounded) median outcome of the session. This procedure allows us to collect beliefs corresponding to a low, medium and high outcome for each participant.

[11]Screen shots of the receivers and dictators' guess table are included in Appendix G, Figure G.4 and Figure G.5, respectively.

[12]Participants are communicated in the instructions that additional earnings are possible based on the accuracy of their guesses. For receivers: they receive 15 tokens if they exactly estimate the amount that the dictator chooses to give them, 10 tokens if their estimate is up to 5 points higher or lower than the actual amount given, 5 tokens if their estimates in up to 10 points higher or lower than the actual amount given and no additional earnings, if their estimate is more than 10 points higher or lower than the actual amount given. For dictators: the computer randomly selects one scenario and compares dictators' guess in that specific case with the corresponding receiver's stated belief. Additional earnings are then calculated following the same scheme implemented for the receiver.

in which the dictator is not informed about the receiver's task outcome, we ask dictators to also guess the hidden task-outcome. This provides us with a direct measure of the inference that dictators make about non-disclosed outcomes.

**Fifth Stage** : The fifth stage requires participants to make decisions in two non-incentivized decision situations.

The first decision situation consists of a hypothetical dictator game in which participants are asked to state their preferences across 10 binary choices (see Table H.1 in the Appendix). Each choice represents a payoff distribution between the participants and another hypothetical subject. The dictator game takes its structure from the game introduced by Kerschbamer (2015), and it provides a non-parametric approach to measure social preferences. The decisions made in the dictator game can, in fact, be transformed in a (X, Y) score tuple and the X score and Y score can be used as indicators of subjects' social preferences [13]. This game provides us with a direct measure of participants' social-preferences and allows us to control for them in our analysis.

In the second decision situation participants are presented with a list of five different categories of outcome from the effort task: lowest number of encoded words in the session, number of words above the lowest, but below the median, number of words equal to the session median number, number of words above the session median, but below the highest number of words encoded in the session and highest number of encoded words in the session[14]. They are therefore asked to guess the rate at which receivers chose to reveal their outcome in each of the five outcome categories. Answers to this question allow us to collect further information regarding dictators' beliefs about receivers' disclosure decisions. Specifically, this information can be used both to assess whether dictators correctly guess the pattern of disclosure decisions of receivers and to better understand how dictators interpret non-disclosed information.

[13] Appendix H.1 provides a detailed description of the game and of the procedure to transform participants' decisions in preferences scores

[14] See Figure G.6 in Appendix for a screen shot of the list of outcome categories presented to participants.

Participants are informed that decisions made in this stage of the experiment do not affect their final payoff. Completing, however, the stage creates an expectation of getting access to information about the final payoff from the experiment and receive the payment.

In addition to the main treatment described above, in which information is endogenous, we run two additional treatments with exogenous information. The difference in the exogenous treatments is determined by the fact that receivers cannot choose whether to reveal or not information about their outcome from the real-effort task. Thus, in one exogenous treatment dictators always receive information about the receiver's outcome, while in the other they are never informed about it. In order to ensure that behaviour is comparable across treatments, however, we maintain the rest of the experiment structure invariant[15]. Differences in behaviour across treatments provide us with some inference about the effect of endogenous information about receivers' deservingness on dictators' behaviour.

Participants' final pay payoff from the experiment is determined by the dictator's decision on how to split the 100 tokens endowment and eventual earnings from the correct guesses made in the beliefs stage.

## 3.3 Subjects pool and procedure

The experiment was programmed in z-Tree (Fischbacher, 2007) and conducted at the EssexLab of the University of Essex between December 2019 and March 2020. Ethical approval was obtained by the Social Sciences Faculty Ethics Sub-Committee of the University of Essex. All participants were recruited using the EssexLab Online Recruitment System. The research involved 326 participants, 62% were females, 74.23% were undergraduate students, 18.71% postgraduate students and 7.06% no-students, the average participants' age was 24.8.

[15]In the two exogenous treatments we do not ask participants to complete the task relative to the guess of the rate at which participants reveal their outcome across different outcome categories. Given the lack of option to reveal, this question, in fact, does not apply to the decision setting. We, however, believe that this change should not affect decisions in the experiment as the question is the last.

Evidence shows that distributive behaviour can be influenced by cultural and demographic characteristics (Alesina and Angeletos, 2005; Engel, 2011). We collect therefore information relative to those background characteristics that have been shown to be relevant for subjects' giving decisions and control for them in our experiment results in order to obtain more reliable estimates. Our data on subjects' characteristics include age, gender, academic background, religion and nationality. We also collect information regarding beliefs about sources of inequalities and distributional preferences. Information about participants' preferences are collected through the study of decisions that participants make in the hypothetical dictator game (Kerschbamer, 2015) [16]. Preferences are determined for 88.96% of participants in our sample. Following inconsistent choices in the hypothetical dictator game, preferences are instead not defined for 11.04% of participants [17]. Table H.3 in the Appendix presents the distribution of preferences measures of our sample. In line with previous results (Kerschbamer, 2015), most participants reveal preferences consistent with inequality aversion (21%) and selfishness (38%). In an open-ended question presented at the end of the questionnaire, we furthermore ask participants to explain in their own words decisions they made in the experiment. Open-ended questions have revealed to be useful instruments to collect insights about underlying motives explaining subjects' behaviour, and can therefore help to better interpret our results [18].

We ran a total of sixteen sessions with 12 to 28 subjects in each session. Sessions lasted on average 65 minutes and the average earnings were £15. In each session, we asked participants to complete the five stages of the experiment and a questionnaire in the same order. Instructions for each of the five stages were presented separately. Participants received instructions for each stage before the start of the specific stage and were given some time to read them and ask questions[19]. Each stage was started

[16] See Appendix H.1 for an explanation of how preference scores are computed from the decisions made in the hypothetical dictator game.

[17] In order to keep all observations of our sample in the analysis, we assign to inconsistent participants an average preference score and control for consistency throughout the analysis.

[18] See Erkal, Gangadharan, and Nikiforakis (2011) for an example.

[19] When first reading the instructions, participants were informed that the experiment consisted of five stages, but they were only provided with instructions for Stage 1. Instructions are presented in the Appendix G.1

only after everyone has finished reading the instructions and was certain to have understood them. Participants assigned to the same role played each stage simultaneously and could move to the next stage only after everyone has made decisions in the previous one. Results about decisions made during the experiment and final payoff were communicated only at the end of the experiment when participants completed the questionnaire.

We implemented three different treatment conditions: the endogenous information treatment, the exogenous information treatment with information and the exogenous information treatment without information. We ran eight sessions for the endogenous information treatment (four pilot sessions and four sessions[20]) and four sessions for each of the exogenous treatments.

## 3.4 Results

We will first analyse behaviour and beliefs of dictators and receivers in the main treatment, then we will provide a synthesis and compare results with the two control treatments.

### 3.4.1 Distribution of outcomes

Previous research in Experimental Economics has shown that in a context where participants earn their endowment giving decision may be affected by a sense for entitlement and concerns for relative status. Cox, Friedman, and Gjerstad (2007), for example, develop a model presenting evidence that relative status increases with a subject's performance and can negatively affect the decision to give to others. Achieving a high relative status may, in fact, affect the sense of entitlement of subjects and induce them to think that they have the right to keep a larger share of the endowment (Engel, 2011;

[20] In the pilot sessions participants were only provided information about the median outcome of the session. In the main sessions we extended the set of information regarding the distribution of outcomes from the task and also provided participants with information concerning the maximum and minimum outcome achieved during the session. The question concerning the share of subjects choosing to report their outcome across different outcome categories was furthermore introduced only in the last pilot session and its wording changed for the main sessions.

Cherry, Frykblom, and Shogren, 2002). In our experiment we could therefore expect subjects encoding a higher number of words to give less if concerned about their relative status. This, however, would also suggest that our analysis regarding dictators' decision to give may be biased if best performing subjects are consistently assigned either to the role of dictator or receiver. We therefore analyse the average and distribution of number of words encoded by dictators and receivers in the first task of the experiment. The average number of words encoded by dictators is equal to 58.97 and it is not significantly different from the average number of words encoded by receivers, 56.65 (two-sided t-test, p-value= 0.1118). Figure 3.1 shows the distribution of the number of words encoded by dictators (right panel) and receivers (left panel). It can be seen that the number of words encoded by participants assigned to each role exhibits similar variation on both sides of the distribution. A two-sample Kolmogorov–Smirnov test (p-value = 0.338) furthermore confirms that there are no differences in the distribution of the number of words encoded by participants assigned to the two roles [21]. We conclude therefore that behaviour of subjects should not be affected by consistent difference in the performance of participants assigned to the two roles and we proceed with the analysis of behaviour in the different stages of the experiment.

[21] We conducted the same analysis also on the sample split by treatment and found that the same results hold.

FIGURE 3.1: Distribution of the number of words encoded by participants' role

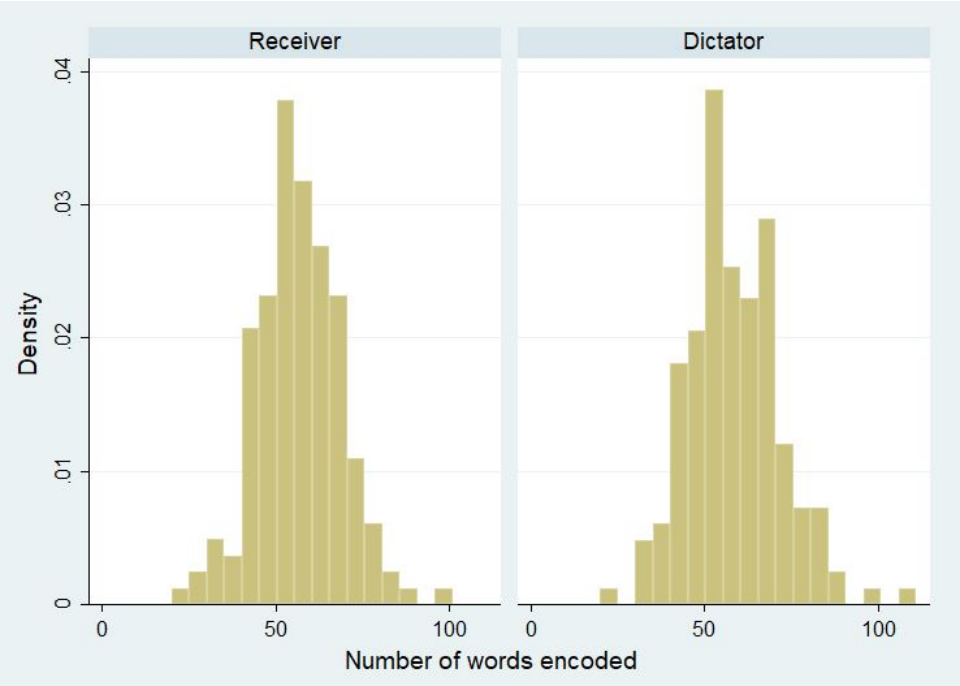


### 3.4.2 Participants' decisions

#### Receivers' disclosure decision

We start the analysis of recipients' decision by looking at the share of recipients who decide to reveal the number of words they encoded during the task in the first stage. Table 3.1 summarizes recipients' choices. The average revealing rate is about 71%, while not revealing is chosen above 28% of the times. This result is, however, in contrast with the unravelling prediction presented in the introduction. According to the unravelling equilibrium, only recipients achieving the lowest score in the task may, in fact, choose non-disclosure. In our experiment, although only 5.56% of the subjects assigned to the role of receiver (5 participants) encoded the lowest number of words in their session and therefore could have chosen to not reveal their outcome.

TABLE 3.1: Share of subjects who decide to reveal their outcome from the task

| Variable | Freq. | Percent |
|---|---|---|
| **Reveal** | 64 | 71.11 |
| **Not Reveal** | 26 | 28.89 |
| **Total** | 90 | 100 |

While the unravelling argument of the disclosure theory allows disclosure decisions to depend solely on the absolute position of the receiver in the task-outcome distribution, research in other contexts suggests that different standards may be also used by subjects to evaluate relative performance. The use of these alternative standards might have affected subjects' choice to disclose in our experiment and could therefore represent a competing explanation for the observed results. Median value, for example, has been observed to represent a relevant parameter that individuals use to compare themselves to others in different situations (Aldrovandi, Brown, and Wood, 2015; Wood, Brown, and Maltby, 2012). Since in our experiment information about session-median outcome has been provided to participants and exploited to introduce deservingness across receivers and assign them to types, we try to gain further insight on the determinants of our result and assess the distribution of disclosure decisions with respect to the median outcome.

While we do not observe difference in share of subjects reporting their outcome above and below the median outcome (p-value= 0.942, Fisher's test on the share of subjects choosing to not-reveal)[22], an interesting pattern relative to the median outcome emerges if we look at the choice to reveal across the complete distribution of task-outcomes. Figure 3.2 shows the distribution of relative rank for subjects who decide to not reveal (left panel) and subjects who choose to disclose their outcome (right panel)[23].

If we look at the right panel of the figure we can observe that the decision to reveal the task-outcome is somewhat uniformly distributed across different rank levels although slightly more common across higher ranks. Conversely, the left panel shows

[22] 14 subjects out of 49 subjects performing below the median outcome of the session decide to not reveal their outcome. Similarly, 12 subjects out of 41 subjects amongst those who encode a number of words above the median outcome choose to not report their score to the matched dictator.

[23] We choose to employ a measure of relative rank instead of absolute rank as it allows us to compare more easily results across sessions and avoid that most of the variation in the analysis is determined by larger sessions (Ejrnæs and Pörtner, 2004). The number of participants per session in our sample is in fact highly heterogeneous and varies from a minimum of 12 to a maximum of 28 participants. The measure of relative rank is defined as (r-1)/(N-1), where r is the absolute rank and N is the number of participants in a session. This measure is constructed such that the best performing participant's relative rank equals to zero and the worst performing participant's equals one. Appendix Figure I.1 shows the results obtained with an alternative measure of rank able to purge the effect of session size from our estimates (Booth and Kee, 2009). The pattern presented in the main text holds also in this case.

FIGURE 3.2: Relative rank distribution by choice to reveal the task outcome

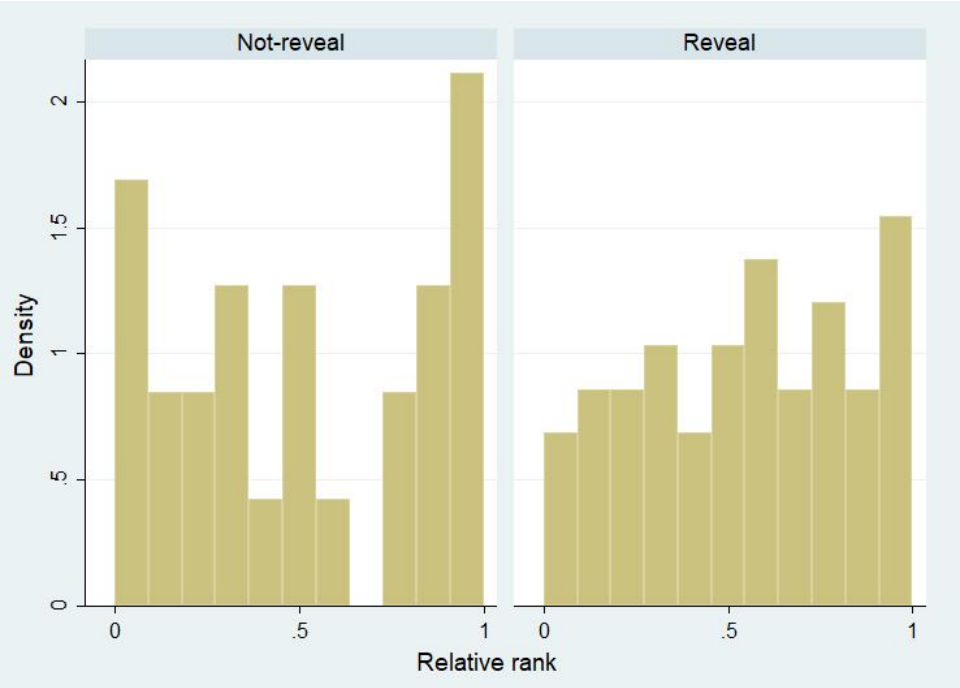


that the decision to withhold information concentrates around two main rank levels which correspond with the extremis of the rank distribution. The latter evidence suggests that there exists a correlation between an individual's position in the outcome distribution and the choice to not reveal the outcome. Specifically, the choice to withhold information from the matched dictator is more common among those subjects that perform either significantly better or worse than other in the session. This impression from Figure 3.2 is furthermore consistent with the detailed analysis of the decision to reveal amongst three main reference groups: participants that achieved the highest, lowest and median outcome in a session. While Figure 3.2 looks at relative rank and presents aggregate data per level of relative rank, a detailed analysis of choices made at the two extremis can provide a direct test of the extent to which the unravelling prediction applies to our scenario. In line with the prediction of the disclosure theory, we find that 80% of subjects that achieve the lowest outcome in a session choose to not disclose it. The data on best performing subjects' decision to disclose show, however, that our initial predictions were incorrect. A high share of participants also in this category (60%) chooses, in fact, to maintain secret its outcome. The non-reporting rates drops instead to 14.29% when we look at the subjects that achieved an outcome

corresponding to the median[24].

All in all, there seem to be both theoretical evidence provided by the literature and hints from the data presented above that suggest that the presence of information about median outcome in our experiment might have affected the standards that participants use to compare their outcomes with others and subsequently their decision to reveal. Specifically, it appears that participants make their decision to disclose not based on the absolute distance of their outcome with respect to the worst outcome in the session (as suggested by the disclosure theory), but focus instead on the distance with respect to the median outcome of the session. We directly test for this hypothesis in Table 3.2, where we assess the effect that the distance of a subject's outcome from the session median outcome has on the probability to reveal. Models 1-2 and 3 present the results of a logit regression where the dependent variable is the probability that a receiver discloses her outcome from the task. The coefficients in the regression are probability derivatives. The median_distance variable is constructed as the absolute difference between a participant's task outcome and the median outcome of the session. The unit of the variable in the regressions is 10 encoded words. Consistent with our hypothesis, Model 1 indicates that receivers who encode a number of words that is far from the median number of words encoded in the session are significantly less likely to reveal their outcome. That is, a 10 words increase in the gap with respect to the median outcome decreases the receiver's likelihood to disclose her outcome by 11.6 percent (p-value<0.01). This negative effect persists also after controlling for other candidate factors able to affect the decision to reveal in our experiment (Model 2) and individuals' characteristics (Model 3).

In order to reinforce our interpretation of the results concerning the decision to not disclose the task outcome, we also analyse receivers' responses to the post-experiment question about their choice to disclose. Although we do not find direct insights for

[24]Since all experiment sessions consisted of an even number of participants and no participant scored an outcome exactly equal to the median, we define median outcome as being equal to median outcome + or − 0.5. The sample of receivers assigned to this group in our analysis consists of 3 subjects who perform above the median and 4 below it. There is only one participant choosing to not reveal amongst those in the median group and she has outcome below the session median.

our explanation concerning the distance from the session median outcome, the questionnaire responses provide us with some evidence for an interpretation of our results. Specifically, responses of subjects choosing to not reveal their outcome seem to suggest that social preferences, and in particular inequality aversion, might play a role in the choice to disclose. The explanation for not reporting the task outcome given by three of the 26 subjects that choose to not disclose is, in fact, that they want to achieve an equal split of the endowment. Other four participants link the choice to not disclose with the idea to achieve a fairer outcome distribution. Four other subjects mention instead to try not to affect the decision of the dictator with their outcome disclosure. If the choice to not reveal the task outcome is motivated by inequality aversion of participants, in line with our results, we could expect high and low performer to be more likely to not reveal. Both types of subjects might in fact reasonably choose to not reveal their outcome in order to avoid the dictator to assign them either a greater or smaller share of endowment than the one she assigns to herself. For evidence of this we examine the impact of X score and Y score on participants' choice to reveal. As mentioned previously, in Kerschbamer (2015)'s task, Y score increases with inequality aversion while X score decreases with inequality aversion. If inequality aversion was a leading explanation for not revealing the task outcome, we would therefore expect Y score and X score to have significant effect on the choice to disclose, but opposite sign. While the significant coefficients on X score and Y score in Model 3 of Table 3.2 seem to suggest that social preferences play a role in the decision to disclose, their equal sign does not support our hypothesis of inequality aversion as a driver of the results. The positive coefficients associated with both variables seem instead to show that disclosure choice is mainly affected by altruism of subjects. The more a receiver cares about the other player's wealth (higher X and Y scores) the more likely she is to disclose her outcome from the task.

To sum up, the results in this section show that the unravelling prediction of the disclosure theory does not explain receivers' behaviour in a context where the receiver has to decide whether to disclose information about her performance. We point instead at the relative position of the receiver' s outcome with respect to the median outcome

TABLE 3.2: Logit regression determinants of choice to reveal task outcome

| Probability to choose to reveal task outcome | | | |
|---|---|---|---|
| | (1) | (2) | (3) |
| Median_distance | -0.116*** | -0.116** | -0.113*** |
| | (0.045) | (0.045) | (0.039) |
| Outcome_above_median | | -0.008 | 0.007 |
| | | (0.099) | (0.089) |
| Type_shock | | -0.001 | 0.021 |
| | | (0.112) | (0.102) |
| Female | | | -0.065 |
| | | | (0.091) |
| Age | | | 0.004 |
| | | | (0.005) |
| Xscore | | | 0.094*** |
| | | | (0.031) |
| Yscore | | | 0.062* |
| | | | (0.035) |
| Log-likelihood | -51.715 | -51.711 | -44.238 |
| Pseudo R2 | 0.0442 | 0.0442 | 0.1823 |
| Observations | 90 | 90 | 90 |
| Controls | No | Yes | Yes+ |

Tables show marginal effects. Robust standard errors clustered on subject level. *** p<0.01, ** p<0.05, * p<0.1
Controls Yes: the dummy for having outcome above the median is one if the receiver has encoded a number of words above the session median number. The dummy for the shock of the type is one if the type of the receiver has changed following the outcome shock. Yes+ includes: a dummy for education, which is one if the receiver is an undergraduate student, religion and a dummy for consistency, which is one if the receiver has made consistent choices in the hypothetical dictator game.

as a main driver of the decision to disclose in our experiment and we provide empirical evidence of this by showing that the likelihood to reveal decreases with the distance of a subject's outcome from the median. We furthermore try to interpret these results with subjects' preferences of inequality aversion, but we find evidence against this hypothesis.

A plausible and alternative interpretation for our results could, however, be found in the literature on in-group bias. A large body of evidence from this field of research has shown that the information that dictators and receivers have about relative group identity can strongly affect respectively the decision to give and expectations about it. Foddy, Platow, and Yamagishi (2009), for example, study in-group-bias in the case where receivers can choose to receive an allocation from an out- or in-group member. The authors observe that subjects prefer the allocation from an in-group dictator only as long as the dictator is informed of the shared group identity. They further show that the result can be explained by the expectation that dictators behave more altruistically and fairly towards same group members only when the group membership is common knowledge.

The mere act of categorizing participants within the outcome distribution might have induced some sort of group identification also in our experiment[25]. Participants with scores closer to the two extremis of the distribution might then have perceived higher the probability to be matched with a dictator belonging to a different outcome category. In line with the above literature, these subjects might have therefore chosen to keep their outcome (group) secret, best responding to the expectation that this would have given them a higher tokens allocation. In Section 3.4.3 we will study receiver's beliefs and we will provide evidence of the importance of expectations about dictators' action for the receivers' choice to not disclose the task outcome.

[25]See (Gerard and Hoyt, 1974; Scheepers et al., 2006; Eckel and Grossman, 2005) for examples of group assignment based on participants' results in a task.

**Dictators' transfer decision**

Having studied the decision to disclose information across receivers, we now turn the analysis to the study of the effect of endogenous information on dictators' decision to give.

Table 3.3 illustrates the average amount given by dictators that have been informed about receiver's outcome and dictators that have not been informed[26]. The table shows that the average amount given to receiver who disclose information is higher than that given to those who do not disclose information on their outcome. The average amount given in the two scenarios, however, is not statistically different (p-value=0.118, two-sample Wilcoxon rank-sum test).

TABLE 3.3: Dictators' decision to give by whether receiver's task outcome is revealed

| | Obs | Mean | Std. Dev. | Min | Max |
|---|---|---|---|---|---|
| Informed | 66 | 36.71 | 17.60049 | 0 | 75 |
| Not informed | 24 | 30.54 | 19.58033 | 0 | 70 |
| Total | 90 | 35.07 | 18.24416 | 0 | 75 |

This result is confirmed by Model 1 in Table 3.4, where we present outcomes of an OLS regression that tests also for other possible factors affecting the decision to give in our setting (Model 2) and controls for individual's characteristics (Model 3). While Table 3.4 confirms previous evidence, it also provides us with some interesting insights about the effect that social preferences have on dictator' decision to give and on the impact of information. First, Model 3 of the regression shows that the coefficient of Y score is positive and strongly statistically significant (p-value<0.01). This implies that dictators who are more inequality averse under favourable inequality (have a higher Y score) choose to give significantly higher amounts to their matched receivers. Second, Model 3 shows that after controlling for the effect of dictators' preferences, the coefficient for being informed becomes statistically significant at 5 percent level. This suggests that preferences might work as confounder for the effect of information and

[26]The change in the receiver's decision to disclose information that occurs in our experiment design with 10% probability has caused the number of dictators that are not informed (24) to be lower than the actual number of receivers that choose to not reveal their outcome (26).

only by taking them into account it becomes possible to appreciate an association between information and giving decision. We therefore split our participants sample according to different preferences types and try to assess whether a different effect of being informed exists across preferences groups. We define "other-regarding" those participants whose Y score is greater than 0.5, and "self-regarding" those participants having an associated Y score that is below or equal 0.5. Last four columns of Table 3.4 present results of an OLS regression for the effect of being informed about receiver's outcome on the decision to give, when the sample is split by preferences types.

TABLE 3.4: OLS estimates of the effect of information on giving decisions

| | Total Sample | | | Other-regarding subjects | | Self-regarding subjects | |
|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Infomed | 6.17 | 5.455 | 7.406** | 1.038 | 2.1 | 11.303** | 11.565* |
| | (4.515) | (4.43) | (3.628) | (5.338) | (4.798) | (5.54) | (6.484) |
| Outcome_above_median | | -2.495 | -0.956 | | -2.864 | | -1.34 |
| | | (4.241) | (3.686) | | (3.722) | | (6.271) |
| Receiver_type_shown | | 3.326 | 0.938 | | 1.171 | | 0.02 |
| | | (4.164) | (3.539) | | (3.827) | | (6.738) |
| Female | | | 2.154 | | -7.122* | | 9.527 |
| | | | (3.697) | | (4.108) | | (6.266) |
| Age | | | 0.308* | | 0.237 | | 0.448* |
| | | | (0.168) | | (0.222) | | (0.227) |
| Xscore | | | -1.11 | | | | |
| | | | (1.344) | | | | |
| Yscore | | | 5.612*** | | | | |
| | | | (1.198) | | | | |
| Constant | 30.542*** | 30.577*** | 22.920** | 41.750*** | 40.161*** | 19.333*** | 12.068 |
| | (3.957) | (5.494) | (9.054) | (4.959) | (9.313) | (4.193) | (12.744) |
| R-squared | 0.023 | 0.039 | 0.343 | 0.0013 | 0.1657 | 0.0654 | 0.1966 |
| Observations | 90 | 90 | 90 | 45 | 45 | 45 | 45 |
| Controls | No | Yes | Yes+ | No | Yes++ | No | Yes++ |

Robust standard errors clustered on subject level. *** p<0.01, ** p<0.05, * p<0.1
Controls Yes: the dummy for having outcome above the median is one if the dictator has encoded a number of words above the session median number. The dummy for the receiver's type shown is one if the type of the receiver observed by the dictator is H-type. Yes+ includes: a dummy for education, which is one if the dictator is an undergraduate student, religion and a dummy for consistency, which is one if the dictator has made consistent choices in the hypothetical dictator game. Yes++ does not include a dummy for consistency.

While the analysis of the fourth and fifth column shows that results for other-regarding participants are very similar to those observed in the overall sample before preferences are controlled for, the last two columns highlight that information has a different effect on the decision of self-regarding subjects. Model 4 and 5 show that receiving information about receiver's outcome does not affect the decision to give of other-regarding subjects. This is in line with evidence from the other-regarding preferences literature which shows that giving across other-regarding subjects is mainly motivated by difference aversion (reduce the inequality in payoff across players) and warm glow (derive utility simply from the act of giving) motives [27]. Other factors,

[27] See: Andreoni (1990), Fehr and Schmidt (1999), and Bolton and Ockenfels (2000)

such as for example being informed or not about receiver's outcome in our experiment, play therefore a small or no role at all in the decision to give of this group of subjects. A different picture emerges instead from model 6 and 7, which report the analysis of behaviour of self-regarding subjects. In this case we see that information has a positive and statistically significant effect on the decision to give of self-regarding subjects (model 6; 11.30 tokens more, p-value<0.05) and the effect persists in magnitude and sign, although with a lower level of statistical significance, after controlling for demographic characteristics and other experimental factors (model 7, p-value<0.10)[28]. Self-regarding subjects seem to value the choice of receivers who choose to disclose information about their outcome and reward it by sharing higher amounts of endowment with them than with receivers who choose instead to withhold information.[29]

Hence, while our initial hypothesis that dictators would give less to receivers who do not reveal their outcome is not confirmed by the analysis of the overall sample, we find some support for this hypothesis when we split the sample according to preferences types and we look at self-regarding subjects only. Behaviour of self-regarding subjects seem indeed to align with the intuition that not disclosed information are viewed with scepticism and cause dictators to be less generous. In Section 3.4.3 we will further study dictator' beliefs about the number of words encoded by receivers who choose to hide their outcome and we will show that expectation of low performance amongst those who do not disclose information represents a plausible explanation for

[28]Results of Table 3.4 in the main text include observations for subjects making inconsistent choices in the dictator game of Kerschbamer (2015) and who are therefore assigned the average Y scores of the sample. In Table I.1 of the Appendix we replicate the analysis of Table 3.4, but we exclude from the sample the 7 observations corresponding to inconsistent subjects. The results show that the difference between self-regarding and other-regarding subjects remains invariant.

[29]We choose to use 0.5 as a threshold for the assignment to one of the two preferences types as it gives us sufficiently large and balanced groups to use throughout the analysis (45 self-regarding subjects and 45 other-regarding subjects, with 13 not informed and 32 informed dictators in each group. Groups are also balanced in terms of the main demographic characteristics). Since most dictators have Y score either positive or close to zero, the use of alternative thresholds or number of groups results instead in more unbalanced groups and reduces the power of the analysis. Table I.2 in the Appendix shows results for an analysis based on the alternative split of subjects in three groups (1st group: subjects with Y scores equal to 2.5 and 1.5, 2nd group: Y scores of 0.5 and -0.5 and 3rd group: Y scores of -1.5 and -2.5). It can be seen that the difference in the effect of information persists between highly other-regarding subjects and the intermediate group of more self-regarding subjects. In this case, though, information does not show an effect on the decision of highly self-regarding subjects (3rd group). It should be noted, however, that this last group consists only of few observations (9 subjects).

the observed decline in transfers of this group. A key results of our analysis is then that preferences play an important role in determining the effect of endogenous information on dictators' decision to give and it is therefore important to control for them in order to be able to account for differences in behaviours across individuals. Finally, looking at other variables in the study of the overall sample we find that except for age, none of our controls have an effect on dictators' decision to give. The positive coefficient associated with age is in line with the results of the meta-analysis by Engel (2011) and suggests that older dictators make higher transfers. The effect of age persists also in the analysis of behaviour in the subsample of self-regarding subjects, but disappears in the study of other-regarding subjects. When analysing the decision of other-regarding subjects we observe instead a negative and significant effect associated with the dummy variable for being female. This result is in contrast with other findings from the literature showing that women are more altruistic and inequality averse (Andreoni and Vesterlund, 2001; Croson and Gneezy, 2009; Kamas and Preston, 2015). A possible explanation could be given by the fact that we analyse a subsample of subjects with specific preferences, whereas other studies look at the general sample.

### 3.4.3 Beliefs

We now study receivers' first-order beliefs and dictators' second-order beliefs in order to complement the analysis about decisions developed in the above sections and try to improve our understanding of the underlying motives driving the observed results.

#### Beliefs - Manipulation check

We use beliefs stated by receivers and dictators to check the effectiveness of introducing deservingness in our experiment by exploiting the collocation of a participants' task outcome in the distribution of outcomes in the session. In the experiment we asked receivers to guess the amount they would receive in three different scenarios: when their real outcome from the task is revealed and when a higher or lower outcome are instead revealed. Similarly, dictators were asked to guess the amount that receivers

expect to receive in the corresponding same three scenarios. Figure 3.3 shows respectively receivers' first-order beliefs and dictators' second-order beliefs for the amount expected in the three different scenarios of revealed outcome.

FIGURE 3.3: Average stated beliefs by task outcome and participant's role

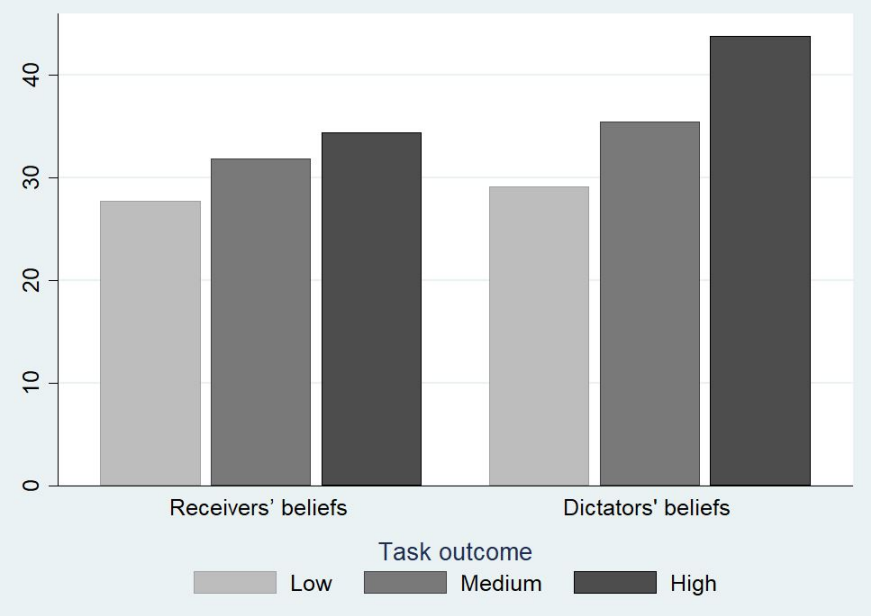


Figure 3.3 shows that beliefs of receivers and dictators are decreasing with the task outcome revealed to the dictator. Receivers expect to receive on average 34.38 tokens when the highest outcome is revealed, 31.83 and 27.74 tokens when instead the intermediate and lowest outcome are respectively revealed to the dictator. The amount that subjects expect to receive across the three different outcome situations is statistically significantly different (p-value<0.0351 across all levels of task outcome, paired t-test). A significant difference can be observed also across dictators' second-order beliefs for the amount expected in the three outcome scenarios (p-value<0.001 across all levels of task outcome, paired t-test). The average dictator's second-order beliefs stated in the case of high outcome is 43.73 tokens, 35.38 tokens in the case of medium outcome and 29.09 tokens in the case the disclosed outcome is low. The decreasing pattern observed across receivers and dictators' beliefs suggests therefore that the experimental design was effective in introducing different perceptions of deservingness in relation to the outcome achieved by a receiver in the initial encoding task.

### Receiver's beliefs

Figure 3.4 shows average beliefs of receivers who choose to disclose their task outcome and of receivers who choose to withhold information. Beliefs of the two groups are presented for four alternative scenarios of task outcomes revealed to the dictator: high, medium, low and no information. Figure 3.4 shows that beliefs of the two groups exhibit the same decreasing pattern already observed when analysing receivers' beliefs across the three levels of outcome (high, medium, low) at the aggregate level. However, while beliefs regarding the choice to not reveal are below beliefs stated for any other level of outcome across subjects who choose to reveal, the beliefs related with no information represent the highest category across subjects choosing to hide their outcome. Table 3.5 complements figure 3.4 and presents average beliefs in the four categories of outcome information for subjects choosing to reveal and subjects choosing to withhold information.

FIGURE 3.4: Average stated beliefs by task outcome and choice to reveal

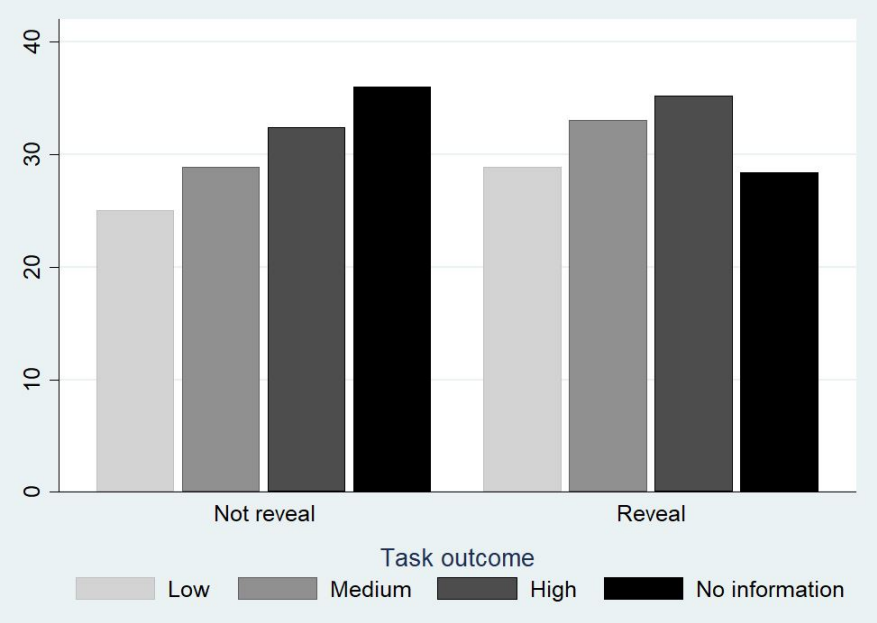


Table 3.5 corroborates the evidence from figure 3.4 and highlights two broad patterns. First, it shows that in all the three outcome scenarios the amount expected by receivers revealing their outcome is higher than the one expected by those who do not

TABLE 3.5: Average stated beliefs by choice to reveal

| | Task outcome | | | | Observations |
|---|---|---|---|---|---|
| | Low | Medium | High | No information | |
| Subjects Not revealing | 25 | 28.88 | 32.42 | 35.96 | 26 |
| Subjects revealing | 28.86 | 33.03 | 35.17 | 28.4 | 64 |
| p-values | 0.4001 | 0.1657 | 0.4307 | 0.3073 | |

reveal. This evidence is in line with the above results of higher other-regarding preferences observed across subjects choosing to reveal and with the phenomenon of false consensus effect often presented in the economic and psychological literature (Engelmann and Strobel, 2000; Butler, Giuliano, and Guiso, 2016). According to this phenomenon subjects tend to believe that their preferences, choices and behaviours match the ones of the larger population. Other-regarding subjects may therefore expect dictators to give higher amount for any level of outcome revealed as this would match their decision to give if they were assigned to the role of dictator. A non-parametric test for the difference in average amount expected by the two groups across the three cases of outcome, however, does not reveal any statistically significant difference (all p-value$\geq$0.1657, two-sample Wilcoxon rank-sum test). Second, Table 3.5 highlights how the average beliefs of subjects not revealing are higher than beliefs of subjects revealing only in the scenario where no information is provided to the dictator. Although beliefs for the amount received when no information is provided do not statistically differ across receivers revealing and not revealing (p-value =0.3073, two-sample Wilcoxon rank-sum test), the latter evidence provides us with some interesting insights about the correlation between the two groups' beliefs and their choice to reveal the task outcome. The data seem, in fact, to suggest that each group chooses the option to reveal associated with highest expectations of amount to receive. We test this hypothesis further by comparing the difference between beliefs stated in the case the real outcome from the task is revealed and beliefs stated in the case no outcome is revealed amongst the two groups. The average difference between beliefs when the real outcome is revealed and no outcome is revealed is 4.23 tokens for subjects choosing to reveal their

outcome and -7.85 tokens for subjects choosing to maintain their outcome secret. The average differential between reveal and not reveal beliefs across the two groups is statistically significantly different (p-value= 0.0219, two-sample Wilcoxon rank-sum test). Thus, this result supports our hypothesis and provide evidence that receivers' decision to reveal incorporates expectations about dictators' action (as discussed in the section above). Receivers in the two groups best respond to their own beliefs and choose the option to reveal that they expect to be associated with the highest premium.

**Dictator's beliefs**

Having studied receivers' beliefs, we now move the analysis to dictators and assess their beliefs in order to understand whether they correctly anticipate receivers' beliefs and their decision to disclose the task outcome. We use dictators' second-order beliefs to obtain a measure of the accuracy with which dictators anticipate recipients' expectations of the amount they will send under the four different information scenarios. Dictators' guess of the number of words encoded by subjects choosing to not reveal and of the rate at which receivers report their task outcome across the four different outcome levels are instead used to assess how well dictators anticipate actual actions of receivers.

**Dictators' second-order beliefs** :

Figure 3.5 presents dictators' average second-order beliefs juxtaposed with receivers' average beliefs in the four different information scenarios. While the comparison of receivers and dictators' beliefs in the low and medium outcome scenarios does not produce any statistically significant difference (two-sided t-test, p-value$\geq$ 0.6182), we observe a statistically significant difference when comparing the two roles' beliefs in the high and no-information scenarios (two-sided t-test, p-value= 0.0017 and p-value= 0.0334 respectively). Specifically, the comparison shows that dictators overestimate the amount expected by recipients by 9.356 tokens in the case of high outcome and by 6.544 tokens in the no-information case. From the results reported above it emerges therefore that there is a discrepancy between the distribution of dictators and receivers' beliefs.

FIGURE 3.5: Dictators and receivers' average stated beliefs by task outcome

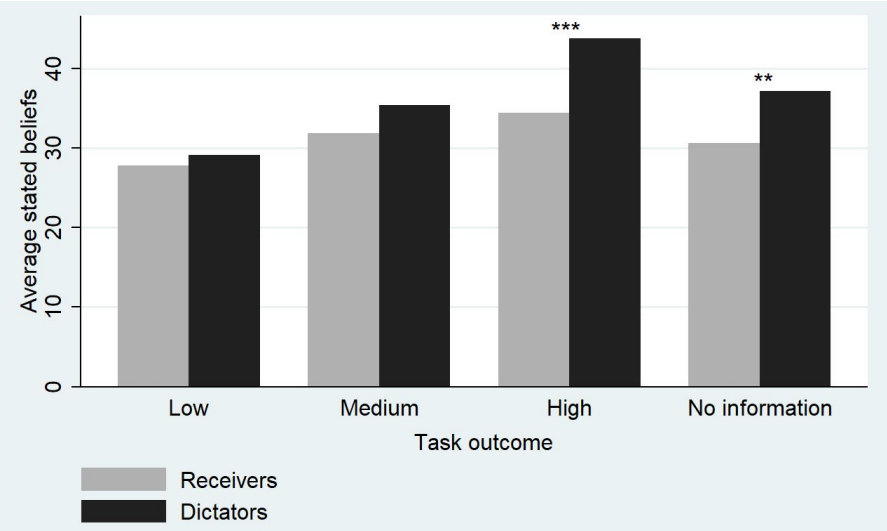


To better understand this result and link it to our findings concerning dictators' behaviour, we next divide our sample of dictators according to the previously used classification of "self-regarding" and "other-regarding" types and separately assess the correctness of dictators' beliefs in the two groups. Previous contributions have reported evidence of the existence of a correlation between dictators' second-order beliefs and behaviour (Guala, Mittone, and Ploner, 2013; Gächter and Renner, 2010). Having previously observed different behaviours amongst self-regarding and other-regarding subjects we could therefore expect to observe similar different patterns also when assessing beliefs of the two groups. Indeed, a comparison of second-order beliefs of the two types of dictators with receiver' firs-order beliefs, shows different results for different types. Other-regarding subjects tend to overestimate the amount expected by receivers across all information levels. The average difference between receivers' beliefs and beliefs of other-regarding dictators is 8.16 tokens and it is statistically significant for any information level (all p-values$\leq$ 0.0887, two-sample Wilcoxon rank-sum test)[30]. The comparison of receivers' average beliefs and beliefs of self-regarding dictators shows instead no significant difference across all level of information, except for the high outcome case (p-value= 0.0405, two-sample Wilcoxon rank-sum test), where

[30]The difference in beliefs for the high outcome is 10.956 tokens, 6.211 tokens for the medium outcome, 5.878 tokens for the low outcome and 9.589 tokens for no-information. The corresponding p-values are: 0.0035, 0.0609, 0.0887 and 0.0045.

also self-regarding subjects over guess the amount that receivers expect by 7.76 tokens. This evidence highlights some interesting insights concerning dictators and receivers' beliefs and behaviours and their interlink in the context of our experiment. Other-regarding type's average beliefs are less correct than those of self-regarding subjects in all outcome scenarios, but in the case of high outcome, where both types overestimate the transfer expected by receivers. While the over guessing of other-regarding subjects could be explained by the phenomenon of false consensus effect described above[31], the fact that the difference in no-information beliefs disappears when self-regarding subjects' beliefs are separately studied, suggests that the overall difference in no-information beliefs reflects the over guessing of the other-regarding type. The persistency instead of the difference in high outcome beliefs across all dictator types highlights an interesting component of inconsistency in beliefs between receivers and dictators. This discrepancy in beliefs amongst the two participants' roles may signal a tendency of receivers to hold too sceptical beliefs for the amount given in the high information scenario and may help to better explain the failure of unravelling of receivers' information in our experiment. As discussed above, low transfers expectations in the high outcome case can, in fact, reasonably result in the choice to not-reveal.

On the other hand, dictators associate high outcomes with expectations of significantly higher transfer and in line with this (and our initial hypothesis), are more likely to link the choice to not reveal with lower task outcomes. The guess of the number of words encoded by subjects choosing to not reveal and the guess of the rate at which receivers report different levels of task outcome show, in fact, that dictators tend to associate the not disclosed outcome with a number of words below the median session number and expect a rate of disclosure that is increasing with the achieved task outcome.

[31] See: Ross, Greene, and House (1977) and Ellingsen et al. (2010) for a discussion on how false consensus effect can explain the correlation between behaviour and beliefs.

**Dictators' guess of not revealed outcome** :

The analysis of the guess of the number of words encoded by receivers choosing to not reveal shows that 84.62% of dictators guess a number of words that is below the median number of the session. However, no dictator guesses a number of words corresponding exactly to the lowest outcome of the session as our initial hypothesis would suggest. On average dictator guess a number that is 13.91 words above the lowest outcome and 5,36 words below the median outcome of the session. An explanation for this result could be given by the specific structure of incentives of this question. Participants may, in fact, try to maximize their likelihood of obtaining additional payoff by exploiting the 10 words interval within which a guess is considered correct, and choosing a value within that interval. Furthermore, the information on the post-shock type of the receiver that is provided to dictators may reasonably cause dictators observing a "High-type" to be more incline to guess a number close to the session median. High type receivers in our experiment could, in fact, be either participants encoding a number of words above the median or participants encoding a number of words below the median, but whose position in the distribution has increased due to the reduction in payoff of other subjects after the shock. Choosing a number equal or above the session median may represent therefore a plausible strategy to try to make a guess closer to the actual outcome of the matched receiver in this information case. Indeed, the only two participants that we observe guessing a number of words above the median were informed that their matched receiver was a high type [32].

[32]The analysis of dictators' guess about the not revealed task outcome is based on a smaller sample (13 observations rather than 24) as in the pilot sessions participants were not given information on the top and bottom outcome. To simplify the comparability of results across sessions we therefore focus the analysis in the main text only on observations from the main sessions. Although the same conclusions do hold when the complete dictators' sample is considered. When we look at the total sample we observe that 79.17% of dictators guess a number of words below the median, the average distance from the bottom outcome becomes 25 and the distance from the median outcome decreases to 4.5. Two out of the three dictators who make a guess above the median were informed that the matched receiver was a High type.

**Dictators' guess of reporting rates** :

The evidence that dictators associate the choice to not reveal with low task outcomes is furthermore corroborated by the dictators' responses to the question concerning the guess of the rate with which receivers report their outcome across five different outcomes categories (lowest session outcome, outcome between the lowest and the median outcome, median outcome, outcome between the median and highest outcome and highest outcome). Figure 3.6 shows the distribution of dictators' guesses across the different outcome categories with weighted markers and a regression line. Larger markers indicate a higher number of guesses made in an interval of 5 percentage points. While dictators' beliefs show some heterogeneity, we can observe from the upward regression line that the reporting rate and the task outcome level are positively correlated with each other. Dictators seem therefore to expect a high reporting rate amongst best performing subjects and instead a significantly lower one amongst those subjects achieving the lowest outcome in the session [33].

FIGURE 3.6: Dictators' guess of receivers' disclosure choice across different outcome levels vs. actual disclosure rates

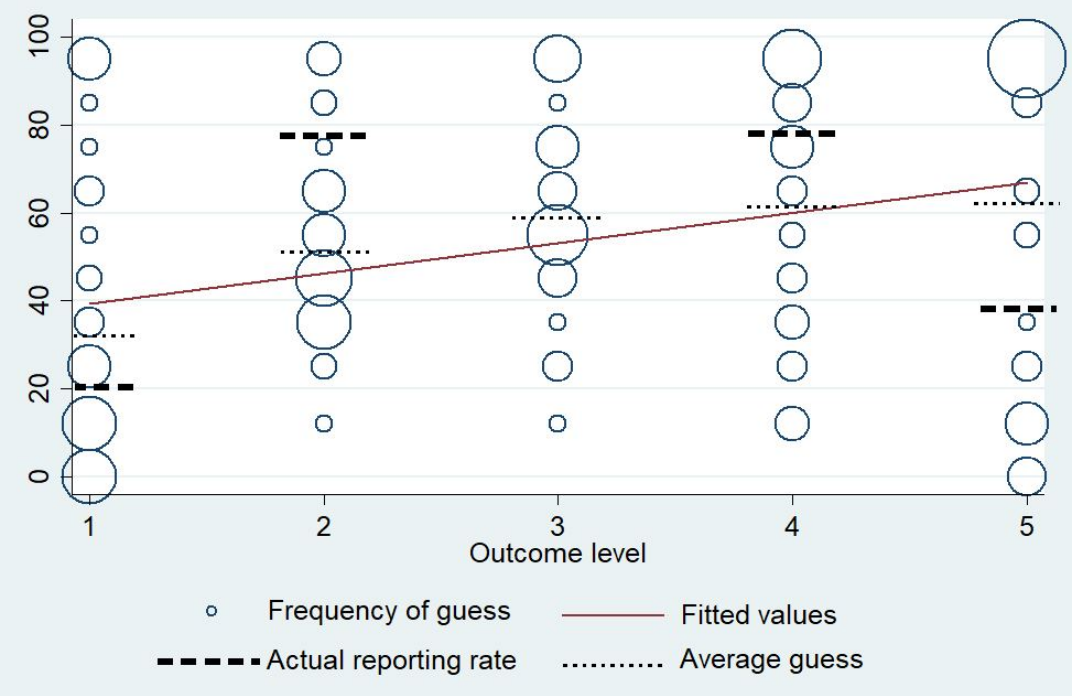


[33]The data reported in the figure do not include guesses made in the pilot sessions as the question has been introduced later in the pilot session and changed in the main sessions due to a difficulty with the wording of the text.

The dashed line and the dotted line in Figure 3.6 represent furthermore the actual reporting rate and the average guess across four outcome levels respectively[34]. Although dictators seem to under guess the revealing rate for the two central categories of outcome and slightly over guess the choice to reveal when the outcome is the lowest, the two lines appear overall to resemble each other in an increasing pattern across this first three outcome levels. A sharp contrast between actual and guessed reporting rate emerges instead in the highest outcome category. Here, in fact, we see the actual reporting rate drastically dropping while dictators' guess remains constant. As a result dictators overestimate the reporting rate by more than 20 percentage points. This result may be interpreted as further evidence of the discrepancy between receivers and dictators' high outcome beliefs that we have already highlighted in different sections of our analysis.

**Receivers' beliefs and dictators' decision**

In the final step of the analysis of participants beliefs, we link receivers' beliefs and dictators' actions in order to gather a measure of correctness of receivers' expectations. Our measure of correctness of receivers' beliefs is obtain by the comparison of the average receivers' first-order beliefs and average amount sent by dictators. Since in the experiment we only observe dictators' decision to give corresponding to the outcome scenario disclosed to them, in this part of the analysis we consider only receivers' beliefs stated with respect to the outcome scenario actually shown to dictators. The average receivers' expected transfer is 33.22 tokens and the average corresponding dictator' decision to give is 35.07 tokens. A mean test fails to reject the null hypothesis of no difference in the number of tokens expected and transferred (two-sided t-test, p-value= p=0.519). Beliefs of receivers seem therefore to be overall correct. The same result holds also if we assess beliefs and actions separately in the scenario of no-information and information. The comparison between beliefs and actions does not produce statistically significant difference in any information scenario (p-value= 0.5893 for no-information

[34]We only consider four outcome levels in this part of the analysis as no receiver in our sample achieved exactly an outcome corresponding to the median session outcome.

scenario and p-value= 0.1734 for informed scenario, two-sample Wilcoxon rank-sum test).

Having observed in previous results that most of the difference in dictators' behaviour occurs when different preferences types are taken into account, we extend the current analysis and also include dictators' preferences classification. We try therefore to understand whether being matched with different categories of dictators has an effect on the correctness of receivers' beliefs. Table 3.6 presents average transfers of other-regarding and self-regarding dictators and the corresponding beliefs stated by the matched receivers in the two information scenarios. A Wilcoxon rank-sum test does not find any significant difference between receivers' beliefs and both other-regarding dictators' choice in the no-information scenario and self-regarding dictators' choice in the information scenario (p-value= 0.360 and p-value=0.493 for no-information and information scenario respectively). A significant difference can be observed instead when receivers' beliefs are compared with both self-regarding dictators' choice in the no-information scenario and other-regarding dictators' choice in the information scenario (p-value= 0.018 and p-value= 0.008 respectively, two-sample Wilcoxon rank-sum test). Specifically, we observe receivers under guessing the amount transferred by other-regarding dictators when information is provided and overestimating instead the amount sent by self-regarding types when no information is disclosed. First of all, these results corroborate our findings concerning dictators' different behaviours across the two preference categories (other-regarding dictators make consistent higher transfers in all information scenarios, while self-regarding subjects significantly decrease their transfers when information on the receiver's task outcome is not provided). Second, the results highlight how and in which direction behaviours of different types of dictators may trigger wrong expectations across receivers.

TABLE 3.6: Average receivers' stated beliefs and dictators' transfer decisions by information scenario and dictators' type

| | No information | | Information | |
|---|---|---|---|---|
| | other-regarding match | self-regarding match | other-regarding match | self-regarding match |
| Receiver belief | 35.83 | 31.28 | 30.58 | 33.97 |
| Dictator transfer | 41.75 | 19.33 | 42.79 | 30.64 |
| p-values Dictator=Receiver | 0.3605 | 0.0175 | 0.0077 | 0.4937 |
| observations | 12 | 12 | 33 | 33 |

### 3.4.4 The effect of endogenous information on dictators' transfer

We now conclude our analysis by directly assessing the effect of endogenous information on the decision to give of dictators. In order to do so, we compare dictators' decisions to give in our main treatment (endogenous treatment) with decisions made in the two exogenous control treatments. In these treatments, information does not depend on receivers' choice to disclose and it is either always or never provided to the dictator. Table 3.7 illustrates average transfer in the exogenous and endogenous treatment both for the scenario when information is provided and when it is not provided. The last two rows of the table report p-values for the test of difference in mean transfers between endogenous and exogenous treatment in different information scenarios. The table shows that overall dictators give on average 3.73 more tokens in the endogenous treatment than in the exogenous treatment. However, the overall difference in average transfer between the two treatments is not statistically different from zero. If we stratify the comparison by information scenario, though, two slightly different patterns emerge. In the no-information scenario dictators give more in the exogenous treatment than in the endogenous one. The difference in average transfer is only 1.91 tokens and it is not statistically significant. When assessing average transfer amongst the two treatments in the scenario with information it can be seen instead that dictators give an average of 7.05 more tokens in the endogenous treatment than in the exogenous treatment and the difference is statistically significant (p-value= 0.0775, two-sample Wilcoxon rank-sum test). These findings are also confirmed if we hold

constant demographics characteristics and experiment conditions across treatments[35].

TABLE 3.7: Transfer decisions by information scenario and origin of information

| Treatment | Information | N | Mean transfer |
|---|---|---|---|
| Endogenous | Yes | 66 | 36.71 |
| | No | 24 | 30.54 |
| | Total | 90 | 35.07 |
| Exogenous | Yes | 29 | 29.66 |
| | No | 44 | 32.45 |
| | Total | 73 | 31.34 |
| Hypotheses and p-value Endogenous=Exogenous | | | |
| | Yes | 0.0775 | |
| | No | 0.4853 | |
| | Total | 0.3429 | |

Our analysis shows that endogeneity of information does not affect dictators' giving decision both when behaviour is studied at the aggregate level and in the no-information scenario. Interestingly, however, we find that endogeneity of information has a significant influence on transfers of dictators in the scenario where information is available. Dictators that receive information as a result of receivers' decision to disclose, in fact, give significantly higher amounts than dictators that receive information exogenously. This seems to suggest that dictators reward receivers for choosing to disclose, giving them more than they would have given if the information was already available in the experiment setting. A potential reason for the increase in transfers in the endogenous treatment might be that some dictators care about supporting specific receiver's types and achieving an outcome distribution that matches their preferences. These dictators might therefore be willing to transfer more to receivers that choose to provide them with the information necessary for their choice. In line with this, Corneo and Fong (2008) use fiscal data to estimate the willingness to pay for distributive justice and find that the achievement of a more just distribution of income is associated with a significant positive value. Similarly, in the context of incomplete information Thunström et al. (2016) show that dictators actively choose to disclose information about

[35]See Table I.3 in the Appendix, which reports results for an OLS regression of average transfer on treatment dummies and other control variables.

receivers' deservingness, while Fong and Oberholzer-Gee (2011) demonstrate that dictators are willing to invest resources to obtain information on their receivers' type.

## 3.5 Conclusion

In this paper we experimentally study distributional decision in the context of a dictator game, where information about the deservingness of the receiver is endogenously determined by the decision of the receiver to disclose the information. We vary deservingness in the eyes of participants by exploiting differences in the investment of effort that subjects make in completing a real-effort task, and by presenting the endowment to split as the reward given for the completion of the task.

We therefore explore whether receivers voluntarily provide information about their outcome from the task (deservingness) to dictators before they decide on the split of the endowment, and assess the effect of information, or lack of it, on dictators' transfer decisions.

Our paper contributes to the understanding of giving decisions in the real-world and of the effect that information about receiver's deservingness has on them. It adds to the existing literature in two ways. First, it offers an analysis of distributional choices in a context of endogenous information about receivers' deservingness that has received little or no attention in the literature - when the information is determined by the receiver choice to disclose the information. Second, it provides a novel analysis of the extent to which receivers value information about their deservingness and are willing to disclose it. Production and dissemination of information about deservingness of the receiver are central questions in the debate of the literature about distributive behaviour in the real-world. To our knowledge, ours is the first paper that assesses how receivers perceive this information and how they decide to reveal it.

We find that subjects achieving outcomes in the task that are close to the median outcome of a session are significantly more likely to disclose their outcome than subjects having instead outcomes distant from the median. A plausible explanation of receivers' decision to disclose information could be found in the literature on in-group

bias. Evidence from this literature shows that information that dictators and receivers have about relative group identities can strongly affect decisions to give and expectations about the amount to receive. If the group membership is common knowledge, subjects have revealed to prefer to be matched with same group members. However, if information about group identity is not common, there is evidence suggesting that subjects are indifferent on the group identity of the matched subjects. All these results have been shown to be in line with expectation of a better treatment within the same group (Foddy, Platow, and Yamagishi, 2009). The categorisation by outcome in our experiment might have induced some sort of group identifications across participants. Participants with outcomes closer to the extremis of the distribution, might then have perceived higher the probability to be matched with a subject with a very different outcome. In order to maximise their outcome from the game they might have therefore chosen to keep their outcome-type secret. Indeed, our results suggest that receivers' decision to disclose can be imputed to the expectation of a better outcome associated with the chosen option. The receiver's decision to disclose information about her outcome is strongly related to the receiver's stated beliefs about dictators' action. Both receivers choosing to reveal and to hide their outcome, in fact, choose the disclosure option that they expect to lead to the highest transfer from the matched dictator. Our results provide interesting insights on whether and how subjects perceive a benefit to revealing information about their deservingness. However, our experimental design does not allow us to directly disentangle across different mechanisms behind the decision to reveal information. One important question would be to further investigate the extent to which common knowledge of receivers and dictators' outcome affect the decision to reveal. We leave this question for future research.

Concerning dictators' behaviour, when we assess behaviour of the total sample, we do not observe any difference in transfer decisions between dictators that receive information and dictators that are instead uninformed. We, however, observe some effect of endogeneity of information in our context, when we assess the behaviour of other-regarding and self-regarding participants separately. In particular, we find that self-regarding subjects significantly decrease their transfers when information is

not provided, while lack of information does not affect the decision to give of other-regarding subjects. Similarly to Fong and Oberholzer-Gee (2011), our experiment provides evidence of a different effect of endogenous information on giving decisions of individuals having different preferences. However, while their experiment documents this effect in the case where endogeneity of information is determined by the dictator decision to disclose information, we extend their finding to the case where the availability of information is determined by the choice of the receiver to reveal. We provide therefore evidence that individuals' preferences represent an essential component that needs to be accounted for in order to better explain subjects' behaviour in contexts of endogenous information.

We furthermore corroborate Fong and Oberholzer-Gee (2011)'s finding concerning beliefs of more selfish subjects in situations of uncertainty of the information about the receiver's deservingness. Fong and Oberholzer-Gee (2011) argue that uncertainty of information and heterogeneity in the deservingness of receivers provide a plausible excuse to give less and to behave as if all the matched receivers are characterized by low deservingness. In our experiment we directly study beliefs of dictators about the not disclosed outcomes and find that a possible explanation for the decrease in transfers of self-regarding subjects is the expectation that a not disclosed outcome is in reality a low outcome.

Finally, we find that endogeneity of information about the receiver's deservingness has a positive and statistically significant effect on dictators' giving decisions compared to decisions made with information exogenously provided. Our finding is in line with the literature and suggests that individuals value information about the deservingness of the receiver and achieving a distribution of resources that matches their preferences. Fong and Oberholzer-Gee (2011) and Corneo and Fong (2008) show that subjects are willing to invest in the information about the deservingness of the receiver, similarly our setting suggests that subjects are willing to sacrifice part of their payoff to reward the receiver who has provided them with the information needed to achieve a preferred allocation.

# Appendix A

# Chapter 1 Appendix A: Experimental Instructions

## General information

This is an experiment in economic decision making which has been funded by an EssexLab Seedcorn Grant and British Academy Leverhulme Small Research Grant. If you read the following instructions carefully you can, depending on your decisions and decisions of other participants, earn a considerable amount of money. This money will be paid to you in cash at the end of the experiment.
These instructions are solely for your private use. During the experiment you are not allowed to communicate with the other participants. If you have any questions, please, raise your hand. The experimenter will come to you and answer them. Any violation of this rule will exclude you from the experiment and from all payments.
During the experiment you will make choices about amounts of tokens. Your earnings will be calculated in tokens and then converted into pounds at the following rate: 10 Tokens= 0.90 pounds.
At the end of the experiment your earnings and an additional show up fee of 4£ will be paid to you in cash.
The confidentiality of your decisions will be maintained during the experiment as none of the other participants will know which decisions you made.

## Experiment instructions

The experiment you are about to participate in consists of **4 different decision environments**, namely, **Game A, Game B, Game C, Game D**.
Below, we will describe each one of these decision environments separately

**Game A Instructions**
Game A will require you to make decisions in **2 different steps: Step 1 and Step 2**. You will earn a separate payoff in each step. Total earnings from Game A will be determined by the sum of Step 1 payoff and Step 2 payoff.
**Step 1**
At the beginning of Step 1 you will be matched with **3 other participants** in this room. Therefore, you will be part of a **group** consisting of **4 people**. No participant will know who is in which group.
You and each other member of the group will be given an **endowment of 20 tokens** and will be asked to decide how much of this to invest in a **Private account** and how

much to contribute to a **Public account**.
Every token you invest in the **Private account** will give you a **return of 1**. One token invested in the Private account will therefore generate 1 token of earnings.
Example: if you decide to invest 10 tokens in the Private account, you will obtain 10 tokens of earnings.
**No one except yourself will benefit from investments made in your Private account.**
Every token contributed to the **Public account** will instead give a **return of 0.4** to each group member. If you decide to contribute 1 token to the Public account, you and each of the other group members, will gain 0.4 tokens of earnings. Similarly, you will obtain 0.4 tokens of earnings for each token that the other three participants will contribute to the Public account. Therefore, **every participant will equally profit from contributions made in the Public account**.
Specifically, your earnings from the Public account will be determined as:

**Your earnings = (sum of all group members' contribution to the Public account) x 0.4**

Example: if you contribute 10 tokens to the Public account and the sum of the contributions of the other three group members is 30 tokens, you (and each of the other group members) will earn 16 tokens.

Your earnings = [(10+30) x 0.4] = 16 tokens

Finally, your **total earnings** will be calculated as the sum of earnings from the Private account and from the Public account:

Earnings from Private account (20 tokens - tokens contributed to the Public account)
+ Earnings from Public account (group contribution to the Public account x 0.4)
= Total earnings

In the decision environment described above you will be asked to make two decisions: one **Unconditional** and one **Conditional**.
1- **Unconditional investment decision**:
You will be asked to decide how many of the 20 tokens of your endowments you want to invest in the Private account and how many of them you want to contribute to the Public account.
2- **Conditional investment decision**:
You will be asked to decide how many of the 20 tokens of your endowments you want to contribute to the **Public account** given **any possible average contribution to the Public account of the other three participants** (rounded to the nearest integer). Therefore, you will be required to state a choice of contribution **conditional on each possible average contribution decision** of the other group members. You will do so by fill in each row of a table reporting the 21 possible average contribution choices of the other three participants.

**Both** your **Unconditional and Conditional decision could be relevant for your final payoff**. It is therefore important that you pay attention to each one of them.
Once all participants have completed their decision tasks, one subject will be randomly selected in each group. The final payoff of each group member will then be calculated taking into account the "Conditional investment decision" of the selected subject and the "Unconditional investment decision" of the other three participants in the group.

**Step 2**

Once you have completed Step 1 you will be asked to make decisions in the following step, Step 2. **Step 2** consists of **10 decision making periods**. At the beginning of the decision process you will be **matched** with **3 other participants** in this room. Therefore, you will be asked to take decisions in a group consisting of 4 people. The **composition of the group will remain unchanged across the 10 periods** of the decision process, i.e. you will play with the same participants for 10 periods.
The decision environment is the same as the one described above for the Unconditional decision in Stage 1 (see page 1).
Now, though, once you have made your contribution decision, you will be **also asked to guess the average** (rounded to the closer integer) **contribution to the Public account** of the **other** three **group members** in that period.

It is important that you try to give correct estimation of others' contribution, as a **correct guess could increase your final earnings**.
At the end of the 10 decision periods, **one period will be randomly selected** and **additional earnings** will be estimated according to the accuracy of your estimation of other's contribution in that period:
- If you **correctly estimate** the average contribution of the other players in that period you will increase your earnings by **3 tokens**.
- If you make an estimation that is **1 token higher or smaller** than the actual others' average contribution, you will increase your earnings of **2 tokens**.
- If you make an estimation that is **2 token higher or smaller** than the actual others' average contribution, you will increase your earnings of **1 tokens**.
- If you make an estimation that is higher or smaller by **more than 2 tokens** than the actual others' average contribution, you **will not increase your earnings**.

The structure of the experiment is the same in all subsequent 9 periods: at the beginning of each period you are given 20 tokens and you are asked to decide how many of these to contribute to the Public account and how many to invest in the Private account. Once your contribution choice is made you are asked to state your guess for the other three participants' contribution decision in that period.
In each period your total earnings are calculated as the sum of the earnings from Private account and earnings from Public account.
Before moving to the next period, **you are always informed about your earnings** in the current period and the actual average contribution of the other group members.

The **sum of the 10 periods earnings plus** additional earnings from the **correct guess** of the randomly selected period, **generate the payoff of Step 2**.
Your total earning from Game A are equal to the payoff obtained from Step 1 plus the payoff obtained from Step 2.
Once you have completed all decisions of Game A, the experiment will require you to move to Game B. Game B is characterized by a different decision environment. We will describe it next.

---

You will receive instructions for Game B, Game C and Game D on your screen at the beginning of these parts.

Your **total earning in the end** of the experiment will be equal to the **sum of the earnings gained in two decision environments**.
The computer will randomly select two decision environments out of the four decision environments (Game A, Game B, Game C, Game D) you made decisions in. The earnings you obtained in each of the two selected decisions environments will be summed up and give your final earnings from the experiment.

At the end of the experiment there will be a short questionnaire for you to fill in.

**Game B Instructions**

**Game B** will require you to make decisions in **12 different decision situations**.
You will take **2 decisions** for **each decision situation**. Therefore, you will be asked to make a **total of 24 decisions**.
**Each decision corresponds to a new Game.**
Each decision you take is independent from the others: decisions that you will take will not affect neither the following situations nor the following payoffs.

**Each decision situation** is characterized by **2 different roles: role A and role B**. In each decision situation, you will be asked to **play** sequentially **both roles** and thus, to take 2 different decisions.
For each role, you will be randomly matched with a different participant than in the previous role.

**In any** case of **decision**, you will be **randomly paired with one other participant** so that your decisions may affect the payoff of the other participant and the other participants' decisions may affect your own payoff.

Some decision situations only have decisions for one role, whereas other decision situations have multiple decisions. In decision situations with multiple decisions, decisions will be made sequentially: player A will make a decision first and, next, player B will make a decision.
However, player B will not be informed of player A decision before making her decision.

You will **not be informed of the outcome of any decision** you made until the end of the experiment.

Once the Game is concluded, you will have a total of **24 outcomes** (2 for each decision situation). Your final earnings will be determined by the **random selection of 2 payoffs** across all the available 24.

**Example of decision situation (decisions for one role)**:

**In this game, you are player A.** You have no choice in this game. Player B's choice determines the outcome. If player B chooses B1, you would receive 80, and player B would receive 20. If player B chooses B2, you would each receive 0.

**In this game, you are player B**. You may choose B1 or B2. Player A has no choice in this game. If you choose B1, you would receive 20 and player A would receive 80. If you choose B2, you would each receive 0.

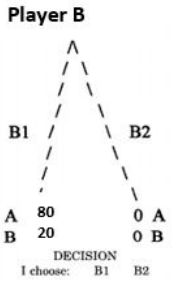


**Example of decision situation (decisions for both roles)**:

Once you have completed all decisions of Game B, the experiment will require you to move to Game C. Game C is characterized by a different decision environment. Below we will describe it.

## Game C Instructions

**Game C** will require you to make decisions in 2 different decision tasks: **Task 1 and Task 2**. For each decision task, you will be **paired** with a **different participant** in this room. No participant will know who is in which pair.
In each pair one person will be assigned to **role A** and one to **role B**.
You will **play both roles**. In one Task (1 or 2) you will be assigned the role of A, while in the other Task you will have the role of B. You will complete the two tasks sequentially.
We first introduce the **basic decision situation** and then explain what the 2 tasks consist of.

**In this Game, you are player A**. You may choose A1 or A2. If you choose A1, you would receive 75 and player B would receive 0. If you choose A2, then player B's choice of B1 or B2 would determine the outcome. If you choose A2 and player B chooses B1, you would each receive 40. If you choose A2 and player B chooses B2, you would receive 75 and he or she would receive 37.5. Player B will make a choice without being informed of your decision. Player B knows that her/his choice only affects the outcome if you choose A2, so that she/he will choose B1 or B2 on the assumption that you have chosen A2 over A1.

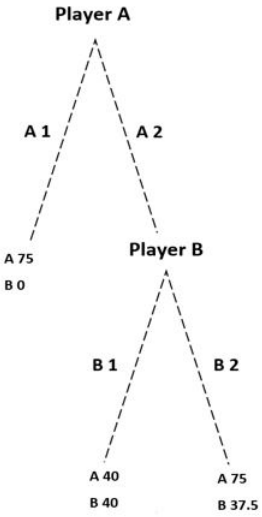


DECISION
I choose: A1 A2

**In this period, you are player B**. You may choose B1 or B2. Player A has already made a choice. If he or she has chosen A1, he or she would receive 75 and you would receive 0. Your decision only affects the outcome if player A has chosen A2. Thus, you should choose B1 or B2 on the assumption that player A has chosen A2 over A1. If player A has chosen A2 and you choose B1, you would each receive 40. If player A has chosen A2 and you choose B2, then player A would receive 75 and you would receive 37.5.

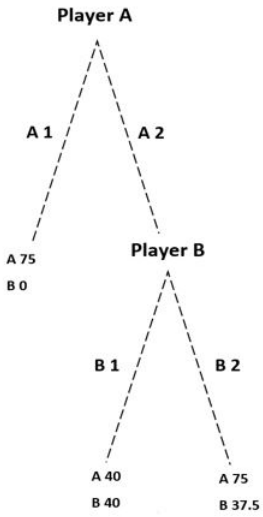


DECISION
I choose: B1 B2

**Decision situation:**
At the beginning of each task **participants** are **paired** and **assigned** each to a different role (**role A or role B**).

The participant having **role A** is **given an endowment of 10 tokens** and is asked to decide whether to **pass some, all or none** of this **to Participant B**. Participant A can pass **any integer** number between **0 and 10**.

Each **token** that is **passed** to Participant B **is tripled** by the experimenter before it reaches Participant B.

Example: if Participant A passes 5 tokens, Participant B receives 15 tokens.

**Participant B** then receives the passed tokens and is asked to decide how many of them to **return to Participant A**. Participant B can **return any integer** number between **0 and 30**.

**Participant A's payoff** is then given by the amount of **tokens** that she/he does **not pass plus** the amount of **tokens** she/he receives **back from Participant B**.

Participant B's payoff is instead given by the amount of tokens that she/he receives and decides not to return to Participant A.

**Tasks structure:**
You will play **2 tasks** (**Task 1** and **Task 2**), each one of these **with** a **different role** (**Role A** or **Role B**) and a different pairing.
- In the **Task** where you are assigned the **role of A**, you will be given an **endowment of 10 tokens** and asked to decide how many tokens to pass to Participant B.
- In the **Task** where you are assigned the **role of B**, you will be asked to decide how many tokens to **return** to Participant A, **given any possible** (integer) **amount** of tokens Participant **A could have passed** to you. However, you will **not be revealed** the **amount** of tokens player **A has actually passed** to you before making your decision.

At the end of **Game C** the computer will randomly select only one Task and your **final earnings** will be given by the payoff you earned in that task.

Once you have completed all decisions of Game C, the experiment will require you to move to Game D. Game D is characterized by a different decision environment. Below we will describe it.

**Game D Instructions**

**Game D** will require you to make a **single decision**.

You will be shown a table reporting **6 different gambles** and you will be asked to **choose** the **one** that you would like to play. **Each gamble** is associated with **2 possible** outcomes: **payoff A** and **payoff B**.

For each gamble, you will be shown the possible outcomes and the associated **probabilities**.

In each gamble each outcome occurs with a **50% chance**.

You will be required to **select** one and **only one gamble** amongst the alternative 6 available.

Your **payoff** will be determined both by the gamble you choose and the outcome of this game.

In particular, at the end of Game D, the computer will select **with 50%** chance one outcome (payoff A or payoff B) for each gamble. Your **final earning** from **Game D** will then be determined by the payoff associated with the **realized outcome** in your **chosen gamble**.

**This is the end of the instructions.**

FIGURE A.1: Screen shot of the Public Goods Game Conditional contribution decision

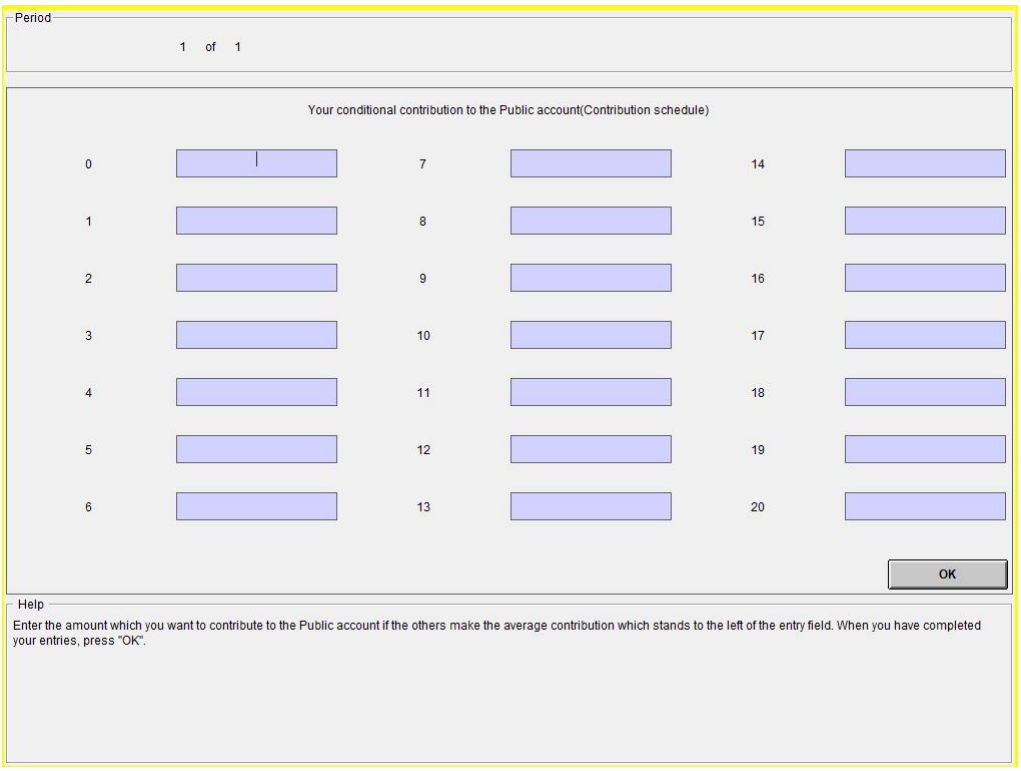

FIGURE A.2: Screen shot of a dictator game from the Response Game

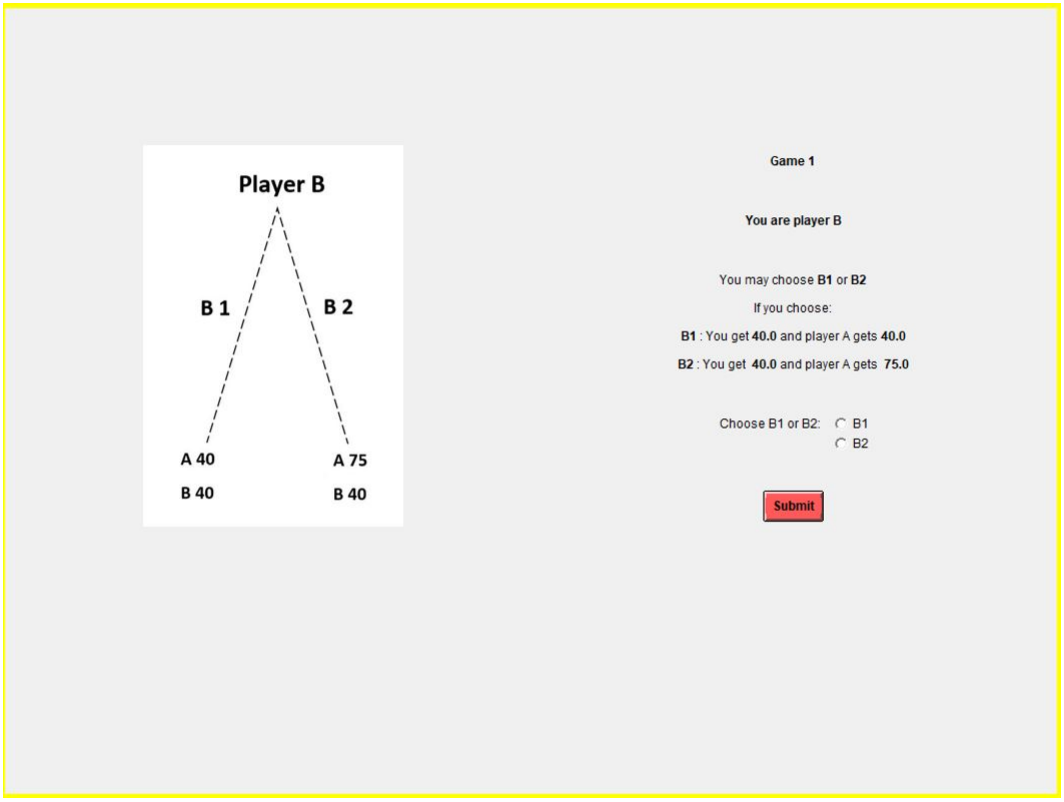


FIGURE A.3: Screen shot of a response game game from the Response Game

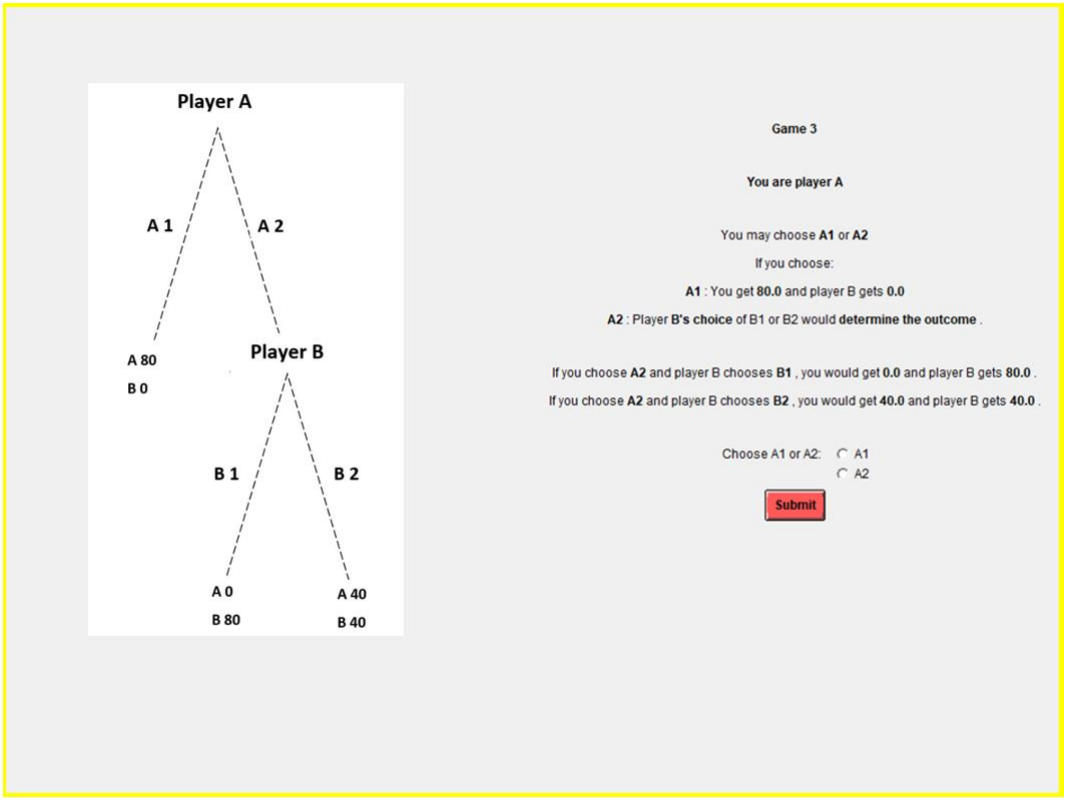

FIGURE A.4: Screen shot of Player A's choice from the Investment Game

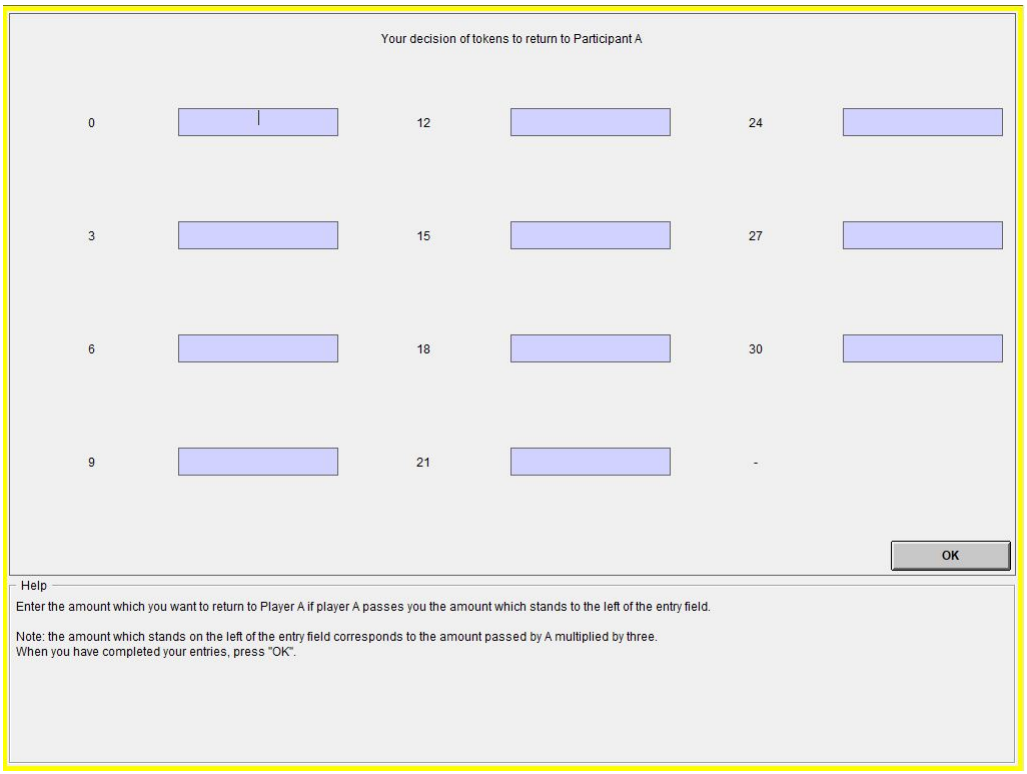


FIGURE A.5: Screen shot of the Risk Task

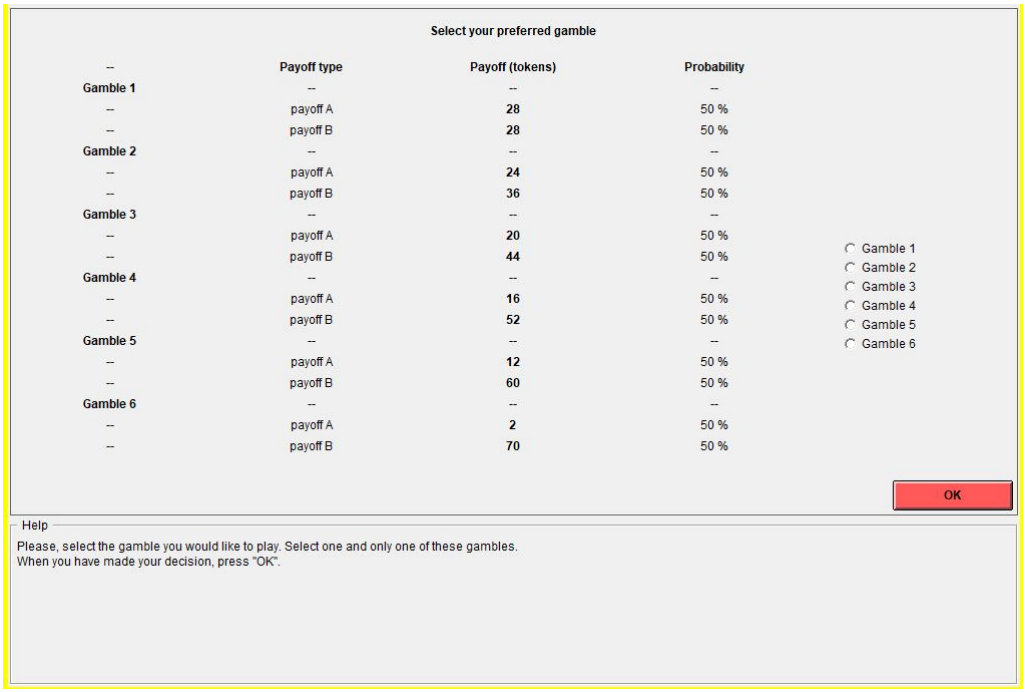

# Appendix B

# Chapter 1 Appendix B: Additional tables

TABLE B.1: Response Game - Games description

| . | A stays out | If A enters, B chooses |
|---|---|---|
| Two person Dictator games | . | . |
| Dict 1 | | (40; 40) vs. (75; 40) |
| Dict 2 | | (40; 40) vs. (75; 37.5) |
| Dict 3 | | (30; 60) vs. (70; 50) |
| Dict 4 | | (20; 70) vs. (60; 60) |
| Dict 5 | | (0; 80) vs. (40; 40) |
| Two person Response games: | . | . |
| B's payoffs identical | . | . |
| Resp 1 | (75; 0) | (40; 40) vs. (75; 40) |
| Resp 2 | (55; 55) | (40; 40) vs. (75; 40) |
| Resp 3 | (10; 100) | (7.5; 12.5) vs. (12.5; 12.5) |
| Resp 4 | (45; 90) | (20; 40) vs. (40; 40) |
| Two person Response games: | . | . |
| B's sacrifice helps A | . | . |
| Resp 5 | (75; 0) | (40; 40) vs. (75; 37.5) |
| Resp 6 | (55; 55) | (40; 40) vs. (75; 37.5) |
| Resp 7 | (75; 10) | (30; 60) vs. (70; 50) |
| Resp 8 | (70; 20) | (20; 70) vs. (60; 60) |
| Resp 9 | (80; 0) | (0; 80) vs. (40; 40) |
| Resp 10 | (0; 80) | (0; 80) vs. (40; 40) |
| Resp 11 | (72.5; 0) | (40; 40) vs. (75; 37.5) |
| Resp 12 | (45; 0) | (35; 45) vs. (45; 35) |
| Two person Response games: | . | . |
| B's sacrifice hurts A | . | . |
| Resp 13 | (37.5; 100) | (40; 40) vs. (35; 35) |
| Resp 14 | (40; 120) | (40; 20) vs. (0; 0) |
| Resp 15 | (37.5; 100) | (40; 40) vs. (25; 35) |
| Resp 16 | (75; 75) | (80; 20) vs. (0; 0) |
| Resp 17 | (75; 75) | (80; 20) vs. (0; 5) |
| Resp 18 | (75; 75) | (80; 20) vs. (0; 10) |
| Resp 19 | (75; 75) | (80; 20) vs. (0; 15) |

Table B.2 shows the logit model we use to estimate the probability that player B rewards player A in games of positive reciprocity.

Models 1 and 2 present the results of logit regressions where the dependent variable is the probability Player B rewards Player A.

Model 1 looks at the effect that factors such as A's benefit from B's choice to reward (which is equal to the gain in A's payoff if B decides to reward), the difference between B and A's payoff if B rewards and the cost of choosing to reward (which is calculated as the difference in B's payoff when B chooses to reward over the alternative action), have on the probability of choosing the action to reward. All these factors have a statistically significant effect on B's likelihood to reward (p-value $< 0.01$).

Model 2 presents instead the results of the logit regression where we add as independent variable the dummy for only-child and we control for demographic and family background characteristics. Results of model 2 suggest no significant effect of being an only-child on the likelihood of positive reciprocity.

TABLE B.2: Logit regression determinant of reciprocity

| Probability B rewards A | | |
|---|---|---|
| | (1) | (2) |
| Only_child | | -0.070 |
| | | (0.057) |
| benefit to A if B rewards | 0.064*** | 0.066*** |
| | (0.013) | (0.013) |
| B's payoff behind of A if rewards | -0.057*** | -0.061*** |
| | (0.013) | (0.012) |
| cost to reward | -0.080*** | -0.085*** |
| | (0.019) | (0.018) |
| Female | | -0.185*** |
| | | (0.046) |
| age | | 0.001 |
| | | (0.004) |
| Observations | 966 | 944 |
| Controls | No | |
| Ethnicity | | Yes |
| Nationality | | Yes |
| Religion | | Yes |
| Time_home | | Yes |
| Par_maritalStat | | Yes |
| AgeMoth_birth | | Yes |
| ParentsEdu | | Yes |
| Log-likelihood | (-598.75) | (-545.63) |
| Pseudo R | 0.024 | 0.093 |

Tables show marginal effects.
Robust standard error clustered on subject level.
*** p<0.01, **p<0.05, * p<0.1

# Appendix C

# Chapter 2 Appendix C: Additional details learning models

In this Appendix we provide additional results for the learning models discussed in Section 2.2.3 as well as more details on the simulations discussed in that Section.

## C.1 Comparative Statics of Choice Probabilities

In a first step consider the reinforcement model without updating of counterfactuals. Assume that agent $i$ chooses $GT$ in match $m$ and consequently earns a payoff of $\pi^{i,m}(GT, s^{-i,m}, T^m)$. Then the change in choice probabilities from match $m$ to match $m+1$ is given by:

$$\Delta p_{GT}^{i,m} = p_{GT}^{i,m+1} - p_{GT}^{i,m} = \frac{\psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m)}{\psi_{GT}^{i,m} + \psi_{AD}^{i,m} + \pi(GT, s^{-i,m}, T^m)} - \frac{\psi_{GT}^{i,m}}{\psi_{GT}^{i,m} + \psi_{AD}^{i,m}}$$

We are interested in comparing the impact in early matches to later matches, keeping the environment constant. Thus, we keep the ratio of propensities fixed by setting $\psi_{AD}^{i,m} = \lambda \psi_{GT}^{i,m}$ where $\lambda > 0$ is a constant. Thus,

$$\Delta p_{GT}^{i,m} = \frac{\lambda \pi(GT, s^{-i,m}, T^m) \psi_{GT}^{i,m}}{\left(\psi_{GT}^{i,m} + \lambda \psi_{GT}^{i,m}\right)\left(\psi_{GT}^{i,m} + \lambda \psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m)\right)}$$

Generally, propensities will be larger in later matches. Thus, in order to compare the impact of earning some payoff early to earning the same payoff later on (for a fixed ratio of propensities), we consider the change in $\Delta p_{GT}^{i,m}$ as the level of propensities $\psi_{GT}^{i,m}$ varies:

$$\frac{\partial \Delta p_{GT}^{i,m}}{\partial \psi_{GT}^{i,m}} = -\frac{\lambda \psi_{GT}^{i,m}}{\left((1+\lambda)\psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m)\right)^2} < 0.$$

Thus, the larger propensities are, the smaller the impact of realized payoffs on choice probabilities will be.

We next turn to reinforcement learning with updating of counterfactuals where $i$ also updates the propensity for $AD$. The change in choice probability from match $m$ to match $m+1$ is now given by:

$$\Delta p_{GT}^{i,m} = p_{GT}^{i,m+1} - p_{GT}^{i,m} = \frac{\psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m)}{\psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m) + \psi_{AD}^{i,m} + \pi(AD, s^{-i,m})} - \frac{\psi_{GT}^{i,m}}{\psi_{GT}^{i,m} + \psi_{AD}^{i,m}}$$

Again we set $\psi_{AD}^{i,m} = \lambda \psi_{GT}^{i,m}$. Thus,

$$\Delta p_{GT}^{i,m} = \frac{\lambda \pi(GT, s^{-i,m}, T^m) - \pi(AD, s^{-i,m}, T^m)}{(1+\lambda)\left(\psi_{GT}^{i,m} + \lambda \psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m) + \pi(AD, s^{-i,m}, T^m)\right)}$$

Note that this change is positive provided that $\pi(GT, s^{-i,m}, T^m)/\pi(AD, s^{-i,m}, T^m) > 1/\lambda = \psi_{GT}^{i,m}/\psi_{AD}^{i,m}$ and is non-positive otherwise. Thus, for an agent to increase her probability of choosing $GT$ in round $m+1$, $GT$ had to earn proportionally more than $AD$ than it previously did (as expressed by the ratio of propensities). We again consider the change in $\Delta p_{GT}^{i,m}$ as the level of propensities $\psi_{GT}^{i,m}$ changes:

$$\frac{\partial \Delta p_{GT}^{i,m}}{\partial \psi_{GT}^{i,m}} = -\frac{\lambda \pi(GT, s^{-i,m}, T^m) - \pi(AD, s^{-i,m}, T^m)}{\left(\psi_{GT}^{i,m} + \lambda \psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m) + \pi(AD, s^{-i,m}, T^m)\right)^2}.$$

Provided that $\pi(GT, s^{-i,m}, T^m)/\pi(AD, s^{-i,m}, T^m) > 1/\lambda$, so that the probability of choosing $GT$ increases in match $m$, this expression is negative. Thus, later experience will lead to smaller changes in choice probabilities.

In the case where $\pi(GT, s^{-i,m}, T^m)/\pi(AD, s^{-i,m}, T^m) \leq 1/\lambda$ we have $\Delta p^{i,m} \leq 0$ and $\partial \Delta p^{i,m}_{GT}/\partial \psi^{i,m}_{GT} \geq 0$. Thus, again later experience leads to less pronounced changes in choice probability.

Finally, consider the case of fictitious play. Unlike the reinforcement learning models propensities are not necessarily increasing over time but are given by the expected payoffs which are determined by beliefs on the play of others, $\sigma^{-i,m}$, and beliefs on match length realization, $\bar{T}^m$). Choice probability is then given by

$$p^{i,m+1}_{GT} = \frac{\hat{\pi}(GT, \sigma^{-i,m}, \bar{T}^m)}{\hat{\pi}(GT, \sigma^{-i,m}, \bar{T}^m) + \hat{\pi}(AD, \sigma - i, m, \bar{T}^m)}$$

One can verify that $\partial p^{i,m+1}_{GT}/\partial \bar{T} > 0$, implying that agents are more likely to choose $GT$ for longer expected match lengths (keeping beliefs on behavior of others fixed). As early matches enter beliefs both directly and indirectly (since they affect what happens in later matches) this model also displays the power law of practice, though, since this is a second order effect, to a lesser extent than the other models.

## C.2 Further details on simulations

The payoff parameters of the underlying PD games were chosen to correspond to the 25th, 50th and 75th percentile of the distribution of SizeBAD. We follow Erev and Roth (1998) and have normalized payoffs of the underlying PD games to ensure payoffs and propensities are positive. Table C.1 reports the corresponding parameter values of Dal Bó and Fréchette (2018) alongside the values for R and P when S and T were normalized to 0 and 1, respectively.

In each simulation run, 16 agents were matched against each other in a round robin tournament and match length realizations were drawn at the session level according to the discount factor $\delta = 0.75$, thus targeting the median group size and the most frequent discount factor in Dal Bó and Fréchette (2018). Our exercise contains 4000 simulated

experiments played over 15 matches for each of the three payoff configuration leading to 960 000 indefinitely repeated PD games for each of the three learning models.

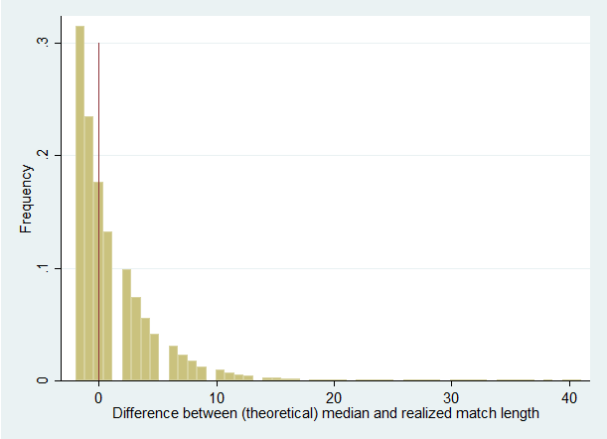


FIGURE C.1: Match Length Realization in our simulations.

The median match length in the simulated data is 3 and the mean 3.99 (as expected). The smallest realization is 1 and the largest 44. Appendix Figure C.1 shows the distribution of match lengths in our simulated data.

| percentile | SizeBAD | $\ell$ | $g$ | R | P |
|---|---|---|---|---|---|
| 25 | 0.1625 | 0.5 | 0.4 | 0.789 | 0.263 |
| 50 | 0.2 | 0.56 | 1 | 0.609 | 0.219 |
| 75 | 0.667 | 1.85 | 2 | 0.588 | 0.381 |

TABLE C.1: Payoff parameters used for simulations.

In the reinforcement learning models, initial propensities for both strategies were set to the expected payoff when the match length corresponds to its mean and the opponent randomizes with equal probability among the two strategies, averaged across the two strategies, i.e. $\hat{\psi}_s^{i,0} = (\hat{\pi}(GT, \frac{1}{2}, T) + \hat{\pi}(AD, \frac{1}{2}, T))/2$. Note that for the reinforcement models this implies initial choice probabilities of $p_{GT}^{i,m} = p_{AD}^{i,m} = \frac{1}{2}$. Correspondingly, in the fictitious play model we started with initial belief on the strategy of the opponent of $\sigma^{-i,0} = 1/2$ and with initial belief on the match length of $\bar{T}^0 = \mathbb{E}[T]$.

# Appendix D

# Chapter 2 Appendix D: Additional Information Experiments

We conducted our own experiment to complete Case Study I (Section 2.3.1). In this Appendix we provide additional information on these experiments. Our experiments were conducted at Essex Lab at the University of Essex. The deterministic, replication and inverse treatments were conducted in February/March 2017 and the MatchStoch treatment in May 2019. Participants were students (and some non-students) who signed up for lab experiment at EssexLab at the University of Essex. They were recruited using recruitment software hroot. Table D.1 shows some demographics of our participants as well their answers to a post experimental questionnaire. We used exactly the same questionnaire as Bigoni et al. (2015), but replaced the question about whether people were born in Italy with whether they were born in the UK.

| | *Sample Characteristics* | | | |
|---|---|---|---|---|
| | Rep | Inverse | MatchStoch | Det |
| Age | 26.54 | 24.25 | 22.76 | 25.27 |
| From UK (1=yes) | 0.38 | 0.23 | 0.37 | 0.44 |
| Gender (1=female) | 0.47 | 0.44 | 0.42 | 0.50 |
| Risk Attitude (0-10) | 6.54 | 5.89 | 6.16 | 6.23 |
| Trust (0-1) | 0.38 | 0.39 | 0.33 | 0.42 |
| Logic1 (0-1) | 0.64 | 0.58 | 0.70 | 0.67 |
| Logic2 (0-1) | 0.33 | 0.31 | 0.52 | 0.40 |

TABLE D.1: Basic Sample Characteristics of Lab Experimental Sample used in Section 2.3.1. Mean Age, fraction of participants born in UK, fraction female, mean risk attitude (0, most risk averse, 10 least risk averse), fraction displaying high trust, fraction answering Logic 1 question correctly and fraction answering Logic 2 question correctly.

In all the experiments we followed exactly the same procedures used by Bigoni et

al. (2015) including using the exact same Instructions and software (translated from Italian to English). The only change made to the software was to change the match length realizations (i) to be drawn at the match level in session MatchStoch and (ii) inversed in treatment "Inverse" (see Figure F.3). While in MatchStoch we draw match lengths within each session as the experiment progresses, we use the draws generated by Bigoni et al. (2015) in the "Replication" and "Inverse" treatments. As the interest of the study is match length realization, there was no other way to conduct an exact replication. Both procedures are in line with what was communicated to participants which is the following: "*How is a period duration established? The period may stop at every tick of 0.16 seconds. This event depends on the results of a random draw...*". No further information was given about when this draw took/will take place.

# Appendix E

# Chapter 2 Appendix E: Additional Tables

This Appendix contains additional tables.

## E.1 Additional Tables for Section 2.2.2

In Table E.1 we consider separately the effect of match length realization of the 1st tenth, ninth, eight,..., half of matches on first-stage cooperation rates in the remaining matches. The table shows that there is a positive effect of match length realization in all these cases. However, the effect is smaller for the first tenth, ninth,...,sixth of matches compared to our baseline specification using the first third of matches. This shows that, while the very first matches are very important there is still learning and match length realizations become more important as more early matches are aggregated. Using the 1st half of matches, however, does not lead to a larger effect than using the 1st third.

Table E.2 compares the impact of match length realization in the first X-th, second X-th, third X-th,... of matches on cooperation in the last X-th of matches, where X ranges from 2,...,10. The coefficient $\Delta^{1st}_{above}$ is statistically significant in seven out of these nine different splits while all other dummies are statistically significant in fewer than 50 percent of the cases where they are included. $\Delta^{1st}_{above}$ also has the biggest coefficient size in all cases $X = 2,...,6$ and the second-biggest coefficient size when $X = 7,..,10$, though it should be noted that these differences are not always statistically significant.

| 1st... | (1) tenth | (2) ninth | (3) eight | (4) seventh | (5) sixth | (6) fifth | (7) fourth | (8) third | (9) half |
|---|---|---|---|---|---|---|---|---|---|
| $\Delta^{1st}_{above}$ | 0.125** | 0.081 | 0.020 | 0.055 | 0.023 | 0.078 | 0.181*** | 0.226*** | 0.165*** |
| | (0.057) | (0.059) | (0.058) | (0.061) | (0.060) | (0.063) | (0.061) | (0.056) | (0.073) |
| SizeBAD | -0.592*** | -0.692*** | -0.781*** | -0.815*** | -0.829*** | -0.744*** | -0.628*** | -0.539*** | -0.688*** |
| | (0.0726) | (0.0769) | (0.0967) | (0.0964) | (0.0867) | (0.116) | (0.120) | (0.0945) | (0.152) |
| SizeBAD × $\Delta^{1st}_{above}$ | -0.189** | -0.056 | 0.053 | -0.066 | -0.105 | -0.028 | -0.212*** | -0.296*** | -0.103 |
| | (0.081) | (0.082) | (0.093) | (0.077) | (0.084) | (0.116) | (0.073) | (0.094) | (0.163) |
| Constant | 0.852*** | 0.916*** | 1.003*** | 1.038*** | 1.051*** | 0.965*** | 0.843*** | 0.747*** | 0.886*** |
| | (0.0574) | (0.0750) | (0.0977) | (0.0992) | (0.0901) | (0.118) | (0.123) | (0.0987) | (0.155) |
| Observations | 45,869 | 45,469 | 44,773 | 43,955 | 43,481 | 41,873 | 38,778 | 34,319 | 25,795 |
| R-squared | 0.188 | 0.188 | 0.189 | 0.195 | 0.196 | 0.199 | 0.212 | 0.223 | 0.237 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.1: Initial (first stage) cooperation rate in remaining matches depending on whether at least $\frac{2}{3}$ of first X-th of matches where above theoretical median length.

Furthermore for all $X = 2,...,9$ the coefficient corresponding to the first $X$-th of matches is bigger than the one corresponding to the most recent (the $X-1$th) $X$-th of matches. As the finer splits often involve only very few matches in each group we consider the cases $X = 2,...,6$ to be more meaningful. The table hence emphasizes the point that early match length realization is at least as important or more important than match length realization in more recent matches.

Table E.3 shows the results of a placebo test, where we regress cooperation in the 1st third of matches on $\Delta^{3rd}_{above}$ (as well as `SizeBAD` and an interaction). We would not expect realized match length of final matches, which have not yet been played, to affect cooperation in the beginning of the experiment. Hence we would expect zero coefficients on $\Delta^{3rd}_{above}$ and the corresponding interaction term. We do indeed find that these coefficients are close to zero and statistically not significant. This shows that the results in Table 2.1 are fundamental and not e.g. driven by correlations of match lengths within sessions or observed or unobserved heterogeneity across papers or treatments, e.g. caused by different ways researchers implement match length draws.

Appendix Table E.4 reproduces our main results reported in Table 2.1 using paper fixed effects.

In Appendix Table E.5 we use the share of matches above theoretical median length instead of a dummy variable.

In Appendix Table E.6 we conduct the same analysis as Table 2.1 with the only

| 1st... | (1) tenth | (2) ninth | (3) eight | (4) seventh | (5) sixth | (6) fifth | (7) fourth | (8) third | (9) half |
|---|---|---|---|---|---|---|---|---|---|
| $\Delta^{1st}_{above}$ | 0.010 | 0.057* | 0.048 | 0.082** | 0.065* | 0.077** | 0.105*** | 0.101** | 0.123*** |
| | (0.037) | (0.031) | (0.033) | (0.035) | (0.034) | (0.039) | (0.036) | (0.041) | (0.044) |
| $\Delta^{2nd}_{above}$ | 0.051 | 0.026 | 0.022 | 0.037 | 0.017 | 0.075** | 0.097** | 0.069* | |
| | (0.03) | (0.041) | (0.037) | (0.034) | (0.045) | (0.036) | (0.038) | (0.039) | |
| $\Delta^{3rd}_{above}$ | -0.014 | -0.015 | 0.019 | 0.087** | 0.048 | 0.076* | 0.074** | | |
| | (0.039) | (0.039) | (0.038) | (0.037) | (0.036) | (0.040) | (0.035) | | |
| $\Delta^{4th}_{above}$ | 0.070* | 0.103*** | 0.017 | 0.003 | 0.050 | 0.075** | | | |
| | (0.038) | (0.035) | (0.037) | (0.034) | (0.037) | (0.036) | | | |
| $\Delta^{5th}_{above}$ | -0.000 | 0.001 | 0.091** | 0.006 | 0.014 | | | | |
| | (0.035) | (0.032) | (0.040) | (0.045) | (0.044) | | | | |
| $\Delta^{6th}_{above}$ | 0.033 | 0.044 | 0.047 | 0.001 | | | | | |
| | (0.033) | (0.043) | (0.039) | (0.039) | | | | | |
| $\Delta^{7th}_{above}$ | -0.004 | 0.012 | 0.022 | | | | | | |
| | (0.041) | (0.040) | (0.034) | | | | | | |
| $\Delta^{8th}_{above}$ | -0.025 | 0.030 | | | | | | | |
| | (0.033) | (0.037) | | | | | | | |
| $\Delta^{9th}_{above}$ | 0.050 | | | | | | | | |
| | (0.037) | | | | | | | | |
| SizeBAD | -0.824*** | -0.814*** | -0.806*** | -0.806*** | -0.809*** | -0.805*** | -0.805*** | -0.809*** | -0.769*** |
| | (0.0808) | (0.0729) | (0.0772) | (0.0875) | (0.0796) | (0.0777) | (0.0734) | (0.0681) | (0.0670) |
| Constant | 0.989*** | 0.966*** | 0.965*** | 0.969*** | 0.977*** | 0.973*** | 0.983*** | 0.994*** | 0.967*** |
| | (0.0940) | (0.0789) | (0.0836) | (0.0926) | (0.0852) | (0.0823) | (0.0798) | (0.0747) | (0.0736) |
| Observations | 5,905 | 7,338 | 7,506 | 8,296 | 9,148 | 11,721 | 13,573 | 18,536 | 25,795 |
| R-squared | 0.257 | 0.273 | 0.266 | 0.269 | 0.262 | 0.268 | 0.265 | 0.251 | 0.236 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.2: Initial (first stage) cooperation rate in remaining matches depending on whether at least $\frac{2}{3}$ of first X-th, second X-th, third X-th,... of matches where above theoretical median length.

| | *Placebo test* | | |
|---|---|---|---|
| | (1) | (2) | (3) |
| $\Delta^{3rd}_{above}$ | 0.038 | -0.054 | -0.053 |
| | (0.055) | (0.061) | (0.060) |
| SizeBAD | | -0.643*** | -0.650*** |
| | | (0.061) | (0.078) |
| SizeBAD × $\Delta^{3rd}_{above}$ | | | 0.009 |
| | | | (0.055) |
| Constant | 0.365*** | 0.908*** | 0.915*** |
| | (0.045) | (0.070) | (0.085) |
| $\delta$ f.e. | NO | YES | YES |
| Observations | 29,467 | 29,467 | 29,467 |
| R-squared | 0.002 | 0.141 | 0.141 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.3: Placebo test. Initial (first stage) cooperation rate in the *1st* third of matches explained by dummy $\Delta^{3rd}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 3rd third of the experiment were longer than the theoretical median match length.

difference that $\Delta^{1st}_{above}$ takes the value 1 if more than half (rather than $\frac{2}{3}$) of early matches are above median. It is hence a much weaker test as match length realizations with more or less than half of matches above median need not differ by much. In the most

| | *Main Result with paper fixed effects* | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\Delta^{1st}_{above}$ | 0.115** | 0.120*** | 0.267*** | 0.120* | 0.119*** | 0.247*** |
| | (0.055) | (0.040) | (0.053) | (0.060) | (0.040) | (0.058) |
| $\Delta^{2nd}_{above}$ | | | | 0.025 | 0.068 | 0.058 |
| | | | | (0.067) | (0.042) | (0.061) |
| `SizeBAD` | | -0.746*** | -0.497*** | | -0.808*** | -0.573*** |
| | | (0.075) | (0.084) | | (0.075) | (0.148) |
| `SizeBAD` $\times$ $\Delta^{1st}_{above}$ | | | -0.328*** | | | -0.291*** |
| | | | (0.084) | | | (0.094) |
| `SizeBAD` $\times$ $\Delta^{2nd}_{above}$ | | | | | | -0.015 |
| | | | | | | (0.132) |
| Constant | 0.826*** | 1.006*** | 0.757*** | 0.227*** | 1.101*** | 0.946*** |
| | (0.0147) | (0.0782) | (0.0869) | (0.0215) | (0.120) | (0.161) |
| $\delta$ f.e. | NO | YES | YES | NO | YES | YES |
| paper f.e. | YES | YES | YES | YES | YES | YES |
| Test $\Delta^{1st}_{above}$=$\Delta^{2nd}_{above}$ | - | - | - | 0.3477 | 0.4220 | 0.0508 |
| Observations | 34,319 | 34,319 | 34,319 | 18,536 | 18,536 | 18,536 |
| R-squared | 0.104 | 0.239 | 0.245 | 0.122 | 0.274 | 0.278 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.4: Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level. Observations stem from 141 sessions spread across 15 papers. With paper fixed effects.

extreme case they might differ by one round only, which is not much if matches are long, i.e. $\delta$ is high. On the other hand it should allow us to better identify the effect of shorter match length when $\delta = 0.5$, as there will be more variation in the dummy in this case. Indeed we do find that the dummy defined in this way is more effective in capturing differences when $\delta = 0.5$. Overall the effect of shorter matches reported in column (1) is now a 38% increase in cooperation rates compared to the 43% increase identified in Table 2.1. Table E.7 goes the opposite direction and shows results when a dummy is used indicating whether more than $\frac{3}{4}$ of early matches were above median. Again results are similar, even though this dummy takes the value 1 much less often.

In Appendix Table E.8 we ask whether unusually short early matches or unusually

| | *Share of matches above median* | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $Share^{1st}_{above}$ | 0.218*** | 0.206** | 0.309*** | 0.072 | 0.210** | 0.310** |
| | (0.068) | (0.079) | (0.113) | (0.112) | (0.080) | (0.120) |
| $Share^{2nd}_{above}$ | | | | 0.235** | 0.151 | 0.111 |
| | | | | (0.112) | (0.094) | (0.130) |
| `SizeBAD` | | -0.764*** | -0.446 | | -0.813*** | -0.573 |
| | | (0.0759) | (0.383) | | (0.0754) | (0.424) |
| `SizeBAD` $\times$ $Share^{1st}_{above}$ | | | -0.441 | | | -0.434 |
| | | | (0.490) | | | (0.498) |
| `SizeBAD` $\times$ $Share^{2nd}_{above}$ | | | | | | 0.104 |
| | | | | | | (0.281) |
| Constant | 0.290*** | 1.023*** | 0.706* | 0.250*** | 1.041*** | 0.801* |
| | (0.031) | (0.079) | (0.384) | (0.030) | (0.078) | (0.425) |
| $\delta$ f.e. | NO | YES | YES | NO | YES | YES |
| Test $Share^{1st}_{above}$=$Share^{2nd}_{above}$ | - | - | - | 0.4465 | 0.6338 | 0.2631 |
| Observations | 34,319 | 34,319 | 34,319 | 18,536 | 18,536 | 18,536 |
| R-squared | 0.023 | 0.236 | 0.237 | 0.044 | 0.271 | 0.271 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.5: Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by share of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by share of matches longer than theoretical median length in 1st and 2nd third of matches. Standard errors clustered at session level. Observations stem from 141 sessions spread across 15 papers.

long early matches have a larger effect. To this end we reproduce Table 2.1 using now a dummy $\Delta^{1st}_{below}$ which takes the value 1 if more $\frac{2}{3}$ of early matches are below the theoretical median length. For the sessions identified by this dummy cooperation rates are 32% lower in specification (1), indicating that the overall effect of match length realizations is broadly symmetric.

| | *Reduced Threshold* | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\Delta^{1st}_{above}$ | 0.120** | 0.142*** | -0.095 | 0.051 | 0.143** | -0.102 |
| | (0.050) | (0.051) | (0.085) | (0.058) | (0.057) | (0.091) |
| $\Delta^{2nd}_{above}$ | | | | 0.157** | 0.064 | 0.076 |
| | | | | (0.062) | (0.054) | (0.094) |
| `SizeBAD` | | -0.792*** | -1.967*** | | -0.826*** | -2.039*** |
| | | (0.066) | (0.286) | | (0.067) | (0.322) |
| `SizeBAD` $\times \Delta^{1st}_{above}$ | | | 1.184*** | | | 1.239*** |
| | | | (0.279) | | | (0.321) |
| `SizeBAD` $\times \Delta^{2nd}_{above}$ | | | | | | -0.018 |
| | | | | | | (0.163) |
| Constant | 0.320*** | 1.000*** | 2.175*** | 0.277*** | 1.011*** | 2.225*** |
| | (0.032) | (0.072) | (0.287) | (0.032) | (0.073) | (0.324) |
| $\delta$ f.e. | NO | YES | YES | NO | YES | YES |
| Test $\Delta^{1st}_{above}$=$\Delta^{2nd}_{above}$ | - | - | - | 0.3178 | 0.2876 | 0.1933 |
| Observations | 34,319 | 34,319 | 34,319 | 18,536 | 18,536 | 18,536 |
| R-squared | 0.014 | 0.219 | 0.223 | 0.036 | 0.249 | 0.254 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.6: Reduced Threshold. Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{1}{2}$ of matches in the 1st third of the experiment were longer than the theoretical median match length.

| | *Increased Threshold* | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\Delta^{1st}_{above}$ | 0.192*** | 0.054 | 0.147** | 0.149** | 0.044 | 0.139** |
| | (0.067) | (0.038) | (0.059) | (0.070) | (0.038) | (0.059) |
| $\Delta^{2nd}_{above}$ | | | | 0.181*** | 0.079** | 0.112* |
| | | | | (0.062) | (0.038) | (0.063) |
| `SizeBAD` | | -0.743*** | -0.671*** | | -0.766*** | -0.661*** |
| | | (0.067) | (0.077) | | (0.068) | (0.090) |
| `SizeBAD` $\times \Delta^{1st}_{above}$ | | | -0.215*** | | | -0.221** |
| | | | (0.081) | | | (0.085) |
| `SizeBAD` $\times \Delta^{2nd}_{above}$ | | | | | | -0.060 |
| | | | | | | (0.092) |
| Constant | 0.346*** | 0.951*** | 0.879*** | 0.297*** | 0.951*** | 0.846*** |
| | (0.030) | (0.073) | (0.082) | (0.029) | (0.074) | (0.095) |
| $\delta$ f.e. | NO | YES | YES | NO | YES | YES |
| Test $\Delta^{1st}_{above}$=$\Delta^{2nd}_{above}$ | - | - | - | 0.3477 | 0.4220 | 0.050 |
| Observations | 34,319 | 34,319 | 34,319 | 18,536 | 18,536 | 18,536 |
| R-squared | 0.031 | 0.215 | 0.218 | 0.064 | 0.249 | 0.252 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.7: Reduced Threshold. Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{3}{4}$ of matches in the 1st third of the experiment were longer than the theoretical median match length.

| | *Short matches* | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\Delta^{1st}_{below}$ | -0.144*** | -0.042 | -0.145** | -0.063 | -0.045 | 0.092 |
| | (0.049) | (0.051) | (0.061) | (0.080) | (0.057) | (0.074) |
| $\Delta^{2nd}_{below}$ | | | | -0.138 | -0.041 | 0.040 |
| | | | | (0.083) | (0.060) | (0.102) |
| `SizeBAD` | | -0.759*** | -0.750*** | | -0.799*** | -0.775*** |
| | | (0.066) | (0.067) | | (0.068) | (0.071) |
| `SizeBAD` $\times$ $\Delta^{1st}_{below}$ | | | -1.093*** | | | -0.768*** |
| | | | (0.188) | | | (0.239) |
| `SizeBAD` $\times$ $\Delta^{2nd}_{below}$ | | | | | | -0.228 |
| | | | | | | (0.152) |
| Constant | 0.442*** | 1.010*** | 1.907*** | 0.468*** | 1.071*** | 1.824*** |
| | (0.036) | (0.097) | (0.150) | (0.039) | (0.127) | (0.185) |
| $\delta$ f.e. | NO | NO | YES | NO | NO | YES |
| Test $\Delta^{1st}_{below}$=$\Delta^{2nd}_{below}$ | - | - | - | 0.6325 | 0.9650 | 0.6977 |
| Observations | 34,319 | 34,319 | 34,319 | 18,536 | 18,536 | 18,536 |
| R-squared | 0.018 | 0.213 | 0.215 | 0.033 | 0.242 | 0.244 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.8: Focus on Short Matches. Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{below}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were shorter than the theoretical median match length.

# E.2 Additional Tables for Section 2.2.3

| | *Simulated Reinforcement Learning* | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\Delta^{1st}_{above}$ | 0.040*** | 0.040*** | 0.042*** | 0.039*** | 0.039*** | 0.039*** |
| | (0.002) | (0.003) | (0.004) | (0.002) | (0.002) | (0.005) |
| $\Delta^{2nd}_{above}$ | | | | 0.027*** | 0.026*** | 0.035*** |
| | | | | (0.003) | (0.002) | (0.004) |
| `SizeBAD` | | -0.139*** | -0.138*** | | -0.161*** | -0.159*** |
| | | (0.003) | (0.003) | | (0.003) | (0.004) |
| $\Delta^{1st}_{above} \times$ `SizeBAD` | | | -0.004 | | | -0.000 |
| | | | (0.011) | | | (0.012) |
| $\Delta^{2nd}_{above} \times$ `SizeBAD` | | | | | | -0.025** |
| | | | | | | (0.010) |
| Constant | 0.415*** | 0.463*** | 0.463*** | 0.419*** | 0.475*** | 0.474*** |
| | (0.001) | (0.001) | (0.001) | (0.001) | (0.001) | (0.001) |
| Test $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$ | - | - | - | 0.0063 | 0.0003 | 0.0989 |
| Observations | 1,920,000 | 1,920,000 | 1,920,000 | 960,000 | 960,000 | 960,000 |
| R-squared | 0.002 | 0.013 | 0.013 | 0.002 | 0.017 | 0.017 |

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.9: Simulated Reinforcement Learning Model. Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level.

| | *Simulated Reinforcement Learning with Counterfactuals* | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\Delta^{1st}_{above}$ | 0.030*** | 0.029*** | 0.037*** | 0.025*** | 0.024*** | 0.029*** |
| | (0.001) | (0.000) | (0.001) | (0.001) | (0.000) | (0.001) |
| $\Delta^{2nd}_{above}$ | | | | 0.016*** | 0.018*** | 0.024*** |
| | | | | (0.001) | (0.001) | (0.002) |
| `SizeBAD` | | -0.181*** | -0.178*** | | -0.184*** | -0.180*** |
| | | (0.001) | (0.001) | | (0.002) | (0.001) |
| $\Delta^{1st}_{above} \times$ `SizeBAD` | | | -0.023*** | | | -0.017*** |
| | | | (0.003) | | | (0.003) |
| $\Delta^{2nd}_{above} \times$ `SizeBAD` | | | | | | -0.017*** |
| | | | | | | (0.005) |
| Constant | 0.399*** | 0.461*** | 0.460*** | 0.400*** | 0.463*** | 0.462*** |
| | (0.000) | (0.000) | (0.000) | (0.000) | (0.000) | (0.000) |
| Test $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$ | - | - | - | 0.0001 | 0.0000 | 0.0000 |
| Observations | 1,920,000 | 1,920,000 | 1,920,000 | 960,000 | 960,000 | 960,000 |
| R-squared | 0.010 | 0.200 | 0.201 | 0.010 | 0.207 | 0.207 |

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.10: Simulated Reinforcement Learning Model with Counterfactuals. Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level.

| | *Simulated Fictitious Play Learning* | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\Delta^{1st}_{above}$ | 0.027*** | 0.029*** | 0.038*** | 0.021*** | 0.023*** | 0.030*** |
| | (0.001) | (0.000) | (0.001) | (0.001) | (0.001) | (0.002) |
| $\Delta^{2nd}_{above}$ | | | | 0.023*** | 0.022*** | 0.029*** |
| | | | | (0.001) | (0.001) | (0.002) |
| `SizeBAD` | | -0.178*** | -0.175*** | | -0.189*** | -0.185*** |
| | | (0.001) | (0.001) | | (0.001) | (0.002) |
| $\Delta^{1st}_{above} \times$ `SizeBAD` | | | -0.025*** | | | -0.019*** |
| | | | (0.003) | | | (0.005) |
| $\Delta^{2nd}_{above} \times$ `SizeBAD` | | | | | | -0.022*** |
| | | | | | | (0.005) |
| Constant | 0.392*** | 0.449*** | 0.448*** | 0.395*** | 0.456*** | 0.454*** |
| | (0.000) | (0.000) | (0.000) | (0.000) | (0.000) | (0.000) |
| Test $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$ | - | - | - | 0.2481 | 0.4042 | 0.9592 |
| Observations | 1,920,000 | 1,920,000 | 1,920,000 | 960,000 | 960,000 | 960,000 |
| R-squared | 0.007 | 0.128 | 0.128 | 0.009 | 0.149 | 0.149 |

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.11: Simulated Fictitious Play Learning Model. Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy $\Delta^{1st}_{above}$ indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies $\Delta^{1st}_{above}$ and $\Delta^{2nd}_{above}$. Standard errors clustered at session level.

## E.3 Additional Tables for Section 2.3.1

| | *Average Cooperation Rates* | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Replication | -14.46*** | -14.72*** | | |
| | (4.722) | (4.955) | | |
| Inverse | -3.142 | -3.582 | | |
| | (3.555) | (3.548) | | |
| Match Stochastic | -6.641 | -6.956 | -6.923 | -7.210 |
| | (5.154) | (5.059) | (5.093) | (4.993) |
| L.Duration | | -0.022 | | 0.013 |
| | | (0.045) | | (0.054) |
| Constant | 54.04*** | 55.31*** | 42.03*** | 42.43*** |
| | (3.454) | (3.497) | (8.895) | (8.445) |
| Demographics | NO | NO | YES | YES |
| Observations | 6,624 | 6,336 | 4,416 | 4,224 |
| Number of id | 288 | 288 | 192 | 192 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.12: Random Effects OLS regression of average cooperation rate on treatment dummies and covariates (age, gender, birthplace, nationality). Standard errors clustered at session level. Columns (3) and (4) only use data from match stochastic and deterministic sessions.

| | *Initial Cooperation Rates* | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Replication | -22.06*** | -22.44*** | | |
| | (3.921) | (3.756) | | |
| Inverse | -0.589 | -0.538 | | |
| | (3.367) | (3.451) | | |
| Match Stochastic | -12.77** | -13.11** | -12.94** | -13.23** |
| | (6.306) | (6.391) | (6.429) | (6.533) |
| L.Duration | | 0.061 | | 0.119* |
| | | (0.041) | | (0.065) |
| Constant | 73.96*** | 73.68*** | 57.61*** | 56.58*** |
| | (3.144) | (3.369) | (12.20) | (11.97) |
| Demographics | NO | NO | YES | YES |
| Observations | 6,624 | 6,336 | 4,416 | 4,224 |
| Number of id | 288 | 288 | 192 | 192 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.13: Random Effects OLS regression of initial cooperation rate on treatment dummies and covariates (age, gender, birthplace, nationality). Standard errors clustered at session level. Columns (3) and (4) only use data from match stochastic and deterministic sessions.

## E.4 Additional Tables for Section 2.3.2

| | *First Round Cooperation Rate* | | | | |
|---|---|---|---|---|---|
| Decision Setting | Finite | Prob All | Prob S1 | Prob S23 | S1 vs S23 |
| $N = 4$, MPCR$= 0.3$ | 24.9 | 28.8 | 16.56 | 34.95 | |
| | | $<$ | $>^{*}$ | $<^{**}$ | $p = 0.002$ |
| $N = 4$, MPCR$= 0.6$ | 44.0 | 47.7 | 38.88 | 51.00 | |
| | | $<$ | $>$ | $<$ | $p = 0.224$ |
| $N = 2$, MPCR$= 0.6$ | 52.4 | 44.3 | 42.56 | 45.46 | |
| | | $>$ | $>$ | $>$ | $p = 0.678$ |
| $N = 2$, MPCR$= 0.6$, Binary | 76.1 | 57.8 | 54.28 | 59.23 | |
| | | $>^{***}$ | $>^{*}$ | $>^{***}$ | $p = 0.613$ |

TABLE E.14: Initial cooperation rates across all rounds in the finite sessions (column (1)) and across all sessions with a probabilistic ending (column (2)) as in Table 3 in Lugovskyy et al. (2017). We further split the sessions with probabilistic ending in those with initially short matches (S1) and those with initially long matches (S23). Below each cooperation rate we show how the finite setting compares to the rate in question (as in Lugovskyy et al. (2017)). The last column shows the p-value when comparing initially short and long sequences. Following Lugovskyy et al. (2017) standard errors are clustered at the participant level in all regressions.

## E.5 Additional Tables for Section 2.3.3

| *with RO* | All groups | Short first matches | Other Groups | Short vs Other |
|---|---|---|---|---|
| Substitutes | 16.5 | 9.0 | 20.1 | |
| Complements | 23.2 | 34.9 | 17.3 | |
| Treatment Effect | - 6.7 | -25.2** | 2.8 | $p = 0.1280$ |
| *without RO* | All groups | Short first matches | Other Groups | Short vs Other |
| Substitutes | 17.6 | 20.9 | 15.9 | |
| Complements | 10.9 | 16.2 | 8.1 | |
| Treatment Effect | 6.7* | 4.7 | 7.8 | $p = 0.8949$ |

TABLE E.15: Efficiency measure from Table 2 (all matches) in Embrey, Mengel, and Peeters (2019). Treatment Effect is the difference between substitutes and complements. The two rightmost columns split out the groups with short initial matches from the rest.

# E.6 Additional Tables for Section 2.4

| | $\delta = 0.5$ | | | | $\delta = 0.75$ | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| mean duration | 0.021 | | -0.001 | 0.022 | 0.015* | | -0.004 | -0.002 |
| | (0.070) | | (0.077) | (0.075) | (0.008) | | (0.016) | (0.012) |
| median duration | | 0.104* | 0.105 | 0.087 | | 0.020** | 0.024* | 0.035** |
| | | (0.063) | (0.082) | (0.105) | | (0.008) | (0.016) | (0.016) |
| SizeBAD | -0.662*** | -0.704*** | -0.705*** | -0.642*** | -0.825*** | -0.847*** | -0.848*** | -0.965*** |
| | (0.113) | (0.113) | (0.103) | (0.089) | (0.071) | (0.065) | (0.064) | (0.077) |
| Constant | 0.665*** | 0.519*** | 0.521*** | 0.527** | 0.776*** | 0.753*** | 0.759*** | 0.575*** |
| | (0.150) | (0.174) | (0.151) | (0.201) | (0.061) | (0.062) | (0.061) | (0.098) |
| paper fixed effects | NO | NO | NO | YES | NO | NO | NO | YES |
| Observations | 10,786 | 10,786 | 10,786 | 10,786 | 12,380 | 12,380 | 12,380 | 12,380 |
| R-squared | 0.151 | 0.160 | 0.160 | 0.173 | 0.185 | 0.189 | 0.189 | 0.225 |

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE E.16: First round cooperation in the 2nd and 3rd third of matches regressed on mean and median match length in the 1st third. We selected the two discounted factors with most papers in the meta-study (8 papers for $\delta = 0.75$ and 4 papers for $\delta = 0.5$).

# Appendix F

# Chapter 2 Appendix F: Additional Figures

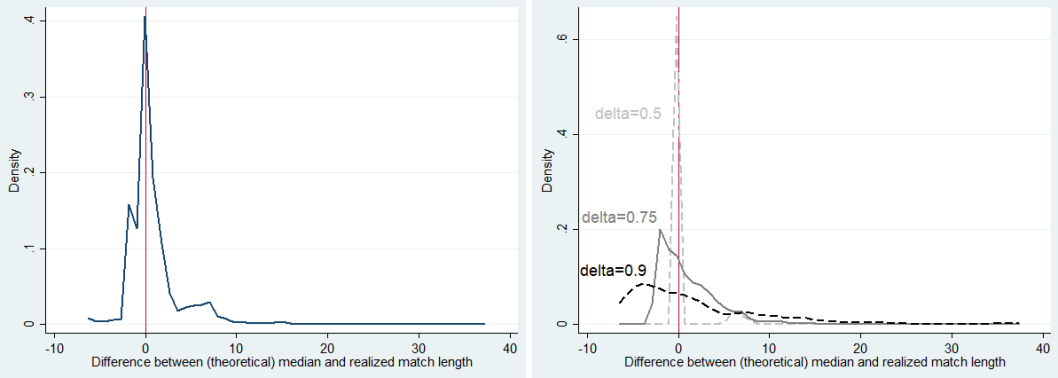


FIGURE F.1: Meta Data used in Section 2.2.2: Kernel density estimates of the difference between theoretical median match length and realized match length overall (left panel) and separately for $\delta = 0.5, 0.75$ and $\delta = 0.9$ (right panel).

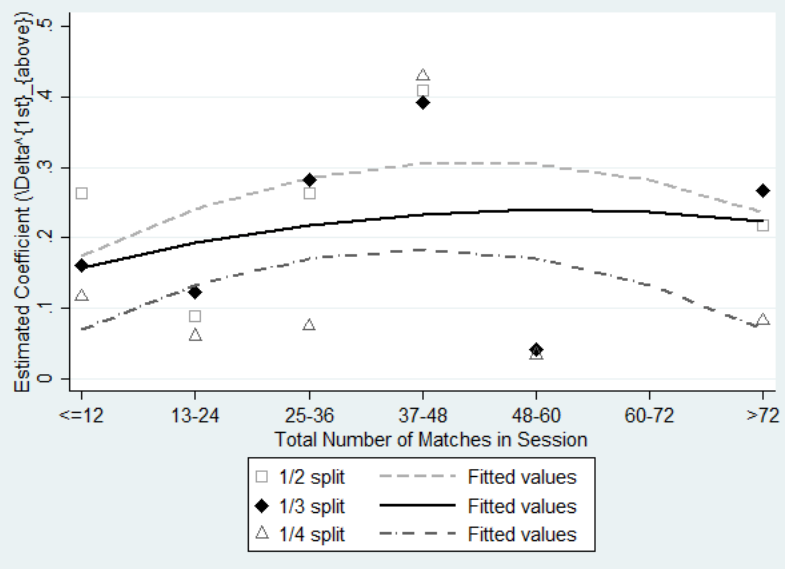


FIGURE F.2: Section 2.2.2 The Effect of Experience (Number of Matches Played in a Session) on Estimated Coefficient $\Delta^{1st}_{above}$. Black diamonds show the estimated coefficient $\Delta^{1st}_{above}$ on cooperation in the last third of matches depending on the total number of matches played in the session. The black line shows fitted values using a square polynomial. Squares and light gray dashed line show the effect of the first half of matches on the second half and triangles and dark gray dashed line the effect of the first quarter on the last quarter of matches.

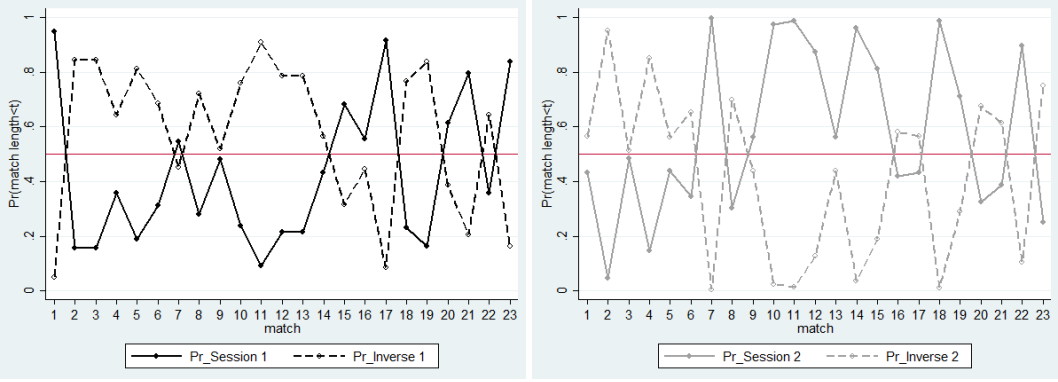


FIGURE F.3: Section 2.3.1: Illustration of how sequence of "inverse" match lengths is generated for the two sessions.

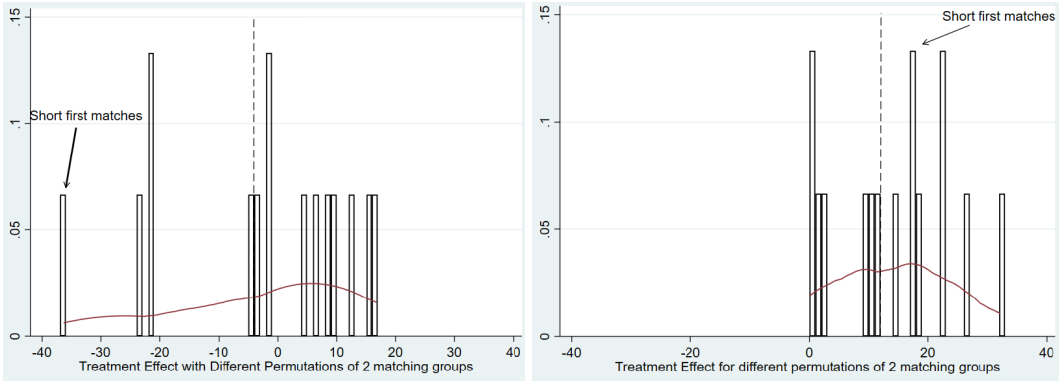


FIGURE F.4: Section 2.3.3: Illustration of treatment effect for different permutations of two matching groups. Histogram and estimated kernel density. Left Panel: treatments with revision opportunities. Right panel: treatments without revision opportunities.

# Appendix G

# Chapter 3 Appendix G

## G.1 Experimental Instructions

### General information

This is an experiment in economic decision making. If you read the following instructions carefully you can, depending on your decisions and decisions of other participants, earn a considerable amount of money. This money will be paid to you in cash at the end of the experiment.
These instructions are solely for your private use. During the experiment you are not allowed to communicate with the other participants. If you have any questions, please, raise your hand. The experimenter will come to you and answer them. Any violation of this rule will exclude you from the experiment and from all payments.
During the experiment you will make choices about amounts of tokens. Your **earnings** will be **calculated in tokens** and then converted into pounds at the following rate: **1 Token= 0.20 pounds**.
At the end of the experiment your earnings and an additional show up fee of £4 will be paid to you in cash.
The confidentiality of your decisions will be maintained during the experiment as none of the other participants will be able to identify which decisions you made.

### Experiment instructions

The experiment you are about to participate in consists of **five stages**. Below you will find the instructions for Stage 1.
Once Stage 1 is completed you will receive instructions for Stage 2, Stage 3, Stage 4 and Stage 5 on your screen at the beginning of each one of these stages.
You will participate in each stage only once.
At the end of the experiment there will be a short questionnaire for you to fill in.

### Stage 1

In Stage 1, you are asked to solve a task which will determine your points at the end of Stage 1.
The task is the same for all participants in the room.
The amount of points you accumulated in the task might affect decisions that will be made in the next two stages by you and other participants.

The task:
You will be presented with a sequence of words and your task will be to encode these words by substituting the letters of the alphabet in each word with numbers using a "letter-number" matching criteria similar to the one shown in Table 1.
Example 1: You are given the word FLAT. The letters in Table 1 show that F=6, L=3, A=8, and T=19. Hence the solution should be: 6 3 8 19.
Every time you encode a word correctly, the computer will immediately show you a new word to encode and so on after every word encoded correctly. This process will continue for **15 minutes** (900 seconds). The time limit, the words and their sequences are the same for all participants. For each word you encode correctly, you will receive 1 point.

**Stage 1 results**
Your final points from the task will be calculated and revealed to you in Stage 2.

TABLE G.1: Table 1

| Letters | Numbers |
|---|---|
| A | 8 |
| B | 12 |
| C | 14 |
| D | 10 |
| E | 9 |
| F | 6 |
| G | 24 |
| H | 22 |
| I | 7 |
| J | 5 |
| K | 11 |
| L | 3 |
| M | 18 |
| N | 1 |
| O | 21 |
| P | 16 |
| Q | 23 |
| R | 2 |
| S | 13 |
| T | 19 |
| U | 25 |
| V | 4 |
| W | 26 |
| X | 17 |
| Y | 20 |
| Z | 15 |

## Stage 2

At the beginning of Stage 2, before your points from Stage 1 task are revealed, you will be randomly paired with one other participant in this room and each of you will be assigned to a different role (**role A** or **role B**). You will therefore be part of a **group of two people**, which will remain the same until the end of the experiment. No participant will know who is in which group.
Depending on the role you have been assigned, you will then receive information as follows:

**If you have been assigned Role A:**

In Stage 2 you will be communicated the number of points you accumulated in Stage 1 task (we will call these points "Actual points"). You will further be provided information about where your "Actual points" from the task fall compared with points of other participants in this experiment session:

- You will be shown the HIGHEST and LOWEST number of "Actual points" accumulated during Stage 1 task of this experiment session.
- You will also be shown the **median of "Actual points"** from the task and whether your "Actual points" from the task are above, equal or below it.

In Stage 2 you will finally be assigned to a **"player type"**.
The player type you will be assigned to will depend on the number of points ("Post -shock points") you will have at the end of the following process:
The computer will flip a 'virtual' coin separately for each participant in role A. If the outcome is Head, then the number of points you accumulated in Stage 1 will remain unaffected. If the outcome is Tail, your number of points from the task will instead be reduced by 50%.

At the end of this process you will be communicated:
- whether your points have been affected by the random 50% reduction shock
- your "Post-shock points", which are the points left after the shock has been applied to your points from the task.

Therefore, you will be communicated whether your "Post-shock points" are above, equal or below the new **median of "Post-shock points"** of this experiment session.
A "player type" will be finally assigned to you depending on the position of your "Post-shock points" with respect to the "Post-shock points" median only.
You will be a **"H" type** if your "Post-shock points" are equal or above the "Post-shock points" median.
You will be a **"L" type** type if your "Post-shock points" are below the "Post-shock points" median.

---

Note: **"Actual points" median and "Post-shock points" median** are the same for all participants in the session.
However, the two medians might be different from each other and where a participants' points fall might therefore vary according to the median that is considered.

Specifically, it can be that:

- A participant affected by the reduction shock, has "Actual points" from the task above the "Actual points" median, but "Post-shock points" below the "Post-shock points" median. This is because his/her points have reduced.

- A participant not affected by the reduction shock, has "Actual points" from the task below the "Actual points" median, but "Post-shock points" above the "Post-shock points" median. This is because the "Post-shock points" median has shifted down (is smaller than the "Actual points" median of the session).
- It might however also be that where a participants' points fall does not vary across the two medians.

What is the median of points in an experiment session? An example

*The MEDIAN of points in the session corresponds to the number of points in the middle of the list of points accumulated in the session, when all participants' points are ordered from the lowest to the highest.*

Example: (Note that the numbers are unrealistic on purpose.)
Assume that in a session with 8 participants, the ordered list of the accumulated "Actual points" is:
2; 2; 4; **5; 6**; 6; 7; 8
The median of "Actual points" in this session is: [(5+6)/2]=5.5

Let's now assume that the points of the participants in the second and sixth position (respectively, 2 points and 6 points) are affected by the 50% reduction shock. The new list of "Post-shock points" is therefore:
1; 2; 3; **4; 5**; 6; 7; 8
The median of "Post-shock points" in this session is: [(4+5)/2]=3.5

---

**If you have been assigned Role B:**
In Stage 2 you will be only communicated the number of points you accumulated in Stage 1 task ("Actual points") and provided information about where your "Actual points" from the task fall compared with points of other participants in this experiment session:

- You will be shown the HIGHEST and LOWEST number of "Actual points" accumulated during Stage 1 task of this experiment session.
- You will also be shown the median of "Actual points" from the task and whether your "Actual points" from the task are above, equal or below it.

If you want, you can take note of the value of the median, your points and where they fall in the distribution. These values might be relevant for your decisions in Stage 3 and 4.

Instructions for Stage 3, 4 and 5 will be shown you on the screen, before the start of each stage.

## Stage 3

In Stage 3, participants assigned role A and participants assigned role B will make two different choices.

**Player A's choice**

You are asked to decide whether to reveal your **"Actual points"** from Stage 1 task to player B, before he/she makes his/her choice.

- If you decide **not to reveal** your "Actual points" from the task, your paired Player B will only be informed of the player type assigned you at the end of Stage 2.
- If you decide **to reveal** your "Actual points" from the task, your paired Player B will be informed of both your player type and the "Actual points" you accumulated in the task before the shock occurred.

Your decision will be implemented with 90% probability. However, there is also 10% of probability that player B will receive different information.

**Final earnings**
Player A and Player B's total final earnings **from Stage 1, Stage 2 and Stage 3** will be **determined solely by Player's B decision** in this stage.

**Player B's choice**

Participants assigned to role B are **always** provided with information about the **"player type"** of the paired player A.
Information about the paired player A's "Actual points" from the task will instead **depend solely on the choice to reveal made by the paired player A**.
Once the information about Player A's "player type" and "Actual points" from the task are disclosed to you, you will be given an endowment.
This endowment has been earned **by you together with your paired player A** by working on the task in Stage 1. Only you will receive the endowment, your paired Player A will not receive the endowment.
The endowment you receive will be equal to **100 tokens** and you will be asked to decide how to split this endowment between you and Player A. You can transfer to Player A any integer from 0 tokens to 100 tokens.
The rest of the endowment stays with you (player B).

**Final earnings**

Player A and Player B's final earnings from Stage 1, Stage 2 and Stage 3 will be **determined solely by your decision on how to split the endowment** received in Stage 3.
Final earnings from Stage 3 will be revealed you only after you have completed the experiment.

Instructions for Stage 4 and 5 will be shown you on the screen, before the start of each stage.

## Stage 4

In Stage 4, you will be asked to make some guesses about choices made by the other player in your group during the experiment.

**Player A's guess**

You will be shown the endowment that Player B received in Stage 3 and you will be asked to guess the number of tokens that she/he decided to give you in
**four different scenarios**.
It is important that you try to give correct estimation of Player B' decision, as a **correct guess** could **increase** your final **earnings**.
Note that all possible scenarios could become relevant as with 10% of probability, your actual choice to reveal your "Actual points" from the task, might be changed and Player B might be provided with different information.
**Additional earnings** can be made according to the accuracy of your estimation of Player's B choice in Stage 3:

- If you **correctly estimate** the number of tokens Player B gave you in Stage 3, you will increase your earnings by **15 tokens**.
- If you make an estimation that is **5 tokens higher or smaller** than the actual Player's B choice of tokens to give you, you will increase your earnings by **10 tokens**.
- If you make an estimation that is **10 tokens higher or smaller** than the actual Player's B choice of tokens to give you, you will increase your earnings by **5 tokens**.
- If you make an estimation that is higher or smaller by more than 10 tokens than the actual Player's B choice of tokens to give you, you **will not increase your earnings**.

**Player B's guess**

You will be asked to guess the number of tokens that Player A expects to receive in **four different scenarios**.
It is important that you try to give correct estimation of Player A' beliefs as a **correct guess** could **increase** your final **earnings**.
It is important that you also try to give a correct guess in each possible scenario as the computer will randomly select one scenario and will compare your guess in that specific case with the beliefs revealed in Stage 4 by Player A.
**Additional earnings** will be estimated according to the accuracy of your estimation of Player's A stated beliefs:

- If you **correctly estimate** the number of tokens Player A believes to receive in Stage 3, you will increase your earnings by **15 tokens**.
- If you make an estimation that is **5 tokens higher or smaller** than the actual number of tokens Player A believes to receive in Stage 3, you will increase your earnings by **10 tokens**.
- If you make an estimation that is **10 tokens higher or smaller** than the actual number of tokens Player A believes to receive in Stage 3, you will increase your earnings by **5 tokens**.
- If you make an estimation that is higher or smaller by more than 10 tokens than the actual number of tokens Player A believes to receive in Stage 3, you **will not increase your earnings**.

If Player A's "Actual points" were not revealed you, you might also be asked to guess the number of "Actual points" Player A accumulated in Stage 1. In this case, the computer will randomly select either one scenario or the "Actual points" guess for additional earnings and the accuracy of your guess will be estimated according to the criteria specified above.

Your **total earnings in the end** of the **experiment** will be equal to the sum of the **earnings** gained at the end of Stage 3 and **earnings** from the **correct guess made in Stage 4**.

## Stage 5

Stage 5 is the last stage of the experiment.
By completing Stage 5 you will complete the experiment and start a short questionnaire at end of which your final earnings will be communicated to you.
Your total earnings at the end of the experiment will not depend on your decisions in this Stage.

Stage 5 consists of two parts: Step 1 and Step 2.

In **Step 1** you will be asked to make **ten decisions. In each of the ten decision problems** you are asked to decide between **two alternatives** which are called **LEFT** and **RIGHT**. Each alternative implies hypothetical earnings for you and another person. The ten decision problems will be presented as rows in a table.

In **Step 2** you will be asked to consider **5 different categories of Actual points** from the task and to make five guesses. For each category of Actual points, you will be asked to guess the percentage of Players A that, having Actual points falling in that category, decided to reveal their Actual points during the experiment.

At the end of Stage 5 there will be a short questionnaire for you to fill in.

FIGURE G.1: Screen shot of the Encryption Real-effort task

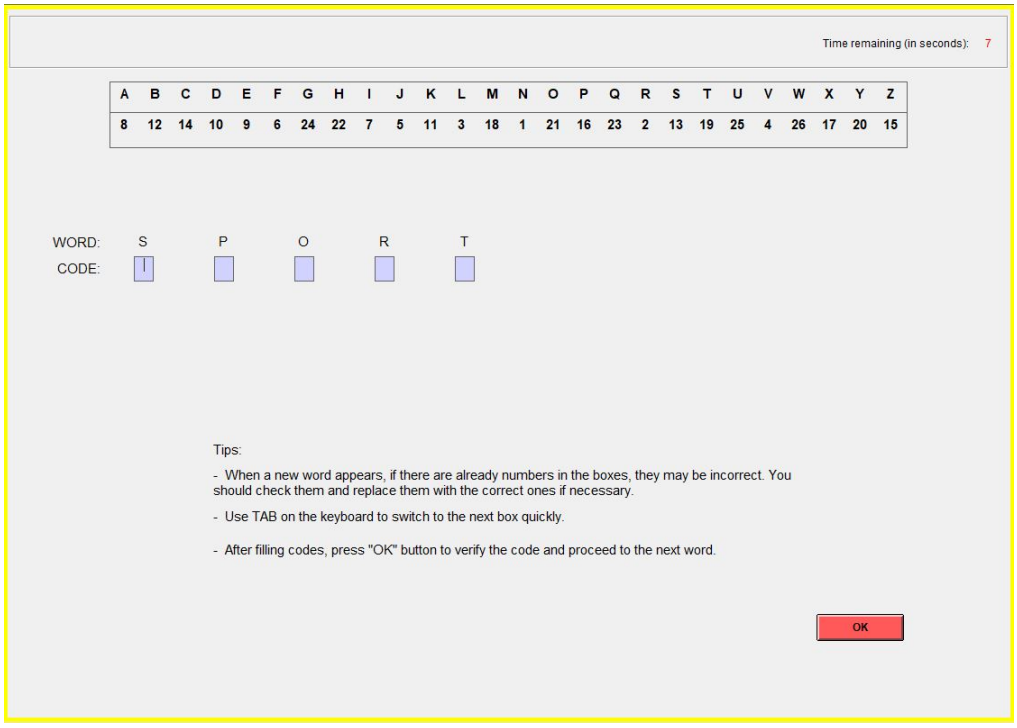


FIGURE G.2: Screen shot of the receiver's decision to reveal the task outcome

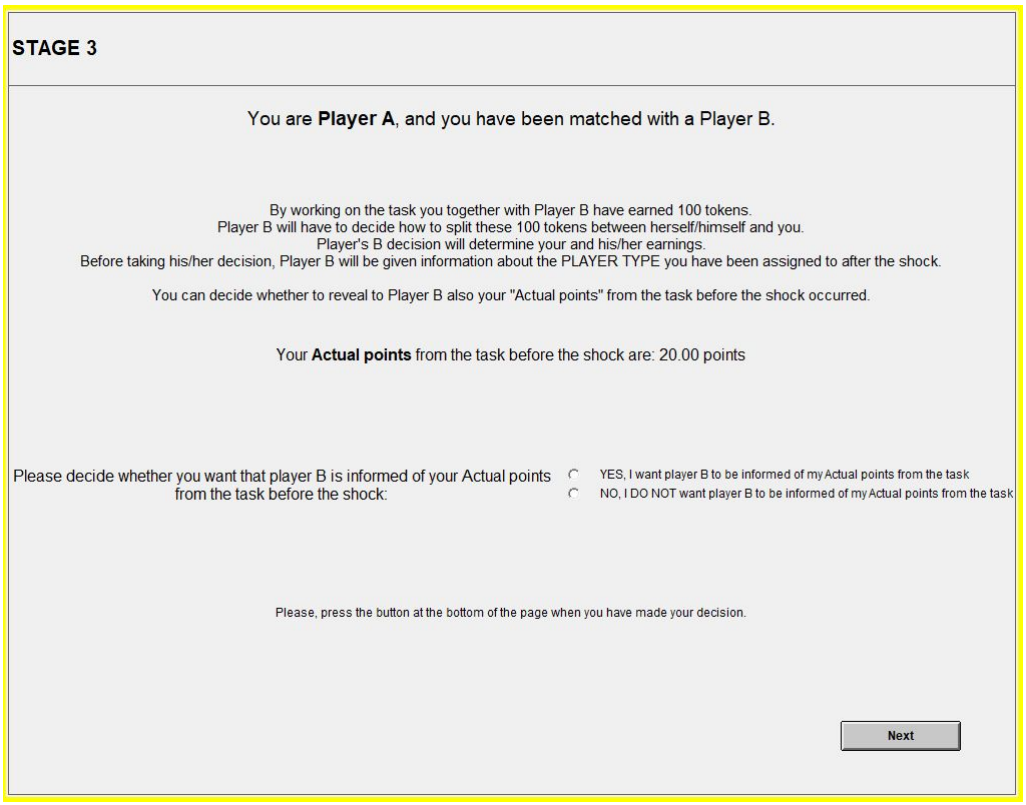

FIGURE G.3: Screen shot of the dictator's decision to give

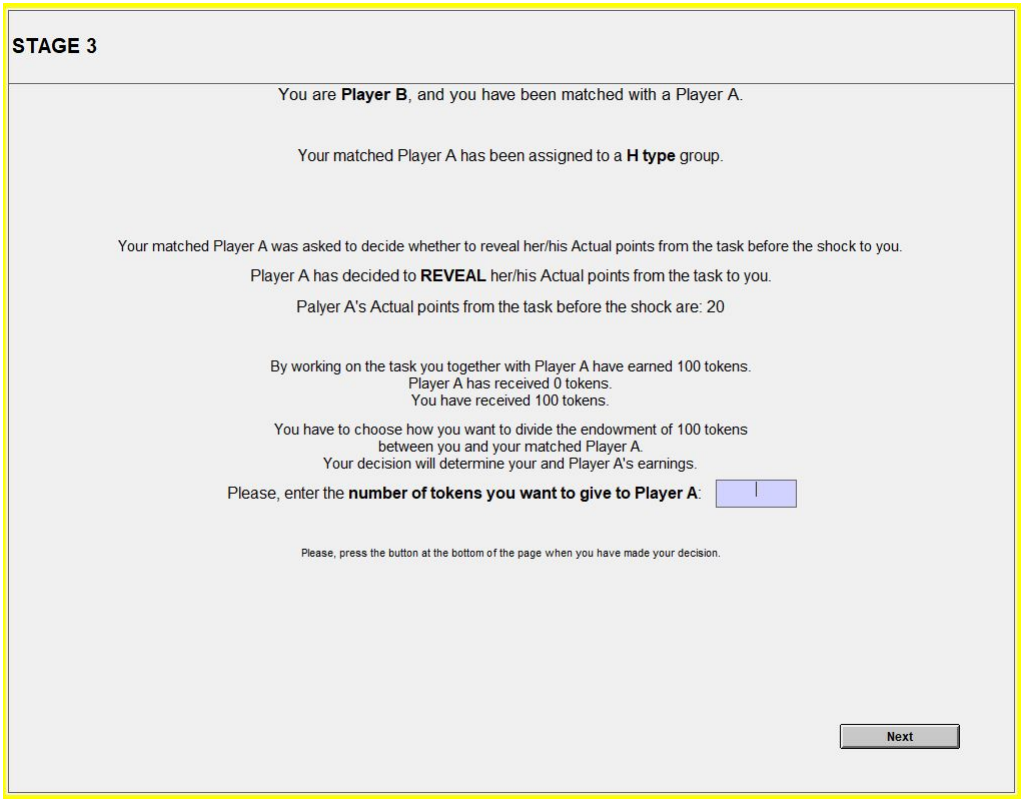


FIGURE G.4: Screen shot of the receiver's guess about the number of tokens received in the four alternative task-outcome scenarios

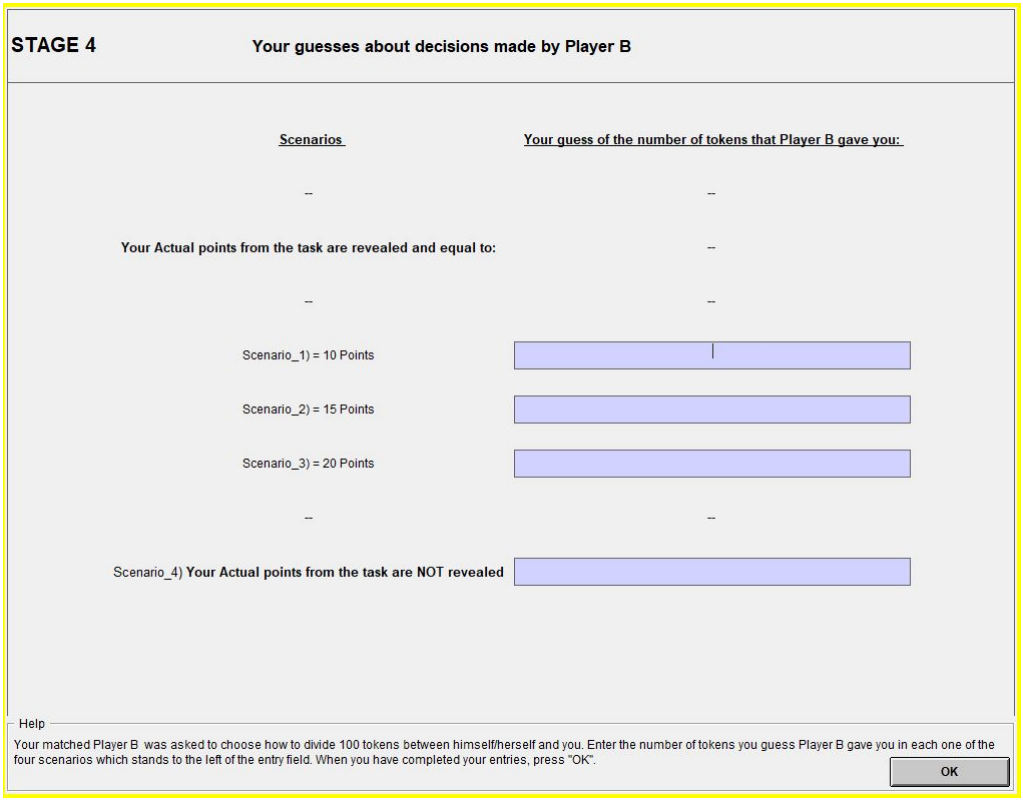

FIGURE G.5: Screen shot of the dictator's guess about the number of tokens expected by the receiver in the four alternative task-outcome scenarios

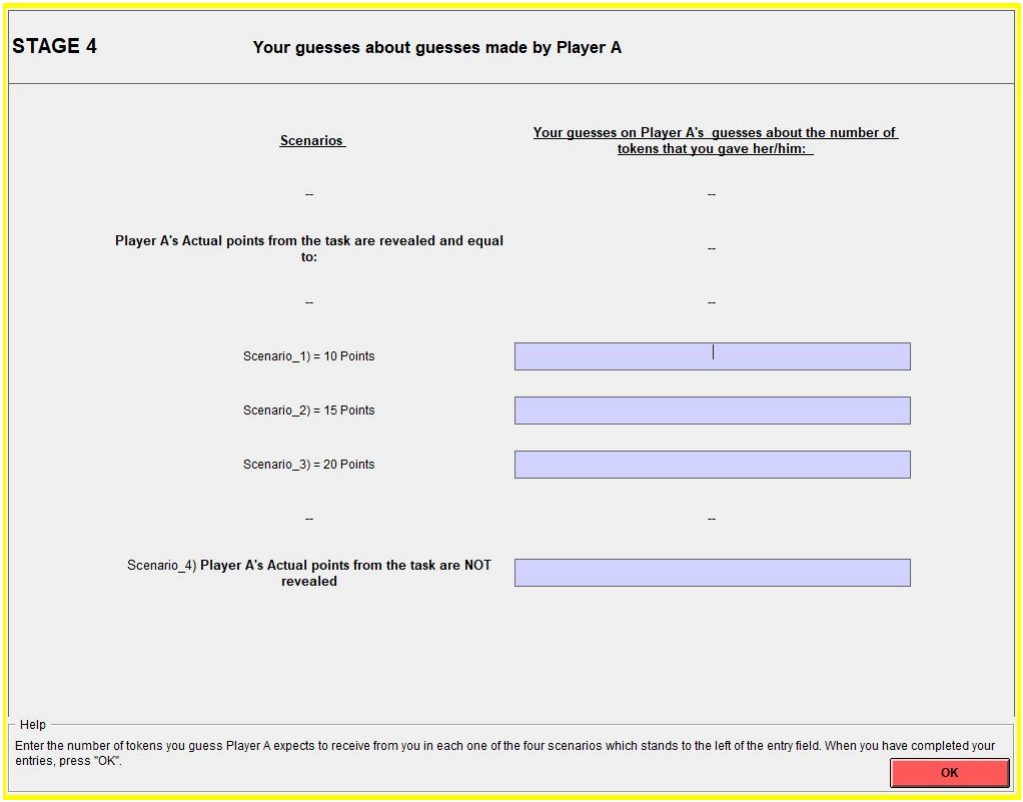


FIGURE G.6: Screen shot of the guess of revealing rates across outcome categories

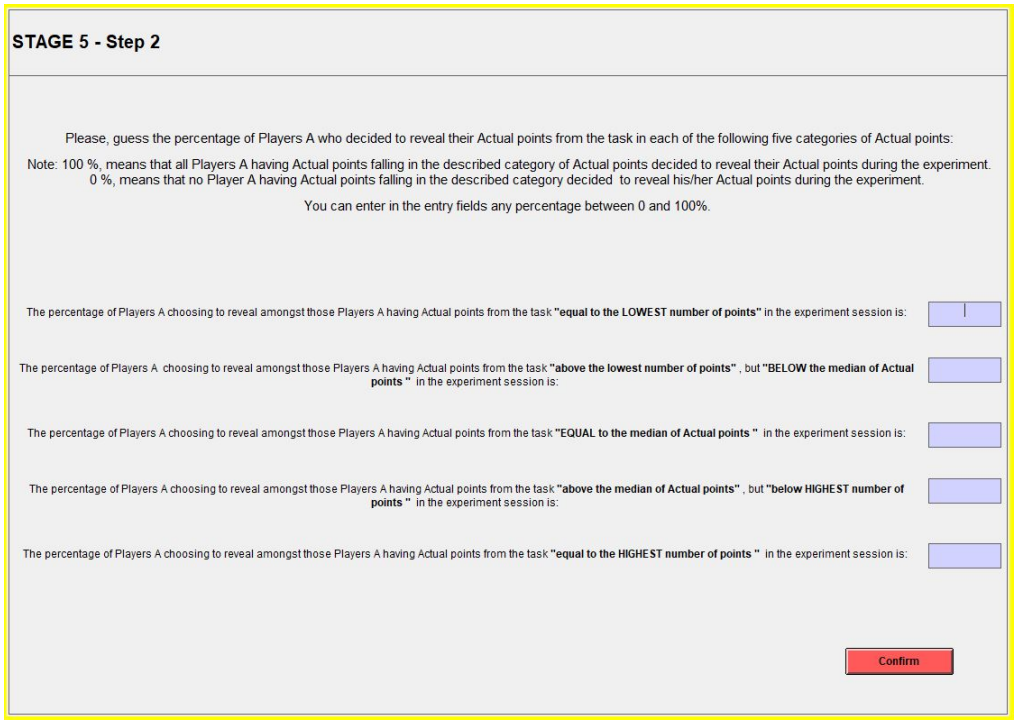

# Appendix H

# Chapter 3 Appendix H

## H.1 Hypothetical dictator game (Stage 5)

Table H.1 reproduces the ten choice problems of the dictator game developed by Kerschbamer (2015). Each choice problem asks participants to make a decision between two alternative payoff distributions, each one of which defines a different payoff allocation for the participant and another subject. Since in our experiment this game is not incentivized, we refer to the other subject as a hypothetical person.

The first five choice problems (X List) require participants to make distributional decisions in situations of disadvantageous inequality, while the last five choice problems (Y List) represent decisions in situations of advantageous inequality. Participants' decisions in the two choice lists provide therefore information about distributional preferences of subjects both in favourable and unfavourable inequality and can be connected both with dictator and receiver's role in our experiment. Furthermore, following the method developed by Kerschbamer (2015), each participant's set of decisions in the game can be summarized in a (X, Y) score tuple. Each score characterises a different aspect of a subject's social preferences and provides a direct measure of individuals' distributional concerns. Using these scores as controls in our analysis enables us to account for heterogeneity of participants' preferences in the results.

The X score, summarises subjects' distributional preferences in situations of disadvantageous inequality and it is determined by the decisions a subject makes in the choice problems of the X List. The X score assigned to a subject is defined by the row

number in which the subject chooses for the first time the "Left" option. Table H.2 shows how X score are determined. The earlier a subject chooses the left option, the higher her X score is. The Y score represents instead subjects' distributional preferences in a situation of favourable inequality. It is determined by choices made in the Y List, and by the row number in which the subject chooses for the first time the "Left" option. The last column of Table H.2 presents the rule used to convert a subject's decisions in Y score. The later a subject chooses the left option, the higher her Y score is.
X and Y scores indicate the weight that subjects put on others' payoff when making decisions respectively in situations of favourable and unfavourable inequality.
Higher X and Y scores, thus, represent a higher weight on the other's payoff and are associated with more altruistic preferences. Selfish preferences, which assume a low weight on others' payoff, are instead identified by X and Y score close to zero (selfish subjects have both X and Y score in {-0.5, 0.5}). The symmetric structure of the choice problems in the scenarios of disadvantageous and advantageous inequality, furthermore, allows us to identify preferences for inequality aversion as the combination of a low X score and a high Y score.

The assignment of X and Y scores to participants' choices in the two scenarios relies on the assumption that an individual's preferences satisfy strict monotonicity. The choice of a participant to switch from the right to the left option only once in each list of the game provides, however, a direct test that this assumption holds. The monotonic increase of the participants' payoff from one choice to the next in the left option, together with the invariant structure of both the other subject's payoff across all choices in the left option and payoffs across all choices in the right option, guarantee, in fact, that if a subject's preferences satisfy monotonicity, she should choose to switch only once.

In our sample, 88.96% of participants make decisions consistent with the monotonicity assumption, while 11.04% switch multiple times between left and right option, in either one or both choice lists. The share of inconsistent subjects is higher than

the one observed by Kerschbamer (2015) (5%). However, this difference could be explained by the fact that in our experiment the game was not incentivised. In order not to exclude observations from the analysis, and since most of the participants make inconsistent decisions only in one of the two choice lists[1], we choose not to eliminate inconsistent subjects from our sample. We instead replace their null X or Y score with the average score and account for this through the analysis, by introducing a control for consistency of choices.

Table H.3 presents the joint distribution of X scores and Y scores of participants making consistent choices. The table shows that most of subjects in our sample exhibit selfish preferences (X and Y score in {-0.5, 0.5}) and preferences for inequality aversion (X score=-2.5 and Y score=2.5). These results are in line with those presented by Kerschbamer (2015).

[1]Only 7 out of 36 inconsistent participants made decisions inconsistent in both choice lists.

TABLE H.1: Kerschbamer (2015)'s hypothetical dictator game

| Dec. Nr. | LEFT | | Your Choice | RIGHT | |
|---|---|---|---|---|---|
| | you receive | other person receives | | you receive | other person receives |
| | | | X List | | |
| 1 | 8 Pounds | 13 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |
| 2 | 9 Pounds | 13 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |
| 3 | 10 Pounds | 13 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |
| 4 | 11 Pounds | 13 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |
| 5 | 12 Pounds | 13 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |
| | | | Y List | | |
| 6 | 8 Pounds | 7 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |
| 7 | 9 Pounds | 7 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |
| 8 | 10 Pounds | 7 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |
| 9 | 11 Pounds | 7 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |
| 10 | 12 Pounds | 7 Pounds | LEFT RIGHT | 10 Pounds | 10 Pounds |

TABLE H.2: Determination of (X, Y)- Score (Kerschbamer, 2015)

| The subject chooses LEFT for the 1st time in row | in the X List (Dec. Nr. 1-5) | in the Y List (Dec. Nr. 6-10) |
|---|---|---|
| | X score | Y score |
| 1 | 2.5 | -2.5 |
| 2 | 1.5 | -1.5 |
| 3 | 0.5 | -0.5 |
| 4 | -0.5 | 0.5 |
| 5 | -1.5 | 1.5 |
| Never | -2.5 | 2.5 |

TABLE H.3: Sample Xscore and Yscore distribution

| | | **Xscore** | | | | | |
|---|---|---|---|---|---|---|---|
| | | **-2.5** | **-1.5** | **-0.5** | **0.5** | **1.5** | **2.5** |
| **Yscore** | **-2.5** | 4.83 | 1.03 | 1.72 | 1.03 | 0 | 0.34 |
| | **-1.5** | 1.03 | 0 | 0 | 0 | 0.34 | 0 |
| | **-0.5** | 2.07 | 0.69 | 3.79 | 4.14 | 0 | 0 |
| | **0.5** | 3.45 | 0.69 | 11.38 | 18.62 | 0.34 | 0 |
| | **1.5** | 1.03 | 1.38 | 2.07 | 3.45 | 1.72 | 0 |
| | **2.5** | 20.69 | 1.03 | 2.76 | 6.90 | 1.38 | 2.07 |

# Appendix I

# Chapter 3 Appendix I

FIGURE I.1: Rank index distribution by choice to reveal the task outcome

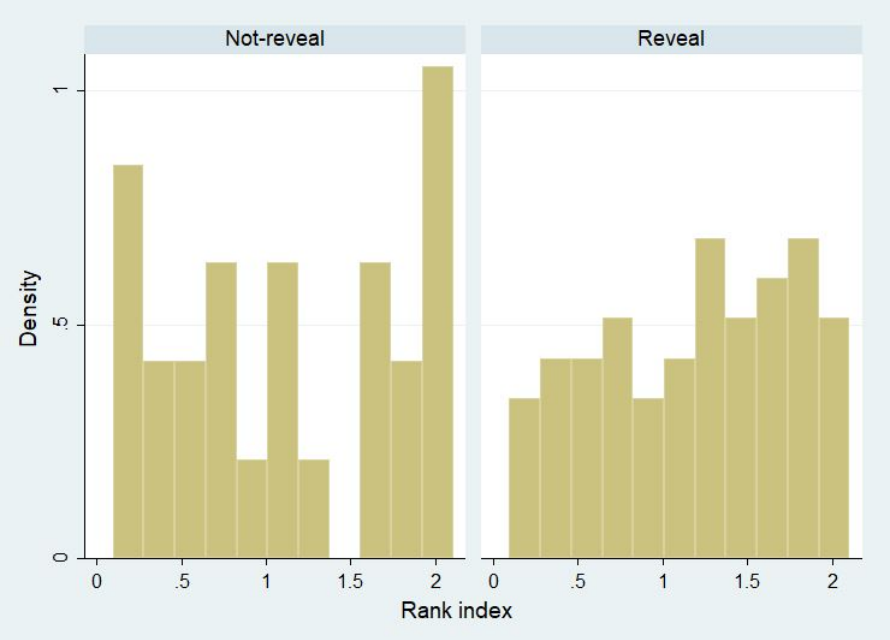


Figure I.1 presents the distribution of a rank index defined as: I=R/A, where R is the absolute rank of a subject in a session, A=(N+1)/2 is the average rank in a session and N is the number of participants in the session (the index is built following a technique developed by Booth and Kee (2009) for the study of birth order). By deflating a subject's rank by the average rank within a session, this index provides us with an alternative measure of rank which is independent from the size of the session and facilitates therefore the comparison of our data across sessions of different size. Figure

I.1 shows the distribution of the rank index for subjects who choose to hide their task outcome and subjects who choose to disclose their outcome in the left and right panel, respectively. The pattern that emerges from the two distributions corroborates the one presented in the main text.

TABLE I.1: OLS estimates of the effect of information on giving decisions - no inconsistent subjects

| | Other-regarding subjects | | Self-regarding subjects | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Infomed | 3.333 | 5.366 | 11.303** | 11.565* |
| | (5.786) | (5.008) | (5.54) | (6.484) |
| Outcome_above_median | | -2.864 | | -1.34 |
| | | (3.722) | | (6.271) |
| Receiver_type_shown | | 1.171 | | 0.02 |
| | | (3.827) | | (6.738) |
| Female | | -9.292** | | 9.527 |
| | | (4.395) | | (6.266) |
| Age | | 0.243 | | 0.448* |
| | | (0.235) | | (0.227) |
| Constant | 41.000*** | 40.778*** | 19.333*** | 12.068 |
| | (5.374) | (9.574) | (4.193) | (12.744) |
| R-squared | 0.013 | 0.229 | 0.0654 | 0.1966 |
| Observations | 38 | 38 | 45 | 45 |
| Controls | No | Yes | No | Yes |

Robust standard errors clustered on subject level. *** p<0.01, ** p<0.05, * p<0.1
Controls Yes: the dummy for having outcome above the median is one if the dictator has encoded a number of words above the session median number. The dummy for the receiver's type shown is one if the type of the receiver observed by the dictator is H-type. The dummy for education is one if the dictator is an undergraduate student and religion is a categorical variable having five categories.

TABLE I.2: OLS estimates of the effect of information on giving decisions - alternative thresholds

| | Yscore>0.5 | | 0.5>=Yscore>=-0.5 | | Yscore<-0.5 | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Infomed | 3.333 | 5.366 | 12.462** | 13.835* | 10.429 | 8.182 |
| | (5.786) | (5.008) | (5.945) | (6.970) | (8.765) | (8.426) |
| Outcome_above_median | | -3.173 | | 7.775 | | -32.225*** |
| | | (4.187) | | (6.181) | | (6.980) |
| Receiver_type_shown | | 1.072 | | 6.326 | | -56.840** |
| | | (4.634) | | (6.719) | | (22.263) |
| Female | | -9.292** | | 7.196 | | 57.548*** |
| | | (4.395) | | (6.918) | | (10.510) |
| Age | | 0.243 | | 0.373** | | -12.066** |
| | | (0.235) | | (0.169) | | (4.681) |
| Constant | 41.000*** | 40.778*** | 22.000*** | -6.730 | 6.000* | 248.742** |
| | (5.374) | (9.574) | (4.556) | (13.640) | (3.207) | (92.195) |
| R-squared | 0.013 | 0.229 | 0.090 | 0.345 | 0.062 | 0.976 |
| Observations | 38 | 38 | 36 | 36 | 9 | 9 |
| Controls | No | Yes | No | Yes | No | Yes |

Robust standard errors clustered on subject level. *** p<0.01, ** p<0.05, * p<0.1
Controls Yes: the dummy for having outcome above the median is one if the dictator has encoded a number of words above the session median number. The dummy for the receiver's type shown is one if the type of the receiver observed by the dictator is H-type.
The dummy for education is one if the dictator is an undergraduate student and religion is a categorical variable having five categories.

TABLE I.3: OLS estimates of treatment effect on transfer decisions

| | Average transfer decision | | |
|---|---|---|---|
| | (1) | (2) | (3) |
| Endogenous_NoInfo | -6.170 | -5.525 | -7.342** |
| | (4.520) | (4.451) | (3.634) |
| Exogenous_Info | -7.057* | -6.712* | -6.956* |
| | (3.895) | (4.036) | (3.689) |
| Exogenous_NoInfo | -4.258 | -4.088 | -2.079 |
| | (3.713) | (3.645) | (3.527) |
| Outcome_above_median | | -5.122 | -4.906* |
| | | (3.188) | (2.865) |
| Receiver_type_shown | | 0.626 | -0.258 |
| | | (3.145) | (2.763) |
| Female | | | 3.592 |
| | | | (2.865) |
| Age | | | 0.203 |
| | | | (0.140) |
| Xscore | | | -1.088 |
| | | | (1.140) |
| Yscore | | | 5.032*** |
| | | | (0.923) |
| Constant | 36.712*** | 38.981*** | 35.168*** |
| | (2.177) | (3.619) | (5.900) |
| R-squared | 0.024 | 0.045 | 0.267 |
| Observations | 163 | 163 | 163 |
| Controls | No | Yes | Yes+ |

Robust standard errors clustered on subject level. *** p<0.01, ** p<0.05, * p<0.1
Controls Yes: the dummy for having outcome above the median is one if the dictator has encoded a number of words above the session median number. The dummy for the receiver's type shown is one if the type of the receiver observed by the dictator is H-type.
Yes+ includes: a dummy for education, which is one if the dictator is an undergraduate student, religion and a dummy for consistency, which is one if the dictator has made consistent choices in the hypothetical dictator game.

Table I.3 shows results of the OLS regression, with robust standard errors clustered on the subject level, for the effect of treatment dummies on dictators' transfer decision. Model 2 controls for dictator's outcome and type of the receiver, while Model 3 adds controls for dictator's demographic characteristics and preference scores. The stars indicate statistically significant difference with respect to the baseline treatment where information about the receiver's outcome is available and endogenous.

# Bibliography


Aldrovandi, S., G. Brown, and A. Wood (2015). "Social norms and rank-based nudging: Changing willingness to pay for healthy food." In: *Journal of Experimental Psychology: Applied* 21.3, pp. 242–254.

Alesina, Alberto and George-Marios Angeletos (2005). "Fairness and redistribution". In: *American economic review* 95.4, pp. 960–980.

Almås, Ingvild, Alexander W Cappelen, and Bertil Tungodden (2020). "Cutthroat capitalism versus cuddly socialism: Are Americans more meritocratic and efficiency-seeking than Scandinavians?" In: *Journal of Political Economy* 128.5, pp. 1753–1788.

Almås, Ingvild et al. (2010). "Fairness and the development of inequality acceptance". In: *Science* 328.5982, pp. 1176–1178.

Andersson, O. and E. Wengstroem (2012). "Credible communication and cooperation: Experimental evidence from multi-stage games". In: *Journal of Economic Behavior and Organization* 81.1, pp. 207–219.

Andreoni, James (1990). "Impure altruism and donations to public goods: A theory of warm-glow giving". In: *The economic journal* 100.401, pp. 464–477.

Andreoni, James and John Miller (2002). "Giving according to GARP: An experimental test of the consistency of preferences for altruism". In: *Econometrica* 70.2, pp. 737–753.

Andreoni, James and Lise Vesterlund (2001). "Which is the fair sex? Gender differences in altruism". In: *The Quarterly Journal of Economics* 116.1, pp. 293–312.

Angerer, Silvia et al. (2015). "Donations, risk attitudes and time preferences: A study on altruism in primary school children". In: *Journal of Economic Behavior & Organization* 115, pp. 67–74.

Barr, Abigail and Pieter Serneels (2009). "Reciprocity in the workplace". In: *Experimental Economics* 12.1, pp. 99–112.

Beck, E., K. Burnet, and J. Vosper (2006). "Birth-order effects on facets of extraversion". In: *Personality and Individual Differences* 40, pp. 953–959.

Becker, Gary (1981). *A Treatise on the Family*. National Bureau of Economic Research, Inc.

Ben-Ner, Avner and Louis Putterman (2001). "Trusting and trustworthiness". In: *BUL Rev.* 81, p. 523.

Berg, Joyce, John Dickhaut, and Kevin McCabe (1995). "Trust, reciprocity, and social history". In: *Games and economic behavior* 10.1, pp. 122–142.

Bernard, Mark, Jack Fanning, and Sevgi Yuksel (2018). "Finding cooperators: Sorting through repeated interaction". In: *Journal of Economic Behavior and Organization* 147, pp. 76–94.

Bicchieri, Cristina (2005). *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press.

Bigoni, M. et al. (2015). "Time Horizon and Cooperation in Continuous Time". In: *Econometrica* 83, pp. 587–616.

Björklund, Anders and Markus Jäntti (1994). "The Impact of the Number of Siblings on Men's Adult Earnings: Evidence from Finland, Sweden and the United States". In: *EALE-conference in Warzaw*.

Black, Sandra E, Paul J Devereux, and Kjell G Salvanes (2005). "The more the merrier? The effect of family size and birth order on children's education". In: *The Quarterly Journal of Economics* 120.2, pp. 669–700.

Blackburn, J.M (1936). "Acquisition of Skill: An Analysis of Learning Curves". In: *IHRB Report* 73.

Blake, Judith (1989). *Family size and achievement*. Vol. 3. Univ of California Press.

Blanco, Mariana, Dirk Engelmann, and Hans Theo Normann (2011). "A within-subject analysis of other-regarding preferences". In: *Games and Economic Behavior* 72.2, pp. 321–338.

Blonski, M., P. Ockenfels, and G. Spagnolo (2011). "Equilibrium selection in the repeated prisoner's dilemma: Axiomatic approach and experimental evidence". In: *American Economic Journal: Microeconomics* 3(3), pp. 164–192.

Blonski, M. and G. Spagnolo (2015). "Prisoners' Other Dilemma". In: *International Journal of Game Theory* 44, pp. 61–81.

Boergers, T. and R. Sarin (1997). "Learning through reinforcement and replicator dynamics". In: *Journal of Economic Theory* 77, pp. 1–14.

Bolton, Gary E and Axel Ockenfels (2000). "ERC: A theory of equity, reciprocity, and competition". In: *American economic review* 90.1, pp. 166–193.

Booth, Alison L and Hiau Joo Kee (2009). "Birth order matters: the effect of family size and birth order on educational attainment". In: *Journal of Population Economics* 22.2, pp. 367–397.

Bortolotti, Stefania et al. (2017). "Too Lucky to be True-Fairness Views under the Shadow of Cheating". In:

Brody, Gene H (1998). "Sibling relationship quality: Its causes and consequences". In: *Annual review of psychology* 49.1, pp. 1–24.

Butler, Jeffrey V, Paola Giuliano, and Luigi Guiso (2016). "The right amount of trust". In: *Journal of the European Economic Association* 14.5, pp. 1155–1180.

Cabral, L., E. Ozbay, and A. Schotter (2014). "Intrinsic and instrumental reciprocity:An experimental study". In: *Games and Economic Behavior* 87, pp. 100–121.

Camera, Gabriele and Marco Casari (2009). "Cooperation among strangers under the shadow of the future". In: *American Economic Review* 99.3, pp. 979–1005.

Camerer, C. and TH Ho (1999). "Experience-Weighted Attraction Learning in Normal Form Games". In: *Econometrica* 67.4, pp. 827–874.

Camerer, Colin F (2003). "Behavioural studies of strategic thinking in games". In: *Trends in cognitive sciences* 7.5, pp. 225–231.

Cameron, Lisa et al. (2013). "Little emperors: behavioral impacts of China's One-Child Policy". In: *Science* 339.6122, pp. 953–957.

Cappelen, Alexander et al. (2019). "The Effect of Early Childhood Education on Social Preferences". In: *Journal of Political Economy*.

Cappelen, Alexander W, Cornelius Cappelen, and Bertil Tungodden (2018). "Second-best fairness under limited information: The trade-off between false positives and false negatives". In: *NHH Dept. of Economics Discussion Paper* 18.

Cappelen, Alexander W, Erik Ø Sørensen, and Bertil Tungodden (2010). "Responsibility for what? Fairness and individual responsibility". In: *European Economic Review* 54.3, pp. 429–441.

Cappelen, Alexander W et al. (2007). "The pluralism of fairness ideals: An experimental approach". In: *American Economic Review* 97.3, pp. 818–827.

Cappelen, Alexander W et al. (2013). "Needs versus entitlements—an international fairness experiment". In: *Journal of the European Economic Association* 11.3, pp. 574–598.

Carpenter, Jeffrey and Erika Seki (2011). "Do social preferences increase productivity? Field experimental evidence from fishermen in Toyama Bay". In: *Economic Inquiry* 49.2, pp. 612–630.

Charness, Gary and Matthew Rabin (2002). "Understanding social preferences with simple tests". In: *The Quarterly Journal of Economics* 117.3, pp. 817–869.

Chen, Yan and Sherry Xin Li (2009). "Group identity and social preferences". In: *The American Economic Review* 99.1, pp. 431–457.

Cherry, Todd L, Peter Frykblom, and Jason F Shogren (2002). "Hardnose the dictator". In: *American Economic Review* 92.4, pp. 1218–1221.

Cherry, Todd L and Jason F Shogren (2008). "Self-interest, sympathy and the origin of endowments". In: *Economics Letters* 101.1, pp. 69–72.

Cheung, Yin-Wong and Daniel Friedman (1997). "Individual learning in normal form games: Some laboratory results". In: *Games and Economic Behavior* 19.1, pp. 46–76.

Cialdini, Robert B, Raymond R Reno, and Carl A Kallgren (1990). "A focus theory of normative conduct: recycling the concept of norms to reduce littering in public places." In: *Journal of personality and social psychology* 58.6, p. 1015.

Conley, Dalton and Rebecca Glauber (2006). "Parental educational investment and children's academic risk estimates of the impact of sibship size and birth order from exogenous variation in fertility". In: *Journal of human resources* 41.4, pp. 722–737.

Cook, K.S. and R.M. Cooper (2003). *Experimental studies of cooperation, trust, and social exchange*, pp. 209–244.

Cooper, D.J. and K.U. Kuehn (2014). "Communication, renegotiation, and the scope for collusion". In: *American Economic Journal: Microeconomics* 6.2, pp. 247–278.

Corneo, Giacomo and Christina M Fong (2008). "What's the monetary value of distributive justice?" In: *Journal of Public Economics* 92.1-2, pp. 289–308.

Cox, James C, Daniel Friedman, and Steven Gjerstad (2007). "A tractable model of reciprocity and fairness". In: *Games and Economic Behavior* 59.1, pp. 17–45.

Croson, Rachel and Uri Gneezy (2009). "Gender differences in preferences". In: *Journal of Economic literature* 47.2, pp. 448–74.

Cunha, Flavio and James Heckman (2007). *The technology of skill formation*. Tech. rep. National Bureau of Economic Research.

Dal Bó, P. and G. Fréchette (2011). "The evolution of cooperation in infinitely repeated games: Experimental evidence". In: *American Economic Review* 101(1), pp. 411–429.

Dal Bó, P. and G. Fréchette (2018). "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey". In: *Journal of Economic Literature* 56.1, pp. 60–114.

Dal Bó, Pedro (2005). "Cooperation under the shadow of the future: experimental evidence from infinitely repeated games". In: *American Economic Review* 95.5, pp. 1591–1604.

Dana, Jason, Daylian M Cain, and Robyn M Dawes (2006). "What you don't know won't hurt me: Costly (but quiet) exit in dictator games". In: *Organizational Behavior and human decision Processes* 100.2, pp. 193–201.

Dana, Jason, Roberto A. Weber, and Jason Xi Kuang (2007). "Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness". In: *Economic Theory* 33.1, pp. 67–80.

Dave, Chetan et al. (2010). "Eliciting risk preferences: When is simple better?" In: *Journal of Risk and Uncertainty* 41.3, pp. 219–243.

Deckers, Thomas et al. (2015). "How does socio-economic status shape a child's personality?" In: *IZA Discussion Paper No. 8977*.

Dohmen, Thomas et al. (2011). "Individual risk attitudes: Measurement, determinants, and behavioral consequences". In: *Journal of the European Economic Association* 9.3, pp. 522–550.

Dohmen, Thomas et al. (2012). "The intergenerational transmission of risk and trust attitudes". In: *The Review of Economic Studies* 79.2, pp. 645–677.

Downey, Douglas B (1995). "When bigger is not better: Family size, parental resources, and children's educational performance". In: *American Sociological Review*, pp. 746–761.

Downey, Douglas B and Dennis J Condron (2004). "Playing well with others in kindergarten: The benefit of siblings at home". In: *Journal of Marriage and Family* 66.2, pp. 333–350.

Eckel, Catherine and Philip Grossman (1996). "Altruism in Anonymous Dictator Games". In: *Games and Economic Behavior* 16.2, pp. 181–191.

— (2002). "Sex differences and statistical stereotyping in attitudes toward financial risk". In: *Evolution and human behavior* 23.4, pp. 281–295.

Eckel, Catherine C and Philip J Grossman (2005). "Managing diversity by creating team identity". In: *Journal of Economic Behavior & Organization* 58.3, pp. 371–392.

Eisenberg, N and R.A. Fabes (1998). *Prosocial development*. National Bureau of Economic Research, Inc.

Ejrnæs, Mette and Claus C Pörtner (2004). "Birth order and the intrahousehold allocation of time and education". In: *Review of Economics and Statistics* 86.4, pp. 1008–1019.

Ellingsen, Tore et al. (2010). "Testing guilt aversion". In: *Games and Economic Behavior* 68.1, pp. 95–107.

Embrey, M., F. Mengel, and R. Peeters (2019). "Strategy Revision Opportunities and Collusion". In: *Experimental Economics* 22 (4), pp. 834–856.

Embrey, Matthew, Guillaume R Fréchette, and Sevgi Yuksel (2017). "Cooperation in the finitely repeated prisoner's dilemma". In: *The Quarterly Journal of Economics* 133.1, pp. 509–551.

Engel, Christoph (2011). "Dictator games: A meta study". In: *Experimental economics* 14.4, pp. 583–610.

Engelmann, Dirk and Martin Strobel (2000). "The false consensus effect disappears if representative information and monetary incentives are given". In: *Experimental Economics* 3.3, pp. 241–260.

Engle-Warnick, Jim and Robert L Slonim (2006). "Learning to trust in indefinitely repeated games". In: *Games and Economic Behavior* 54.1, pp. 95–114.

Erev, I. and A. Roth (1998). "Predicting How People Play in Games: Reinforcement Learning in Experimental Games with Unique Mixed Strategy Equiloibria". In: *The American Economic Review* 88.4, pp. 848–881.

Erkal, Nisvan, Lata Gangadharan, and Nikos Nikiforakis (2011). "Relative earnings and giving in a real-effort experiment". In: *American Economic Review* 101.7, pp. 3330–48.

Falbo, Toni and Denise F Polit (1986). "Quantitative review of the only child literature: Research evidence and theory development." In: *Psychological Bulletin* 100.2, p. 176.

Falk, Armin and Urs Fischbacher (2006). "A theory of reciprocity". In: *Games and Economic Behavior* 54.2, pp. 293–315.

Fehr, Ernst (2009). "On the economics and biology of trust". In: *Journal of the european economic association* 7.2-3, pp. 235–266.

Fehr, Ernst, Helen Bernhard, and Bettina Rockenbach (2008). "Egalitarianism in young children". In: *Nature* 454.7208, pp. 1079–1083.

Fehr, Ernst and Klaus M Schmidt (1999). "A theory of fairness, competition, and cooperation". In: *The quarterly journal of economics* 114.3, pp. 817–868.

Fischbacher, Urs (2007). "z-Tree: Zurich toolbox for ready-made economic experiments". In: *Experimental economics* 10.2, pp. 171–178.

Fischbacher, Urs and Simon Gächter (2010). "Social preferences, beliefs, and the dynamics of free riding in public goods experiments". In: *The American economic review* 100.1, pp. 541–556.

Fischbacher, Urs, Simon Gächter, and Ernst Fehr (2001). "Are people conditionally cooperative? Evidence from a public goods experiment". In: *Economics Letters* 71.3, pp. 397–404.

Foddy, Margaret, Michael J Platow, and Toshio Yamagishi (2009). "Group-based trust in strangers: The role of stereotypes and expectations". In: *Psychological Science* 20.4, pp. 419–422.

Fong, Christina (2001). "Social preferences, self-interest, and the demand for redistribution". In: *Journal of Public economics* 82.2, pp. 225–246.

Fong, Christina M (2007). "Evidence from an experiment on charity to welfare recipients: Reciprocity, altruism and the empathic responsiveness hypothesis". In: *The Economic Journal* 117.522, pp. 1008–1024.

Fong, Christina M and Felix Oberholzer-Gee (2011). "Truth in giving: Experimental evidence on the welfare effects of informed giving to the poor". In: *Journal of Public Economics* 95.5-6, pp. 436–444.

Fréchette, G. and S. Yuksel (2017). "Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination". In: *Experimental Economics* 20.2, pp. 279–308.

Friedman, Daniel and Ryan Oprea (2012). "A continuous dilemma". In: *American Economic Review* 102.1, pp. 337–63.

Fudenberg, Drew and David M Kreps (1993). "Learning mixed equilibria". In: *Games and Economic Behavior* 5.3, pp. 320–367.

Fudenberg, Drew and David K Levine (1998). *The Theory of Learning in Games*. Vol. 2. MIT press.

Gächter, Simon and Elke Renner (2010). "The effects of (incentivized) belief elicitation in public goods experiments". In: *Experimental Economics* 13.3, pp. 364–377.

Gerard, Harold B and Michael F Hoyt (1974). "Distinctiveness of social categorization and attitude toward ingroup members." In: *Journal of Personality and Social Psychology* 29.6, p. 836.

Glaeser, Edward L et al. (2000). "Measuring trust". In: *The Quarterly Journal of Economics* 115.3, pp. 811–846.

Golsteyn, Bart HH, Hans Grönqvist, and Lena Lindahl (2014). "Adolescent time preferences predict lifetime outcomes". In: *The Economic Journal* 124.580, F739–F761.

Grossman, Sanford J (1981). "The informational role of warranties and private disclosure about product quality". In: *The Journal of Law and Economics* 24.3, pp. 461–483.

Grossman, Sanford J and Oliver D Hart (1980). "Takeover bids, the free-rider problem, and the theory of the corporation". In: *The Bell Journal of Economics*, pp. 42–64.

Guala, Francesco, Luigi Mittone, and Matteo Ploner (2013). "Group membership, team preferences, and expectations". In: *Journal of Economic Behavior & Organization* 86, pp. 183–190.

Güth, Werner and Reinhard Tietz (1990). "Ultimatum bargaining behavior: A survey and comparison of experimental results". In: *Journal of Economic Psychology* 11.3, pp. 417–449.

Harbaugh, William et al. (2002). "Trust in Children". In: *University of Oregon Economics Department Working Papers*.

Hardin, Russell (2001). "Conceptions and Explanations of Trust". In: *Trust in Society*.

Hauser, Robert M and Hsiang-Hui Daphne Kuo (1998). "Does the gender composition of sibships affect women's educational attainment?" In: *Journal of Human Resources*, pp. 644–657.

Heckman, James J and Stefano Mosso (2014). "The economics of human development and social mobility". In: *Annu. Rev. Econ.* 6.1, pp. 689–733.

Herrmann, Benedikt and Christian Thöni (2003). "Understanding Determinants of Social Capital: Trust, Cooperation, and Informal Sanctions in a Cross-Societal Perspective". In:

Herrmann, Benedikt, Christian Thöni, and Simon Gächter (2008). "Antisocial punishment across societies". In: *Science* 319.5868, pp. 1362–1367.

Hopkins, E. (2002). "Two competing models of how people learn in games". In: *Econometrica* 70.6, pp. 2141–2166.

Jin, Ginger Zhe, Michael Luca, and Daniel Martin (2015). *Is no news (perceived as) bad news? An experimental investigation of information disclosure*. Tech. rep. National Bureau of Economic Research.

Juhn, Chinhui, Yona Rubinstein, and C Andrew Zuppann (2015). *The quantity-quality trade-off and the formation of cognitive and non-cognitive skills*. Tech. rep. National Bureau of Economic Research.

Kamas, Linda and Anne Preston (2015). "Can social preferences explain gender differences in economic behavior?" In: *Journal of Economic Behavior & Organization* 116, pp. 525–539.

Kerschbamer, Rudolf (2015). "The geometry of distributional preferences and a non-parametric identification approach: The Equality Equivalence Test". In: *European economic review* 76, pp. 85–103.

Keser, Claudia and Frans van Winden (2000). "Conditional Cooperation and Voluntary Contributions to Public Goods". In: *Scandinavian Journal of Economics* 102.1, pp. 23–39.

Kitzmann, Katherine M, Robert Cohen, and Rebecca L Lockwood (2002). "Are only children missing out? Comparison of the peer-related social competence of only children and siblings". In: *Journal of Social and Personal Relationships* 19.3, pp. 299–316.

Knoef, M and P Kooreman (2011). "The effect of cooperation: A structural model of Sibling's caregiving-interactions". In: 5733.

Konow, James (2000). "Fair shares: Accountability and cognitive dissonance in allocation decisions". In: *American economic review* 90.4, pp. 1072–1091.

Krupka, Erin L and Roberto A Weber (2013). "Identifying social norms using coordination games: Why does dictator game sharing vary?" In: *Journal of the European Economic Association* 11.3, pp. 495–524.

Lampi, Elina and Katarina Nordblom (2010). "Money and success–Sibling and birth-order effects on positional concerns". In: *Journal of Economic Psychology* 31.1, pp. 131–142.

Lefgren, Lars J, David P Sims, and Olga B Stoddard (2016). "Effort, luck, and voting for redistribution". In: *Journal of Public Economics* 143, pp. 89–97.

Lugovskyy, V. et al. (2017). "An experimental study of finitely and infinitely repeated linear public goods games". In: *Games and Economic Behavior* 102, pp. 286–302.

Mengel, F. (2018). "Risk and Temptation: A Meta-Study on Prisoner's dilemma games". In: *Economic Journal* 128, pp. 3182–3209.

Milgrom, Paul R (1981). "Good news and bad news: Representation theorems and applications". In: *The Bell Journal of Economics*, pp. 380–391.

Mitzkewitz, Michael and Rosemarie Nagel (1993). "Experimental results on ultimatum games with incomplete information". In: *International Journal of Game Theory* 22.2, pp. 171–198.

Mollerstrom, Johanna, Bjørn-Atle Reme, and Erik Ø Sørensen (2015). "Luck, choice and responsibility—An experimental study of fairness views". In: *Journal of Public Economics* 131, pp. 33–40.

Mookherjee, D. and B. Sopher (1997). "Learning and Decision Costs in Experimental Constant Sum Games". In: *Games and Economic Behavior* 19, pp. 97–132.

Piaget, J (1973). *Moral judgment of the child*. Free Press Paperbacks.

Rosenzweig, Mark R and Junsen Zhang (2009). "Do population control policies induce more human capital investment? Twins, birth weight and China's "one-child" policy". In: *The Review of Economic Studies* 76.3, pp. 1149–1174.

Ross, Lee, David Greene, and Pamela House (1977). "The "false consensus effect": An egocentric bias in social perception and attribution processes". In: *Journal of experimental social psychology* 13.3, pp. 279–301.

Roth, A. E. and J. K. Murnighan (1978). "Equilibrium behavior and repeated play of the prisoner's dilemma". In: *Journal of Mathematical Pyschology* 17.2, pp. 189–198.

Roth, Alvin E and Ido Erev (1995). "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term". In: *Games and Economic Behavior* 8.1, pp. 164–212.

Roth, Alvin E and J Keith Murnighan (1982). "The role of information in bargaining: An experimental study". In: *Econometrica: Journal of the Econometric Society*, pp. 1123–1142.

Ruffle, Bradley J (1998). "More is better, but fair is fair: Tipping in dictator and ultimatum games". In: *Games and Economic Behavior* 23.2, pp. 247–265.

Rustichini, Aldo (1999). "Optimal properties of stimulus—Response learning models". In: *Games and Economic Behavior* 29.1-2, pp. 244–273.

Sabater-Grande, G. and N. Georgantzis (2002). "Accounting for risk aversion in repeated prisoners' dilemma games: an experimental test". In: *Journal of Economic Behavior and Organization* 48.1, pp. 37–50.

Saroglou, V. and L. Fiasse (2003). "Birth order, personality, and religion: A study among young adults from a three-siblings family". In: *Personality and Individual Differences* 35, pp. 19–29.

Scheepers, Daan et al. (2006). "Diversity in in-group bias: Structural factors, situational features, and social functions." In: *Journal of personality and social psychology* 90.6, p. 944.

Selten, R. (1967). "Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments". In: *Beiträge Zur Experimentellen Wirtschaftsforschung*, pp. 136–168.

Shang, Jen and Rachel Croson (2005). "Field experiments in charitable contribution: the impact of social influence on the voluntary provision of public goods. Working paper". In:

Solomon, Daniel et al. (1988). "Enhancing children's prosocial behavior in the classroom". In: *American Educational Research Journal* 25.4, pp. 527–554.

Staub, Ervin (1979). "Positive social behavior and morality, Volume 2: Socialization and development". In:

Sugden, Robert (1984). "Reciprocity: The Supply of Public Goods through Voluntary Contributions". In: *Economic Journal* 94.376, pp. 772–87.

Sutter, Matthias et al. (2013). "Impatience and uncertainty: Experimental decisions predict adolescents' field behavior". In: *The American Economic Review* 103.1, pp. 510–531.

Thaler, Richard H (1988). "Anomalies: The ultimatum game". In: *Journal of economic perspectives* 2.4, pp. 195–206.

Thunström, Linda et al. (2016). "Endogenous context in a dictator game". In: *Journal of Behavioral and Experimental Economics* 65, pp. 117–120.

Tinghög, Gustav, David Andersson, and Daniel Västfjäll (2017). "Are Individuals Luck Egalitarians?–An experiment on the influence of brute and option luck on social preferences". In: *Frontiers in psychology* 8, p. 460.

Tversky, Amos and Itamar Simonson (1993). "Context-dependent preferences". In: *Management science* 39.10, pp. 1179–1189.

Vespa, E. (2019). "An experimental investigation of cooperation in the dynamic common pool game". In: *International Economic Review* 61.1, pp. 417–440.

Viscusi, W. Kip (1978). "A Note on "Lemons" Markets with Quality Certification". In: *The Bell Journal of Economics* 9.1, pp. 277–279.

Vriend, Nicolaas J (1997). "Will reasoning improve learning?" In: *Economics Letters* 55.1, pp. 9–18.

Wood, A.M, Gordon D.A. Brown, and John Maltby (2012). "Social Norm Influences on Evaluations of the Risks Associated with Alcohol Consumption: Applying the Rank-Based Decision by Sampling Model to Health Judgments". In: *Alcohol and Alcoholism* 47.1, pp. 57–62.