



# Behavioral Game Theory

HEINRICH HARALD NAX  
ETH ZÜRICH

A THESIS SUBMITTED FOR HABILITATION AT THE PROFESSORSHIP OF  
COMPUTATIONAL SOCIAL SCIENCE IN THE DEPARTMENT OF HUMANITIES,  
SOCIAL AND POLITICAL SCIENCES (D-GESS).

2015

Grüße an Oma Käthi, Elisabeth Eiche und Wilhelm und Margarete Busch.

# Behavioral Game Theory

HH NAX, 2015

HABILITATION COMMITTEE:

PROF. DIRK HELBING

(Computational Social Science, ETH Zürich)

PROF. GUILLAUME HOLLARD

(Economics, École polytechnique)

PROF. ANDREAS DIEKMANN

(Sociology, ETH Zürich)

PROF. TATSUYOSHI SAIJO

(Experimental Economics, Kochi University of Technology)

## Thesis outline

In one of his last publications, “The Agencies Method for Modeling Coalitions and Cooperation in Games” (*IGTR* 10: 539-564, 2008), John F. Nash wrote *I feel, personally, that the study of experimental games is the proper route of travel for finding “the ultimate truth” in relation to games as played by human players. But in practical game theory the players can be corporations or states; so the problem of usefully analyzing a game does not, in a practical sense, reduce to a problem only of the analysis of human behavior. It is apparent that actual human behavior is guided by complex human instincts promoting cooperation among individuals and that moreover there are also the various cultural influences acting to modify individual human behavior and at least often to influence the behavior of humans toward enhanced cooperativeness.*

To most, this quote will come as a surprise because the name John Nash is associated, not with behavioral or experimental research, but rather with the mathematics of game theory and, in particular, with the ‘Nash equilibrium’. The Nash equilibrium is an interactive solution concept mostly applied to idealized decision-making situations amongst infallible optimizers in the sense of pure material self-interest. But before we turn to Nash’s contributions and what the scope of this present thesis in addressing some of the issues raised above will be, let us first turn to the origins and broad evolution of the fields of game theory and of the study of behavior in games.

Beginning with the publication of “The Theory of Games and Economic Behavior” by John von Neumann and Oskar Morgenstern in 1944 (Princeton University Press), the study of human interactions by what since is known as “game theory” has revolutionized the sciences, especially the social sciences and biology. In economics, for instance, game theory changed the way equilibrium concepts are understood and eleven Nobel Prizes have since gone to game theorists. Game theory provides a sharp language to formulate mathematical models of underlying interactions that promise clean predictions, now integral parts of the social sciences toolbox.

A game is defined by a mapping from various combinations of “strategies” taken by the involved “players” into resulting consequences in terms of “pay-offs”. A “solution” predicts which outcomes of the game are to be expected. A major issue with traditional/neoclassical game theory, however, has been that its solution concepts, such as the Nash equilibrium (John F. Nash, 1950) or the strong equilibrium (Robert Aumann, 1959), rely on four rather extreme behavioral and informational assumptions. These are:

1. The joint strategy space is common knowledge.
2. The payoff structure is common knowledge.
3. Players have correct beliefs about other players’ behaviors and beliefs.
4. Players optimize their behavior so as to maximize their own material payoffs.

In the real-world, the whole ensemble of these assumptions is often untenable. Players often do not behave like infallible optimizers in the sense of pure material self-interest, and it would be negligent to think of the resulting consistent and structured deviations as inexplicable irrationalities. “Behavioral game theory”, the title and subject of this thesis, seeks to model these. Broadly speaking, behavioral game theory is separable into two strands of models.

The first strand of models presumes that the mismatch between theory and real-world behaviors may be the result of capacity and/or informational constraints. Hence, decisions are not best described by strictly maximizing behavior. In particular, maximization fails when players have incomplete information about the structure of the game and/or about the payoff consequences of different actions taken by themselves and others for the other players. In a repeated game setting, moreover, players may be unable -or only imperfectly able- to observe other players’ actions and payoffs as the game goes on. Hence, to describe more realistic human behaviors in complex game settings, models of boundedly rational behavior, possibly allowing for learning dynamics, are necessary.

The second strand of explanations for behaviors that consistently contradict equilibrium predictions based on the standard assumptions of self-interest and unbounded rationality may be that players are guided by alternative preferences. In other words, provided the assumption of material self-interest is flawed and that, instead, higher-order motives such as altruism or social norms guide a player’s actions, then self-interest predictions are misguided, even if players follow maximizing behavior. It is not that players’ behavior is not maximizing, rather their maximand is something else than narrow self-interest.

The two strands of explanations both have their respective appeals, and which model is preferable will depend on the context of the application. To describe the trading behavior of agents on financial markets, for example, one may favor the first type of explanation; intention but failure to maximize own material payoff. Similarly, such an approach may be preferable to describe behavior in traffic/congestion games. By contrast, richer preference formulations allowing for, for example, reciprocal considerations may be suitable to model voluntary contributions in situations such as community effort tasks or collective bargaining. Of course, in reality, we would expect an admixture of both explanations to matter in most situations, and their relative degrees to depend on the precise context and setting of the game.

“Behavioral game theory”, with the subtitle “Experiments in strategic interaction”, is also the title of one of the first and best-known textbooks that introduce this area of research to a broader audience (Camerer 2003; Princeton University Press). Therein, the two strands of explanations sketched above are expertly summarized and reviewed. This thesis builds on this body of work, its aim being to synthesize the two approaches. New methods to combine and

to disentangle the two are proposed and applied to different games, illustrating the need for a more nuanced theory, allowing for context-dependent behavior in games. The thesis consists of theoretical and behavioral studies. Moreover, the thesis also proposes a theoretical model of the complex decision-making of coalitions of individuals, and not just of individuals. The thesis is structured as follows:

## **Chapter 1. Introduction**

Chapter 1 provides an introduction to behavioral game theory. It focusses on the context of public goods games in general, and on social preference explanations in particular. A novel, unbiased method to estimate players social preferences is proposed, and used to disentangle social preferences from other factors such as learning. Data from several laboratory experiments is used. Interactions in preferences are discovered and assessed.

The author is the first and main author of the materials contained in this chapter. Underlying the chapter are two papers. One is joint work with Maxwell Burton-Chellew (Department of Zoology and Magdalen College, University of Oxford) and Stuart West (Department of Zoology and Nuffield College, University of Oxford), the other with Ryan Murphy and Kurt Ackermann (Department of Humanities, Social and Political Sciences, ETH Zürich). One paper is currently under review, the other published in *Economics Letters* (“Interactive preferences”, *Economics Letters* 135: 133-136, 2015).

## **Chapter 2. Learning**

Chapter 2 turns to learning behavior. It introduces a model of directional learning and proposes a novel solution concept that bridges Nash (1950) and Aumann (1959)-strong equilibria in the context of public goods games. Directional learning is shown to offer an alternative explanation of the behavioral regularities commonly observed in public goods experiments.

This chapter is joint work with Matjaž Perc (Faculty of Natural Sciences and Mathematics, University of Maribor, Slovenia & Department of Physics, Faculty of Science, King Abdulaziz University, Jeddah, Saudi Arabia). The paper is published in *Scientific Reports* (“Directional learning and the provisioning of public goods”, *Scientific Reports* 5: 8010, 2015). Matjaž Perc and the author are joint first authors.

## **Chapter 3. Social preferences versus learning**

Chapter 3 creates a horse-race between social preference explanations and directional learning in a repeated public goods game setting. Learning comes

out as the winner, but explanations related to social preferences, especially conditional cooperation, continue to matter.

This chapter is joint work with Maxwell Burton-Chellew (Department of Zoology and Magdalen College, University of Oxford) and Stuart West (Department of Zoology and Nuffield College, University of Oxford). It is published in the Proceedings of the Royal Society B (“Payoff-based learning explains the decline in cooperation in public goods games”, *Proceedings of the Royal Society of London B* 282: 20142678, 2015). Maxwell Burton-Chellew and the author are joint first authors.

#### **Chapter 4. Evolution of market equilibria**

Chapter 4 analyzes an evolutionary model of directional learning in a cooperative one-to-one matching game (Shapley and Shubik 1972) as used to model buyer-seller/firm-worker markets. It turns out that random interactions driven by directional adjustments of players’ bids and offers, over time, lead to optimality and stability of the cooperative market outcomes (competitive equilibria). Moreover, market equilibria with equity features are favored in the long-run because they turn out to be more stable.

This chapter is joint work with Bary Pradelski (Oxford-Man Institute of Quantitative Finance, University of Oxford). It is published in the International Journal of Game Theory (“Evolutionary dynamics and equitable core selection in assignment games”, *International Journal of Game Theory*, accepted and in print). Bary Pradelski and the author are joint first authors.

#### **Chapter 5. Complex cooperation**

Chapter 5 assesses the stability of complex cooperative outcomes when contracts amongst individual actors are written across multiple spheres of interaction. The definition of the Core solution concept is generalized from standard cooperative games (von Neumann and Morgenstern 1944), and possibilities for cross-cutting of contractual arrangements are illustrated and discussed.

This chapter is single-authored and published in *Games* (“A note on the Core of TU-cooperative games with multiple membership externalities”, *Games* 5: 191-203, 2014).

#### **Chapter 6. Dynamics of financial expectations**

Chapter 6 turns to the study of dynamics of risk expectations as captured by options data from the S&P 500 over the decade 2003 to 2013, separable into pre-crisis, crisis and post-crisis regimes. Particular focus is on the directedness of option-implied (expected) returns and on the causal structure of realized

and option-implied returns. The analysis reveals super-exponential growth expectations leading up to the Global Financial Crisis.

This chapter is joint work with Matthias Leiss (Department of Humanities, Social and Political Sciences, ETH Zürich) and Didier Sornette (Department of Management, Technology and Economics, ETH Zürich & Swiss Finance Institute, University of Geneva). It is published in the Journal of Economic Dynamics and Control (“Super-exponential growth expectations and the Global Financial Crisis”, *Journal of Economic Dynamics and Control* 55: 1-13, 2015).

## **Chapter 7. Meritocratic mechanism design**

Chapter 7 addresses the issue of mechanism design in the context of public goods games, both theoretically and experimentally. A mechanism based on “meritocratically” matching contributors-with-contributors and freeriders-with-freeriders is studied. Theory predicts that efficiency increases with meritocracy, but at the cost of growing inequality, so that a social optimum, when efficiency and equality are traded off, is commonly reached at intermediate levels of meritocracy. Experimental evidence suggests that higher levels of meritocracy increase both efficiency and equality. Fairness considerations explain this departure from the theoretical predictions.

This chapter is based on two papers that are joint work with Dirk Helbing and Ryan Murphy, and the experimental paper is furthermore co-authored with Stefano Balmelli (all authors are at the Department of Humanities, Social and Political Sciences, ETH Zürich). Both papers underlying the chapter are currently under review. The author is first author on both papers.

## **Conclusion**

Finally, a conclusion summarizes the main points and overarching findings of the thesis, and sketches the scope of the wider research program.



## Acknowledgements

This habilitation thesis has been written with the help and support of Dirk Helbing and Ryan Murphy while at ETH Zurich, and largely profiting from continued guidance of my former supervisor Peyton Young. I thank all of them for their help and support.

First and foremost, I am indebted to my co-authors who inspired many thoughts for further research beyond our joint projects. I would also like to thank Jean-Paul Carvalho, Gabrielle Demange, Jan Dörrie, Stefano Duca, Françoise Forges, Edoardo Gallo, Tom Norman, Heiko Rauhut, and, last but not least, Michael Mäs.

I am also thankful for comments and discussions from participants at various conferences, seminar series and workshops including at King's College, LSE, Maastricht, Cambridge, Oxford, Tokyo, Kyoto, Kochi, Stony Brook, Paris 1, Paris School of Economics, Monte Verità, and ETH Zürich.

I am grateful to have been supported by the European Commission through the ERC Advanced Investigator Grant 'Momentum' (Grant No. 324247). Finally, I thank Stefan Karlen and Dietmar Huber for guidance through the habilitation process.

All remaining errors are my own.

# Contents

Table of Contents . . . . .	i
<b>1 Introduction:</b>	
<b>Estimating social preferences</b>	<b>1</b>
1.1 Introduction . . . . .	5
1.2 Preference estimation . . . . .	13
1.3 Discussion . . . . .	21
1.4 Interactive preferences . . . . .	26
<b>2 Learning:</b>	
<b>Directional learning and the provisioning of public goods</b>	<b>41</b>
2.1 Introduction . . . . .	44
2.2 Results . . . . .	48
2.3 Discussion . . . . .	54
2.4 Methods . . . . .	57
<b>3 Social preferences versus learning:</b>	
<b>Learning and the contribution decline in public goods games</b>	<b>65</b>
3.1 Introduction . . . . .	68
3.2 Material and methods . . . . .	72
3.3 Results and discussion . . . . .	78
<b>4 Evolution of market equilibria:</b>	
<b>Equity dynamics in matching markets</b>	<b>89</b>
4.1 Introduction . . . . .	92
4.2 Related literature . . . . .	93
4.3 Matching markets with transferable utility . . . . .	96
4.4 Evolving play . . . . .	103
4.5 Core stability – absorbing states of the unperturbed process . . . . .	110
4.6 Core selection . . . . .	119
4.7 Conclusion . . . . .	129
<b>5 Complex cooperation:</b>	
<b>Agreements with multiple spheres of cooperation</b>	<b>138</b>
5.1 Introduction . . . . .	141
5.2 A worked example . . . . .	142
5.3 The model . . . . .	146
5.4 Coalitional stability and the core . . . . .	153
5.5 Concluding remarks . . . . .	158

<b>6</b>	<b>Dynamics of financial expectations:</b>	
	<b>Super-exponential growth expectations and crises</b>	<b>164</b>
6.1	Introduction . . . . .	166
6.2	Materials and methods . . . . .	169
6.3	Results . . . . .	174
6.4	Conclusion . . . . .	183
<b>7</b>	<b>Meritocratic mechanism design:</b>	
	<b>Theory and experiments</b>	<b>200</b>
7.1	Motivation . . . . .	204
7.2	Related literature . . . . .	205
7.3	Meritocratic matching . . . . .	208
7.4	Theoretical predictions . . . . .	225
7.5	The efficiency-equality tradeoff . . . . .	237
7.6	The experiment . . . . .	242
7.7	Experimental results . . . . .	246
7.8	Discussion . . . . .	262
<b>8</b>	<b>Conclusion</b>	<b>292</b>

# Chapter 1

**Introduction:**

**Estimating social preferences**

## Abstract

When players are involved in a voluntary contributions game, rich evidence shows that many agents often contribute substantially even when free-riding is the strictly dominant strategy. Assuming that agents maximize utility functions with a sociality parameter measuring their concern for other agents, this suggests a population bias toward pro-sociality. Indeed, there seems to be a widespread belief that contribution behavior in such contexts is adequately explained by pro-sociality. In this paper, we argue that the consensus has settled on this explanation too quickly. Our argument is backed by evidence from recent experiments that vary the strategic incentives of the game so that, in half of the games played by each agent, free-riding ceases to be a dominant strategy, and contributing fully is instead either a dominant or an equilibrium strategy. Applying the same logic, less-than-full contributions in these games would mean anti-sociality. Based on balanced within-subject comparisons, we identify a relatively symmetric distribution of pro- and anti-social preferences. Moreover, we reveal substantial inconsistencies at the individual level, that is, players whose behavior is suggestive of pro-sociality in the standard game often appear to act anti-socially in the game variation. This casts doubt on unconditional (pro-)sociality explanations, especially since most players whose actions are in line with the pursuit of pure material self-interest are found to do this consistently across treatments and adjust their actions accordingly when playing the two different games. Hence, learning and social motivations appear to coexist even in very simple games.

**Acknowledgements.** This research was supported by the ERC, Nuffield College and the Calleva Research Centre, Magdalen College. We thank Dan Wood for several important discussions and Stefano Balietti for the experimental design of the ‘meritocracy’ project (<http://nodegame.org/games/merit/>). We are also grateful to Yoshi Saijo and members of the Research Center for Social Design Engineering at Kochi University of Technology for helpful comments.

## Motivation.

*“The reader will find that the public goods environment is a very sensitive one. Many factors interact with each other in unknown ways. Nothing is known for sure. [...] There appear to be three types of players: dedicated Nash players who act pretty much as predicted by game theory with possibly a small number of mistakes, a group of subjects who will respond to self-interest as will Nash players if the incentives are high enough but who also make mistakes and respond to decision costs, fairness, altruism, etc., and a group of subjects who behave in an inexplicable (irrational?) manner. Casual observation suggests that the proportions are 50 percent, 40 percent, 10 percent in many subject pools. Of course, we need a lot more data before my outrageous conjectures can be tested.”*

(Ledyard, [1995](#))

## 1.1 Introduction

An important interaction studied in game theory is that of a population's joint effort to provide a public good by means of voluntary contributions. In its simplest form (the *voluntary contributions mechanism*; Isaac, McCue, and Plott, 1985; Isaac and Walker, 1988), the game can be split into three steps: first, each individual privately decides how much to contribute; second, the contribution total is multiplied by a rate of return  $r > 1$ ; third, the resulting public good is shared evenly amongst a group of  $s$  players. Assuming  $1 < r < s$ , this simple game thus succinctly captures potential conflict between individual and collective interests: on the one hand, each individual maximizes his own payoff by contributing nothing, but the sum of payoffs is highest if everyone contributes fully.

Research from over three decades has gone into understanding how people behave in these situations (see an early review by Ledyard, 1995). Typically, the design of such experiments has meant that the strictly dominant strategy is to contribute nothing (at least in the one-shot game and/or in the final stage of a repeated game), providing they conform to the predictions of models based on rationality and material self-interest (*homo oeconomicus*). The evidence from laboratory experiments, however, has been that many players contribute substantially and consistently, and therefore their behaviour contradicts the rational *homo oeconomicus* model. Subsequently, the consensus seems to have settled on explanations of this phenomenon according to which many people are not purely self-interested but prefer to consciously sacrifice own material payoff to increase the welfare of others instead (see a more recent review by Chaudhuri, 2011). An important question is whether this social preference explanation captures the actual thought process of agents or is more of an 'as if' kind?



Importantly, the social preference explanation, whilst rejecting one assumption of the *homo oeconomicus* model, that of pure self-interest, relies upon another assumption, that of perfect rationality. Rational choice theory assumes individuals to be fully rational and thus capable of expressing their preferences perfectly through the consequences of their actions (Becker, 1976). Therefore, a player who does not maximize his income in such games must have different preferences. For example, he willingly pays a price in terms of his own income in order to increase the income of others, which overall increases his utility.

An alternative explanation of behavior that contradicts the *homo oeconomicus* model is a lack of perfect rationality instead of a lack of pure self-interest, or even both. Andreoni, 1995a was amongst the first to try and disentangle these two elements, which turns out to be hard in general (see e.g. Andreoni, 1993; Andreoni, 1995a; Andreoni, 1995b; Palfrey and Prisbrey, 1996; Palfrey and Prisbrey, 1997; Goeree, Holt, and Laury, 2002; Houser and Kurzban, 2002; Gunnthorsdottir, Houser, and McCabe, 2007; Ferraro and Vossler, 2010; Bayer, Renner, and Sausgruber, 2013; Burton-Chellew and West, 2013). We refer the reader again to Ledyard, 1995 and Chaudhuri, 2011, the two principal reviews of this literature. We would like to point out here how pro-social interpretations seem to have become increasingly accepted in Chaudhuri, 2011 despite the unaddressed critical remarks regarding these interpretations present throughout Ledyard, 1995. We add to the literature concerned with disentangling intentions and learning/confused behavior by performing a ‘consistency/sanity check’ on preferences based on strategic variations of the same game played by the same players, thus proposing a novel approach to disentangle intentions and confusion, which we shall detail shortly.

Before we turn to details regarding our approach, let us first turn to typical

results as recorded in public goods experiments, which can be summarized by the following four regularities (see Ledyard, 1995; Chaudhuri, 2011; Anderson, Goeree, and Holt, 1998; Laury and Holt, 2008). First, aggregate contribution levels are increasing with the rate of return, even if the Nash equilibrium remains unchanged (at no-contribution). Second, contributions increase with the population size. Third, average contributions lie between the Nash equilibrium level and half the budget. Fourth, the population bias toward above-equilibrium contributions is increasingly neutralized (or even reversed) if the game is modified so that the Nash equilibrium becomes an interior (or high contribution) solution. The ensemble of these regularities can actually not be explained by most pro-social preference models, but Goeree and Holt, 2005 show their relation with quantal-response equilibrium (Palfrey and Prisbrey, 1997; Anderson, Goeree, and Holt, 1998); the logic being that off-equilibrium decisions occur with probabilities that are decreasing in their costliness vis-à-vis the reply that would maximize material self-interest.<sup>1</sup>

Despite the fact that social preference theory is yet to provide a fully conclusive account of experimental regularities, it seems that pro-social preference explanations have become predominant in the experimental economics literature, and whole areas of economics are now psychologically loaded with this tendency.<sup>2</sup>

The contribution of our paper is two-fold. First, we estimate social preferences explicitly under the assumption of full rationality by analysis of individual actions in games that vary the nature of the game (and of the Nash equilibrium).

---

<sup>1</sup>A related explanation is directional group (mis-)learning (Nax and Perc, 2015), which occurs more often when fewer members are needed to be able to mutually reinforce off-equilibrium paths.

<sup>2</sup>In psychology, social preference models were popular long before they reached economics (e.g. social value orientation; Griesinger and Livingston, 1973a). However, well-known issues with this approach deriving from non-reproducibility and confusion have since become widely accepted in psychology and there has since been substantial distancing from such models (Wicherts et al., 2006).

Second, we propose a novel way of discerning whether behavior is actually due to (fully rational) other-regarding preferences or due to bounded rationality by investigating the consistency within-subject. We compare how people play the standard public goods game with play of variants thereof where individual and collective interests are aligned. We shall refer to the class of variants as ‘profitable public goods games’. One way of creating a ‘profitable public goods game’ is to group players by contributions so that contributing more is rewarded by being matched with others doing likewise (Gunnthorsdottir et al., 2010; Nax et al., 2014). Hence, players who still do not contribute in these variants not only hurt others but may also not maximize their own payoffs. Another way to align individual and group interests is to make the public good sufficiently valuable in that the benefits of contributing outweigh the costs, not just at the group level (when contributions are multiplied by  $1 < r < s$ ), but also at the individual contributor level (when contributions are multiplied by  $r > s$ ). In such situations, a purely self-interested and rational player (*homo oeconomicus*) will contribute fully, as will any pro-social and rational player. Less-than-full contributions may either be due to an agent’s sub-rationality, thus hurting himself and others, in which case he may learn to contribute more with experience. Or the agent is rational but anti-social, in which case he may consistently contribute less than fully.

We use data from two experiments.<sup>3</sup> Half of the data from each experiment corresponds to the standard setting whereby rational self-interest predicts non-contributing behavior (‘standard public goods game’). The other half corresponds to a ‘profitable public goods game’. The two experiments differ with respect to the implementations of the two treatment types. In our analysis,

---

<sup>3</sup>Experimental instructions for one, based on data first reported in Burton-Chellew and West, 2013, can be found at <http://www.pnas.org/content/suppl/2012/12/14/1210960110.DCSupplemental/pnas.201210960SI.pdf>; instructions for the other, first reported in Nax et al., 2014, are at <http://nodegame.org/games/merit/>.

we first examine which preferences are being expressed in a given game under the assumption of rational choice, allowing players to have varying degrees of concerns for the payoffs of the other players. Similar efforts have previously focused on standard public goods games with/without punishment (e.g. Fehr and Gächter, 2000; Fehr and Gächter, 2002). The advantage of our study is that we are not restricted to inferring only utility functions that are either self-interested (and rational) or pro-social (and rational). By also using the data from profitable public goods, we are furthermore able *(i)* to infer anti-social (and rational) preferences, *(ii)* to examine within-player (in)consistencies between standard and profitable public goods, and *(iii)* to assess if the total population is on average pro-social, anti-social, or neither/both. If our players appear to have inconsistent social preferences, then it may be that our assumption of full rationality is unjustified instead. We can therefore use our findings to classify people into the three types as proposed by Ledyard, 1995; dedicated Nash players (ca. 50%), social players (ca. 35%, of which pro-social ca. 20%/anti-social ca. 15%), and inexplicable/irrational players (ca. 15%).

This paper is a companion to our related papers on ‘learning’ in public goods games (Burton-Chellew and West, 2013; Nax et al., 2013; Burton-Chellew, Nax, and West, 2015; Nax et al., 2014). In different ways, these papers explored the learning dynamics in the underlying experiments (separately) and compared aggregate patterns across treatments. None of the previous papers addressed the estimation of pro/anti-social preference parameters explicitly, which is the novelty of this study. The original data underlying our analysis is first recorded in Burton-Chellew and West, 2013 and Nax et al., 2014; please also refer to these papers for detailed experimental instructions. In this paper, our analysis proceeds as follows. First, we assume a non-linear functional form (Cobb and Douglas, 1928) to express individual preferences that allows agents to have varying degrees of concern for the payoffs of the other

players. Compared to most existing studies (see Chaudhuri, 2011; Saijo, 2014) ours is different with regards to the non-linearity of the utility function which allows us to rationalize (in terms of degrees of pro/anti-sociality) intermediate contribution decisions.<sup>4</sup> This is important as ca. 30% of all contributions are intermediate even in the final period of play. A linear utility function makes ‘bang-bang’ predictions of extreme decisions, and explains intermediate contribution decisions only at the knife-edge of indifference (e.g. Levine, 1998; Saijo, 2014). Previously, the Cobb-Douglas function has been used to generate the payoffs of the game (Andreoni, 1993; Chan et al., 2002; Cason, Saijo, and Yamato, 2002) but not as a utility function to measure players’ altruistic concerns regarding game payoffs as we do here.<sup>5</sup> Both data sets used in this study contain one ‘dilemma treatment’ (corresponding to a ‘standard public good’) and one ‘provision treatment’ (corresponding to a ‘profitable public good’). Burton-Chellew and West, 2013’s design follows Andreoni, 1988, and complements their design with a high-rate-of-return game for the ‘provision treatment’ as pioneered in Saijo and Nakamura, 1995 (see also Cason, Saijo, and Yamato, 2002). Nax et al., 2014 also follow Andreoni, 1988 in the ‘dilemma treatment’ (with a different rate of return), but adopt Gunthorsdottir et al., 2010’s mechanism in the ‘provision treatment’.<sup>6</sup>

Our results summarize as follows. In the dilemma treatments, there exists a sizeable fraction of *homo oeconomicus* and in addition an array of hetero-

---

<sup>4</sup>Saijo, 2014 shows that substantial instabilities are associated with the voluntary contribution mechanism in terms of best reply dynamics when the payoff functions underlying the game. (This is a different on-linearity than the one concerning individuals’ utility functions as we assume.)

<sup>5</sup>There are a number of hybrid (linear/non-linear) payoff functions that have been used; either using a linear own-payoff component and a non-linear other-payoff component (e.g. Isaac, McCue, and Plott, 1985; Isaac, Walker, and Thomas, 1991; Isaac and Walker, 1998; Laury, Walker, and Williams, 1999), or using a non-linear own-payoff component and a linear other-payoff component (e.g. Sefton and Steinberg, 1996; Keser, 1996; Falkinger et al., 2000).

<sup>6</sup>More detail will be provided later in this chapter, and chapter 7 of this thesis will contain full details.

geneously pro-social players. The distributions implied by the two datasets are similar and consistent with many previous studies. Diametrically opposed results come out of the provision treatments. In the provision treatment by Burton-Chellew and West, 2013, the social preference story is literally turned upside-down by failure to play Nash equilibrium, suggesting that there exists a fraction of *homo oeconomicus* and, in addition, an array of heterogeneously anti-social players. This is a result of incentive structure and contribution decisions being mirror images of the dilemma treatments. The provision treatment by Nax et al., 2014 reveals an even larger fraction of players indistinguishable from *homo oeconomicus* and, in addition, behaviors inexplicable by social preferences. Combining the two treatments and checking for within-player consistencies, we find roughly half of the population to be consistent with the *homo oeconomicus* model, one third consistent with social preferences (characterized by either pro-sociality or anti-sociality), and roughly 15% of players to be inexplicable in terms of (social) preferences. Our findings show that social preference estimations may be highly sensitive to equilibrium-relevant parameter changes even within the same class of games played by the same population, to the extent that implied signs of pro/anti-sociality may actually be reversed. The current status of social preference theory is far from being predictive as regards such phenomena. In light of other factors that have previously been shown to affect social preference distributions such as framing, stakes, beliefs, learning, or preference conditionality, it appears unlikely that this will change soon outside very specific classes of games, where the reason for predictability may be due to data patterns being reliably similar rather than pro-sociality actually accurately explaining the nature of man. This may imply that the proclaimed generality and solidity of findings regarding pro-social preferences (e.g. Fehr and Camerer, 2007; Bowles and Gintis, 2011) was exaggerated, perhaps even grossly.

Apart from our aforementioned companion papers on learning (Burton-Chellew and West, 2013; Nax et al., 2013; Burton-Chellew, Nax, and West, 2015; Nax et al., 2014), the paper that is closest to ours is Saijo and Nakamura, 1995 (see also Saijo and Yamato, 1999; Cason, Saijo, and Yamato, 2002; Brandts, Saijo, and Schram, 2004; Saijo, 2008). Similarly to Burton-Chellew and West, 2013, Saijo and Nakamura, 1995 have an experimental design with one high-rate-of-return game and one low-rate-of-return game, leading to off-equilibrium behavior that is ‘kind’ in one (low rate) and ‘spiteful’ (high rate) in the other.<sup>7</sup> The part of our analysis based on Burton-Chellew and West, 2013’s data can be seen as a formal extension of their study beyond the two-by-two case which allows us to infer degrees of kindness/spite using a nonlinear utility representation. Saijo and Nakamura, 1995 classify people without formal utility assumptions. In addition, we compare our estimates with the data from alternative matching mechanisms (the Nax et al., 2014 data). Also related to our study is Levine, 1998 who estimates spite and kindness in a linear (negative or positive) altruism framework. Most of Levine’s analysis is based on the ultimatum game, but he also considers the low-rate-of-return contributions game data of Isaac and Walker, 1988, where he finds estimates consistent with our provision treatments. Again, the key difference with respect to the parameter estimation as compared to his model is our choice of a non-linear utility function, hence our model does not have a ‘bang-bang’ solution, that is, we do not explain intermediate contributions by randomization at indifference, but by intermediate levels of kindness/spite. Moreover, we complement Levine, 1998’s analysis (which across games yields a distribution of spite and kindness that is qualitatively somewhat similar to ours) by a balanced within-subject rather than between-subject design in the context of public goods games. Finally, we investigate the presence of contradictory social motives across games.

---

<sup>7</sup>An important difference is that Saijo and Nakamura, 1995’s experimental design is between-subject, not within-subject.

In the remainder of this note, we shall dive straight into the analysis next, and come back to the existing literature in our concluding discussion.

## 1.2 Preference estimation

### 1.2.1 Experimental design

*Contribution game.* Population  $N = \{1, 2, \dots, n\}$  plays the following game repeatedly during periods  $t = \{1, 2, \dots, T\}$ . Each  $i \in N$  chooses to contribute a certain amount,  $c_i = \{0, 1, 2, \dots, B\}$ , where  $B$  is the budget. Write  $c$  for a resulting strategy vector. Given the rate of return  $r > 0$ ,  $i$ 's resulting payoff (but not necessarily final utility!) is

$$\phi_i(c) = (1 - c_i) + \sum_{j \in S} \frac{r}{s} c_j,$$

where  $S$  is the group (of fixed size  $s$ ) of which  $i$  is a member. Write  $\phi$  for a resulting payoff vector. We shall call  $\frac{r}{s}$  the game's marginal per-capita rate of return (mpcr).

*Information.* All players get full instructions about the game (see Burton-Chellew and West, 2013; Nax et al., 2014 for details).<sup>8</sup> After the experiment, total earnings are paid out according to a known exchange rate. Each round, players get information about other players' contributions (about the players in one's own group only in Burton-Chellew and West, 2013, about all players in Nax et al., 2014).

*Group re-matching.* Each round, players are re-matched into groups of a fixed size of four players in all cases. In all but the provision treatment of Nax et al.,

---

<sup>8</sup>For details, please see the Supporting Information (201210960SI) for Burton-Chellew and West, 2013, and <http://nodegame.org/games/merit/> for Nax et al., 2014.



2014, this occurs randomly as in Andreoni, 1988: that is, group formation is independent of contribution decisions. In Nax et al., 2014’s provision treatment, players are matched according to the *group-based mechanism* (Gunnthorsdottir et al., 2010), that is, groups form in order of their contributions; the highest contributors form group one, etc. (with random tie-breaking).

Table 1.1 summarizes the treatments.

Table 1.1: Summary of treatments.

	Burton-Chellew and West, 2013		Nax et al., 2014	
<b>Treatment</b>	<b>‘Dilemma’</b>	<b>‘Provision’</b>	<b>‘Dilemma’</b>	<b>‘Provision’</b>
<b>Population size</b> $n =$	12 or 16	12 or 16	16	16
<b>Group size</b> $s =$	4	4	4	4
<b>Budget</b> $B =$	40	40	40	40
<b>Rate of return</b> $r =$	1.6	6.4	2	2
<b>Repetitions</b> $T =$	20	20	40	40
<b>Group re-matching</b>	Random	Random	Random	Contribution-based
<b>Dominant strategy</b>	$c_i = 0$	$c_i = 40$	$c_i = 0$	n/a

## 1.2.2 Preference assumptions

*Utility.* Agent  $i$ ’s utility is assumed to be Cobb-Douglas of the form


$$u_i(c) = \phi_i^{1-\alpha_i} * \phi_{-i}^{\alpha_i},$$

where  $\phi_{-i}^{\alpha_i}$  is the average payoff to players  $j \neq i$  about whom  $i$  learns in the relevant treatment.

*Concern for others.*  $\alpha_i$  measures player  $i$ ’s concern for others. He has no concern for others when  $\alpha_i = 0$  (*homo oeconomicus*). He has anti-social concern for others when  $\alpha_i < 0$  in which case he is willing to sacrifice own payoff if others’ suffer even greater payoff losses (*competitive*). He has pro-social concern for others when  $0 < \alpha_i \leq 10$  (*other-regarding*): when  $0 < \alpha_i < 0.5$  he cares

more for himself (*moderate altruism*), when  $\alpha_i = 0.5$  he cares equally (*impartial altruism*), when  $0.5 < \alpha_i < 1$  he cares more for others (*strong altruism*), when  $\alpha_i = 1$  he cares only for others (*pure altruism*). When  $\alpha_i > 1$ , he is willing to sacrifice others' payoff only when he loses even more (such behavior is somewhat *anti-competitive* or simply *irrational*). These latter irrational motivations (when  $\alpha_i > 1$ ) are so strange that we shall consider them as evidence for (weak) inconsistency. Note that such an agent would prefer burning own and others' payoffs as long as he makes himself worse-off than others as a result. Figure 1.1 illustrates the types of concerns for others.

Figure 1.1: Types of concern.



Of course, many alternative utility functions could be assumed. We choose this utility function for two main reasons. First, it is governed by a single additional parameter, elegantly interpretable by social preferences as outlined above. Second, as opposed to, for example, a linear weight on others' payoff (e.g. Levine, 1998), the Cobb-Douglas representation can explain intermedi-

ate contributions and typically associates them with a unique level of concern for others. A linear function in the dilemma treatment, for example, by contrast would imply ‘bang-bang’ behavior, i.e. an insufficiently pro-social agent would always free-ride, while any agent with pro-sociality above some threshold would contribute fully. Intermediate contributions could only be rationalized by randomization at the threshold (see Levine, 1998).

### 1.2.3 Estimation technique

In our main analysis, we focus on the last period of each experiment since earlier decisions may be rationalizable in other ways, for which alternative preference representation such as Fehr-Schmidt, Bolton-Ockenfels or Charness-Rabin preferences may be preferable. (We will also compare with early-game evidence later on.) Focus on the final period also has the advantage that conditional cooperation and other phenomena are likely to be ‘over’ in the sense that dynamics of such determinants are likely to have settled and/or led to equilibrium. We estimate, for each individual separately, the concern for others as implied by his action taking as given the others’ penultimate-period average contributions. If his action coincides with what *homo oeconomicus* would do, we set  $\alpha_i = 0$ . Otherwise, we assume an interior solution and obtain an expression for  $\alpha_i$  using the first-order condition  $\frac{\partial u_i(c)}{\partial c_i} = 0$ , where  $c_{-i}$  is taken to be  $c_{-i}^{t-1}$ . Solving this toward  $\alpha_i$  for random re-matching gives

$$\alpha_i = (1 - mpcr) * \phi_{-i}(c) / (mpcr * \phi_i(c) + (1 - mpcr) * \phi_{-i}(c)).$$

Note that, by setting  $\alpha_i = 0$ , i.e. by assuming the agent is *homo oeconomicus*, when his action coincides with what *homo oeconomicus* would do, we create a gap in estimating concerns around  $\alpha_i = 0$  in the sense that concerns for

others that are too tiny to actually ever matter are not estimated. In other words, we do not distinguish between an agent with concerns for others (either pro-social or anti-social ones) that are too small to ever matter and an actual *homo oeconomicus*.

*Example 1.* Take the dilemma treatment by Nax et al., 2014. Suppose every other player, from the viewpoint of player  $i$ , made a contribution of zero in the penultimate period ( $t = 39$ ). If player  $i$  decides to contribute half his budget in the last period then this implies a concern for others of  $\alpha_i = 0.625$ , i.e. he cares relatively more about others than about himself (is pro-social of type ‘strong altruist’). If he contributes zero, then we assume he is *homo oeconomicus* and set  $\alpha_i = 0$ .

*Example 2.* Take the provision treatment by Burton-Chellew and West, 2013. Suppose every other player made a full contribution in the penultimate period ( $t = 19$ ). If player  $i$  decides to contribute half his budget in the last period then this implies a concern for others of  $\alpha_i = -0.525$ , i.e. he has a negative concern for others (is anti-social/‘competitive’). If he contributes fully, then we assume he is *homo oeconomicus* and set  $\alpha_i = 0$ .

Obtaining an expression for  $\alpha_i$  in the provision treatment for Nax et al., 2014 where re-matching follows the group-based mechanism (Gunnthorsdottir et al., 2010) is more complicated because no strategy is a priori dominated (see chapter 7 for more detail on equilibrium structure; see also Gunnthorsdottir et al., 2010 and Nax, Murphy, and Helbing, 2014 for details of equilibrium analysis). We illustrate at hand of an example how we obtain expressions for  $\alpha_i$ .

*Example 3.* Take the provision treatment by Nax et al., 2014. Suppose every player but player  $i$  made a full contribution in the penultimate period ( $t = 39$ ). If player  $i$  decides to free-ride in the last period this coincides with what *homo*


oeconomicus would do and we set  $\alpha_i = 0$ . If he decides to contribute say 10, then we conclude that his concern for others is pro-social of order  $\alpha_i = 0.802$  (is strongly altruistic).

## 1.2.4 Estimation results

Figure 1.2 shows the final-round contributions, Figure 1.3 the implied social preferences.

Figure 1.2: Distributions of final-round contributions (final round).

(a) Burton-Chellew and West, 2013: contributions in period  $t = 20$



(b) Nax et al., 2014: contributions in period  $t = 40$


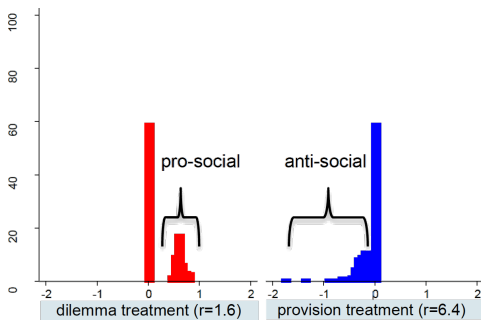


Figure 1.3: Distributions of preferences (final period).

(a) Burton-Chellew and West, 2013: preferences



(b) Nax et al., 2014: preferences

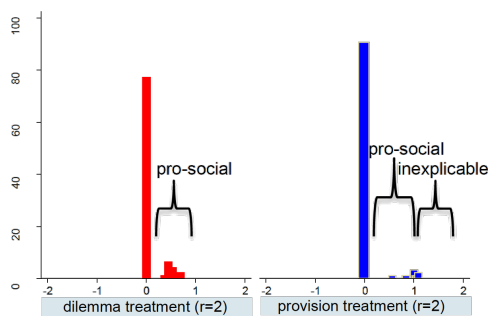



Figure 1.4 combines Figures 1.2 and 1.3. Table 1.2 summarizes the moments of Figure 1.4. The shaded areas in Figure 1.4 indicate the ranges of preferences over which actions by *homo oeconomicus* and those by agents with preferences from this range coincide.

Figure 1.4: Distributions combined (final period).



**Preference consistency.** Of the 236 individuals in Burton-Chellew and West, 2013, 88 (37.3%) are consistent with pure *homo oeconomicus* preferences in both treatments. Another 102 (43.2%) are consistent in the following sense: 51 (21.6%) are *homo oeconomicus* in the dilemma treatment and anti-social in the provision treatment; another 51 (21.6%) are *homo oeconomicus* in the provision treatment and pro-social in the dilemma treatment. 46 (19.5%) individuals are strongly inconsistent, i.e. pro-social in the dilemma treatment and anti-social in the contribution treatment.

Of the 96 individuals in Nax et al., 2014, 67 (69.8%) are consistent with pure *homo oeconomicus* preferences in both treatments. Another 20 (20.1%) are consistent in the sense that *homo oeconomicus* in the provision treatment and pro-social in the dilemma treatment. 7 (7.3%) individuals are weakly inconsistent in that they are *homo oeconomicus* in the dilemma treatment but act anti-competitively in the provision treatment (burning others' and

their own payoffs). 2 (2.1%) individuals are strongly inconsistent in the sense of pro-social in the dilemma treatment and anti-social in the contribution treatment.

Overall, in terms of implied social preferences from both experiments combined, there are 46.7% players consistent with *homo oeconomicus*, 15.4% consistent and anti-social, 21.4% consistent and pro-social, and the remaining 16.5% are inconsistent.

Table 1.2: Distribution summaries for “ $\alpha$ ” (final round).

	Combined	Burton-Chellew and West, 2013		Nax et al., 2014	
		<b>Dilemma</b>	<b>Provision</b>	<b>Dilemma</b>	<b>Provision</b>
<b>mean</b>	0.08	0.26	−0.13	0.12	0.09
<b>s.d.</b>	0.31	0.32	0.21	0.22	0.29
<b>skewness</b>	0.51	0.45	−3.37	1.41	2.91
<b>kurtosis</b>	6.03	1.36	21.37	3.20	9.60
<b>median</b>	0	0	0	0	0
<b>#observations</b>	664	236	236	96	96

**Comparison with early-round play.** In round 1, 85% contributed a positive amount in the dilemma treatment, and 58% contributed less than their full budget in the provision treatment in the Burton-Chellew and West, 2013 data. There are no significant differences in play between initial and final rounds in Nax et al., 2014. Hence, let us talk about Burton-Chellew and West, 2013 data for now. Note first that we cannot apply our preference estimation technique to judge the initial-round contribution decisions because our estimation relies on information about previous-period decisions. Moreover, our technique is, strictly speaking, uninformative regarding inconsistencies and social preference motivations in early rounds of the game as any action is potentially *homo oeconomicus*-rationalizable by a more complex repeated game strategy in initial rounds, especially if we model beliefs (Kreps et al., 1982). But what we can do is check whether players who, in the final round, contribute positive in the dilemma treatment or less-than-full in the provision treatment already did

so in the initial period. Indeed, over 90% percent did (98% in the dilemma treatment, 82% in the provision treatment), suggesting a high degree of internal consistency over time.<sup>9</sup>

If we apply our preference estimation technique to period 2, which is the first period in which this is possible (taking period-1 contributions as inputs), we obtain only a slightly different picture of the combined distribution compared with those obtained from analysis of the final period. First and foremost, there appear to be fewer individuals of the *homo oeconomicus* type, reflecting the fact that we see more intermediate decisions. Note that intermediate decisions, by the end of the experiment, have largely disappeared, possibly through learning or reciprocity dynamics. As in the final period, the median of the combined distribution in the second period again corresponds to *homo oeconomicus* (the median is 0). The reason is that the quantity of additional above-equilibrium contributions in the dilemma treatments is similar to that of additional below-equilibrium contributions in the provision treatment.

### 1.3 Discussion

A large part of experimental economics has focused on what can be called the ‘subjective expected utility correction project’ (Gigerenzer and Selten, 2001). The most famous experiments in economics (Allais, 1953; Ellsberg, 1961; Ainslie, 1975; Kahneman and Tversky, 1979) challenge the axioms of standard decision theory and with it the notion of man as a perfect expected utility maximizer (Ramsey, 1931; Von Neumann and Morgenstern, 1944; Savage, 1954), proposing behavioral corrections of the model. More recently and almost as famously, laboratory experiments on interactive decisions were con-

---

<sup>9</sup>Further confirmation of internal consistency stems from the fact that over two thirds of the inexplicable/irrational agents, as judged by our consistency check in the final period, also already violate Nash equilibrium play in both treatments in the initial period.



ducted, and they have been interpreted as evidence for the fact that humans not only care about their own material payoffs (pure self-interest) but also, in a pro-social way, about those of others (e.g. Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Fischbacher and Gächter, 2010; see Murphy and Ackermann, 2014 for a recent review that includes psychology references). Quite obviously this challenges the *homo oeconomicus* model very fundamentally. The important question is whether, for the purpose of economics, these alternative models of man are ultimately better and more useful for applications. For that, the first testing ground should probably again be the economics laboratory.

A sizeable substrand of experimental economics (e.g. Fischbacher and Gächter, 2010, see Chaudhuri, 2011 for a review) has studied voluntary contributions games as introduced by Isaac, McCue, and Plott, 1985; Isaac and Walker, 1988. These games summarize succinctly the possible conflict between private interests and collective interests. Typically, the game is such that, if indeed individuals care only about their own material payoffs, then universal non-provision of the public good would result, i.e. the least efficient outcome. However, substantial levels of contributions are consistently observed across different experiments, and this has been taken as evidence that individuals are pro-social, and agents' utility functions are adapted to account for this (e.g. Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002). There are a number of problems with this approach in general, some of which are discussed in Binmore and Shaked, 2010a (and, published alongside that article, in Fehr and Schmidt, 2010; Eckel and Gintis, 2010; Binmore and Shaked, 2010b; see also West, El Mouden, and Gardner, 2011). One is that learning or 'erroneous' play, as an alternative explanation of the observed data, may not be easy to distinguish from social preferences (e.g. Andreoni, 1993; Andreoni, 1995a; Andreoni, 1995b; Palfrey and Prisbrey, 1996; Palfrey and Prisbrey, 1997; Goeree, Holt, and Laury, 2002; Ferraro and Vossler,

2010; Bayer, Renner, and Sausgruber, 2013; Burton-Chellew and West, 2013; Burton-Chellew, Nax, and West, 2015).

The issue of non-distinguishability of social motivations and ‘erroneous’ play is perhaps best illustrated by the existing experiments on non-linear variants of the standard voluntary contributions game that yield interior Nash equilibria (e.g. Walker, Gardner, and Ostrom, 1990; Smith and Walker, 1993; Keser, 1996; Van Dijk, Sonnemans, and Winden, 2002; Sefton and Steinberg, 1996; Isaac and Walker, 1998; Laury, Walker, and Williams, 1999; Chan et al., 2002; as reviewed in Laury and Holt, 2008; Saijo, 2014). These results indicate that the population sociality bias depends on the relative location of equilibrium vis-à-vis the social optimum in a particular setting.

However, whether decisions are taken intentionally, that is, corresponding to maximization of any utility at all, or whether decisions are subject to ‘error’ or learning, is at first sight orthogonal to another important issue which has not received the attention it deserves. Namely, there is a salient design issue resulting from the fact that public goods experiments have focused on games where free-riding is the strictly dominant strategy from the vantage point of pure material self-interest. Any deviation from Nash, in the sense of the ‘subjective expected utility correction project’ when sociality parameters are added to utility functions, is therefore evidence for pro-sociality (e.g. Kümmerli et al., 2010). The average man is therefore biased to be pro-social, because *homo oeconomicus* and anti-sociality are collapsed into behavior corresponding to the free-riding boundary of the strategy space. By studying simple variants of the public goods game where contributing fully is (at least, part of) equilibrium, it is possible to estimate not only ‘kind’ motives (such as pro-sociality, altruism or fairness), but also the flipside (such as anti-sociality, envy, or spite). In such variants, the conflict between ‘kind’ motives and ‘*homo oeconomicus*’

dissolves: contributing fully is both individually rational and beneficial to others. Hence, not contributing fully becomes indicative of ‘spiteful’ motives in the same way as positive contributions indicate pro-sociality in the standard settings. We have estimated social preferences from both classes of games and identified an almost-symmetrical distribution of anti-social, *homo oeconomicus* and pro-social preferences with a slight skew towards pro-sociality. Moreover, we reveal another category of players who are inconsistent with respect to their pro/anti-social motivations, and thus inexplicable/irrational in terms of preferences.

Apart from the studies discussed in the introduction (Saijo and Nakamura, 1995; Levine, 1998, etc.) and the ones whose data we use (Burton-Chellew and West, 2013; Nax et al., 2014), we are aware of the following experiments that are also related. The high-rate-of-return variation of the voluntary contributions game is an approach previously taken in Kümmerli et al., 2010. Group-based mechanisms similar to Gunnthorsdottir et al., 2010 (see also Gunnthorsdottir, Vragov, and Shen, 2010; Gunnthorsdottir and Thorsteinsson, 2014) are also proposed in Rabanal and Rabanal, 2014. The novelty of this present paper is the exploitation of the balanced within-subject design, featuring both ‘dilemma’ and ‘provision’ treatments to check for individual-level (in)consistencies.

In terms of results, first and foremost, our analysis validates the use of heterogeneous agent models instead of representative agent models. Agents seem to vary substantially in terms of their social concerns and/or rationality levels, and this may have consequences for the (in)stability of equilibrium, which is not a priori guaranteed under non-linear utilities (Saijo, 2014). Moreover, our findings also cast doubt on whether the typical (median and mean) agents in a population are really as pro-social as previous experiments suggested.

The image our work depicts of the population is rather one of an equilibrium-dependent distribution of pro/anti-sociality around the *homo oeconomicus* median, with the overall mean agent not lying far off either. To extrapolate this finding, we need to repeat similar analyses for other datasets and for other classes of games. How social preferences change with equilibrium properties of the game is an avenue left open for future research. Whilst we shall not attempt to extrapolate our findings to other games, we would like to point out that Bardsley, 2008, when considering generalizations of the *dictator game* that allow a balanced view of pro/anti-sociality in that context, concludes that altruism may well be an artefact of experimentation. In summary, there is evidence suggesting the need for future work in this direction for many more games, motivated by the fact that there are many contradictions that have been revealed even in those games about which we thought we knew a lot. A ‘ban’ (Camerer, 2003) on games like the ultimatum game or the public goods game, as has been recently proposed by Camerer because (supposedly) we know what is going on and why, appears premature.

Finally, we compare our findings with Ledyard, 1995’s “outrageous conjecture” (pp. 172–173) that there typically are 50% “dedicated Nash players”, 40% who respond to altruism or other social motivations (such as spite), and 10% “inexplicable”/“irrational” players. We find 47%, 37%, 16% respectively, with limited evidence for an altruistic tendency in the population average unless we were to re-label the spiteful as selfish and/or the inexplicable as altruists. In light of this evidence, it appears questionable whether social preference theory (alone) will successfully predict human behavior outside classes of games where we know how people will deviate from *homo oeconomicus* and behave instead based on many experiments (without needing any theory). Nevertheless, the image of man as *homo oeconomicus* remains an imperfect description of human behavior, especially given the heterogeneities we have observed. Yet, in light

of our findings, *homo oeconomicus* would probably still be the one to choose amongst all other types when forced to pick only one as the representative agent to make a prediction about a game drawn from a distribution over games that is ex ante unbalanced. Better, however, would be to use models with heterogeneous populations, consisting of learning types, *homo oeconomicus*, and various socially motivated agents.

## 1.4 Interactive preferences

The above findings lead to the pursuit, jointly with Kurt Ackermann and Ryan Murphy, of the question of whether preferences may be interactive, and not unresponsive as standard game theory presumes. The resulting analysis is reported in the below note.

Game theory presumes that agents have unique preference orderings over outcomes that prescribe unique preference orderings over actions in response to other players' actions, independent of other players' preferences. This independence assumption is necessary to permit game-theoretic best response reasoning, but at odds with introspection, because preferences towards one another often dynamically depend on each other. In this note, we propose a model of interactive preferences. The model is validated with data from a laboratory experiment. The main finding of our study is that pro-sociality diminishes over the course of the interactions.

### Introduction

Mother Teresa does not defect in prisoners' dilemmas, because she cares for her opponents in ways that transform the games' mixed motives into other games where her and common motives are aligned (e.g., harmony). Cooperation thus

emerges as a dominant strategy. The experimental economics literature is concerned with ‘subjective expected utility corrections’ (Gigerenzer and Selten, 2001) that modify players’ utility representations to account for such other-regarding concerns. Numerous corrections have been proposed (e.g., Rabin 1993; Levine 1998; Fehr and Schmidt 1999; Bolton and Ockenfels 2000 in light of laboratory evidence that manifests systematic deviations from narrow self-interest predictions (see Ledyard 1995 and Chaudhuri 2011 for reviews).<sup>10</sup> This route of enquiry is bothersome for many theoretical game theorists who question how these findings generalize beyond the laboratory.<sup>11</sup>

Missing from most alternative utility formulations are interactive components that meaningfully alter the game-theoretic analysis. Standard theory (Neumann and Morgenstern, 1944) equips players with preferences that prescribe actions vis-à-vis others’ actions, independent of others’ preferences. Here, inspired by Rabin, 1993 and Levine, 1998, we propose a model of interactive preferences among players that depend on each other and investigate their dynamic interdependence. The model is validated with laboratory studies involving repeated voluntary contributions games (VCM; Isaac, McCue, and Plott 1985) sandwiched by two sets of dictator games (DG; Kahneman, Knetsch, and Thaler 1986) used to evaluate individuals’ social value orientation (SVO; Murphy, Ackermann, and Handgraaf 2011). Our results show that, independent of unintended behavioral deviations, the proportion of behavior associated with pro-sociality diminishes over the course of the interactions and is replaced by individualism. These patterns carry over between VCMs and DGs. Our model fares predictively well.

---

<sup>10</sup>With some exceptions (e.g., Saijo and Nakamura 1995; Saijo 2008, many analytical set-ups have been biased as discussed in, for example, Burton-Chellew and West, 2013; Burton-Chellew, Nax, and West, 2015.

<sup>11</sup>See controversies in JEBO 73, 2010.

## Methods

**Experimental setup.** Experiments were run at ETH’s Decision Science Laboratory during February 2013 involving 128 subjects in 6 sessions (4\*20+2\*24). Subjects were informed in detail and in advance of each stage of the experiment using standard instructions.<sup>12</sup> Every decision was monetarily incentivized, and subjects earned over 40CHF>40US\$ on average. The experiment lasted roughly 90 minutes.

The experiment had the following three stages:

**Stage 1: *Initial SVO*.** Subjects played 6 DGs choosing allocations in different ranges representing different himself-versus-other tradeoffs; for example, between 100 for himself and 50 for the other, (100,50), and 50 for himself and 100 for the other, (50,100).<sup>13</sup> The 6 decisions are represented as angles in the classical SVO ring (Griesinger and Livingston, 1973b), and an individual’s initial SVO is taken as the average angle, representing a compact indicator of his *ex ante* SVO.<sup>14</sup>

**Stage 2: *VCM*.** Subjects played 10 VCMs in groups of 4 that were randomly formed in round 1 and then remained fixed for the remainder of rounds. In each round subjects made contributions and guessed others’ average contributions (with incentives for accuracy). Before each round, players were informed of the previous-period contributions. (More detail will be provided shortly.)

**Stage 3: *Final SVO*.** Stage 1 is repeated, thus measuring individuals’ *ex post* SVOs.

---

<sup>12</sup>See Murphy and Ackermann, 2013 for details.

<sup>13</sup>The remaining 5 choices are amongst linear combinations in the ranges [(100, 50), (85, 85)], [(50, 100), (85, 15)], [(50, 100), (85, 85)], [(85, 85), (85, 15)], and [(85, 15), (100, 50)].

<sup>14</sup>Angles close to 0° represent individualistic preferences in the the sense of material self-interest, angles  $\geq 22.5^\circ$  indicate pro-sociality.

Our analysis focusses on 22 data points p.p., namely his 2 – initial and final – SVOs, plus his 10 contributions and 10 guesses about others’ contributions from the VCM, yielding a total of 2,816 data points.

## The model

**Static model.** Population  $N = \{1, 2, 3, 4\}$  plays a VCM with marginal per capita rate of return  $r = 0.4$  and budget  $B = 20$ . Each  $i \in N$  sets a private contribution  $c_i \in B$  which, jointly with the others’ average contribution,  $c_{-i}$ , results in payoff

$$\phi_i = 20 - c_i + 0.4(c_i + 3c_{-i}).$$

We assume  $i$ ’s utility depends on payoffs in Cobb-Douglas form

$$u_i(c) = \phi_i^{1-\alpha_i} * \phi_{-i}^{\alpha_i}, \tag{1.1}$$

where  $\alpha_i \in [0, 1]$  measures player  $i$ ’s concern for others. The nonlinearity of expression 1.1 distinguishes it from most representations, including Levine, 1998, thus rationalizing intermediate contributions in terms of intermediate concerns. We obtain the following expression for  $\alpha_i$  by assuming  $c_i$  is chosen optimally given his guess about  $c_{-i}$  (expressed as  $\widehat{c}_{-i}$ ):

$$\alpha_i = \frac{0.6\phi_{-i}(c_i, \widehat{c}_{-i})}{0.4\phi_i(c_i, \widehat{c}_{-i}) + 0.6\phi_{-i}(c_i, \widehat{c}_{-i})} \tag{1.2}$$

Note that  $\partial\alpha_i/\partial c_i > 0$  and  $\partial\alpha_i/\partial\widehat{c}_{-i} < 0$ , that is, higher own contributions (holding beliefs about others constant) indicate more concern for others, and higher beliefs regarding others’ contributions (keeping own contributions fixed) indicate less concern for others.

The interdependence of preferences results from imposing that, in static equi-



librium,  $\alpha_i = \hat{\alpha}_{-i}$ , where  $\hat{\alpha}_{-i}$  is  $i$ 's belief about  $\alpha_{-i}$ .<sup>15</sup> The resulting set of equilibria, the general structure of which is under investigation in an ongoing study, contains the standard case (when  $\alpha_i = \alpha_{-i} = 0$ ) but also new ones when  $\alpha_i = \alpha_{-i} > 0$  as in fairness equilibria (Rabin, 1993).

**Dynamic components.** The above game repeats with revelation of past outcomes. Each period  $t$ , suppose  $i$  contributes to maximize expression 1.1 so that expression 1.2 implies  $\alpha_i^t$  given  $(c_i^t, \tilde{c}_{-i}^t)$ . We assume  $\alpha_i^t$  is updated in light of evidence by

$$\alpha_i^t = (1 - \beta_i^t)\alpha_i^{t-1} + \beta_i^t\tilde{\alpha}_{-i}^{t-1}, \quad (1.3)$$

where  $\tilde{\alpha}_{-i}^{t-1}$  is  $i$ 's deduction of  $\alpha_{-i}^{t-1}$  from previous-period evidence, and  $\beta_i^t \in [0, 1]$  measures  $i$ 's period- $t$  degree of belief responsiveness.

### Estimation strategy

**Classification.** Initial SVOs are used to classify individuals as ‘individualistic’ and ‘pro-social’. An individual is pro-social (individualistic) according to the SVO measure if his SVO-angle is  $\geq 22.5$  ( $< 22.5$ ) degree.<sup>16</sup> The initial SVO classifications are used to predict initial VCM contributions

‘Responsive’ and ‘unresponsive’ types are classified based on the VCM data. Individual  $i$  is said to be responsive (unresponsive) if the estimation of expression 1.3 in light of his VCM decisions from rounds 2-10 yields an average coefficient for  $\beta_i^t$  which is positive (not positive).

**Prediction.** We use our estimated  $2 \times 2$  typology (from initial SVO and VCM) to make predictions regarding final SVO classifications, which we shall assess in light of the recorded final SVOs. We shall use the following terminology: an individual is associated with a VCM group matching that is said to be

---

<sup>15</sup>A weaker assumption in the same spirit would be to weigh this dependence by some parameter as in Levine, 1998, something we shall introduce via ‘responsiveness’ instead.

<sup>16</sup>See Murphy and Ackermann 2013 for a more fine-grained categorization.

‘individualistic’ (‘pro-social’) if those players he is matched with, on average, contribute less (more) than himself.

We predict unresponsive types (pro-social and individualistic alike) not to change their preferences. We predict responsive types to change their types in the direction of their interaction partners as matched with during the VCM group matching. Hence, a responsive pro-social (individualist) in a VCM group matching that is pro-social (individualistic) will remain pro-social (individualistic). A responsive pro-social (individualist) matched with individualistic (pro-social) others, however, may become individualistic (pro-social), dependant on the action/payoff difference between himself and his opponents. In particular, whichever payoff difference is larger we shall assume will be associated with a preference-change flow of probability one, and the lesser payoff-difference to be proportional to that flow depending on the relative payoff difference.

Table 1.3: Regressions 1 and 2 (standard errors adjusted for 128 individual clusters)

Regression 1		Regression 2	
‘Contribution’ (VCM, $t=1$ )		‘Responsiveness’ (VCM, $t=1-10$ )	
Initial pro-sociality	3.54* (1.19)	$\alpha^{t-1}$	-0.35* (0.04)
Constant	10.76* (2.72)	$\tilde{\alpha}_{-i}^{t-1}$	0.44* (0.15)
Controls	not listed	Controls	not listed
$N$	128	$N$	1,152
$R^2$	0.13	$R^2$	0.20

\* : significance level < 0.01

## Results

Initial SVOs classify 53% of players as individualistic and 47% as pro-social, and pro-socials give over 30% more in period 1 of the VCM (regression 1). Expression 1.3 is structurally confirmed at the population level in the VCM data (regression 2). Re-running regressions for expression 1.3 at individual

levels (omitted) for the VCM, we find 71% responsiveness (34% pro-socials, 37% individualists); 14% (20%) are responsive pro-socials (individualists) matched by chance in individualistic (pro-social) groups.

It is amongst those 34% matched in opposite groups where we expect preference interactions to materialize. An average of 2.3 coins less was earned by the 14% responsive pro-socials in individualistic groups versus 0.6 more by responsive individualists in pro-social groups. Hence, flowing from (responsive) individualistic to pro-social, we expect ca.  $1/4$  ( $\approx 0.6/2.3$ ) of the flow from (responsive) pro-social to individualistic.

Predictions compare with the data as follows. Final SVOs categorize 64% individualists and 36% pro-socials (62% and 38% predicted). 47% are individualistic in initial and final SVOs, which means that 6% individualists turned pro-socials (5% predicted). 30% were pro-social in both, hence 17% pro-socials turned individualists (14% predicted). The model made two types of errors. First, 7% changed preferences whom we classified unresponsive. Second, we predicted 1% (3%) too few individualists turning pro-socials (vice versa), thus incorrectly predicting flow of 3% responsiveness. Overall, our model was therefore accurate in predicting global preferences (95%), less in individualizing flow (90%).

## Conclusion

Individuals become less (more) pro-social when interacting with individualists (pro-socials). On average, there is a trend toward individualism over the course of the VCM, independent of the contribution decay. Our result is therefore not a byproduct of learning. Even in the sterile and anonymous context of the laboratory we found evidence for interactive preferences among players that depend on each other and evolve over time. Our model explains indirect reci-

procuity (Alexander, 1987; Fischbacher and Gächter, 2010) as driven by natural dynamics governing the interactions of preferences. Since stakes and intentions of players certainly matter more outside the laboratory, such phenomena are likely not to be artifacts. Preference dynamics should therefore be studied further, as the long-run predictions of models without preference interactions are potentially misguided.

## References

- Ainslie, George (1975). “Specious reward: A behavioral theory of impulsiveness and impulse control.” In: *Psychological Bulletin* 82.4, pp. 463–496.
- Alexander, RD (1987). *The Biology of Moral Systems*. De Gruyter.
- Allais, Maurice (1953). “Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’école américaine”. In: *Econometrica: Journal of the Econometric Society* 21, pp. 503–546.
- Anderson, Simon P, Jacob K Goeree, and Charles A Holt (1998). “A theoretical analysis of altruism and decision error in public goods games”. In: *Journal of Public Economics* 70.2, pp. 297–323.
- Andreoni, James (1988). “Why free ride?: Strategies and learning in public goods experiments”. In: *Journal of Public Economics* 37.3, pp. 291–304.
- (1993). “An experimental test of the public goods crowding-out hypothesis”. In: *American Economic Review* 83, pp. 1317–1327.
- (1995a). “Cooperation in public-goods experiments: kindness or confusion?”. In: *American Economic Review* 85, pp. 891–904.
- (1995b). “Warm-glow versus cold-prickle: The effects of positive and negative framing on cooperation in experiments”. In: *Quarterly Journal of Economics* 110, pp. 1–21.

- Bardsley, Nicholas (2008). “Dictator game giving: altruism or artefact?” In: *Experimental Economics* 11.2, pp. 122–133.
- Bayer, Ralph-C, Elke Renner, and Rupert Sausgruber (2013). “Confusion and learning in the voluntary contributions game”. In: *Experimental Economics* 16.4, pp. 478–496.
- Becker, Gary S (1976). *The economic approach to human behavior*. University of Chicago press.
- Binmore, Ken and Avner Shaked (2010a). “Experimental Economics: Where Next?” In: *Journal of Economic Behavior & Organization* 73.1, pp. 87–100.
- (2010b). “Experimental Economics: Where Next? Rejoinder”. In: *Journal of Economic Behavior & Organization* 73.1, pp. 120–121.
- Bolton, GE and A Ockenfels (2000). “ERC: A theory of equity, reciprocity, and competition”. In: *AER* 90, pp. 166–193.
- Bowles, Samuel and Herbert Gintis (2011). *A cooperative species: Human reciprocity and its evolution*. Princeton University Press.
- Brandts, J, T Saijo, and A Schram (2004). “How universal is behavior? A four country comparison of spite and cooperation in voluntary contribution mechanisms”. In: *Public Choice* 119.3-4, pp. 381–424.
- Burton-Chellew, MN, HH Nax, and SA West (2015). “Payoff-based learning explains the decline in cooperation in public goods games”. In: *Proc. Roy. Soc. B* 282.1801, p. 20142678.
- Burton-Chellew, MN and SA West (2013). “Prosocial preferences do not explain human cooperation in public-goods games”. In: *PNAS* 110.1, pp. 216–221.
- Camerer, Colin F (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

- Cason, TN, T Saijo, and T Yamato (2002). “Voluntary participation and spite in public good provision experiments: An international comparison”. In: *EE* 5.2, pp. 133–153.
- Chan, Kenneth S et al. (2002). “Crowding-out voluntary contributions to public goods”. In: *Journal of Economic Behavior & Organization* 48.3, pp. 305–317.
- Charness, G and M Rabin (2002). “Understanding social preferences with simple tests”. In: *QJE* 117, pp. 817–869.
- Chaudhuri, A (2011). “Sustaining cooperation in laboratory public goods experiments: A selective survey of the literature”. In: *EE* 14.1, pp. 47–83.
- Cobb, Charles W and Paul H Douglas (1928). “A theory of production”. In: *The American Economic Review* 18, pp. 139–165.
- Eckel, Catherine and Herbert Gintis (2010). “Blaming the messenger: Notes on the current state of experimental economics”. In: *Journal of Economic Behavior & Organization* 73.1, pp. 109–119.
- Ellsberg, Daniel (1961). “Risk, ambiguity, and the Savage axioms”. In: *Quarterly Journal of Economics* 75, pp. 643–669.
- Falkinger, Josef et al. (2000). “A simple mechanism for the efficient provision of public goods: Experimental evidence”. In: *American Economic Review* 90, pp. 247–264.
- Fehr, E and KM Schmidt (1999). “A theory of fairness, competition, and cooperation”. In: *QJE* 114, pp. 817–868.
- (2010). “On inequity aversion: A reply to Binmore and Shaked”. In: *JEBO* 73.1, pp. 101–108.
- Fehr, Ernst and Colin F Camerer (2007). “Social neuroeconomics: the neural circuitry of social preferences”. In: *Trends in Cognitive Sciences* 11.10, pp. 419–427.

- Fehr, Ernst and Simon Gächter (2000). “Cooperation and punishment in public goods experiments”. In: *American Economic Review* 90.4, pp. 980–994.
- Fehr, Ernst and Simon Gächter (2002). “Altruistic punishment in humans”. In: *Nature* 415.6868, pp. 137–140.
- Ferraro, Paul J and Christian A Vossler (2010). “The source and significance of confusion in public goods experiments”. In: *The BE Journal of Economic Analysis & Policy* 10.1, p. 53.
- Fischbacher, U and S Gächter (2010). “Social preferences, beliefs, and the dynamics of free riding in public goods experiments”. In: *AER* 100.1, pp. 541–556.
- Gigerenzer, G and R Selten (2001). *Bounded Rationality: The Adaptive Toolbox*. MIT Press.
- Goeree, Jacob K and Charles A Holt (2005). “An explanation of anomalous behavior in models of political participation”. In: *American Political Science Review* 99.2, pp. 201–213.
- Goeree, Jacob K, Charles A Holt, and Susan K Laury (2002). “Private costs and public benefits: unraveling the effects of altruism and noisy behavior”. In: *Journal of Public Economics* 83.2, pp. 255–276.
- Griesinger, Donald W and James W Livingston (1973a). “Toward a model of interpersonal motivation in experimental games”. In: *Behavioral Science* 18.3, pp. 173–188.
- Griesinger, DW and JW Livingston (1973b). “Toward a model of interpersonal motivation in experimental games”. In: *Behavioral Science* 18.3, pp. 173–188.
- Gunthorsdottir, Anna, Daniel Houser, and Kevin McCabe (2007). “Disposition, history and contributions in public goods experiments”. In: *Journal of Economic Behavior & Organization* 62.2, pp. 304–315.

- Gunthorsdottir, Anna and Palmar Thorsteinsson (2014). “Tacit coordination and equilibrium selection in a merit-based grouping mechanism: A cross-cultural validation study”. In: *mimeo*.
- Gunthorsdottir, Anna, Roumen Vragov, and Jianfei Shen (2010). “Tacit coordination in contribution-based grouping with two endowment levels”. In: *Research in Experimental Economics* 13, pp. 13–75.
- Gunthorsdottir, Anna et al. (2010). “Near-efficient equilibria in contribution-based competitive grouping”. In: *Journal of Public Economics* 94.11, pp. 987–994.
- Houser, Daniel and Robert Kurzban (2002). “Revisiting kindness and confusion in public goods experiments”. In: *American Economic Review* 92, pp. 1062–1069.
- Isaac, R Mark, James Walker, and S Thomas (1991). “On the suboptimality of voluntary public goods provision: Further experimental evidence”. In: *Research in Experimental Economics* 4, pp. 211–221.
- Isaac, R Mark and James M Walker (1998). “Nash as an organizing principle in the voluntary provision of public goods: Experimental evidence”. In: *Experimental Economics* 1.3, pp. 191–206.
- Isaac, RM, KF McCue, and CR Plott (1985). “Public goods provision in an experimental environment”. In: *J. Pub. Econ.* 26.1, pp. 51–74.
- Isaac, RM and JM Walker (1988). “Group size effects in public goods provision: The voluntary contributions mechanism”. In: *Quarterly Journal of Economics* 103, pp. 179–199.
- Kahneman, A, JL Knetsch, and RH Thaler (1986). “Fairness and the assumptions of economics”. In: *Journal of Business* 59.4, pp. 285–300.
- Kahneman, D and A Tversky (1979). “Prospect theory: An analysis of decision under risk”. In: *Econometrica* 47, pp. 263–291.



- Keser, Claudia (1996). “Voluntary contributions to a public good when partial contribution is a dominant strategy”. In: *Economics Letters* 50.3, pp. 359–366.
- Kreps, David M et al. (1982). “Rational cooperation in the finitely repeated prisoners’ dilemma”. In: *Journal of Economic Theory* 27.2, pp. 245–252.
- Kümmerli, R et al. (2010). “Resistance to extreme strategies, rather than prosocial preferences, can explain human cooperation in public goods games”. In: *PNAS* 107.22, pp. 10125–10130.
- Laury, Susan K and Charles A Holt (2008). “Voluntary provision of public goods: experimental results with interior Nash equilibria”. In: *Handbook of Experimental Economics Results* 1, pp. 792–801.
- Laury, Susan K, James M Walker, and Arlington W Williams (1999). “The voluntary provision of a pure public good with diminishing marginal returns”. In: *Public Choice* 99.1-2, pp. 139–160.
- Ledyard, JO (1995). “Public goods: A survey of experimental research”. In: *The Handbook of Experimental Economics*. Ed. by JH Kagel and AE Roth. Princeton University Press, pp. 111–194.
- Levine, DK (1998). “Modeling altruism and spitefulness in experiments”. In: *Review of Economic Dynamics* 1.3, pp. 593–622.
- Murphy, RO and KA Ackermann (2013). “Explaining Behavior in Public Goods Games: How Preferences and Beliefs Affect Contribution Levels”. In: *SSRN* 2244895.
- (2014). “Explaining behavior in public goods games: How preferences and beliefs affect contribution levels”. In: *mimeo*.
- Murphy, RO, KA Ackermann, and MJJ Handgraaf (2011). “Measuring social value orientation”. In: *Journal of Judgment and Decision Making* 6, pp. 771–781.

- Nax, Heinrich H, Ryan O Murphy, and Dirk Helbing (2014). “Stability and welfare of meritocratic matching in voluntary contribution games”. In: *mimeo*.
- Nax, Heinrich H and Matjaž Perc (2015). “Directional learning and the provisioning of public goods”. In: *Scientific reports* 5, p. 8010.
- Nax, Heinrich H et al. (2013). “Learning in a black box”. In: *Department of Economics WP, University of Oxford* 653.
- Nax, Heinrich H et al. (2014). “How assortative re-matching regimes affect cooperation levels in public goods interactions”. In: *mimeo*.
- Neumann, J von and O Morgenstern (1944). *Theory of Games and Economic Behaviour*. Princeton University Press.
- Palfrey, Thomas R and Jeffrey E Prisbrey (1997). “Anomalous behavior in public goods experiments: How much and why?” In: *American Economic Review* 87, pp. 829–846.
- Palfrey, TR and JE Prisbrey (1996). “Altruism, reputation and noise in linear public goods experiments”. In: *Journal of Public Economics* 61.3, pp. 409–427.
- Rabanal, Jean Paul and Olga A Rabanal (2014). “Efficient investment via assortative matching: A laboratory experiment”. In: *mimeo*.
- Rabin, M (1993). “Incorporating fairness into game theory and economics”. In: *AER* 83.5, pp. 1281–1302.
- Ramsey, Frank Plumpton (1931). “Foundations: Essays in Philosophy, Logic, Mathematics and Economics”. In: *Humanities Press*.
- Saijo, T (2008). “Spiteful behavior in voluntary contribution mechanism experiments”. In: *Handbook of Experimental Economics Results*. Ed. by CR Plott and VL Smith. Vol. 1, pp. 802–816.
- Saijo, T and H Nakamura (1995). “The “spite” dilemma in voluntary contribution mechanism experiments”. In: *Journal of Conflict Resolution* 39.3, pp. 535–560.

- Saijo, T and T Yamato (1999). “A voluntary participation game with a non-excludable public good”. In: *JET* 84.2, pp. 227–242.
- Saijo, Tatsuyoshi (2014). “The instability of the voluntary contribution mechanism”. In: *mimeo*.
- Savage, Leonard J (1954). “The Foundations of Statistics”. In: *Wiley*.
- Sefton, Martin and Richard Steinberg (1996). “Reward structures in public good experiments”. In: *Journal of Public Economics* 61.2, pp. 263–287.
- Smith, Vernon L and James M Walker (1993). “Monetary rewards and decision cost in experimental economics”. In: *Economic Inquiry* 31.2, pp. 245–261.
- Van Dijk, Frans, Joep Sonnemans, and Frans van Winden (2002). “Social ties in a public good experiment”. In: *Journal of Public Economics* 85.2, pp. 275–299.
- Von Neumann, John and Oskar Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Walker, James M, Roy Gardner, and Elinor Ostrom (1990). “Rent dissipation in a limited-access common-pool resource: Experimental evidence”. In: *Journal of Environmental Economics and Management* 19.3, pp. 203–211.
- West, Stuart A, Claire El Mouden, and Andy Gardner (2011). “Sixteen common misconceptions about the evolution of cooperation in humans”. In: *Evolution and Human Behavior* 32.4, pp. 231–262.
- Wicherts, Jelte M et al. (2006). “The poor availability of psychological research data for reanalysis.” In: *American Psychologist* 61.7, p. 726.

## Chapter 2

**Learning:**

**Directional learning and the  
provisioning of public goods**

## Abstract

We consider an environment where players are involved in a public goods game and must decide repeatedly whether to make an individual contribution or not. However, players lack strategically relevant information about the game and about the other players in the population. The resulting behavior of players is completely uncoupled from such information, and the individual strategy adjustment dynamics are driven only by reinforcement feedbacks from each player's own past. We show that the resulting "directional learning" is sufficient to explain cooperative deviations away from the Nash equilibrium. We introduce the concept of  $k$ -strong equilibria, which nest both the Nash equilibrium and the Aumann-strong equilibrium as two special cases, and we show that, together with the parameters of the learning model, the maximal  $k$ -strength of equilibrium determines the stationary distribution. The provisioning of public goods can be secured even under adverse conditions, as long as players are sufficiently responsive to the changes in their own payoffs and adjust their actions accordingly. Substantial contribution levels can thus be explained without arguments involving selflessness or social preferences, solely on the basis of uncoordinated directional (mis)learning.

**Acknowledgements.** This research was supported by the European Commission through the ERC Advanced Investigator Grant ‘Momentum’ (Grant 324247), by the Slovenian Research Agency (Grant P5-0027), and by the Deanship of Scientific Research, King Abdulaziz University (Grant 76-130-35-HiCi).

## 2.1 Introduction

Cooperation in sizable groups has been identified as one of the pillars of our remarkable evolutionary success. While between-group conflicts and the necessity for alloparental care are often cited as the likely sources of the other-regarding abilities of the genus *Homo* (Bowles and Gintis, 2011; Hrdy, 2011), it is still debated what made us the “supercooperators” that we are today (Nowak and Highfield, 2011; Rand and Nowak, 2013). Research in the realm of evolutionary game theory (Maynard Smith, 1982; Weibull, 1995; Hofbauer and Sigmund, 1998; Mesterton-Gibbons, 2001; Nowak, 2006a; Myatt and Wallace, 2008) has identified a number of different mechanisms by means of which cooperation might be promoted (Mesterton-Gibbons and Dugatkin, 1992; Nowak, 2006b), ranging from different types of reciprocity and group selection to positive interactions (Rand et al., 2009), risk of collective failure (Santos and Pacheco, 2011), and static network structure (Santos, Santos, and Pacheco, 2008; Rand et al., 2013).

The public goods game (Isaac, McCue, and Plott, 1985), in particular, is established as an archetypical context that succinctly captures the social dilemma that may result from a conflict between group interest and individual interests (Ledyard, 1997; Chaudhuri, 2011). In its simplest form, the game requires that players decide whether to contribute to a common pool or not. Regardless of the chosen strategy by the player himself, he receives an equal share of the public good which results from total contributions being multiplied by a fixed rate of return. For typical rates of return it is the case that, while the individual temptation is to free-ride on the contributions of the other players, it is in the interest of the collective for everyone to contribute. Without additional mechanisms such as punishment (Fehr and Gächter, 2000), contribution decisions in such situations (Ledyard, 1997; Chaudhuri, 2011) approach the free-riding

Nash equilibrium (Nash, 1950) over time and thus lead to a “tragedy of the commons” (Hardin, 1968). Nevertheless, there is rich experimental evidence that the contributions are sensitive to the rate of return (Fischbacher, Gächter, and Fehr, 2001) and positive interactions (Rand et al., 2009), and there is evidence in favor of the fact that social preferences and beliefs about other players’ decisions are at the heart of individual decisions in public goods environments (Fischbacher and Gächter, 2010).

In this paper, however, we shall consider an environment where players have no strategically relevant information about the game and/ or about other players, and hence explanations in terms of social preferences and beliefs are not germane. Instead, we shall propose a simple learning model, where players may mutually reinforce learning off the equilibrium path. As we will show, this phenomenon provides an alternative and simple explanation for why contributions rise with the rate of return, as well as why, even under adverse conditions, public cooperation may still prevail. Previous explanations of this experimental regularity (Ledyard, 1997) are based on individual-level costs of ‘error’ (Palfrey and Prisbey, 1997; Goeree and Holt, 2005).

Suppose each player knows neither who the other players are, nor what they earn, nor how many there are, nor what they do, nor what they did, nor what the rate of return of the underlying public goods game is. Players do not even know whether the underlying rate of return stays constant over time (even though in reality it does), because their own payoffs are changing due to the strategy adjustments of other players, about which they have no information. Without any such knowledge, players are unable to determine ex ante whether contributing or not contributing is the better strategy in any given period, i.e., players have no strategically relevant information about how to respond best. As a result, the behavior of players has to be *completely uncoupled* (Foster



and Young, 2006; Young, 2009), and their strategy adjustment dynamics are likely to follow a form of *reinforcement* (Roth and Erev, 1995; Erev and Roth, 1998) feedback or, as we shall call it, *directional learning* (Selten and Stoecker, 1986; Selten and Buchta, 1994). We note that, in our model, due to the one-dimensionality of the strategy space, reinforcement and directional learning are both adequate terminologies for our learning model. Since reinforcement applies also to general strategy spaces and is therefore more general we will prefer the terminology of directional learning. Indeed, such directional learning behavior has been observed in recent public goods experiments (Bayer, Renner, and Sausgruber, 2013; Young et al., 2013). The important question is how *well* will the population learn to play the public goods game despite the lack of strategically relevant information. Note that *well* here has two meanings due to the conflict between private and collective interests: on the one hand, how close will the population get to playing the Nash equilibrium, and, on the other hand, how close will the population get to playing the socially desirable outcome.

The learning model considered in this paper is based on a particularly simple “directional learning” algorithm which we shall now explain. Suppose each player plays both cooperation (contributing to the common pool) and defection (not contributing) with a mixed strategy and updates the weights for the two strategies based on their relative performances in previous rounds of the game. In particular, a player will increase its weight on contributing if a previous-round switch from not contributing to contributing led to a higher realized payoff or if a previous-round switch from contributing to not contributing led to a lower realized payoff. Similarly, a player will decrease its weight on contributing if a previous-round switch from contributing to not contributing led to a higher realized payoff or if a previous-round switch from not contributing to contributing led to a lower realized payoff. For simplicity,

we assume that players make these adjustments at a fixed incremental step size  $\delta$ , even though this could easily be generalized. In essence, each player adjusts its mixed strategy directionally depending on a Markovian performance assessment of whether a previous-round contribution increase/decrease led to a higher/lower payoff.

Since the mixed strategy weights represent a well-ordered strategy set, the resulting model is related to the directional learning/ aspiration adjustment models (Sauermann and Selten, 1962; Selten and Stoecker, 1986; Selten and Buchta, 1994), and similar models have previously been proposed for bid adjustments in assignment games (Nax, Pradelski, and Young, 2013), as well as in two-player games (Laslier and Walliser, 2014). In Nax, Pradelski, and Young, 2013 the dynamic leads to stable cooperative outcomes that maximize total payoffs, while Nash equilibria are reached in Laslier and Walliser, 2014. The crucial difference between these previous studies and our present study is that our model involves more than two players in a voluntary contributions setting, and, as a result, that there can be interdependent directional adjustments of groups of players including more than one but not all the players. This can lead to uncoordinated (mis)learning of subpopulations in the game.

Consider the following example. Suppose all players in a large standard public goods game do not contribute to start with. Then suppose that a group of players in a subpopulation uncoordinatedly but by chance simultaneously all decide to contribute. If this group is sufficiently large (the size of which depends on the rate of return), then this will result in higher payoffs for all players including those in the ‘contributors group’, despite the fact that not contributing is the dominant strategy in terms of unilateral replies. In our model, if indeed this generates higher payoffs for all players including the freshly-turned contributors, then the freshly-turned contributors would continue to increase

their probability to contribute and thus increase the probability to trigger a form of stampede or herding effect, which may thus lead away from the Nash equilibrium and towards a socially more beneficial outcome.

Our model of uncoordinated but mutually reinforcing deviations away from Nash provides an alternative explanation for the following regularity that has been noted in experiments on public goods provision (Ledyard, 1997). Namely, aggregate contribution levels are higher the higher the rate of return, despite the fact that the Nash equilibrium remains unchanged (at no-contribution). This regularity has previously been explained only at an individual level, namely that ‘errors’ are less costly – and therefore more likely – the higher the rate of return, following quantal-response equilibrium arguments (Palfrey and Prisbey, 1997; Goeree and Holt, 2005). By contrast, we provide a group-dynamic argument. Note that the alternative explanation in terms of individual costs is not germane in our setting, because we have assumed that players have no information to make such assessments. It is in this sense that our explanation perfectly complements the explanation in terms of costs.

In what follows, we present the results, where we first set up the model and then deliver our main conclusions. We discuss the implications of our results in section 3. Further details about the applied methodology are provided in the Methods section.

## 2.2 Results

### Public goods game with directional learning

In the public goods game, each player  $i$  in the population  $N = 1, 2, \dots, n$  chooses whether to contribute ( $c_i = 1$ ) or not to contribute ( $c_i = 0$ ) to the

common pool. Given a fixed *rate of return*  $r > 0$ , the resulting payoff of player  $i$  is then  $u_i = (1 - c_i) + (r/n) * \sum_{j \in N} c_j$ . We shall call  $r/n$  the game's *marginal per-capita rate of return* and denote it as  $R$ . Note that for simplicity, but without loss of generality, we have assumed that the group is the whole population. In the absence of restrictions on the interaction range of players (Perc et al., 2013), i.e., in well-mixed populations, the size of the groups and their formation can be shown to be of no relevance in our case, as long as  $R$  rather than  $r$  is considered as the effective rate of return.

The directional learning dynamics is implemented as follows. Suppose the above game is infinitely repeated at time steps  $t = 0, 1, 2, \dots$ , and suppose further that  $i$ , at time  $t$ , plays  $c_i^t = 1$  with probability  $p_i^t \in [\delta, 1 - \delta]$  and  $c_i^t = 0$  with probability  $(1 - p_i^t)$ . Let the vector of contribution probabilities  $p^t$  describe the state of the game at time  $t$ . We initiate the game with all  $p_i^0$  lying on the  $\delta$ -grid between 0 and 1, while subsequently individual mixed strategies evolve randomly subject to the following three “directional bias” rules:

**upward:** if  $u_i(c_i^t) > u_i(c_i^{t-1})$  and  $c_i^t > c_i^{t-1}$ , or if  $u_i(c_i^t) < u_i(c_i^{t-1})$  and  $c_i^t < c_i^{t-1}$ , then  $p_i^{t+1} = p_i^t + \delta$  if  $p_i^t < 1$ ; otherwise,  $p_i^{t+1} = p_i^t$ .

**neutral:** if  $u_i(c_i^t) = u_i(c_i^{t-1})$  and/or  $c_i^t = c_i^{t-1}$ , then  $p_i^{t+1} = p_i^t$ ,  $p_i^t + \delta$ , or  $p_i^t - \delta$  with equal probability if  $0 < p_i^t < 1$ ; otherwise,  $p_i^{t+1} = p_i^t$ .

**downward:** if  $u_i(c_i^t) > u_i(c_i^{t-1})$  and  $c_i^t < c_i^{t-1}$ , or if  $u_i(c_i^t) < u_i(c_i^{t-1})$  and  $c_i^t > c_i^{t-1}$ , then  $p_i^{t+1} = p_i^t - \delta$  if  $p_i^t > 0$ ; otherwise,  $p_i^{t+1} = p_i^t$ .

Note that the second, neutral rule above allows random deviations from any intermediate probability  $0 < p_i < 1$ . However,  $p_i = 0$  and  $p_i = 1$  for all  $i$  are absorbing state candidates. We therefore introduce perturbations to this directional learning dynamics and study the resulting stationary states. In particular, we consider perturbations of order  $\epsilon$  such that, with probability  $1 - \epsilon$ , the dynamics is governed by the original three “directional bias” rules.

However, with probability  $\epsilon$ , either  $p_i^{t+1} = p_i^t$ ,  $p_i^{t+1} = p_i^t - \delta$  or  $p_i^{t+1} = p_i^t + \delta$  happens equally likely (with probability  $\epsilon/3$ ) but of course obeying the  $p_i^{t+1} \in [0, 1]$  restriction.

## Provisioning of public goods

We begin with a formal definition of the  $k$ -strong equilibrium. In particular, a pure strategy imputation  $s^*$  is a  $k$ -strong equilibrium of our (symmetric) public goods game if, for all  $C \subseteq N$  with  $|C| \leq k$ ,  $u_i(s_C^*; s_{N \setminus C}^*) \geq u_i(s'_C; s_{N \setminus C}^*)$  for all  $i \in C$  for any alternative pure strategy set  $s'_C$  for  $C$ . As noted in the previous section, this definition bridges, on the one hand, the concept of the Nash equilibrium in pure strategies (Nash, 1950) in the sense that any  $k$ -strong equilibrium with  $k > 0$  is also a Nash equilibrium, and, on the other hand, that of the (Aumann-)strong equilibrium (Aumann, 1974; Aumann, 1987) in the sense that any  $k$ -strong equilibrium with  $k = n$  is Aumann strong. Equilibria in between (for  $1 < k < n$ ) are “more stable” than a Nash equilibrium, but “less stable” than an Aumann-strong equilibrium.

The maximal  $k$ -strengths of the equilibria that still exist in our public goods game as a function of  $r$  are depicted in Fig. 2.1 for  $n = 16$ . The cyan-shaded region indicates the “public bad game” region for  $r < 1$  ( $R < 1/n$ ), where the individual and the public motives in terms of the Nash equilibrium of the game are aligned towards defection. Here  $c_i = 0$  for all  $i$  is the unique Aumann-strong equilibrium, or in terms of the definition of the  $k$ -strong equilibrium,  $c_i = 0$  for all  $i$  is  $k$ -strong for all  $k \in [1, n]$ . The magenta-shaded region indicates the typical public goods game for  $1 < r < n$  ( $1/n < R < 1$ ), where individual and public motives are conflicting. Here there exists no Aumann-strong equilibria. The outcome  $c_i = 0$  for all  $i$  is the unique Nash equilibrium, and that outcome is also  $k$ -strong equilibrium for some  $k \in [1, n)$ , where the

size of  $k$  depends on  $r$  and  $n$  in that  $\partial k/\partial r \leq 0$  while  $\partial k/\partial n \geq 0$ . Finally, the gray-shaded region indicates the unconflicted public goods game for  $r > n$  ( $R > 1$ ), where individual and public motives are again aligned, but this time towards cooperation. Here  $c_i = 1$  for all  $i$  abruptly becomes the unique Nash and Aumann-strong equilibrium, or equivalently the unique  $k$ -strong equilibrium for all  $k \in [1, n]$ .

If we add perturbations of order  $\epsilon$  to the unperturbed public goods game with directional learning that we have introduced in section 2, there exist stationary distributions of  $p_i$  and the following proposition can be proven. In the following, we denote by “ $k$ ” the maximal  $k$ -strength of an equilibrium.

**Proposition:** As  $t \rightarrow \infty$ , starting at any  $p^0$ , the expectation with respect to the stationary distribution is  $E[p^t] > 1/2$  if  $R \geq 1$  and  $E[p^t] < 1/2$  if  $R < 1$ .  $\partial E[p^t]/\partial \epsilon < 0$  if  $R \geq 1$ , and  $\partial E[p^t]/\partial \epsilon > 0$  if  $R < 1$ . Moreover,  $\partial E[p^t]/\partial \delta > 0$ , and  $\partial E[p^t]/\partial \delta < 0$  if  $R \geq 1$ . Finally,  $\partial E[p^t]/\partial k < 0$  if  $R < 1$ .

We begin the proof by noting that the perturbed process given by our dynamics results in an irreducible and aperiodic Markov chain, which has a unique stationary distribution. When  $\epsilon = 0$ , any absorbing state must have  $p_i^t = 0$  or 1 for all players. This is clear from the positive probability paths to either extreme from intermediate states given by the unperturbed dynamics. We shall now analyze whether  $p_i^t = 0$  or 1, given that  $p_j^t = 0$  or 1 for all  $j \neq i$ , has a larger attraction given the model’s underlying parameters.

If  $R \geq 1$ , the probability path for any player to move from  $p_i^t = 0$  to  $p_i^{t+T} = 1$  in some  $T = 1/\delta$  steps requires a single perturbation for that player and is therefore of the order of a single  $\epsilon$ . By contrast, the probability for any player to move from  $p_i^t = 1$  to  $p_i^{t+T} = 0$  in  $T$  steps is of the order  $\epsilon^3$ , because at least two other players must increase their contribution in order for that player to

experience a payoff increase from his non-contribution. Along any other path or if  $p^t$  is such that there are not two players  $j$  with  $p_j^t = 0$  to make this move, then the probability for  $i$  to move from  $p_i^t = 1$  to  $p_i^{t+T} = 0$  in  $T$  steps requires even more perturbations and is of higher order. Notice that, for any one player to move from  $p_i^t = 0$  to  $p_i^{t+T} = 1$  we need at least two players to move away from  $p_i^t = 0$  along the least-resistance paths. Because contributing 1 is a best reply for all  $R \geq 1$ , those two players will also continue to increase if continuing to contribute 1. Notice that the length of the path is  $T = 1/\delta$  steps, and that the path requires no perturbations along the way, which is less likely the smaller  $\delta$ .

If  $R < 1$ , the probability for any player to move from  $p_i^t = 1$  to  $p_i^{t+T} = 0$  in some  $T = 1/\delta$  steps requires a single perturbation for that player and is therefore of the order of a single  $\epsilon$ . By contrast, the probability for any player to move from  $p_i^t = 0$  to  $p_i^{t+T} = 1$  in some  $T$  steps is at least of the order  $\epsilon^k$ , because at least  $k$  players (corresponding to the maximal  $k$ -strength of the equilibrium) must contribute in order for all of these players to experience a payoff increase. Notice that  $k$  decreases in  $R$ . Again, the length of the path is  $T = 1/\delta$  steps, and that path requires no perturbations along the way, which is less likely the smaller  $\delta$ . With this, we conclude the proof of the proposition. However, it is also worth noting a direct corollary of the proposition; namely, as  $\epsilon \rightarrow 0$ ,  $E[p^t] \rightarrow 1$  if  $R \geq 1$ , and  $E[p^t] \rightarrow 0$  if  $R < 1$ .

Lastly, we simulate the perturbed public goods game with directional learning and determine the actual average contribution levels in the stationary state. Color encoded results in dependence on the normalized rate of return  $R$  and the responsiveness of players to the success of their past actions  $\delta$  (alternatively, the sensitivity of the individual learning process) are presented in Fig. 2.2 for  $\epsilon = 0.1$ . Small values of  $\delta$  lead to a close convergence to the respective

Nash equilibrium of the game, regardless of the value of  $R$ . As the value of  $\delta$  increases, the pure Nash equilibria erode and give way to a mixed outcome. It is important to emphasize that this is in agreement, or rather, this is in fact a consequence of the low  $k$ -strengths of the non-contribution pure equilibria (see Fig 2.1). Within intermediate to large  $\delta$  values the Nash equilibria are implemented in a zonal rather than pinpoint way. When the Nash equilibrium is such that all players contribute ( $R > 1$ ), then small values of  $\delta$  lead to more efficient aggregate play (recall any such equilibrium is  $n$ -strong). Conversely, by the same logic, when the Nash equilibrium is characterized by universal free-riding, then larger values of  $\delta$  lead to more efficient aggregate play. Moreover, the precision of implementation also depends on the rate of return in the sense that uncoordinated deviations of groups of players lead to more efficient outcomes the higher the rate of return. In other words, the free-riding problem is mitigated if group deviations lead to higher payoffs for every member of an uncoordinated deviation group, the minimum size of which (that in turn is related to the maximal  $k$ -strength of equilibrium) is decreasing with the rate of return.

Simulations also confirm that the evolutionary outcome is qualitatively invariant to: i) The value of  $\epsilon$  as long as the latter is bounded away from zero, although longer convergence times are an inevitable consequence of very small  $\epsilon$  values (see Fig. 2.3); ii) The replication of the population (i.e., making the whole population a group) and the random remixing between groups; and iii) The population size, although here again the convergence times are the shorter the smaller the population size. While both ii and iii are a direct consequence of the fact that we have considered the public goods game in a well-mixed rather than a structured population (where players would have a limited interaction range and where thus pattern formation could play a decisive role; Perc et al., 2013), the qualitative invariance to the value of  $\epsilon$  is elucidated further



in Fig. 2.3. We would like to note that by “qualitative invariance” it is meant that, regardless of the value of  $\epsilon > 0$ , the population always diverges away from the Nash equilibrium towards a stable mixed stationary state. But as can be observed in Fig. 2.3, the average contribution level and its variance both increase slightly as  $\epsilon$  increases. This is reasonable if one perceives  $\epsilon$  as an exploration/mutation rate. More precisely, it can be observed that, the lower the value of  $\epsilon$ , the longer it takes for the population to move away from the Nash equilibrium where everybody contributes zero in the case that  $1/n < R < 1$  (which was also the initial condition for clarity). However, as soon as initial deviations (from  $p_i = 0$  in this case) emerge (with probability proportional to  $\epsilon$ ), the neutral rule in the original learning dynamics takes over, and this drives the population towards a stable mixed stationary state. Importantly, even if the value of  $\epsilon$  is extremely small, the random drift sooner or later gains momentum and eventually yields similar contribution levels as those attainable with larger values of  $\epsilon$ . Most importantly, note that there is a discontinuous jump towards staying in the Nash equilibrium, which occurs only if  $\epsilon$  is exactly zero. If  $\epsilon$  is bounded away from zero, then the free-riding Nash equilibrium erodes unless it is  $n$ -strong (for very low values of  $R \leq 1/n$ ).

## 2.3 Discussion

We have introduced a public goods game with directional learning, and we have studied how the level of contributions to the common pool depends on the rate of return and the responsiveness of individuals to the successes and failures of their own past actions. We have shown that directional learning alone suffices to explain deviations from the Nash equilibrium in the stationary state of the public goods game. Even though players have no strategically relevant information about the game and/ or about each others’ actions, the population

could still end up in a mixed stationary state where some players contributed at least part of the time although the Nash equilibrium would be full free-riding. Vice versa, defectors emerged where cooperation was clearly the best strategy to play. We have explained these evolutionary outcomes by introducing the concept of  $k$ -strong equilibria, which bridge the gap between the Nash equilibrium and Aumann-strong equilibrium concepts. We have demonstrated that the lower the maximal  $k$ -strength and the higher the responsiveness of individuals to the consequences of their own past strategy choices, the more likely it is for the population to (mis)learn what is the objectively optimal unilateral (Nash) play.

These results have some rather exciting implications. Foremost, the fact that the provisioning of public goods even under adverse conditions can be explained without any sophisticated and often lengthy arguments involving selflessness or social preference holds promise of significant simplifications of the rationale behind seemingly irrational individual behavior in sizable groups. It is simply enough for a critical number (depending on the size of the group and the rate of return) of individuals to make a “wrong choice” at the same time once, and if only the learning process is sufficiently fast or naive, the whole subpopulation is likely to adopt this wrong choice as their own at least part of the time. In many real-world situations, where the rationality of decision making is often compromised due to stress, propaganda or peer pressure, such “wrong choices” are likely to proliferate. As we have shown in the context of public goods games, sometimes this means more prosocial behavior, but it can also mean more free-riding, depending only on the rate of return.

The power of directional (mis)learning to stabilize unilaterally suboptimal game play of course takes nothing away from the more traditional and established explanations, but it does bring to the table an interesting option

that might be appealing in many real-life situations, also those that extend beyond the provisioning of public goods. Fashion trends or viral tweets and videos might all share a component of directional learning before acquiring mainstream success and recognition. We hope that our study will be inspirational for further research in this direction. The consideration of directional learning in structured populations (Szabó and Fáth, [2007](#); Perc and Szolnoki, [2010](#)), for example, appears to be a particularly exciting future venture.

## 2.4 Methods

For the characterization of the stationary states, we introduce the concept of  $k$ -strong equilibria, which nests both the Nash equilibrium (Nash, 1950) and the Aumann-strong equilibrium (Aumann, 1974; Aumann, 1987) as two special cases. While the Nash equilibrium describes the robustness of an outcome against unilateral (1-person) deviations, the Aumann-strong equilibrium describes the robustness of an outcome against the deviations of any subgroup of the population. An equilibrium is said to be (Aumann-)strong if it is robust against deviations of the whole population or indeed of any conceivable subgroup of the population, which is indeed rare. Our definition of the  $k$ -strong equilibrium bridges the two extreme cases, measuring the size of the group  $k \geq 1$  (at or above Nash) and hence the degree to which an equilibrium is stable. We note that our concept is related to coalition-proof equilibrium (Bernheim, Peleg, and Whinston, 1987; Moreno and Wooders, 1996). In the public goods game, the free-riding Nash equilibrium is typically also more than 1-strong but never  $n$ -strong. As we will show, the maximal strength  $k$  of an equilibrium translates directly to the level of contributions in the stationary distribution of our process, which is additionally determined by the normalized rate of return  $R$  and the responsiveness of players to the success of their past actions  $\delta$ , i.e., the sensitivity of the individual learning process.

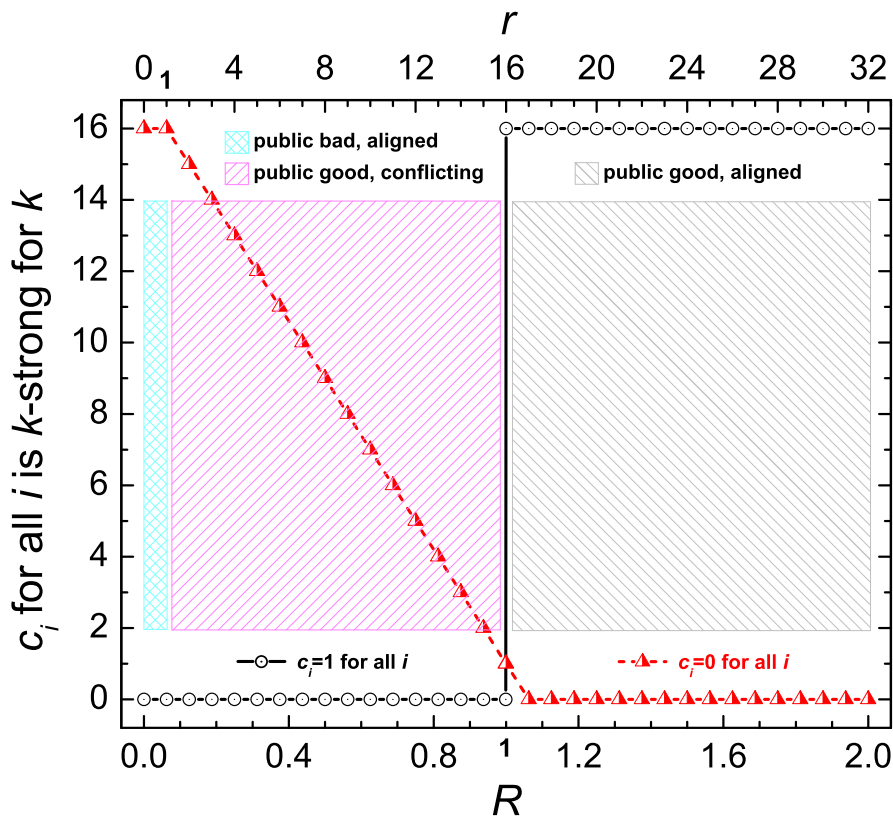


Figure 2.1: The maximal  $k$ -strength of equilibria in the studied public goods game with directional learning. As an example, we consider the population size being  $n = 16$ . As the rate of return  $r$  increases above 1, the Aumann-strong ( $n$ -strong)  $c_i = 0$  for all  $i$  (full defection) equilibrium loses strength. It is still the unique Nash equilibrium, but its maximal strength is bounded by  $k = 17 - r$ . As the rate of return  $r$  increases further above  $n$  ( $R > 1$ ), the  $c_i = 1$  for all  $i$  (full cooperation) equilibrium suddenly becomes Aumann-strong ( $n$ -strong). Shaded regions denote the public bad game ( $r < 1$ ), and the public goods games with conflicting ( $1 < r < n$ ) and aligned ( $R > 1$ ) individual and public motives in terms of the Nash equilibrium of the game (see main text for details). We note that results for other population and/or group sizes are the same over  $R$ , while  $r$  and the slope of the red line of course scale accordingly.




Figure 2.2: Color-encoded average contribution levels in the unperturbed public goods game with directional learning. Simulations confirm that, with little directional learning sensitivity (i.e. when  $\delta$  is zero or very small), for the marginal per-capita rate of return  $R > 1$  the outcome  $c_i = 1$  for all  $i$  is the unique Nash and Aumann-strong equilibrium. For  $R = 1$  (dashed horizontal line), any outcome is a Nash equilibrium, but only  $c_i = 1$  for all  $i$  is Aumann-strong while all other outcomes are only Nash equilibria. For  $R < 1$ ,  $c_i = 0$  for all  $i$  is the unique Nash equilibrium, and its maximal  $k$ -strength depends on the population size. This is in agreement with results presented in Fig. 2.1. Importantly, however, as the responsiveness of players increases, contributions to the common pool become significant even in the defection-prone  $R < 1$ -region. In effect, individuals' (mis)learn what is best for them and end up contributing even though this would not be a unilateral best reply. Similarly, in the  $R > 1$  region free-riding starts to spread despite of the fact that it is obviously better to cooperate. For both these rather surprising and counterintuitive outcomes to emerge, the only thing needed is directional learning.




Figure 2.3: Time evolution of average contribution levels, as obtained for  $R = 0.7$ ,  $\delta = 0.1$  and different values of  $\epsilon$  (see legend). If only  $\epsilon > 0$ , the Nash equilibrium erodes to a stationary state where at least some members of the population always contribute to the common pool. There is a discontinuous transition to complete free-riding (defection) as  $\epsilon \rightarrow 0$ . Understandably, the lower the value of  $\epsilon$  (the smaller the probability for the perturbation), the longer it may take for the drift to gain on momentum and for the initial deviation to evolve towards the mixed stationary state. Note that the time horizontally is in logarithmic scale.

## References

- Aumann, Robert J (1974). “Subjectivity and correlation in randomized strategies”. In: *Journal of Mathematical Economics* 1, pp. 67–96.
- (1987). “Correlated equilibrium as an expression of Bayesian rationality”. In: *Econometrica* 55, pp. 1–18.
- Bayer, Ralph-C., Elke Renner, and Rupert Sausgruber (2013). “Confusion and learning in the voluntary contributions game”. In: *Experimental Economics* 16, pp. 478–496.
- Bernheim, B. D., B. Peleg, and M. D. Whinston (1987). “Coalition-Proof Equilibria I. Concepts”. In: *Journal of Economic Theory* 42, pp. 1–12.

- Bowles, Samuel and Herbert Gintis (2011). *A Cooperative Species: Human Reciprocity and Its Evolution*. Princeton, NJ: Princeton University Press.
- Chaudhuri, Ananish (2011). “Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature”. In: *Experimental Economics* 14, pp. 47–83.
- Erev, Ido and Alvin E Roth (1998). “Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria”. In: *American Economic Review* 88, pp. 848–881.
- Fehr, Ernst and Simon Gächter (2000). “Cooperation and Punishment in Public Goods Experiments”. In: *Am. Econ. Rev.* 90, pp. 980–994.
- Fischbacher, U., S. Gächter, and E. Fehr (2001). “Are people conditionally cooperative? Evidence from a public goods experiment”. In: *Econ. Lett.* 71, pp. 397–404.
- 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, pp. 541–556.
- Foster, Dean P and H Peyton Young (2006). “Regret testing: learning to play Nash equilibrium without knowing you have an opponent”. In: *Theoretical Economics* 1, pp. 341–367.
- Goeree, Jacob K. and Charles A. Holt (2005). “An Explanation of Anomalous Behavior in Models of Political Participation”. In: *American Political Science Review* 99, pp. 201–213.
- Hardin, Gerrett (1968). “The Tragedy of the Commons”. In: *Science* 162, pp. 1243–1248.
- Hofbauer, Josef and Karl Sigmund (1998). *Evolutionary Games and Population Dynamics*. Cambridge, U.K.: Cambridge University Press.
- Hrdy, Sarah Blaffer (2011). *Mothers and Others: The Evolutionary Origins of Mutual Understanding*. Cambridge, MA: Harvard University Press.



- Isaac, Mark R., Kenneth F. McCue, and Charles R. Plott (1985). “Public goods provision in an experimental environment”. In: *Journal of Public Economics* 26, pp. 51–74.
- Laslier, J.-F. and B. Walliser (2014). “Stubborn Learning”. In: *Theory and Decision* forthcoming.
- Ledyard, J. O. (1997). “Public Goods: A Survey of Experimental Research”. In: *The Handbook of Experimental Economics*. Ed. by J. H. Kagel and A. E. Roth. Princeton, NJ: Princeton University Press, pp. 111–194.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge, U.K.: Cambridge University Press.
- Mesterton-Gibbons, M. (2001). *An Introduction to Game-Theoretic Modelling, 2nd Edition*. Providence, RI: American Mathematical Society.
- Mesterton-Gibbons, M. and L. A. Dugatkin (1992). “Cooperation among unrelated individuals: Evolutionary factors”. In: *The Quarterly Review of Biology* 67, pp. 267–281.
- Moreno, D. and J. Wooders (1996). “Coalition-proof equilibrium”. In: *Games and Economic Behavior* 17, pp. 82–112.
- Myatt, D. P. and C. Wallace (2008). “An evolutionary analysis of the volunteer’s dilemma”. In: *Games Econ. Behav.* 62, pp. 67–76.
- Nash, John (1950). “Equilibrium points in n-person games”. In: *Proc. Natl. Acad. Sci. USA* 36, pp. 48–49.
- Nax, Heinrich H., Bary S. R. Pradelski, and H. Peyton Young (2013). “Decentralized dynamics to optimal and stable states in the assignment game”. In: *Proc. IEEE* 52, pp. 2398–2405.
- Nowak, Martin A. (2006a). *Evolutionary Dynamics*. Cambridge, MA: Harvard University Press.
- (2006b). “Five Rules for the Evolution of Cooperation”. In: *Science* 314, pp. 1560–1563.

- Nowak, Martin A. and Roger Highfield (2011). *SuperCooperators: Altruism, Evolution, and Why We Need Each Other to Succeed*. New York: Free Press.
- Palfrey, Thomas R. and Jeffrey E. Prisbey (1997). “Anomalous Behavior in Public Goods Experiments: How Much and Why?” In: *The American Economic Review* 87, pp. 829–846.
- Perc, M. and A. Szolnoki (2010). “Coevolutionary games – a mini review”. In: *BioSystems* 99, pp. 109–125. DOI: [10.1016/j.biosystems.2009.10.003](https://doi.org/10.1016/j.biosystems.2009.10.003).
- Perc, M. et al. (2013). “Evolutionary dynamics of group interactions on structured populations: a review”. In: *J. R. Soc. Interface* 10, p. 20120997.
- Rand, D. A. and M. A. Nowak (2013). “Human cooperation”. In: *Trends in Cognitive Sciences* 17, pp. 413–425.
- Rand, D. G. et al. (2009). “Positive Interactions Promote Public Cooperation”. In: *Science* 325, pp. 1272–1275.
- Rand, D. G. et al. (2013). “Evolution of fairness in the one-shot anonymous Ultimatum Game”. In: *Proc. Natl. Acad. Sci. USA* 110, pp. 2581–2586.
- 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, pp. 164–212.
- Santos, F. C. and J. M. Pacheco (2011). “Risk of collective failure provides an escape from the tragedy of the commons”. In: *Proc. Natl. Acad. Sci. USA* 108, pp. 10421–10425. DOI: [10.1073/pnas.1015648108](https://doi.org/10.1073/pnas.1015648108).
- Santos, F. C., M. D. Santos, and J. M. Pacheco (2008). “Social diversity promotes the emergence of cooperation in public goods games”. In: *Nature* 454, pp. 213–216.
- Sauermann, Heinz and Reinhard Selten (1962). “Anspruchsanpassungstheorie der unternehmung”. In: *Journal of Institutional and Theoretical Economics* 118, pp. 577–597.

- Selten, R. and J. Buchta (1994). *Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions*. Discussion Paper Series B.
- Selten, R. and R. Stoecker (1986). “End behavior in sequences of finite Prisoner’s Dilemma supergames A learning theory approach”. In: *Journal of Economic Behavior & Organization* 7, pp. 47–70.
- Szabó, György and Gábor Fáth (2007). “Evolutionary games on graphs”. In: *Phys. Rep.* 446, pp. 97–216.
- Weibull, J. W. (1995). *Evolutionary Game Theory*. Cambridge, MA: MIT Press.
- Young, H Peyton (2009). “Learning by trial and error”. In: *Games and Economic Behavior* 65, pp. 626–643.
- Young, H. Peyton et al. (2013). *Learning in a Black Box*. Economics Series Working Papers.

## Chapter 3

Social preferences versus  
learning:

Learning and the contribution  
decline in public goods games

## Abstract

Economic games such as the public goods game are increasingly being used to measure social behaviours in humans and non-human primates. The results of such games have been used to argue that people are pro-social, and that humans are uniquely altruistic, willingly sacrificing their own welfare in order to benefit others. However, an alternative explanation for the empirical observations is that individuals are mistaken, but learn, during the game, how to improve their personal payoff. We test between these competing hypotheses, by comparing the explanatory power of different behavioural rules, in public goods games, where individuals are given different amounts of information. We find: (i) that individual behaviour is best explained by a learning rule that is trying to maximize personal income; (ii) that conditional cooperation disappears when the consequences of cooperation are made clearer; and (iii) that social preferences, if they exist, are more anti-social than pro-social.

**Acknowledgements.** We thank Jay Biernaskie, Innes Cuthill, Claire El Mouden, Nichola Raihani and two anonymous referees for comments. We thank the ERC, Nuffield College and the Calleva Research Centre, Magdalen College, for funding.

### 3.1 Introduction

The results from economic games have been used to argue that humans are altruistic in a way that differs from most if not all other organisms (Fehr and Gächter, 2002; Gintis et al., 2003; Fehr and Fischbacher, 2003; Henrich, 2006). In public goods games experiments, participants have to choose how much of their monetary endowment they wish to keep for themselves and how much to contribute to a group project (Ledyard, 1995; Chaudhuri, 2011). Contributions to the group project are automatically multiplied by the experimenter before then being shared out equally among all group members regardless of their relative contributions (Isaac and Walker, 1988b; Isaac and Walker, 1988a). The multiplication is usually less than the group size, so that a contributor receives back less from her contribution than she contributed. In this case, participants have to choose between retaining their full endowment and thus maximizing their personal income, or sacrificing some of their earnings to the benefit of the group. Hundreds of experiments have shown that most people partially contribute to the group project and thus fail to maximize personal income (Ledyard, 1995; Chaudhuri, 2011). It has been argued that this robust result demonstrates that humans have a unique regard for the welfare of others, termed pro-social preferences, which cannot be explained by kin selection (Hamilton, 1964), reciprocity (Trivers, 1971) and/or via improved reputation (Alexander, 1987; Nowak and Sigmund, 1998; Wedekind and Milinski, 2000; Nowak and Sigmund, 2005). Consequently, economic games are also increasingly being used in non-human primates in attempts to explore the evolutionary origins of such puzzling social behaviours (Brosnan and Waal, 2003; Jensen, Call, and Tomasello, 2007; Proctor et al., 2013).

The conclusion that humans are especially, perhaps uniquely, altruistic has relied on the assumption that individuals play ‘perfectly’ in experiments such as

the public goods game. Specifically, that individuals have a full understanding of the game, in terms of the consequences of their behaviour for themselves and others, such that their play reflects how they value the welfare of others (social preferences) as in Fehr and Gächter, 2002; Fehr and Schmidt, 1999. This results in the inference that the costly decisions that players make knowingly inflict a personal cost in order to benefit others (Fehr and Fischbacher, 2003). Consequently the typical decline in contributions when players are made to play the game repeatedly (Ledyard, 1995; Chaudhuri, 2011); see figure 3.1, is argued to be a withdrawal of cooperation in response to a minority of non-cooperators (Fischbacher, Gächter, and Fehr, 2001; Fischbacher and Gächter, 2010; Camerer, 2013).

An alternative explanation for the data is that individuals are trying to maximize their financial gain, but they are not playing the game ‘perfectly’ (Kümmerli et al., 2010; Burton-Chellew and West, 2013). This hypothesis predicts individuals initially cooperate to some degree, because they are uncertain and bet-hedge (Burton-Chellew and West, 2013), or they are mistaken about how the payoffs operate (Kümmerli et al., 2010; Houser and Kurzban, 2002; Andreoni, 1995), or perhaps they operate a heuristic from every-day life that starts off cooperating without calculating the consequences (Rand, Greene, and Nowak, 2012). This hypothesis consequently predicts a decline in cooperation over time as individuals learn, albeit imperfectly, how behaviour influences payoffs. Consistent with this alternative hypothesis, individuals have been found to contribute similar amounts over time to the group project (as observed in standard experiments) even in low-information environments, that is, even when they do not know they are playing the public goods game with others (Burton-Chellew and West, 2013; Bayer, Renner, and Sausgruber, 2013). However, this alternate hypothesis has been argued against, with the suggestion that the decline in cooperation is better explained by pro-social individuals conditionally






Figure 3.1: We analyse the data from Burton-Chellew & West Burton-Chellew and West, 2013. Participants played a public goods game for 20 repeated rounds, with random group composition each round. There were three different information treatments (see text for details). The results conform to the stereotypical results of public goods games, in that contributions commence at intermediate values and decline steadily with repetition of the game.

cooperating depending upon the behaviour of others, rather than individuals learning how to better play the game (Camerer, 2013).

We explicitly test these competing hypotheses, by examining the rules that individuals use to vary their behaviour when playing the public goods game (Camerer, 2003; Erev and Haruvy, 2013) (figure 3.2). Our first rule assumes that individuals are trying to maximize their own income, but are uncertain or mistaken as to how to do this. They thus subsequently use information from game play to try and improve their earnings. For example, if contributing less over time to the public good coincided with an increase in such an individual’s financial reward, then this individual would contribute even less next time, and vice versa; behavior known as ‘directional learning’ (Bayer, Renner, and

Sausgruber, 2013; Erev and Haruvy, 2013; Cross, 1983; Selten and Stoecker, 1986; Sauermann and Selten, 1962; Selten and Buchta, 1998)). Our second and third rules are based on two forms of pro-social behaviour that have been previously argued to lead to altruistic behaviour in public goods games (Fischbacher, Gächter, and Fehr, 2001; Fischbacher and Gächter, 2010; Croson, Fatas, and Neugebauer, 2005; Croson, 2007). Our second rule assumes that individuals are trying to maximize a weighted function of their own income and that of their group-mates (Croson, 2007). This also allows directional learning, but in a way that takes account of the consequences of behaviour for others. Our third rule is conditional cooperation, in response to the cooperation of others (Fischbacher, Gächter, and Fehr, 2001; Fischbacher and Gächter, 2010; Croson, Fatas, and Neugebauer, 2005; Böhm and Rockenbach, 2013). For example, if the average contributions of one's group-mates increase from one round to the next, then one will respond by contributing more in the next round.

We analysed data from three public goods games, all with the same payoff-structure, but which differ in the amount of information that the players are given about the consequences of their behaviour for others. Specifically, individuals had no knowledge that their behaviour even benefited others (black box), or were told at the start how their behaviour benefited others (standard), or were also shown after each round of play that contributions benefited others (enhanced).<sup>1</sup> By comparing behaviour in these different games, we could explicitly examine the extent to which behaviour was influenced by consequences for the actor himself/herself (the only concern in the black box), and consequences for others (increasingly highlighted in the standard and enhanced treatments). In addition, we told players in the standard and enhanced treatment the decisions of their group-mates after each round. This allows us to

---

<sup>1</sup>See Burton-Chellew and West, 2013 for further instruction details.

test whether players are attempting to condition their cooperation and whether this depends on how clear the benefits of contributing are for others.

## 3.2 Material and methods

### 3.2.1 Data collection

We analysed the dataset from our previously published study, where the experimental methods are described in detail (Burton-Chellew and West, 2013); see figure 3.1. This experiment examined the behaviour of 236 individuals, distributed among 16 sessions. Here, we provide a brief summary of the parts of the experimental design relevant to this study.

We tested three versions of the public goods game and used an identical set-up and payoff matrix, but provided different levels of social information, each time. In each session, we had 12 or 16 participants and we grouped them into groups of four and had them play the public good game, before repeating the game again and again for a total of 20 rounds. Groups were randomly created every round. In all treatments, we gave our participants a fresh endowment of 40 monetary units (MU), or 40 coins (for the black box), per round, and multiplied the contributions of players by 1.6 before sharing them out equally among all four group members. This meant that the marginal-per-capita-return (MPCR) for each unit contributed was 0.4. Consequently, contributions were always personally costly and to not contribute was the payoff-maximizing (strictly dominant) strategy in each round.

Our most extreme condition was an entirely asocial set-up, with no social framing, and where instead of allowing participants to contribute to a group-project, we let them contribute to a ‘black box’, even though they were in

reality playing a standard inter-connected public goods game. We told the participants that the black box ‘performs a mathematical function that converts the number of coins inputted into a number of coins to be outputted.’ This allowed us to deliberately create participants that would not know the payoff-maximizing strategy, and are also unconcerned by other-regarding preferences. In such a condition, the participants could only be motivated to try and adjust their behaviour so as to maximize their own income, as much as participants are ever so motivated.

Our other two treatments were revealed public goods games, where we told our participants they could either contribute each MU to a group project (the public good) or keep it for themselves. We told our players how the game works, specifically that contributions are multiplied by 1.6 before being shared out equally among all four players. In both of these ‘revealed’ versions of the game, we gave our participants the exact same instructions, but we gave more information after each round of play in one treatment than the other. Specifically, in the ‘standard’ set-up, we told participants after each round what their own payoffs were, and also what the decisions of their three group-mates were. This is the most typical information content of public goods game studies (e.g. Fehr and Gächter, 2002), which has provided the template for many subsequent studies. In our ‘enhanced’ treatment, we also informed our participants what their groupmates individual returns from the group project were and their subsequent individual earnings. Note that in this enhanced treatment, there is strictly speaking no new information relative to the standard treatment, if players (i) understood the game and (ii) were calculating the earnings of their group-mates from their contributions.

Methodologically, in each session, we had our participants play two ‘game-frames’, i.e. both a black box game and a revealed public goods game, in order

to enable a within-participant analysis. We presented the two games as two entirely separate experiments to minimize spill-over effects: in one they could ‘input’ ‘coins’ into a ‘black box’, in the other they could ‘contribute’ ‘MU’ to a ‘group project’, and the order of play of these games was counter-balanced across sessions.

### 3.2.2 Statistical analysis

We tested three learning rules (figure 3.2). In all cases, we assumed that players adjusted their behaviour according to whether previous behavioural adjustments lead to positive or negative consequences for the proposed underlying utility function. For example, if players derive utility only from their personal income, and a previous reduction (or increase) in their contributions led to an increase in their personal income, then in the next time step they would gravitate towards the lesser (or greater), more successful, level of contribution. Similarly, if players value the payoffs to others, then *ceteris paribus*, others’ changes in income would be responded to in an equivalent way. The three underlying utility functions that we examine were as follows:

- (I) payoff-based learning: individuals set contributions,  $c_i$ , in response to their own income,  $\varphi_i(c)$  and the resulting utility is simply  $u_i(c) = \varphi_i$ ;
- (II) pro-social learning: individuals set contributions,  $c_i$ , in response to both their own income,  $\varphi_i(c)$ , and the income of the other members of their group,  $\varphi_j(c)$  and the resulting utility, a weighted function of the two, is  $u_i(c) = (1 - \alpha_i)\varphi_i(c) + \alpha_i \sum_{j \neq i} \varphi_j(c)$ , where  $\alpha_i$  measures the agent’s concern for others’ payoffs. Pro-sociality implies  $\alpha_i > 0$ ; and
- (III) conditional cooperation: individuals set contributions,  $c_i$ , in response to their own income,  $\varphi_i(c)$  and to the contributions of their groupmates,

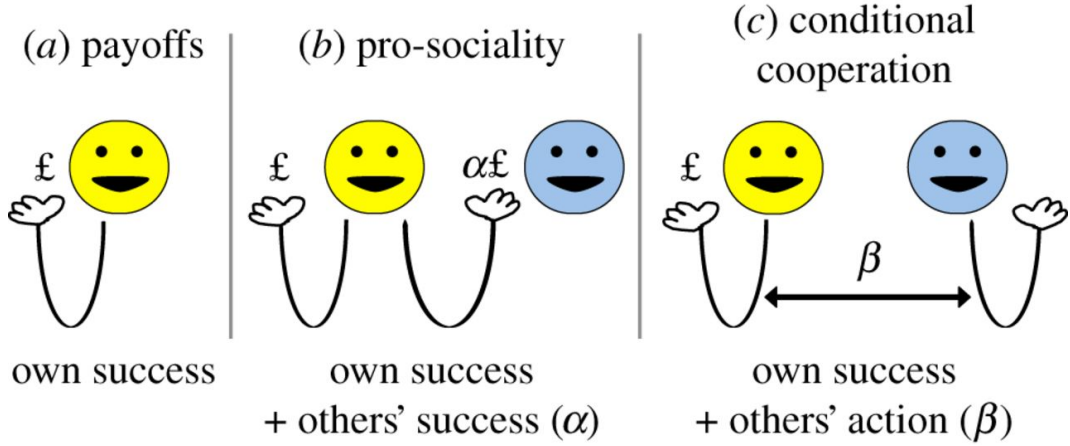


Figure 3.2: We considered the explanatory power of three behavioural response rules: (a) payoff-based learning based on increasing own income; (b) pro-social directional learning, based on own income and the income of others (weighted by  $\alpha$ ); and (c) conditional cooperation, based on own income and a desire to equalize incomes (weighted by  $\beta$ ).

such that the resulting utility is  $u_i(c) = \varphi_i(c) + -\beta_i \left| \sum_{j \neq i} (c_j - c_i) \right|$ , where  $\beta_i$  measures the agent's concern to match others' contributions.

We chose these utility functions because of their relationship to the utility functions already discussed in the literature, and because they allow a clear comparison between cases with and without pro-social preferences. Different, and potentially more elaborate behavioural rules could be favoured in different scenarios allowing more behavioural flexibility (Hauert et al., 2002; Semmann, Krambeck, and Milinski, 2003).

We perform ordinary linear regressions with individual-level clustering of the form  $f(c_i^{t+1}) = \beta X_i^t + e^{t+1}$ , where  $f(c_i^{t+1})$ , measuring a contribution adjustment by player  $i$ , is the response variable and  $X_i^t$  is the vector of the predictor variables including those for the three hypotheses.  $\beta$  is the vector of parameters to be estimated and  $\beta_i$  is the estimator of predictor variable  $x_i$ 's positive effect on the response variable for a unit change in  $x_i$ .  $e^{t+1}$  represents the standard (normally distributed) error term for this model. We focus on adjustments in periods 1 – 10 because median contributions, having reached zero

in the enhanced treatment, and near zero otherwise (5/6 for black box and 4 for standard), change little after this and we are interested in modelling how cooperative behaviour changes over time.

Our response variable records an individual's directional changes in contributions over time:  $f(c_i^{t+1})$ , and takes the value +1 when representing an increase in contributions (relative to the average of the previous two periods), -1 when representing a decrease and 0 otherwise. Our predictor variables specify the directional change in contributions that should occur in line with the relevant utility function or learning rule.

The predictor variables  $x_i$  represent the three different learning rules above by encoding the previous relationship between an agent's contributions and (I) their payoffs, (II) their group-mates' payoffs or (III) their group-mates' actions, respectively. They take integer values from -1 to 1. Specifically, for utility function (I), payoff-based learning, if a player's contribution increased across the two rounds (if  $c_i^t > c_i^{t-1}$ ) along with their payoff ( $\varphi_i^t \geq \varphi_i^{t-1}$ ), then we predict that this coupling of increased contributions with 'success' (increased payoff) will lead to a contribution increase (relative to the mean of the two previous rounds). We therefore encode this as +1. Likewise, following a contribution decrease and 'failure' (if  $c_i^t < c_i^{t-1}$  and  $\varphi_i^t < \varphi_i^{t-1}$ ) we also predict a contribution increase and encode +1. By contrast, following a contribution decrease and 'success' (if  $c_i^t < c_i^{t-1}$  and  $\varphi_i^t \geq \varphi_i^{t-1}$ ) or a contribution increase and 'failure' (if  $c_i^t > c_i^{t-1}$  and  $\varphi_i^t < \varphi_i^{t-1}$ ), we predict a contribution decrease (relative to the mean of the two previous rounds) and encode -1, and we predict 0 for all other cases.

For utility function (II), pro-social learning, we likewise encode the value +1 following either a contribution increase and 'other-regarding success' (if  $c_i^t > c_i^{t-1}$  and  $\sum_{j \neq i} \varphi_j^t \geq \sum_{j \neq i} \varphi_j^{t-1}$ ) or a contribution decrease and 'other-

Table 3.1: Summary of results from testing the three different learning rules together. (The table details the statistical significance of the three learning rules (payoff-based learning, pro-social learning and conditional cooperation) for the three information treatments (black box, standard and enhanced).  $\checkmark$ , estimators significantly support direction of hypothesis in this treatment.  $\times$ , estimators significantly contradict direction of hypothesis in this treatment, n.s., non-significant. The values represent the estimate of the effects of unit changes in the hypothesis-specific predictor variables on the response variable; positive (negative) parameter estimators support (contradict) the respective hypothesis. Table 3.2 details the regressions fully.)

	black box	standard	enhanced
payoff-based learning	$\checkmark$ 0.30*	$\checkmark$ 0.25*	$\checkmark$ 0.14*
pro-social learning <sup>a</sup>	$\times$ -0.13*	$\times$ -0.23*	$\times$ -0.29*
conditional cooperation <sup>a</sup>	n.s. 0.05	$\checkmark$ 0.21*	n.s. -0.001

\* significance < 0.001.

<sup>a</sup> Controlling for payoff-based learning.

regarding failure' (if  $c_i^t < c_i^{t-1}$  and  $\sum_{j \neq i} \varphi_j^t < \sum_{j \neq i} \varphi_j^{t-1}$ ); -1 following either a contribution decrease coupled with 'other-regarding success' (if  $c_i^t < c_i^{t-1}$  and  $\sum_{j \neq i} \varphi_j^t \geq \sum_{j \neq i} \varphi_j^{t-1}$ ) or a contribution increase with 'other-regarding failure' (if  $c_i^t > c_i^{t-1}$  and  $\sum_{j \neq i} \varphi_j^t < \sum_{j \neq i} \varphi_j^{t-1}$ ), and 0 otherwise. Thus this variable, along with the payoff-based learning variable, is also positive if the prior directional changes in contributions were maintained after success or reversed after failure, but success and failure are now judged in terms of others' payoffs instead of own payoffs. For our third utility function, (III), conditional cooperation, we encode +1 when there has been an increase in the mean contribution of group-mates across the previous two rounds (if  $\sum_{j \neq i} c_j^t \geq \sum_{j \neq i} c_j^{t-1}$ ) and 0 otherwise.

Positive estimators of the  $\beta_i$ , mean a positive correlation between the learning rule and the subsequent changes in contributions, and thus support the



respective hypothesis, whereas negative estimators, meaning a negative correlation between the learning rule and the subsequent changes in contributions, contradict the respective hypothesis. For pro-social learning, the coefficient indicates whether the average of weights  $\alpha_i$  on others' income is supportive of pro-sociality (positive) or not. Table 3.1 summarizes the results according to their implications for the various hypotheses. Table 3.2 provides full details of the parameter estimates for all models on all the data. The electronic supplementary material provides the parameter estimates for models that analysed sub-sets of the data according to which game-frame order they belonged to (see Material and methods, data collection). We also provide a table detailing the utility functions and their quantitative relationship to the data (electronic supplementary material, table 3.2).

### 3.3 Results and discussion

We found that our payoff-based learning rule was significant for all three versions of the public goods game, in contrast to both our pro-social and conditional-cooperation rules which were typically non-significant or significant in the wrong direction (tables 3.1 and 3.2; electronic supplementary material).

#### 3.3.1 Learning in a black box

In the black box treatment, the behaviour of individuals could best be explained by payoff-based responses, with players significantly learning to improve their income (tables 3.1 and 3.2). Figure 3.1 confirms that, this leads to behaviour at the group level which is strikingly similar to play in standard public goods games. By contrast, the pro-social response rule estimate was sig-

Table 3.2: A comparison of the different behavioural rules, plus one combining them all together, across three different information treatments. (PBL, payoff-based learning (own success); PSL, pro-social learning (own success and others' success); CC, conditional cooperation (own success and others' actions). All, a combination of all the components from the three rules (own success, others' success, and others' actions). The parameters in the first three rows estimate the effects of unit changes in the predictor variables that act as components in the three learning rules; positive (negative) parameter estimators support (contradict) the respective hypothesis.)

	black box estimate (significance)			standard estimate (significance)			enhanced estimate (significance)					
	PBL	PSL	CC	All	PBL	PSL	CC	All	PBL	PSL	CC	All
own	0.31	0.29	0.30	0.30	0.28	0.22	0.30	0.25	0.22	0.14	0.19	0.14
success	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
others'	-0.12		-0.13		-0.16			-0.23		-0.29		-0.29
success	(0.001)		(0.001)		(0.001)			(0.001)		(0.001)		(0.001)
others'		-0.04	0.05			0.09	0.21			-0.11	-0.001	
actions		(0.241)	(0.178)			(0.038)	(0.001)			(0.012)	(0.95)	
$r^2$	0.09	0.10	0.09	0.10	0.07	0.09	0.07	0.10	0.04	0.11	0.05	0.11
no.	1888	1888	1888	1888	928	928	928	928	960	960	960	960
obs.												

nificantly negative, attributing a negative weight to the welfare of other players. This would represent anti-social preferences if it were not for the asocial frame of the black box treatment and provides a baseline estimate for the anti-social nature of payoff-based learning. The conditional-cooperation response rule was not significant when payoff-based learning is controlled for.

### 3.3.2 Learning in public goods games

We found that the behaviour of individuals in public goods games could, as in the black box, be significantly explained by payoff-based learning, but not by pro-sociality (tables 3.1 and 3.2). Again the pro-social learning rule estimated a significantly negative weight to the other players ( $\alpha$ ), implying ‘anti-social’ behaviour in this socially framed game (tables 3.1 and 3.2). The coefficient was considerably larger in the enhanced treatment, than in the standard treatment, and considerably larger in the standard treatment than in the black box, suggesting that providing players with more information on how contributions benefit others but are personally costly has anti-social consequences. This would not be the case if players understood the game and were willingly sacrificing in order to benefit others.

Conditional cooperation was significant in the standard version but not the enhanced version of the game, which has identical instructions and game structure, but where individuals were explicitly shown the returns to the other group members from the group project. This enhanced information could of course in principle be calculated by participants in the standard treatment as they knew the decisions of their group-mates. In the standard version, the conditional cooperation rule was not so significant unless controlling for anti-social responses to others’ success (table 3.2).

Conditional cooperation is proposed to explain the typical decline in contri-

butions over time (Fischbacher, Gächter, and Fehr, 2001; Fischbacher and Gächter, 2010), but contributions declined faster in the enhanced treatment where conditional cooperation was either non-significant (combined model, table 3.2) or significantly negative (non-combined model, table 3.2). This suggests that the conditional cooperation in the standard treatment is more to do with social learning than social preferences, as the reduced uncertainty in the enhanced treatment may reduce uncertain participants' reliance upon imitation (Carpenter, 2004). In addition, if some participants have incorrect beliefs about how the payoffs are determined and choose to match others in the standard treatment, they may be less likely to do so in the enhanced treatment as they revise their mistaken beliefs.

The dataset we used also contains three additional experimental treatments, where the contributions were multiplied by 6.4 instead of 1.6 and thus the resulting MPCR was 1.6 instead of 0.4 (Burton-Chellew and West, 2013). In these treatments, the MPCR  $> 1.0$ , which means that contributing fully was both the income-maximizing (strictly dominant) strategy for any particular round and also the social optimum. We do not analyse the data from these treatments here, because in such treatments it is impossible to differentiate our first and second behavioural rules, as individual and pro-social outcomes are aligned in these settings (there is no conflict between individual and group outcomes). However, the fact that contributions were significantly below full contribution in all three treatments, even after 20 rounds, but increased over time in both the black box and the standard games (Burton-Chellew and West, 2013), is also consistent with the payoff-based learning hypothesis.

However, such payoff-based learning does not require that people realize that the dominant strategy is independent of their group-mates' actions. Therefore the re-start phenomenon (Andreoni, 1988; Croson, 1996) whereby average

cooperation levels temporarily increase from a previous decline when the experiment is ‘re-started’, while challenging, does not falsify learning hypotheses, and may also be partly owing to selfish players attempting to manipulate others (Andreoni, 1988; Kreps et al., 1982; Ambrus and Pathak, 2011).

### 3.3.3 Cooperation in public goods games

Overall, our analyses suggest that changes in behaviour over time in public goods games are largely explained by participants learning how to improve personal income. We found conflicting support for conditional cooperation as such behaviour disappeared when the consequences of contributing were made clearer. This suggests that conditional cooperation is largely due to confusion/error and not pro-sociality. This is reinforced by our lack of evidence of a desire to help others (pro-sociality). Indeed, we found that, if anything, the benefits to others are weighted negatively, with individuals adjusting their behaviour to better reduce the income of others. We are not suggesting that humans are anti-social, nor that they are never pro-social — pro-sociality is found across the tree of life from genes to cells to vertebrates (West, Griffin, and Gardner, 2007) — rather, that public goods games do not demonstrate that humans are uniquely altruistic.

Our conclusions contradict a widely accepted paradigm in the field of human behaviour, that the results of public goods games reflect a uniquely human regard for the welfare of others (Fehr and Fischbacher, 2003; Fehr and Schmidt, 1999; Fischbacher and Gächter, 2010). We suggest that the acceptance of this human pro-sociality hypothesis was based on two things. First, there has perhaps been a lack of control treatments where imperfect behaviour would not always lead to higher than expected levels of cooperation (Kümmerli et al., 2010), and null hypotheses, such as that provided by the black box treatment (Burton-

Chelley and West, 2013). Second, there has been an implicit assumption that humans behave as utility-maximizers, such that their costly choices reliably reveal their (social) preferences (Fehr and Schmidt, 1999).

However, there is an increasing range of evidence that individuals do not play games as perfect maximizing machines (Kümmerli et al., 2010; Burton-Chelley and West, 2013; Houser and Kurzban, 2002; Andreoni, 1995), that they instead exhibit bounded-rationality, and can be influenced by a variety of ‘irrelevant’ factors that do not influence payoffs in the game (Burnham and Hare, 2007; Nettle et al., 2013; Burton-Chelley and West, 2012). This is in accord with one of the revolutionary findings of behavioural economics, that people are predictably irrational, and make systematic errors that limit their own welfare (Camerer, 2003). Yet paradoxically, the behavioural economics approach is routinely used to ‘measure’ pro-sociality, using methods that rely upon the assumption of rational choice and revealed preferences.

**Data accessibility.** All the data have been submitted to Dryad and are available at [doi:10.5061/dryad.cr829](https://doi.org/10.5061/dryad.cr829).

## References

- Alexander, Richard D (1987). *The Biology of Moral Systems*. Hawthorne, New York: Aldine de Gruyter.
- Ambrus, Attila and Parag A Pathak (2011). “Cooperation over finite horizons: A theory and experiments”. In: *Journal of Public Economics* 95.7-8, pp. 500–512.
- Andreoni, James (1988). “Why free ride: strategies and learning in public-goods experiments”. In: *Journal of Public Economics* 37.3, pp. 291–304.

- Andreoni, James (1995). “Cooperation in public-goods experiments: kindness or confusion?” In: *The American Economic Review* 85.4, pp. 891–904.
- Bayer, Ralph-C., Elke Renner, and Rupert Sausgruber (2013). “Confusion and learning in the voluntary contributions game”. In: *Experimental Economics* 16.4, pp. 478–496.
- Böhm, Robert and Bettina Rockenbach (2013). “The Inter-Group Comparison – Intra-Group Cooperation Hypothesis: Comparisons between Groups Increase Efficiency in Public Goods Provision”. In: *PloS ONE* 8.2, e56152.
- Brosnan, Sarah F and Frans B M de Waal (2003). “Monkeys reject unequal pay”. In: *Nature* 425.6955, pp. 297–299.
- Burnham, Terence C and Brian Hare (2007). “Engineering Human Cooperation”. In: *Human Nature* 18.2, pp. 88–108.
- Burton-Chellew, Maxwell N and Stuart A West (2012). “Pseudocompetition among groups increases human cooperation in a public-goods game”. In: *Animal Behaviour* 84.4, pp. 947–952.
- (2013). “Prosocial preferences do not explain human cooperation in public-goods games”. In: *Proceedings of the National Academy of Sciences* 110.1, pp. 216–221.
- Camerer, Colin F (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.
- (2013). “Experimental, cultural, and neural evidence of deliberate prosociality”. In: *Trends in Cognitive Sciences* 17.3, pp. 106–108.
- Carpenter, Jeffrey P (2004). “When in Rome: conformity and the provision of public goods”. In: *The Journal of Socio-Economics* 33.4, pp. 395–408.
- Chaudhuri, Ananish (2011). “Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature”. In: *Experimental Economics* 14.1, pp. 47–83.

- Croson, Rachel T A (1996). “Partners and strangers revisited”. In: *Economics Letters* 53.1, pp. 25–32.
- (2007). “Theories of commitment, altruism and reciprocity: evidence from linear public goods games.” In: *Economic Inquiry* 45.2, pp. 199–216.
- Croson, Rachel T A, Enrique Fatas, and Tibor Neugebauer (2005). “Reciprocity, matching and conditional cooperation in two public goods games”. In: *Economics Letters* 87.1, pp. 95–101.
- Cross, John G (1983). *A Theory of Adaptive Economic Behavior*. Cambridge, UK: Cambridge University Press.
- Erev, Ido and Ernan Haruvy (2013). “Learning and the economics of small decisions”. In: *The handbook of experimental economics*. Ed. by John H Kagel and Alvin E Roth. Vol. 2. Princeton, NJ: Princeton University Press.
- Fehr, Ernst and Urs Fischbacher (2003). “The nature of human altruism”. In: *Nature* 425.6960, pp. 785–791.
- Fehr, Ernst and Simon Gächter (2002). “Altruistic punishment in humans”. In: *Nature* 415.6868, pp. 137–140.
- 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 and Simon Gächter (2010). “Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments”. In: *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.
- Gintis, Herbert et al. (2003). “Explaining altruistic behavior in humans”. In: *Evolution and Human Behavior* 24.3, pp. 153–172.



- Hamilton, William D (1964). “The genetical evolution of social behaviour. I & II”. In: *Journal of Theoretical Biology* 7.1, pp. 1–52.
- Hauert, Christoph et al. (2002). “Volunteering as Red Queen Mechanism for Cooperation in Public Goods Games”. In: *Science* 296.5570, pp. 1129–1132.
- Henrich, Joseph (2006). “Costly Punishment Across Human Societies”. In: *Science* 312.5781, pp. 1767–1770.
- Houser, Daniel and Robert Kurzban (2002). “Revisiting Kindness and Confusion in Public Goods Experiments”. In: *American Economic Review* 92.4, pp. 1062–1069.
- Isaac, R Mark and James M. Walker (1988a). “Communication and free-riding behavior: The voluntary contribution mechanism”. In: *Economic Inquiry* 26.4, pp. 585–608.
- (1988b). “Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism”. In: *The Quarterly Journal of Economics* 103.1, p. 179.
- Jensen, Keith, Josep Call, and Michael Tomasello (2007). “Chimpanzees Are Rational Maximizers in an Ultimatum Game”. In: *Science* 318.5847, pp. 107–109.
- Kreps, David M et al. (1982). “Rational cooperation in the finitely repeated prisoners’ dilemma”. In: *Journal of Economic Theory* 27.2, pp. 245–252.
- Kümmerli, Rolf et al. (2010). “Resistance to extreme strategies, rather than prosocial preferences, can explain human cooperation in public goods games”. In: *Proceedings of the National Academy of Sciences* 107.22, pp. 10125–10130.
- Ledyard, John O (1995). “Public goods: A survey of experimental research”. In: *Handbook of experimental economics*. Ed. by John H Kagel and Alvin E Roth. Princeton, NJ: Princeton University Press, pp. 253–279.

- Nettle, Daniel et al. (2013). “The watching eyes effect in the Dictator Game: it’s not how much you give, it’s being seen to give something”. In: *Evolution and Human Behavior* 34.1, pp. 35–40.
- Nowak, Martin A and Karl Sigmund (1998). “Evolution of indirect reciprocity by image scoring”. In: *Nature* 393.6685, pp. 573–577.
- (2005). “Evolution of indirect reciprocity”. In: *Nature* 437.7063, pp. 1291–1298.
- Proctor, Darby et al. (2013). “Chimpanzees play the ultimatum game”. In: *Proceedings of the National Academy of Sciences* 110.6, pp. 2070–2075.
- Rand, David G., Joshua D. Greene, and Martin A. Nowak (2012). “Spontaneous giving and calculated greed”. In: *Nature* 489.7416, pp. 427–430.
- Sauermann, Heinz and Reinhard Selten (1962). “Anspruchsanpassungstheorie der Unternehmung”. In: *Zeitschrift für die gesamte Staatswissenschaft/Journal of Institutional and Theoretical Economics*, pp. 577–597.
- Selten, Reinhard and Joachim Buchta (1998). “Experimental sealed bid first price auctions with directly observed bid functions”. In: *Games and human behavior: essays in honor of Amnon Rapoport*. Ed. by Amnon Rapoport et al. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 79–104.
- Selten, Reinhard and Rolf Stoecker (1986). “End behavior in sequences of finite Prisoner’s Dilemma supergames: A learning theory approach”. In: *Journal of Economic Behavior & Organization* 7.1, pp. 47–70.
- Semmann, Dirk, Hans-Jürgen Krambeck, and Manfred Milinski (2003). “Volunteering leads to rock-paper-scissors dynamics in a public goods game”. In: *Nature* 425.6956, pp. 390–393.
- Trivers, Robert L (1971). “The Evolution of Reciprocal Altruism”. In: *The Quarterly Review of Biology* 46.1, pp. 35–57.
- Wedekind, Claus and Manfred Milinski (2000). “Cooperation Through Image Scoring in Humans”. In: *Science* 288.5467, pp. 850–852.

West, Stuart A, Ashleigh S Griffin, and Andy Gardner (2007). “Evolutionary Explanations for Cooperation”. In: *Current Biology* 17.16, R661–R672.

## Chapter 4

Evolution of market  
equilibria:

Equity dynamics in matching  
markets

## **Abstract**

We study evolutionary dynamics in assignment games where many agents interact anonymously at virtually no cost. The process is decentralized, very little information is available and trade takes place at many different prices simultaneously. We propose a completely uncoupled learning process that selects a subset of the core of the game with a natural equity interpretation. This happens even though agents have no knowledge of other agents' strategies, payoffs, or the structure of the game, and there is no central authority with such knowledge either. In our model, agents randomly encounter other agents, make bids and offers for potential partnerships and match if the partnerships are profitable. Equity is favored by our dynamics because it is more stable, not because of any ex ante fairness criterion.

**Acknowledgements.** Foremost, we thank Peyton Young for his guidance. He worked with us throughout large parts of this project and provided invaluable guidance. Further, we thank Itai Arieli, Peter Biró, Gabrielle Demange, Gabriel Kreindler, Jonathan Newton, Tom Norman, Tamás Solymosi and anonymous referees for suggesting a number of improvements to earlier versions. We are also grateful for comments by participants at the 23rd International Conference on Game Theory at Stony Brook, the Paris Game Theory Seminar, the AFOSR MUIR 2013 meeting at MIT, and the 18th CTN Workshop at the University of Warwick. The research was supported by the United States Air Force Office of Scientific Research Grant FA9550-09-1-0538 and the Office of Naval Research Grant N00014-09-1-0751.

## 4.1 Introduction

Many matching markets are decentralized and agents interact repeatedly with very little knowledge about the market as a whole. Examples include online markets for bringing together buyers and sellers of goods, matching workers with firms, matching hotels with clients, and matching men and women. In such markets matchings are repeatedly broken, reshuffled, and restored. Even after many encounters, however, agents may still have little information about the preferences of others, and they must experiment extensively before the market stabilizes.

In this paper we propose a simple adaptive process that reflects the participants' limited information about the market. Agents have aspiration levels that they adjust from time to time based on their experienced payoffs. Matched agents occasionally experiment with higher bids in the hope of extracting more from another match, while single agents occasionally lower their bids in the hope of attracting a partner. There is no presumption that market participants or a central authority know anything about the distribution of others' preferences or that they can deduce such information from prior rounds of play. Instead they follow a process of trial and error in which they adjust their bids and offers in the hope of increasing their payoffs. Such aspiration adjustment rules are rooted in the psychology and learning literature.<sup>1</sup> A key feature of the rule we propose is that an agent's behavior does not require any information about other agents' actions or payoffs: the rule is completely uncoupled.<sup>2</sup>

---

<sup>1</sup>There is an extensive literature in psychology and experimental game theory on trial and error and aspiration adjustment. See in particular the learning models of Thorndike, 1898, Hoppe, 1931, Estes, 1950, Bush and Mosteller, 1955, Herrnstein, 1961, and aspiration adjustment and directional learning dynamics of Heckhausen, 1955, Sauermaun and Selten, 1962, Selten and Stoecker, 1986, Selten, 1998.

<sup>2</sup>This idea was introduced by Foster and Young, 2006 and is a refinement of the concept of uncoupled learning due to Hart and Mas-Colell, 2003; Hart and Mas-Colell, 2006. Recent work has shown that there exist completely uncoupled rules that lead to Nash equilibrium in generic noncooperative games (Germano and Lugosi 2007, Marden et al. 2009, Young 2009, Pradelski and Young 2012).

It is therefore particularly well-suited to environments such as decentralized online markets where players interact anonymously and trades take place at many different prices. We shall show that this simple adaptive process leads to equitable solutions inside the core of the associated *assignment game* (Shapley and Shubik 1972). In particular, core stability and equity are achieved even though agents have no knowledge of the other agents' strategies or preferences, and there is no *ex ante* preference for equity.

The paper is structured as follows. The next section discusses the related literature on matching and core implementation. Section 3 formally introduces assignment games and the relevant solution concepts. Section 4 describes the process of adjustment and search by individual agents. In sections 5 and 6 we show that the stochastically stable states of the process lie in the least core. Section 7 concludes with several open problems.

## 4.2 Related literature

Our results fit into a growing literature showing how cooperative game solutions can be implemented via noncooperative dynamic learning processes (Agastya 1997; Agastya 1999, Arnold and Schwalbe 2002, Newton 2010; Newton 2012, Sawa 2011, Rozen 2013). A particularly interesting class of cooperative games are assignment games, in which every potential matched pair has a cooperative 'value'. Shapley and Shubik, 1972 showed that the core of such a game is always nonempty.<sup>3</sup> Subsequently various authors have explored refinements of the assignment game core, including the kernel (Rochford 1984) and the nucleolus (Huberman 1980, Solymosi and Raghavan 1994, Nunez 2004, Llerena, Nunez, and Rafels 2012). To the best of our knowledge, however,

---

<sup>3</sup>Important subsequent papers include Crawford and Knoer, 1981, Kelso and Crawford, 1982, Demange and Gale, 1985, and Demange, Gale, and Sotomayor, 1986.



there has been no prior work showing how a core refinement is selected via a decentralized learning process, which is the subject of the present paper.

This paper establishes convergence to the core of the assignment game for a class of natural dynamics and selection of a core refinement under payoff perturbations. We are not aware of prior work comparable with our selection result. There are, however, several recent papers that also address the issue of core convergence for a variety of related processes (Chen, Fujishige, and Yang 2011, Biró et al. 2012, Klaus and Payot 2013, Bayati et al. 2014). These processes are different from ours, in particular they are not aspiration-adjustment learning processes, and they do not provide a selection mechanism for a core refinement as we do here. The closest relative to our paper is the concurrent paper by Chen, Fujishige, and Yang, 2011, which demonstrates a decentralized process where, similarly as in our process, pairs of players from the two market sides randomly meet in search of higher payoffs. This process also leads almost surely to solutions in the core. Chen, Fujishige, and Yang, 2011 and our paper are independent and parallel work. They provide a constructive proof based on their process which is similar to ours for the proof of the convergence theorem. Thus, theirs as well as our algorithm (proof of Theorem 1) can be used to find core outcomes. Biró et al., 2012 generalizes Chen, Fujishige, and Yang, 2011 to transferable-utility roommate problems. In contrast to Chen, Fujishige, and Yang, 2011 and our proof, Biró et al., 2012 use a target argument which cannot be implemented to obtain a core outcome. Biró et al., 2012's proof technique is subsequently used in Klaus and Payot, 2013 to prove the result of Chen, Fujishige, and Yang, 2011 for continuous payoff space in the assignment game. A particularity in this case is the fact that the assignment may continue to change as payoffs approximate a core outcome. Finally, Bayati et al., 2014 study the rate of convergence of a related bargaining process for the roommate problem in which players know their best alternatives at each

stage. The main difference of this process to ours is that agents best reply (i.e. they have a lot of information about their best alternatives), the order of activation is fixed, not random, and matches are only formed once a stable outcome is found.<sup>4</sup> An important feature of our learning process is that it is explicitly formulated in terms of random bids of workers and random offers of firms (as in Shapley and Shubik 1972), which allows a completely uncoupled set-up of the dynamic.

There is also a related literature on the marriage problem (Gale and Shapley 1962). In this setting the players have ordinal preferences for being matched with members of the other population, and the core consists of matchings such that no pair would prefer each other to their current partners.<sup>5</sup> Typically, many matchings turn out to be stable. Roth and Vande Vate, 1990 demonstrate a random blocking pair dynamic that leads almost surely to the core in such games. Chung, 2000, Diamantoudi, Xue, and Miyagawa, 2004 and Inarra, Larrea, and Molis, 2008; Inarra, Larrea, and Molis, 2013 establish similar results for nontransferable-utility roommate problems, while Klaus and Klijn, 2007 and Kojima and Uenver, 2008 treat the case of many-to-one and many-to-many nontransferable-utility matchings. Another branch of the literature considers stochastic updating procedures that place high probability on core solutions, that is, the stochastically stable set is contained in the core of the game (Jackson and Watts 2002, Klaus, Klijn, and Walzl 2010, Newton and Sawa 2013).

The key difference between marriage problems and assignment games is that the former are framed in terms of nontransferable (usually ordinal) utility, whereas in the latter each potential match has a transferable ‘value’. The core of the assignment game consists of outcomes such that the matching is optimal

---

<sup>4</sup>In a recent paper Pradelski, 2014 discusses the differences to our set-up in more detail. He then investigates the convergence rate properties of a process closely related to ours.

<sup>5</sup>See Roth and Sotomayor, 1992 for a text on two-sided matching.

and the allocation is pairwise stable. Generically, the optimal matching is unique and the allocations supporting it infinite. On the face of it one might suppose that the known results for marriage games would carry over easily to assignment games but this is not the case. The difficulty is that in marriage games (and roommate games) a payoff-improving deviation is determined by the players' current matches and their preferences, whereas in an assignment game it is determined by their matches, the value created by these matches, and by how they currently split the value of the matches. Thus the core of the assignment game tends to be significantly more constrained and paths to the core are harder to find than in the marriage game.

The contribution of the present paper is to demonstrate a simple completely uncoupled adjustment process that has strong selection properties for assignment games. Using a proof technique introduced by Newton and Sawa, [2013](#) (the one-period deviation principle), we show that the stochastically stable solutions of our process lie in a subset of the core of the assignment game. These solutions have a natural equity interpretation: namely, every pair of matched agents splits the difference between the highest and lowest payoffs they could get without violating the core constraints.

### **4.3 Matching markets with transferable utility**

In this section we shall introduce the conceptual framework for analyzing matching markets with transferable utility; in the next section we introduce the learning process itself.

### 4.3.1 The assignment game

The population  $N = F \cup W$  consists of firms  $F = \{f_1, \dots, f_m\}$  and workers  $W = \{w_1, \dots, w_n\}$ .<sup>6</sup> They interact by making bids and offers to randomly encountered potential partners. We assume matches form only if these bids are mutually profitable for both agents.

**Willingness to pay.** Each firm  $i$  has a *willingness to pay*,  $p_{ij}^+ \geq 0$ , for being matched with worker  $j$ .

**Willingness to accept.** Each worker  $j$  has a *willingness to accept*,  $q_{ij}^- \geq 0$ , for being matched with firm  $i$ .

We assume that these numbers are specific to the agents and are not known to the other market participants or to a central market authority.

**Match value.** Assume that utility is linear and separable in money. The *value* of a match  $(i, j) \in F \times W$  is the potential surplus

$$\alpha_{ij} = (p_{ij}^+ - q_{ij}^-)_+. \quad (4.1)$$

It will be convenient to assume that all values  $p_{ij}^+$ ,  $q_{ij}^-$ , and  $\alpha_{ij}$  can be expressed as multiples of some minimal unit of currency  $\delta$ , for example, “dollars”.

We shall introduce time at this stage to consistently develop our notation. Let  $t = 0, 1, 2, \dots$  be the time periods.

**Assignment.** For all pairs of agents  $(i, j) \in F \times W$ , let  $a_{ij}^t \in \{0, 1\}$ .

$$\text{If } (i, j) \text{ is } \begin{cases} \textit{matched} & \text{then } a_{ij}^t = 1, \\ \textit{unmatched} & \text{then } a_{ij}^t = 0. \end{cases} \quad (4.2)$$

---

<sup>6</sup>The two sides of the market could also, for example, represent buyers and sellers, or men and women in a (monetized) marriage market.

If for a given agent  $i \in N$  there exists  $j$  such that  $a_{ij}^t = 1$  we shall refer to that agent as *matched*; otherwise  $i$  is *single*. An assignment  $\mathbf{A} = (a_{ij}^t)_{i \in F, j \in W}$  is such that if  $a_{ij}^t = 1$  for some  $(i, j)$ , then  $a_{ik}^t = 0$  for all  $k \neq j$  and  $a_{lj}^t = 0$  for all  $l \neq i$ .

**Matching market.** The *matching market* is described by  $[F, W, \alpha, \mathbf{A}]$ :

- $F = \{f_1, \dots, f_m\}$  is the set of  $m$  firms (or men or sellers),
- $W = \{w_1, \dots, w_n\}$  is the set of  $n$  workers (or women or buyers),
- $\alpha = \begin{pmatrix} \alpha_{11} & \dots & \alpha_{1n} \\ \vdots & \alpha_{ij} & \vdots \\ \alpha_{m1} & \dots & \alpha_{mn} \end{pmatrix}$  is the matrix of match values.
- $\mathbf{A} = \begin{pmatrix} a_{11} & \dots & a_{1n} \\ \vdots & a_{ij} & \vdots \\ a_{m1} & \dots & a_{mn} \end{pmatrix}$  is the assignment matrix with 0/1 values and row/column sums at most one.

The set of all possible assignments is denoted by  $\mathcal{A}$ .

Note, that the game at hand is a cooperative game:

**Cooperative assignment game.** Given  $[F, W, \alpha]$ , the *cooperative assignment game*  $G(v, N)$  is defined as follows. Let  $N = F \cup W$  and define  $v : S \subseteq N \rightarrow \mathbb{R}$  such that

- $v(i) = v(\emptyset) = 0$  for all singletons  $i \in N$ ,
- $v(S) = \alpha_{ij}$  for all  $S = (i, j)$  such that  $i \in F$  and  $j \in W$ ,
- $v(S) = \max\{v(i_1, j_1) + \dots + v(i_k, j_k)\}$  for every  $S \subseteq F \times W$ ,

where the maximum is taken over all sets  $\{(i_1, j_1), \dots, (i_k, j_k)\}$  consisting of disjoint pairs that can be formed by matching firms and workers in  $S$ . The number  $v(N)$  specifies the value of an optimal assignment.

### 4.3.2 Dynamic components

**Aspiration level.** At the end of any period  $t$ , a player has an *aspiration level*,  $d_i^t$ , which determines the minimal payoff at which he is willing to be matched. Let  $\mathbf{d}^t = \{d_i^t\}_{i \in F \cup W}$ .

**Bids.** In any period  $t$ , one pair of players is drawn at random and they make bids for each other. We assume that the two players' bids are such that the resulting payoff to each player is at least equal to his aspiration level, and with positive probability is exactly equal to his aspiration level.

Formally, firm  $i \in F$  encounters  $j \in W$  and submits a random bid  $b_i^t = p_{ij}^t$ , where  $p_{ij}^t$  is the maximal amount  $i$  is currently willing to pay if matched with  $j$ . Similarly, worker  $j \in W$  submits  $b_j^t = q_{ij}^t$ , where  $q_{ij}^t$  is the minimal amount  $j$  is currently willing to accept if matched with  $i$ . A bid is separable into two components; the current (deterministic) aspiration level and a random variable that represents an exogenous shock to the agent's aspiration level. Specifically let  $P_{ij}^t, Q_{ij}^t$  be independent random variables that take values in  $\delta \cdot \mathbb{N}_0$  where 0 has positive probability.<sup>7</sup> We thus have for all  $i, j$

$$p_{ij}^t = (p_{ij}^+ - d_i^{t-1}) - P_{ij}^t \quad \text{and} \quad q_{ij}^t = (q_{ij}^- + d_j^{t-1}) + Q_{ij}^t \quad (4.3)$$

Consider, for example, worker  $j$ 's bid for firm  $i$ . The amount  $q_{ij}^-$  is the minimum that  $j$  would ever accept to be matched with  $i$ , while  $d_j^{t-1}$  is his previous aspiration level over and above the minimum. Thus  $Q_{ij}^t$  is  $j$ 's attempt to get even more in the current period. Note that if the random variable is zero, the agent bids exactly according to his current aspiration level.

**Prices.** When  $i$  is matched with  $j$  they trade at a unique *price*,  $\pi_{ij}^t$ .

---

<sup>7</sup>Note that  $\mathbb{P}[P_{ij}^t = 0] > 0$  and  $\mathbb{P}[Q_{ij}^t = 0] > 0$  are reasonable assumptions, since we can adjust  $p_{ij}^+$  and  $q_{ij}^-$  in order for it to hold. This would alter the underlying game but then allow us to proceed as suggested.

**Payoffs.** Given  $[\mathbf{A}^t, \mathbf{d}^t]$  the *payoff* to firm  $i$  / worker  $j$  is

$$\phi_i^t = \begin{cases} p_{ij}^+ - \pi_{ij}^t & \text{if } i \text{ is matched to } j, \\ 0 & \text{if } i \text{ is single.} \end{cases}, \quad \phi_j^t = \begin{cases} \pi_{ij}^t - q_{ij}^- & \text{if } j \text{ is matched to } i, \\ 0 & \text{if } j \text{ is single.} \end{cases} \quad (4.4)$$

Note that, players' payoffs can be deduced from the aspiration levels and the assignment matrix.

**Profitability.** A pair of bids  $(p_{ij}^t, q_{ij}^t)$  is *profitable* if both players, in expectation, receive a higher payoff if the match is formed.

Note that if two players' bids are at their aspiration levels and  $p_{ij}^t = q_{ij}^t$  they are only profitable if both players are currently single. Also note that a pair of players  $(i, j)$  with  $\alpha_{ij} = 0$  will never match.

**Re-match.** At each moment in time, a pair  $(i, j)$  that randomly encounters each other matches if their bids are profitable. The resulting price,  $\pi_{ij}^t$ , is set anywhere between  $q_{ij}^t$  and  $p_{ij}^t$ . (Details about how players are activated are specified in the next section.)

To summarize, when a new match forms that is profitable, both agents receive a higher payoff in expectation due to the full support of the resulting price.<sup>8</sup>

**States.** The *state* at the end of period  $t$  is given by  $Z^t = [\mathbf{A}^t, \mathbf{d}^t]$  where  $\mathbf{A}^t \in \mathcal{A}$  is an assignment and  $\mathbf{d}^t$  is the aspiration level vector. Denote the set of all states by  $\Omega$ .

### 4.3.3 Solution concepts

**Optimality.** An assignment  $\mathbf{A}$  is *optimal* if  $\sum_{(i,j) \in F \times W} a_{ij} \cdot \alpha_{ij} = v(N)$ .

---

<sup>8</sup>In this sense any alternative match that may block a current assignment because it is profitable (as defined earlier) is a strict blocking pair.

**Pairwise stability.** An aspiration level vector  $\mathbf{d}^t$  is *pairwise stable* if  $\forall i, j$  and  $a_{ij}^t = 1$ ,

$$p_{ij}^+ - d_i^t = q_{ij}^- + d_j^t, \quad (4.5)$$

and  $p_{i'j}^+ - d_{i'}^t \leq q_{i'j}^- + d_j^t$  for every alternative firm  $i' \in F$  with  $i' \neq i$  and  $q_{ij'}^- + d_{j'}^t \geq p_{ij'}^+ - d_i^t$  for every alternative worker  $j' \in W$  with  $j' \neq j$ .

**Core (Shapley and Shubik 1972).** The *core* of any assignment game is always non-empty and consists of the set  $\mathbf{C} \subseteq \Omega$  of all states  $Z$  such that  $\mathbf{A}$  is an optimal assignment and  $\mathbf{d}$  is pairwise stable.

Subsequent literature has investigated the structure of the assignment game core, which turns out to be very rich.<sup>9</sup> In order to investigate the constraints of pairwise stability in more detail the concept of ‘payoff excess’ will be useful:

**Excess.** Given state  $Z^t$ , the *excess* for a player  $i$  who is matched with  $j$  is

$$e_i^t = \phi_i^t - \max_{k \neq j} (\alpha_{ik} - \phi_k^t)_+. \quad (4.6)$$

The excess for player  $i$  describes the gap to his next-best alternative, that is, the smallest amount he would have to give up in order to profitably match with some other player  $k \neq j$ . If a player has negative excess, pairwise stability is violated. In a core allocation, therefore, all players have nonnegative excess. For the analysis of absorbing core states, note that the excess in payoff can be equivalently expressed in terms of the excess in aspiration level. This is the case since in absorbing core states aspiration levels are directly deducible from payoffs.

---

<sup>9</sup>See, for example, Roth and Sotomayor, 1992, Balinski and Gale, 1987, Sotomayor, 2003.



**Minimal excess.** Given state  $Z^t$ , the *minimal excess* is

$$e_{\min}^t(Z^t) = \min_{i: i \text{ matched}} e_i^t. \quad (4.7)$$

Based on the minimal excess of a state, we can define the kernel (Davis and Maschler 1965). For assignment games, the kernel coincides with the solution concept proposed by Rochford, 1984, which generalizes a pairwise equal split solution à la Nash, 1950.

**Kernel (Davis and Maschler 1965, Rochford 1984).** The *kernel*  $\mathbf{K}$  of an assignment game is the set of states such that the matching is optimal and for all matched pairs  $(i, j)$ ,

$$e_i^t =_{\delta} e_j^t, \quad (4.8)$$

where  $=_{\delta}$  means “equality up to  $\delta$ ”. (This is necessary given that we operate on the discrete grid.)

Given  $Z^t$ , extend the definition of excess to any coalition  $S \subseteq N$ ;  $e^t(S) = \sum_{i \in S} \phi_i^t - v(S)$ . Now let  $\mathbf{E}(\phi^t) \in \mathbb{R}^{m+n}$  be the vector of excesses for all  $S \subseteq N$ , ordered from smallest to largest. Say  $\mathbf{E}(\phi)$  is lexicographically larger than  $\mathbf{E}(\phi')$  for some  $k$ , if  $\mathbf{E}_i(\phi) = \mathbf{E}_i(\phi')$  for all  $i < k$  and  $\mathbf{E}_k(\phi) < \mathbf{E}_k(\phi')$ .<sup>10</sup>

**Nucleolus (Schmeidler 1969).** The *nucleolus*  $\mathbf{N}$  of the assignment game is the unique solution that minimizes the lexicographic measure. (See also Huberman 1980, Solymosi and Raghavan 1994.)

For an analysis of the welfare properties and of the links between the kernel and the nucleolus of the assignment game see Nunez, 2004 and Llerena, Nunez, and Rafels, 2012.

**Least core (Maschler, Peleg, and Shapley 1979).** The *least core*  $\mathbf{L}$  of an

---

<sup>10</sup>Note that the excess for coalitions,  $e^t(S)$ , is usually defined with a reversed sign. In order to make it more concurrent in light of definition (4.6) we chose to reverse the sign.

assignment game is the set of states  $Z$  such that the matching is optimal and the minimum excess is maximized, that is,

$$e_{\min}(Z) = \max_{Z' \in \mathbf{C}} e_{\min}(Z'). \quad (4.9)$$

Note that our definition of excess applies to essential coalitions only (that is, for the case of the assignment game, to two-player coalitions involving exactly one agent from each market side). Hence, the least core generalizes the nucleolus of the assignment game in the following sense. Starting with the nucleolus, select any player with minimum excess (according to equation (4.6)): the least core contains all outcomes with a minimum excess that is not smaller.<sup>11</sup>

The following inclusions are known for the assignment game:<sup>12</sup>

$$\mathbf{N} \in (\mathbf{K} \cap \mathbf{L}), \quad \mathbf{K} \subseteq \mathbf{C}, \quad \mathbf{L} \subseteq \mathbf{C}. \quad (4.10)$$

## 4.4 Evolving play

A fixed population of agents,  $N = F \cup W$ , plays the assignment game  $G(v, N)$ . Repeatedly, a randomly activated agent encounters another agent, they make bids for each other and match if profitable. The distinct times at which one agent becomes active will be called *periods*. Agents are activated by independent Poisson clocks.<sup>13</sup> Suppose that an active agent randomly encounters one agent from the other side of the market drawn from a distribution with full

---

<sup>11</sup>See Shapley and Shubik, 1963; Shapley and Shubik, 1966 for the underlying idea of the least core, the *strong  $\epsilon$ -core*. See Maschler, Peleg, and Shapley, 1979, Driessen, 1999, Llerena and Nunez, 2011 for geometric interpretations of these concepts.

<sup>12</sup> $\mathbf{N} \in \mathbf{K}$  is shown by Schmeidler, 1969 for general cooperative games. Similarly  $\mathbf{N} \in \mathbf{L}$  is shown by Maschler, Peleg, and Shapley, 1979. Driessen, 1998 shows for the assignment game that  $\mathbf{K} \subseteq \mathbf{C}$ .  $\mathbf{L} \subseteq \mathbf{C}$  follows directly from the definitions.

<sup>13</sup>The Poisson clocks' arrival rates may depend on the agents' themselves or on their position in the game. Single agents, for example, may be activated faster than matched agents.

support. The two players enter a new match if their match is profitable, which they can see from their current bids, offers and their payoffs. If the two players are already matched with each other, they remain so.

#### 4.4.1 Behavioral dynamics

The essential steps and features of the learning process are as follows. At the start of period  $t + 1$ :

1. The activated agent,  $i$ , makes a random encounter,  $j$ .
- 2a. If the encounter is profitable given their current bids and assignment, the pair matches.
- 2b. If the match is not profitable, both agents return to their previous matches (or remain single).
- 3a. If a new match  $(i, j)$  forms, the price is set anywhere between bid and offer. The aspiration levels of  $i$  and  $j$  are set to equal their realized payoffs.
- 3b. If no new match is formed, the active agent, if he was previously matched, keeps his previous aspiration level and stays with his previous partner. If he was previously single, he remains single and lowers his aspiration level with positive probability.

Our rules have antecedents in the psychology literature (Thorndike 1898, Hoppe 1931, Estes 1950, Bush and Mosteller 1955, Herrnstein 1961). To the best of our knowledge, however, such a framework has not previously been used in the study of matching markets in cooperative games. The approach seems especially well-suited to modeling behavior in large decentralized assignment markets, where agents have little information about the overall game

and about the identity of the other market participants. Following aspiration adjustment theory (Sauermann and Selten 1962, Selten 1998) and related bargaining experiments on directional and reinforcement learning (e.g., Tietz and Weber 1972, Roth and Erev 1995), we shall assume a simple directional learning model: matched agents occasionally experiment with higher offers if on the sell-side (or lower bids if on the buy-side), while single agents, in the hope of attracting partners, lower their offers if on the sell-side (or increase their bids if on the buy-side).

We shall now describe the process in more detail, distinguishing the cases where the active agent is currently *matched* or *single*. Let  $Z^t$  be the state at the end of period  $t$  (and the beginning of period  $t + 1$ ), and let  $i \in F$  be the unique *active* agent which for ease of exposition we assume to be a firm.

*I. The active agent is currently matched and meets  $j$*

If  $i, j$  are profitable (given their current aspiration levels) they match. As a result,  $i$ 's former partner is now single (and so is  $j$ 's former partner if  $j$  was matched in period  $t$ ). The price governing the new match,  $\pi_{ij}^{t+1}$ , is randomly set between  $p_{ij}^{t+1}$  and  $q_{ij}^{t+1}$ .

At the end of period  $t + 1$ , the aspiration levels of the newly matched pair  $(i, j)$  are adjusted according to their newly realized payoffs:

$$d_i^{t+1} = p_{ij}^+ - \pi_{ij}^{t+1} \quad \text{and} \quad d_j^{t+1} = \pi_{ij}^{t+1} - q_{ij}^- \quad (4.11)$$

All other aspiration levels and matches remain fixed. If  $i, j$  are not profitable,  $i$  remains matched with his previous partner and keeps his previous aspiration level. See Figure 4.1 for an illustration.

II. *The active agent is currently single and meets  $j$*

If  $i, j$  are profitable (given their current aspiration levels) they match. As a result,  $j$ 's former partner is now single if  $j$  was matched in period  $t$ . The price governing the new match,  $\pi_{ij}^{t+1}$ , is randomly set between  $p_{ij}^{t+1}$  and  $q_{ij}^{t+1}$ .

At the end of period  $t+1$ , the aspiration levels of the newly matched pair  $(i, j)$  are adjusted to equal their newly realized payoffs:

$$d_i^{t+1} = p_{ij}^+ - \pi_{ij}^{t+1} \quad \text{and} \quad d_j^{t+1} = \pi_{ij}^{t+1} - q_{ij}^-. \quad (4.12)$$

All other aspiration levels and matches remain as before. If  $i, j$  are not profitable,  $i$  remains single and, with positive probability, reduces his aspiration level,

$$d_i^{t+1} = (d_i^t - X_i^{t+1})_+, \quad (4.13)$$


where  $X_i^{t+1}$  is an independent random variable taking values in  $\delta \cdot \mathbb{N}_0$  and  $\delta$  occurs with positive probability.<sup>14</sup> See Figure 4.2 for an illustration.

#### 4.4.2 Example

Let  $N = F \cup W = \{f_1, f_2\} \cup \{w_1, w_2, w_3\}$ ,  $p_{1j}^+ = (40, 31, 20)$  and  $p_{2j}^+ = (20, 31, 40)$  for  $j = 1, 2, 3$ , and  $q_{i1}^- = (20, 30)$ ,  $q_{i2}^- = (20, 20)$  and  $q_{i3}^- = (30, 20)$  for  $i = 1, 2$ .

---

<sup>14</sup>Note that  $X_i^{t+1}$  may depend on time.




Then one can compute the match values:  $\alpha_{11} = \alpha_{23} = 20$ ,  $\alpha_{12} = \alpha_{22} = 11$ , and  $\alpha_{ij} = 0$  for all other pairs  $(i, j)$ . Let  $\delta = 1$ .

**period  $t$ :**     *Current state*

Suppose that, at the end of some period  $t$ ,  $(f_1, w_1)$  and  $(f_2, w_2)$  are matched and  $w_3$  is single.

The current aspiration level is shown next to the name of that agent, and the values  $\alpha_{ij}$  are shown next to the edges (if positive). Bids will be shown to the right of the aspiration level. Solid edges indicate matched pairs, and dashed edges indicate unmatched pairs. (Edges with value zero are not shown.) Note that no player can see the bids or the status of the players on the other side of the market.


Note that some matches can never occur. For example  $f_1$  is never willing to pay more than 20 for  $w_3$ , but  $w_3$  would only accept a price above 30 from  $f_1$ .



Note that the aspiration levels satisfy  $d_i^t + d_j^t \geq \alpha_{ij}$  for all  $i$  and  $j$ , but the assignment is not optimal (firm 2 should match with worker 3).


**period  $t + 1$ :** *Activation of single agent  $w_3$  and encounter of  $f_2$*

$w_3$ 's current aspiration level is too high in order to be profitable with  $f_2$ . Hence, independent of the specific bid he makes, he remains single and, with positive probability, reduces his aspiration level by 1.




**period  $t + 2$ :** *Activation of matched agent  $f_2$  and encounter of  $w_3$*

$f_2$  and  $w_3$  are profitable. With positive probability  $f_2$  bids 30 for  $w_3$  and  $w_3$  bids 29 for  $f_2$  (hence the match is profitable), and the match forms. The price is set at random to either 29 such that  $f_2$  raises his aspiration level by one unit (11) and  $w_3$  keeps his aspiration level (9), or to 30 such that  $f_2$  keeps his aspiration level (10) and  $w_3$  raises his aspiration level by one unit (10). (Thus in expectation the agents get a higher payoff than before.)




**period  $t + 3$ :**      *Activation of single agent  $w_2$  and encounter of  $f_2$*

$w_2$ 's current aspiration level is too high in the sense that he has no profitable matches and thus in particular is not profitable with  $f_2$ . Hence he remains single and, with positive probability, reduces his aspiration level by 1.



The resulting state is in the core:<sup>15</sup>



<sup>15</sup>Note that the states  $Z^{t+2}$  and  $Z^{t+3}$  are both in the core, but  $Z^{t+3}$  is absorbing whereas  $Z^{t+2}$  is not.



## 4.5 Core stability – absorbing states of the unperturbed process

Recall that a state  $Z^t$  is defined by an assignment  $\mathbf{A}^t$  and aspiration levels  $\mathbf{d}^t$  that jointly determine the payoffs.  $\mathbf{C}$  is the set of core states; let  $\mathbf{C}_0$  be the set of core states such that singles' aspiration levels are zero.

**Theorem 1.** *Given an assignment game  $G(v, N)$ , from any initial state  $Z^0 = [\mathbf{A}^0, \mathbf{d}^0] \in \Omega$ , the process is absorbed into the core in finite time with probability 1. The set of absorbing states consists of  $\mathbf{C}_0$ . Further, starting from  $\mathbf{d}^0 = 0$  any absorbing state is attainable.*

Throughout the proof we shall omit the time superscript since the process is time-homogeneous. The general idea of the proof is to show a particular path leading into the core which has positive probability. The proof uses integer programming arguments (Kuhn 1955, Balinski 1965) but no single authority ‘solves’ an integer programming problem. It will simplify the argument to restrict our attention to a particular class of paths with the property that the realizations of the random variables  $P_{ij}^t, Q_{ij}^t$  are always 0 and the realizations of  $X_i^t$  are always  $\delta$ .  $P_{ij}^t, Q_{ij}^t$  determine the gaps between the bids and the aspiration levels, and  $X_i^t$  determines the reduction of the aspiration level by a single agent. One obtains from equation (4.3) for the bids:

$$\text{for all } i, j, \quad p_{ij}^t = p_{ij}^+ - d_i^{t-1} \quad \text{and} \quad q_{ij}^t = q_{ij}^- + d_j^{t-1} \quad (4.14)$$

Recall that any two agents encounter each other in any period with positive probability. It shall be understood in the proof that the relevant agents in any period encounter each other. Jointly with equation (4.3), we can then say that

a pair of aspiration levels  $(d_i^t, d_j^t)$  is *profitable* if

$$\text{either } d_i^t + d_j^t < \alpha_{ij} \quad \text{or} \quad d_i^t + d_j^t = \alpha_{ij} \text{ and both } i \text{ and } j \text{ are single.} \quad (4.15)$$

Restricting attention to this particular class of paths will permit a more transparent analysis of the transitions, which we can describe solely in terms of the aspiration levels.

We shall proceed by establishing the following two claims.

*Claim 1.* There is a positive probability path to aspiration levels  $\mathbf{d}$  such that  $d_i + d_j \geq \alpha_{ij}$  for all  $i, j$  and such that, for every  $i$ , either there exists a  $j$  such that  $d_i + d_j = \alpha_{ij}$  or else  $d_i = 0$ .

Any aspiration levels satisfying Claim 1 will be called *good*. Note that, even if aspiration levels are good, the assignment does not need to be optimal and not every agent with a positive aspiration level needs to be matched. (See the period- $t$  example in the preceding section.)

*Claim 2.* Starting at any state with good aspiration levels, there is a positive probability path to a pair  $(\mathbf{A}, \mathbf{d})$  where  $\mathbf{d}$  is good,  $\mathbf{A}$  is optimal, and all singles' aspiration levels are zero.<sup>16</sup>

*Proof of Claim 1.*

*Case 1.* Suppose the aspiration levels  $\mathbf{d}$  are such that  $d_i + d_j < \alpha_{ij}$  for some  $i, j$ . Note that this implies that  $i$  and  $j$  are not matched with each other since otherwise the entire surplus is allocated and  $d_i + d_j = \alpha_{ij}$ . With positive probability, either  $i$  or  $j$  is activated and  $i$  and  $j$  become matched. The new aspiration levels are set equal to the new payoffs. Thus the sum of the

---

<sup>16</sup>Note that this claim describes an absorbing state in the core. It may well be that the core is reached while a single's aspiration level is more than zero. The latter state, however, is transient and will converge to the corresponding absorbing state.

aspiration levels is equal to the match value  $\alpha_{ij}$ . Therefore, there is a positive probability path along which  $\mathbf{d}$  increases monotonically until  $d_i + d_j \geq \alpha_{ij}$  for all  $i, j$ .

*Case 2.* Suppose the aspiration levels  $\mathbf{d}$  are such that  $d_i + d_j \geq \alpha_{ij}$  for all  $i, j$ .

We can suppose that there exists a single agent  $i$  with  $d_i > 0$  and  $d_i + d_j > \alpha_{ij}$  for all  $j$ , else we are done. With positive probability,  $i$  is activated. Since no profitable match exists, he lowers his aspiration level by  $\delta$ . In this manner, a suitable path can be constructed along which  $\mathbf{d}$  decreases monotonically until the aspiration levels are good. Note that at the end of such a path, the assignment does not need to be optimal and not every agent with a positive aspiration level needs to be matched. (See the period- $t$  example in the preceding section.) □

*Proof of Claim 2.*

Suppose that the state  $(\mathbf{A}, \mathbf{d})$  satisfies Claim 1 ( $\mathbf{d}$  is good) and that some single exists whose aspiration level is positive. (If no such single exists, the assignment is optimal and we have reached a core state.) Starting at any such state, we show that, within a bounded number of periods and with positive probability (bounded below), one of the following holds:

The aspiration levels are good, the number of single agents with positive aspiration level decreases, and the sum of the aspiration levels remains constant.

(4.16)

The aspiration levels are good, the sum of the aspiration levels decreases by  $\delta > 0$ , and the number of single agents with a positive aspiration level does not increase.

(4.17)

In general, say an edge is *tight* if  $d_i + d_j = \alpha_{ij}$  and *loose* if  $d_i + d_j = \alpha_{ij} - \delta$ . Define a *maximal alternating path*  $P$  to be a path that starts at a single player with positive aspiration level, and that alternates between unmatched tight edges and matched tight edges such that it can not be extended (hence maximal). Note that, for every single with a positive aspiration level, at least one maximal alternating path exists. Figure 3 (left panel) illustrates a maximal alternating path starting at  $f_1$ . Unmatched tight edges are indicated by dashed lines, matched tight edges by solid lines and loose edges by dotted lines.

Without loss of generality, let  $f_1$  be a single firm with positive aspiration level.

*Case 1.* Starting at  $f_1$ , there exists a maximal alternating path  $P$  of odd length.

*Case 1a.* All firms on the path have a positive aspiration level.

We shall demonstrate a sequence of adjustments leading to a state as in (16).


Let  $P = (f_1, w_1, f_2, w_2, \dots, w_{k-1}, f_k, w_k)$ . Note that, since the path is maximal and of odd length,  $w_k$  must be single. With positive probability,  $f_1$  is activated. Since no profitable match exists, he lowers his aspiration level by  $\delta$ . With positive probability,  $f_1$  is activated again next period, he snags  $w_1$  and with positive probability he receives the residual  $\delta$ . At this point the aspiration levels are unchanged but  $f_2$  is now single. With positive probability,  $f_2$  is activated. Since no profitable match exists, he lowers his aspiration level by  $\delta$ . With positive probability,  $f_2$  is activated again next period, he snags  $w_2$  and with positive probability he receives the residual  $\delta$ . Within a finite number

of periods a state is reached where all players on  $P$  are matched and the aspiration levels are as before. (Note that  $f_k$  is matched with  $w_k$  without a previous reduction by  $f_k$  since  $w_k$  is single and thus their bids are profitable.)

*In summary, the number of matched agents has increased by two and the number of single agents with positive aspiration level has decreased by at least one.*

*The aspiration levels did not change, hence they are still good.*

*Transition diagram for Case 1a.*



*Case 1b.* At least one firm on the path has aspiration level zero.


We shall demonstrate a sequence of adjustments leading to a state as in (16).

Let  $P = (f_1, w_1, f_2, w_2, \dots, w_{k-1}, f_k, w_k)$ . There exists a firm  $f_i \in P$  with current aspiration level zero ( $f_2$  in the illustration), hence no further reduction by  $f_i$  can occur. (If multiple firms on  $P$  have aspiration level zero, let  $f_i$  be the first such firm on the path.) Apply the same sequence of transitions as in Case 1a up to firm  $f_i$ . At the end of this sequence the aspiration levels are as before. Once  $f_{i-1}$  snags  $w_{i-1}$ ,  $f_i$  becomes single and his aspiration level is still zero.

*In summary, the number of single agents with a positive aspiration level has decreased by one because  $f_1$  is no longer single and the new single agent  $f_i$  has aspiration level zero. The aspiration levels did not change, hence they are still*

good.

Transition diagram for Case 1b.



Case 2. Starting at  $f_1$ , all maximal alternating paths are of even length.

Case 2a. All firms on all maximal alternating paths starting at  $f_1$  have a positive aspiration level.

We shall demonstrate a sequence of adjustments leading to a state as in (17). Note that, given aspiration levels will have changed by the end of the sequence of transitions, it does not suffice to only consider players along one maximal alternating path. Instead, we need to consider all alternating paths starting at  $f_1$ .

With positive probability  $f_1$  is activated. Since no profitable match exists, he lowers his aspiration level by  $\delta$ . Hence, all previously tight edges starting at  $f_1$  are now loose.

We shall describe a sequence of transitions under which a given loose edge is eliminated (by making it tight again), the matching does not change and the sum of aspiration levels remains fixed. Consider a loose edge between a firm, say  $f'_1$ , and a worker, say  $w'_1$ . Since all maximal alternating paths starting at  $f_1$  are of even length, the worker has to be matched to a firm, say  $f'_2$ . With positive probability  $w'_1$  is activated, snags  $f'_1$ , and with positive probability

$f'_1$  receives the residual  $\delta$ . (Such a transition occurs with strictly positive probability whether or not  $f'_1$  is matched because aspiration levels are strictly below the match value of  $(w'_1, f'_1)$ .) Note that  $f'_2$  and possibly  $f'_1$ 's previous partner, say  $w''_1$ , are now single. With positive probability  $f'_2$  is activated. Since no profitable match exists, he lowers his aspiration level by  $\delta$ . (This occurs because all firms on any maximal alternating path starting at  $f_1$  have an aspiration level at least  $\delta$ .) With positive probability,  $f'_2$  is activated again, snags  $w'_1$ , and with positive probability  $w'_1$  receives the residual  $\delta$ . Finally, with positive probability  $f'_1$  is activated. Since no profitable match exists, he lowers his aspiration level by  $\delta$ . If previously matched,  $f'_1$  is activated again in the next period and matches with the single  $w''_1$  (note that there is no additional surplus to be split). At the end of this sequence the matching is the same as at the beginning. Moreover,  $w'_1$ 's aspiration level went up by  $\delta$  while  $f'_2$ 's aspiration level went down by  $\delta$  and all other aspiration levels stayed the same. The originally loose edge between  $f'_1$  and  $w'_1$  is now tight.


We iterate the latter construction for  $f'_1 = f_1$  until all loose edges at  $f'_1$  have been eliminated. However, given  $f'_2$ 's reduction by  $\delta$  there may be new loose edges connecting  $f'_2$  to workers (possibly on several alternating paths). In this case we repeat the preceding construction for  $f'_2$  until all of the loose edges at  $f'_2$  have been eliminated. If any agents still exist with loose edges we repeat the construction again. This iteration eventually terminates given the following observation. Any worker on a maximal alternating path who previously increased his aspiration level cannot still be connected to a firm by a loose edge. Similarly, any firm that previously reduced its aspiration level cannot now be matched to a worker with a loose edge because such a worker increased his aspiration level. Therefore the preceding construction involves any given firm (or worker) at most once. It follows that, in a finite number of periods, all firms on any maximal alternating paths starting at  $f_1$  have reduced

their aspiration level by  $\delta$  and all workers have increased their aspiration level by  $\delta$ . (Again, note that it is necessary to use this construction on *all* maximal alternating paths starting at  $f_1$ .)

*In summary, the number of aspiration level reductions outnumbers the number of aspiration level increases by one (namely by the firm  $f_1$ ), hence the sum of the aspiration levels has decreased. The number of single agents with a positive aspiration level has not increased. Moreover the aspiration levels are still good.*

Note that the  $\delta$ -reductions may lead to new tight edges, resulting in new maximal alternating paths of odd or even lengths.

*Transition diagram for Case 2a.*



*Case 2b.* At least one firm on a maximal alternating paths starting at  $f_1$  has aspiration level zero.


We shall demonstrate a sequence of adjustments leading to a state as in (16).



Let  $P = (f_1, w_1, f_2, w_2, \dots, w_{k-1}, f_k)$  be a maximal alternating path such that a firm has aspiration level zero. There exists a firm  $f_i \in P$  with current aspiration level zero ( $f_2$  in the illustration), hence no further reduction by  $f_i$  can occur. (If multiple firms on  $P$  have aspiration level zero, let  $f_i$  be the first such firm on the path.) With positive probability  $f_1$  is activated. Since no profitable match exists, he lowers his aspiration level by  $\delta$ . With positive probability,  $f_1$  is activated again next period, he snags  $w_1$  and with positive probability he receives the residual  $\delta$ . Now  $f_2$  is single. With positive probability  $f_2$  is activated, lowers, snags  $w_2$ , and so forth. This sequence continues until  $f_i$  is reached, who is now single with aspiration level zero.

*In summary, the number of single agents with a positive aspiration level has decreased. The aspiration levels did not change, hence they are still good.*

Transition diagram for Case 2b.



Let us summarize the argument. Starting in a state  $[\mathbf{A}, \mathbf{d}]$  with good aspiration levels  $\mathbf{d}$ , we successively (if any exist) eliminate the odd paths starting at firms/workers followed by the even paths starting at firms/workers, while maintaining good aspiration levels. This process must come to an end because at each iteration either the sum of aspiration levels decreases by  $\delta$  and the number of single agents with positive aspiration levels stays fixed, or the sum

of aspiration levels stays fixed and the number of single agents with positive aspiration levels decreases. The resulting state must be in the core and is absorbing because single agents cannot reduce their aspiration level further and no new matches can be formed. Since an aspiration level constitutes a lower bound on a player's bids we can conclude that the process  $Z^t$  is absorbed into the core in finite time with probability 1. Finally note that, starting from  $\mathbf{d}^0 = 0$  we can trivially reach any state in  $\mathbf{C}_0$ .  $\square$

## 4.6 Core selection

In this section, we investigate the effects of random perturbations to the adjustment process. Suppose that players occasionally experience shocks when in a match and that larger shocks are less likely than smaller shocks. The effect of such a shock is that a player receives more or less payoff than anticipated given the current price he agreed to with his partner. We shall formalize these perturbations and investigate the resulting selection of stochastically stable states as the probability of shocks becomes vanishingly small (Foster and Young 1990, Kandori, Mailath, and Rob 1993, Young 1993). It turns out that the set of stochastically stable states is contained in the *least core*; moreover there are natural conditions under which it coincides with the least core.

Given a player  $i$  who is matched in period  $t$ . Suppose his unperturbed payoff  $\phi_i^t$  is subject to a shock. Denote the new payoff by  $\hat{\phi}_i^t$  and define:

$$\hat{\phi}_i^t = \begin{cases} \phi_i^t + \delta \cdot R_i^t & \text{with probability } 0.5, \\ \phi_i^t - \delta \cdot R_i^t & \text{with probability } 0.5, \end{cases} \quad (4.18)$$

where  $R_i^t$  is an independent geometric random variable with  $\mathbb{P}[R_i^t = k] =$

$\epsilon^k \cdot (1 - \epsilon)$  for all  $k \in \mathbb{N}_0$ .<sup>17</sup> Note that for  $\epsilon = 0$  the process is unperturbed.

The immediate result of a given shock is that players receive a different payoff than anticipated. We shall assume that players update their aspiration levels to their new perturbed payoff if positive and zero if negative. If, in a given match, one of the players experiences a negative payoff the match breaks and both players become single. Note that if the partnership remains matched the price does not change.

#### 4.6.1 Stochastic stability

We are interested in the long-run behavior of the process when  $\epsilon$  becomes small. We shall employ the concept of *stochastic stability* developed by Foster and Young, 1990, Kandori, Mailath, and Rob, 1993 and Young, 1993. In particular, we shall conduct the analysis along the lines of Newton and Sawa, 2013 who introduce the notion of ‘one-shot stability’. Note that the perturbed process is ergodic for  $\epsilon > 0$  and thus has a unique stationary distribution, say  $\Pi_\epsilon$  over the state space  $\Omega$ . We are thus interested in  $\lim_{\epsilon \rightarrow 0} \Pi_\epsilon = \Pi_0$ .

**Stochastic stability.** A state  $Z \in \Omega$  is *stochastically stable* if  $\Pi_0(Z) > 0$ . Denote the set of stochastically stable states by  $\mathbf{S}$ .

For a given parameter  $\epsilon$  denote the probability of transiting from  $Z$  to  $Z'$  in  $k$  periods by  $\mathbb{P}_\epsilon^k[Z, Z']$ . The *resistance* of a one-period transition  $Z \rightarrow Z'$  is the unique real number  $r(Z, Z') \geq 0$  such that  $0 < \lim_{\epsilon \rightarrow 0} \mathbb{P}_\epsilon^1[Z, Z'] / \epsilon^{r(Z, Z')} < \infty$ . For completeness let  $r(Z, Z') = \infty$  if  $\mathbb{P}_\epsilon^1[Z, Z'] = 0$ . Hence a transition with resistance  $r$  has probability of the order  $O(\epsilon^r)$ . We shall call a transition (possibly in multiple periods)  $Z \rightarrow Z'$  a *least cost transition* if it exhibits the lowest order of resistance, that is, let  $Z = Z_0, Z_1, \dots, Z_k = Z'$  ( $k$  finite) be

<sup>17</sup>For simplicity we propose this specific distribution. But note that, any probability distribution can be assumed as long as there exists a parameter  $\epsilon$  such that  $\mathbb{P}[x + 1] = \epsilon \cdot O(\mathbb{P}[x])$  for all  $x \in \mathbb{N}_0$ .

a path of one-period transitions from  $Z$  to  $Z'$ . Then a least cost transition minimizes  $\sum_{l=0}^{k-1} r(Z_l, Z_{l+1})$  over all such paths. For a core state  $Z \in \mathbf{C}$  we shall say that a transition out of the core is a *least cost transition* if it minimizes the resistance among all transitions from  $Z$  to any non-core state.

Young, 1993 shows that the computation of the stochastically stable states can be reduced to an analysis of rooted trees on the set of recurrent classes of the unperturbed dynamic. Define the *resistance* between two recurrent classes  $Z$  and  $Z'$ ,  $r(Z, Z')$  to be the sum of resistances of a least cost transition that starts in  $Z$  and ends in  $Z'$ . Now identify the recurrent classes with the nodes of a graph. Given a node  $Z$ , a collection of directed edges  $T$  forms a *Z-tree* if from every node  $Z' \neq Z$  there exists a unique outgoing edge in  $T$ ,  $Z$  has no outgoing edge, and the graph has no cycles.

**Stochastic potential.** The resistance  $r(T)$  of a  $Z$ -tree  $T$  is the sum of the resistances of its edges. The *stochastic potential* of  $Z$ ,  $\rho(Z)$ , is given by

$$\rho(Z) = \min\{r(T) : T \text{ is a } Z\text{-tree}\}. \quad (4.19)$$

Theorem 4 in Young, 1993 states that the stochastically stable states are precisely those states where  $\rho$  is minimized.

## 4.6.2 Analysis

With this machinery at hand we shall show that the stochastically stable states are contained in the least core. To establish this result we shall adapt a proof technique due to Newton and Sawa, 2013 and show that the least core is the set of states which is most stable against one-shot deviations. We shall also provide conditions on the game under which the stochastically stable set is identical with the least core.

Recall that the least core consists of states that maximize the following term:

$$e_{\min}^t = \min_{i: i \text{ matched}} \{ \phi_i^t - \max_{j: a_{ij}^t=0} (\alpha_{ij} - \phi_j^t)_+ \} \quad (4.20)$$

$$= \min \left\{ \underbrace{\min_{i,j: a_{ij}^t=0, i \text{ matched}} (\phi_i^t + \phi_j^t - \alpha_{ij})}_{=: \text{case A}} ; \underbrace{\min_{i: i \text{ matched}} \phi_i^t}_{=: \text{case B}} \right\} \quad (4.21)$$

Case *A* holds if the minimal cost deviation is such that two players who are currently not matched experience shocks such that they become profitable. Case *B*, on the other hand, is the case where the minimal cost deviation is such that a matched agent experiences a shock that leads to a negative payoff and thus to him breaking up his relationship.

Given two states  $Z$  and  $Z^*$ , let the *distance* between them be

$$D(Z, Z^*) = \sum_{i \in F \cup W} |\phi_i - \phi_i^*|. \quad (4.22)$$

**Lemma 2.** *Given  $Z^* \in \mathbf{L}$  and  $Z \in \mathbf{C} \setminus \mathbf{L}$ . Let  $Z'$  be a state not in the core which is reachable from  $Z$  by a least cost transition. Then there exists  $Z_1 \in \mathbf{C}$  such that  $D(Z^*, Z_1) < D(Z^*, Z)$  and  $\mathbb{P}_0^t[Z', Z_1] > 0$  for some  $t \geq 0$ .*

*Proof.* By theorem 1, the recurrent classes consist of all singleton states in  $\mathbf{C}_0 \subseteq \mathbf{C}$ . Thus it suffices to limit our analysis to  $Z^* \in \mathbf{L} \cap \mathbf{C}_0$  and  $Z \in \mathbf{C}_0 \setminus \mathbf{L}$  since other core states have zero-resistance paths to the states in  $\mathbf{C}_0$ .

*Case A.* Suppose that the least-cost transition to a non-core state is such that two (currently not matched) players experience trembles such that their match becomes profitable. That is, there exists  $i$ , matched to  $j$ , and a nonempty set  $J'$  such that  $i, j'$  is least costly to destabilize for any  $j' \in J'$ . Note that  $d_i + d_{j'} - \alpha_{ij'}$  is minimal for all  $j' \in J'$  and thus constant and also the difference is non-negative since we are in a core state.

*Case A.1.*  $d_i > d_i^*$ .

*Case A.1a.* For all  $j' \in J'$ ,  $d_i + d_{j'} > \alpha_{ij'}$ .

We can construct a sequence of transitions such that  $i$  reduces his aspiration level by  $\delta$ ,  $j$  increases his aspiration level by  $\delta$  (note that we have  $d_j < d_j^*$ ), and all other aspiration levels stay the same. Note that  $D$  then decreased and the resulting state is again a core state given that for all  $j' \neq j$  we started out with  $d_i + d_{j'} > \alpha_{ij'}$ .

Now we shall explain the sequence in detail. Suppose the tremble occurs such that  $i$  reduces his aspiration level by at least  $\delta$  and  $i$  and  $j'$  match at a price such that  $i$ 's aspiration level does not increase. Consequently  $j$  and  $i'$  ( $j'$ 's former partner, if he is matched in the core assignment) are now single. In the following period  $i$  and  $j$  are profitable. With positive probability, they match at a price such that  $d_i$  decreases by  $\delta$ . Now  $i'$  and  $j'$  are both single. With positive probability they reduce their aspiration levels and rematch at their previous price, returning to their original aspiration levels. Thus, with positive probability the prices are set such that  $d_i$  decreases by  $\delta$ ,  $d_j$  increases by  $\delta$ , and all other aspiration levels do not change. Hence  $D$  decreased and given the earlier observation the resulting state is again in the core, since now for all  $j'$ ,  $d_i + d_{j'} \geq \alpha_{ij'}$  and all other inequalities still hold.

In the subsequent cases we shall omit a description of the period by period transitions since they are conceptually similar.

*Case A.1b.* For all  $j' \in J'$ ,  $d_i + d_{j'} = \alpha_{ij'}$ .

It follows that  $d_{j'} < d_{j'}^*$ , hence an aspiration level reduction by  $\delta$  by  $i$  and a  $\delta$ -increase by  $j$  and all  $j' \in J'$  yields a reduction in  $D$  and leads to a core state.

*Case A.2.*  $d_i = d_i^*$ .

Since  $Z \notin \mathbf{L}$  we must have  $d_{j'} < d_{j'}^*$ . For otherwise, given  $(i, j')$  is least

costly to destabilize, we would have  $Z \in \mathbf{L}$ . But then  $j'$  must be matched in the core assignment and we have for  $j'$ 's partner  $i'$  that  $d_{i'} > d_{i'}^*$ . Hence an aspiration level reduction by  $\delta$  by  $i'$  and a  $\delta$ -increase by  $j'$  and all  $j''$  for whom  $d_{i'} + d_{j''} = \alpha_{i'j''}$  yields a reduction in  $D$  and leads to a core state.

*Case A.3.*  $d_i < d_i^*$ .

We have  $d_j > d_j^*$  and again a similar argument applies. An aspiration level reduction by  $\delta$  by  $j$  and a  $\delta$ -increase by  $i$  and all  $i'$  for whom  $d_{i'} + d_j = \alpha_{i'j}$  yields a reduction in  $D$  and leads to a core state.

*Case B.* Suppose that the least cost deviation to a non-core state is such that one player experiences a shock and therefore wishes to break up. That is, there exists  $i$  such that  $d_i$  is least costly to destabilize.

It follows that  $d_i < d_i^*$ , for otherwise  $Z \in \mathbf{L}$  would constitute a contradiction.

*Case B.1.* For all  $i' \neq i$ ,  $d_{i'} + d_j > \alpha_{i'j}$ .

We again can construct a sequence of transitions such that  $i$  increases his aspiration level by  $\delta$ , and  $j$  reduces his aspiration level by  $\delta$ . Note that  $D$  then decreased and the resulting state is again in the core given that for all  $i' \neq i$  we started out with  $d_{i'} + d_j > \alpha_{i'j}$ .

Now we shall explain the sequence in detail. Suppose the tremble occurs such that  $i$  turns single. Consequently  $j$  is now single too and, given that we are in a core state  $(i', j)$  is not profitable for any  $i' \neq i$ . Therefore if  $j$  encounters any  $i' \neq i$  he will reduce his aspiration level. Now  $i$  can rematch with his optimal match  $j$  at a new price such that  $i$  can increase his aspiration level by  $\delta$  while  $d_j$  decreases his by  $\delta$ . (Note that, for the latter transition it is crucial that any matched couple has match value at least  $\delta$ .) Hence  $D$  decreased and given the earlier observation the resulting state is again in the core, since now for all  $i'$ ,  $d_{i'} + d_j \geq \alpha_{i'j}$  and all other inequalities still hold.

*Case B.2.* There exists  $I' \neq \emptyset$  and  $i \notin I'$  such that for all  $i' \in I'$ ,  $d_{i'} + d_j = \alpha_{i'j}$ .

Similar to case B.1 we can construct a sequence such that  $i$  increased his aspiration level by  $\delta$ ,  $j$  reduced his by  $\delta$ , and all  $i' \in I'$  increased their aspiration level by  $\delta$  (which will only reduce  $D$  further). The resulting state is in the core.  $\square$

**Theorem 3.** *The stochastically stable states are maximally robust to one-period deviations, and hence  $\mathbf{S} \subseteq \mathbf{L}$ .*

*Proof.* We shall prove the theorem by contradiction. Suppose there exists  $Z^* \in \mathbf{S} \setminus \mathbf{L}$ . Let  $T^*$  be a minimal cost tree rooted at  $Z^*$  and suppose that  $\rho(Z^*)$  is minimal. Let  $Z^{**} \in \mathbf{L}$ . By lemma 2 together with the fact that the state space is finite, we can construct a finite path of least cost transitions between different core states such that their distance to a core state in  $\mathbf{L}$  is decreasing:

$$Z^* \rightarrow Z_1 \rightarrow Z_2 \rightarrow \dots \rightarrow Z_k = Z^{**} \quad (4.23)$$

Now we perform several operations on the tree  $T^*$  to construct a tree  $T^{**}$  for  $Z^{**}$ . First add the edges  $Z_1 \rightarrow Z_2, \dots, Z_{k-1} \rightarrow Z_k$  and remove the previously exiting edges from  $Z_1, \dots, Z_{k-1}$ . Note that since the newly added edges are all minimal cost edges the sum of resistances does not increase. Next, let us add the edge  $Z^* \rightarrow Z_1$  and delete the exiting edge from  $Z_k$ . Since  $Z^* \notin \mathbf{L}$  it follows that  $r(Z^* \rightarrow Z_1) < r(Z_k \rightarrow \cdot)$  and hence

$$\rho(Z^{**}) \leq \rho(Z^*) + r(Z^* \rightarrow Z_1) - r(Z_k \rightarrow \cdot) < \rho(Z^*) \quad (4.24)$$

This constitutes a contradiction.  $\square$

We can formulate natural conditions under which the stochastically stable set coincides with the least core:



**Well-connected.** An assignment game is *well-connected* if for any non-core state and for any player  $i \in F \times W$  there exists a sequence of transitions in the unperturbed process such that  $i$  is single at its end.

**Rich.** An assignment game with match values  $\alpha$  is *rich* if for every player  $i \in F$  there exists a player  $j \in W$  such that  $(i, j)$  is never profitable, that is  $\alpha_{ij} = 0$ .

**Corollary 4.** *Given a well-connected and rich assignment game with a unique core matching<sup>18</sup>, the set of stochastically stable states coincides with the least core, that is  $\mathbf{S} = \mathbf{L}$ .*

*Proof.* Given two recurrent classes of the process  $Z^*, Z^{**} \in \mathbf{C}_0$ , and a non-core state  $Z \notin \mathbf{C}$  which is reachable from  $Z^*$  by a least cost transition, we shall show that  $r(Z, Z^{**}) = 0$ . Suppose that  $Z^{**}$  has aspiration levels  $\mathbf{d}^{**}$ .

The idea of the proof is to construct a finite family of well-connected sequences, such that after going through all transitions players have aspiration levels less than or according to  $\mathbf{d}^{**}$ . Once we are in such a state  $Z^{**}$  can be reached easily.

Given a well-connected sequence in  $Z$  which makes  $j_k$  single at its end (note that the sequence naturally needs to alternate between firms and workers in order to make players single along the way; also it needs to start with a single),  $(i_1, j_1), \dots, (i_k, j_k)$ , we shall show inductively that for any neighboring couples  $(i_l, j_l), (i_{l+1}, j_{l+1}), l = 1, \dots, k - 1$ , such that  $i_l$  and  $j_l$  are currently single, there exists a zero-resistance path such that  $i_l$  and  $j_l$  rematch at any price and  $i_{l+1}$  and  $j_{l+1}$  are both single.

---

<sup>18</sup>Generically the optimal matching is unique. In particular this holds if the weights of the edges are independent, continuous random variables. Then, with probability 1, the optimal matching is unique.

By richness  $i_l$  and  $j_l$  reduce their aspiration levels to zero with positive probability. Suppose that, by well-connectedness,  $i_l$  can match with  $j_{l+1}$  and therefore make  $i_{l+1}$  single. Suppose that the price is set such that  $i_l$  keeps aspiration level 0. Then in the next period  $i_l$  and  $j_l$  can match at any price. Suppose they match such that  $d_i \leq d_i^{**}$  and  $d_j \leq d_j^{**}$ . (This can actually occur since otherwise the aspiration level vector  $\mathbf{d}^{**}$  would not be pairwise stable, contradicting that  $Z^{**}$  is a core state. Also note that  $(i_l, j_l)$  are profitable since both have aspiration level zero and  $\alpha_{i_l j_l} > 0$  given they are on a well-connected sequence.) Now  $i_{l+1}$  and  $j_{l+1}$  are both single. Thus we can apply the same argument for any two subsequent couples in the sequence to conclude that any couple in the sequence can rematch at any aspiration levels such that for all  $i$  along the sequence,  $d_i \leq d_i^{**}$ .

Now by the well-connectedness assumption we know that from any non-core state and for every player there exists such a sequence. Hence successively applying sequences such that each player is at the end once, we can conclude that there exists a path to a state such that for all  $i$ ,  $d_i \leq d_i^{**}$ .

Next we have to show how  $Z^{**}$  is reached from the latter state. Successively match all  $(i, j)$  who are matched in  $Z^{**}$ , who are not matched yet, and for whom  $d_i + d_j < \alpha_{ij}$  (they are profitable) at a price such that their new aspiration levels are  $d_i^{**}, d_j^{**}$  (note that it is here that we need the fact that the core matching is unique). This leads to a state where aspiration levels are at  $\mathbf{d}^{**}$ . Note that these aspiration levels are good. Further note that a reduction of the sum of aspiration levels will lead to a state which is not good. Cases 1a,b and 2b of the proof of Claim 2 of Theorem 1 can now be applied iteratively (Case 2a can not hold, given that otherwise aspiration levels will no longer be good). These cases do not change the aspiration levels but only the matchings. Hence eventually the desired core state  $Z^{**}$  is going to be reached.

We showed that once the process is in a non-core state any core state can be reached. Hence the analysis of stochastic stability reduces to the resistance of exiting a core state. But this resistance is uniquely maximized by the states in the least core which thus coincides with the set of stochastically stable states.  $\square$

### 4.6.3 Example

We shall illustrate the predictive power of our result for the  $3 \times 3$  game studied by Shapley and Shubik, 1972. Let three sellers  $(w_1, w_2, w_3)$  and three buyers  $(f_1, f_2, f_3)$  trade houses. Their valuations are as follows:

These prices lead to the following match values,  $\alpha_{ij}$  (units of 1,000), where sellers are occupying rows and buyers columns:

$$\boldsymbol{\alpha} = \begin{pmatrix} 5 & \mathbf{8} & 2 \\ 7 & 9 & \mathbf{6} \\ \mathbf{2} & 3 & 0 \end{pmatrix} \quad (4.25)$$

The unique optimal matching is shown in bold numbers. Shapley and Shubik, 1972 note that it suffices to consider the 3-dimensional imputation space spanned by the equations  $d_{w_1} + d_{f_2} = 8$ ,  $d_{w_2} + d_{f_3} = 6$ ,  $d_{w_3} + d_{f_1} = 2$ . Figure 4.3 shows the possible core allocations.

We shall now consider the least core,  $\mathbf{L}$ . Note that the particular states in  $\mathbf{L}$  depend on the step size  $\delta$ . We shall consider  $\delta \rightarrow 0$  to best illustrate the core selection. By an easy calculation one finds that the states which are least vulnerable to one-period deviations are such that

$$d_{w_1} \in [11/3, 13/3], d_{w_2} = 17/3, d_{w_3} = 1/3. \quad (4.26)$$

The minimal excess in the least core is  $e_{\min} = 1/3$ . The bold line in figure 4.4 shows the set  $\mathbf{L}$ . The nucleolus,  $d_{w_1} = 4$ ,  $d_{w_2} = 17/3$ ,  $d_{w_3} = 1/3$ , is indicated by a cross. (One can verify, that here the kernel coincides with the nucleolus.)

## 4.7 Conclusion

In this paper we have shown that agents in large decentralized matching markets can learn to play equitable core outcomes through simple trial-and-error learning rules. We assume that agents have no information about the distribution of others' preferences, their past actions and payoffs, or the value of different matches. The unperturbed process leads to the core with probability one but no authority 'solves' an optimization problem. Rather, a path into the core is discovered in finite time by a random sequence of adjustments by the agents themselves. This result is similar in spirit to that of Chen, Fujishige, and Yang, 2011, but in addition our process selects equitable outcomes within the core. In particular, the stochastically stable states of the perturbed process are contained in the least core, a subset of the core that generalizes the nucleolus for assignment games. This result complements the stochastic stability analysis of Newton and Sawa, 2013 in ordinal matching and of Newton, 2012 in more general coalitional games. It is an open problem to extend the analysis to more general classes of cooperative games and matching markets.

## References

- Agastya, M. (1997). “Adaptive Play in Multiplayer Bargaining Situations”. In: *The Review of Economic Studies* 64.3, pp. 411–426.
- (1999). “Perturbed Adaptive Dynamics in Coalition Form Games”. In: *Journal of Economic Theory* 89, pp. 207–233.
- Arnold, T. and U. Schwalbe (2002). “Dynamic coalition formation and the core”. In: *Journal of Economic Behavior and Organization* 49, pp. 363–380.
- Balinski, M. L. (1965). “Integer Programming: Methods, Uses, Computation”. In: *Management Science* 12, pp. 253–313.
- Balinski, M. L. and D. Gale (1987). “On the Core of the Assignment game”. In: *Functional Analysis, Optimization and Mathematical Economics* 1, pp. 274–289.
- Bayati, M. et al. (2014). *Bargaining dynamics in exchange networks*. *Journal of Economic Theory*, forthcoming.
- Biró, P. et al. (2012). “Solutions for the stable roommates problem with payments”. In: *Lecture Notes in Computer Science* 7551, pp. 69–80.
- Bush, R. and F. Mosteller (1955). *Stochastic Models of Learning*. Wiley.
- Chen, B., S. Fujishige, and Z. Yang (2011). *Decentralized Market Processes to Stable Job Matchings with Competitive Salaries*. KIER Working Papers 749, Kyoto University.
- Chung, K.-S. (2000). “On the existence of stable roommate matching”. In: *Games and Economic Behavior* 33, pp. 206–230.
- Crawford, V. P. and E. M. Knoer (1981). “Job Matching with Heterogeneous Firms and Workers”. In: *Econometrica* 49, pp. 437–540.
- Davis, M. and M. Maschler (1965). “The kernel of a cooperative game”. In: *Naval Research Logistic Quarterly* 12, pp. 223–259.
- Demange, G. and D. Gale (1985). “The strategy of two-sided matching markets”. In: *Econometrica* 53, pp. 873–988.

- Demange, G., D. Gale, and M. Sotomayor (1986). “Multi-item auctions”. In: *Journal of Political Economy* 94, pp. 863–872.
- Diamantoudi, E., L. Xue, and E. Miyagawa (2004). “Random paths to stability in the roommate problem”. In: *Games and Economic Behavior* 48.1, pp. 18–28.
- Driessen, T. S. H. (1998). “A note on the inclusion of the kernel in the core for the bilateral assignment game”. In: *International Journal of Game Theory* 27.2, pp. 301–303.
- (1999). “Pairwise-Bargained consistency and game theory: the case of a two-sided firm”. In: *Fields Institute Communications* 23, pp. 65–82.
- Estes, W. (1950). “Towards a statistical theory of learning”. In: *Psychological Review* 57, pp. 94–107.
- Foster, D. and H. P. Young (1990). “Stochastic evolutionary game dynamics”. In: *Theoretical Population Biology* 38, pp. 219–232.
- Foster, D. P. and H. P. Young (2006). “Learning to play Nash equilibrium without knowing you have an opponent”. In: *Theoretical Economics* 1, pp. 341–367.
- Gale, D. and L. S. Shapley (1962). “College admissions and stability of marriage”. In: *American Mathematical Monthly* 69, pp. 9–15.
- Germano, F. and G. Lugosi (2007). “Global Nash convergence of Foster and Young’s regret testing”. In: *Games and Economic Behavior* 60, pp. 135–154.
- Hart, S. and A. Mas-Colell (2003). “Uncoupled Dynamics Do Not Lead to Nash Equilibrium”. In: *American Economic Review* 93, pp. 1830–1836.
- (2006). “Stochastic uncoupled dynamics and Nash equilibrium”. In: *Games and Economic Behavior* 57, pp. 286–303.
- Heckhausen, H. (1955). “Motivationsanalyse der Anspruchsniveau-Setzung”. In: *Psychologische Forschung* 25.2, pp. 118–154.

- Herrnstein, R. J. (1961). “Relative and absolute strength of response as a function of frequency of reinforcement”. In: *Journal of Experimental Analysis of Behavior* 4, pp. 267–272.
- Hoppe, F. (1931). “Erfolg und Mißerfolg”. In: *Psychologische Forschung* 14, pp. 1–62.
- Huberman, G. (1980). “The nucleolus and the essential coalitions”. In: *Analysis and Optimization of Systems, Lecture Notes in Control and Information Systems* 28, pp. 417–422.
- Inarra, E., C. Larrea, and E. Molis (2008). “Random paths to P-stability in the roommate problem”. In: *International Journal of Game Theory* 36, pp. 461–471.
- (2013). “Absorbing sets in roommate problem”. In: *Games and Economic Behavior* 81, pp. 165–178.
- Jackson, M. O. and A. Watts (2002). “The evolution of social and economic networks”. In: *Journal of Economic Theory* 106, pp. 265–295.
- Kandori, M., G. J. Mailath, and R. Rob (1993). “Learning, mutation, and long run equilibria in games”. In: *Econometrica* 61.1, pp. 29–56.
- Kelso, A. S. and V. P. Crawford (1982). “Job Matching, Coalition Formation, and Gross Substitutes”. In: *Econometrica* 50, pp. 1483–1504.
- Klaus, B. and F. Klijn (2007). “Paths to stability for matching markets with couples”. In: *Games and Economic Behavior* 58.1, pp. 154–171.
- Klaus, B., F. Klijn, and M. Walzl (2010). “Stochastic stability for roommates markets”. In: *Journal of Economic Theory* 145, pp. 2218–2240.
- Klaus, B. and F. Payot (2013). *Paths to Stability in the Assignment Problem*. working paper.
- Kojima, F. and M. U. Uenver (2008). “Random paths to pairwise stability in many-to-many matching problems: a study on market equilibrium”. In: *International Journal of Game Theory* 36, pp. 473–488.

- Kuhn, H. W. (1955). “The Hungarian Method for the assignment problem”. In: *Naval Research Logistic Quarterly* 2, pp. 83–97.
- Llerena, F. and M. Nunez (2011). “A geometric characterization of the nucleolus of the assignment game”. In: *Economics Bulletin* 31.4, pp. 3275–3285.
- Llerena, F., M. Nunez, and C. Rafels (2012). *An axiomatization of the nucleolus of the assignment game*. Working Papers in Economics 286, Universitat de Barcelona.
- Marden, J. R. et al. (2009). “Payoff-based dynamics for multi-player weakly acyclic games”. In: *SIAM Journal on Control and Optimization* 48, pp. 373–396.
- Maschler, M., B. Peleg, and L. S. Shapley (1979). “Geometric properties of the kernel, nucleolus, and related solution concepts”. In: *Mathematics of Operations Research* 4.4, pp. 303–338.
- Nash, J. (1950). “The bargaining problem”. In: *Econometrica* 18, pp. 155–162.
- Newton, J. (2010). “Non-cooperative convergence to the core in Nash demand games without random errors or convexity assumptions”. PhD thesis. University of Cambridge.
- (2012). “Recontracting and stochastic stability in cooperative games”. In: *Journal of Economic Theory* 147.1, pp. 364–381.
- Newton, J. and R. Sawa (2013). *A one-shot deviation principle for stability in matching markets*. Economics Working Papers 2013-09, University of Sydney, School of Economics.
- Nunez, M. (2004). “A note on the nucleolus and the kernel of the assignment game”. In: *International Journal of Game Theory* 33.1, pp. 55–65.
- Pradelski, B. S. R. (2014). *Evolutionary dynamics and fast convergence in the assignment game*. Department of Economics Discussion Paper Series 700, University of Oxford.



- Pradelski, B. S. R. and H. P. Young (2012). “Learning Efficient Nash Equilibria in Distributed Systems”. In: *Games and Economic Behavior* 75, pp. 882–897.
- Rochford, S. C. (1984). “Symmetrically pairwise-bargained allocations in an assignment market”. In: *Journal of Economic Theory* 34.2, pp. 262–281.
- Roth, A. E. and I. Erev (1995). “Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term”. In: *Games and Economic Behavior* 8, pp. 164–212.
- Roth, A. E. and M. Sotomayor (1992). “Two-sided matching”. In: *Handbook of Game Theory with Economic Applications* 1, pp. 485–541.
- Roth, A. E. and H. Vande Vate (1990). “Random Paths to Stability in Two-Sided Matching”. In: *Econometrica* 58, pp. 1475–1480.
- Rozen, K. (2013). “Conflict Leads to Cooperation in Nash Bargaining”. In: *Journal of Economic Behavior and Organization* 87, pp. 35–42.
- Sauermann, H. and R. Selten (1962). “Anspruchsanpassungstheorie der Unternehmung”. In: *Zeitschrift fuer die gesamte Staatswissenschaft* 118, pp. 577–597.
- Sawa, R. (2011). *Coalitional stochastic stability in games, networks and markets*. Working Paper, Department of Economics, University of Wisconsin-Madison.
- Schmeidler, D. (1969). “The nucleolus of a characteristic function game”. In: *SIAM Journal of Applied Mathematics* 17, pp. 1163–1170.
- Selten, R. (1998). “Aspiration Adaptation Theory”. In: *Journal of Mathematical Psychology* 42.2-3, pp. 191–214.
- Selten, R. and R. Stoecker (1986). “End Behavior in Sequences of Finite Prisoner’s Dilemma Supergames: A Learning Theory Approach”. In: *Journal of Economic Behavior and Organization* 7, pp. 47–70.

- Shapley, L. S. and M. Shubik (1963). “The core of an economy with nonconvex preferences”. In: *The Rand Corporation* 3518.
- (1966). “Quasi-cores in a monetary economy with nonconvex preferences”. In: *Econometrica* 34.4.
- (1972). “The Assignment Game 1: The Core”. In: *International Journal of Game Theory* 1.1, pp. 111–130.
- Solymosi, T. and T. E. S. Raghavan (1994). “An algorithm for finding the nucleolus of assignment games”. In: *International Journal of Game Theory* 23, pp. 119–143.
- Sotomayor, M. (2003). “Some further remarks on the core structure of the assignment game”. In: *Mathematical Social Sciences* 46, pp. 261–265.
- Thorndike, E. (1898). “Animal Intelligence: An Experimental Study of the Associative Processes in Animals”. In: *Psychological Review* 8, pp. 1874–1949.
- Tietz, A. and H. J. Weber (1972). “On the nature of the bargaining process in the Kresko-game”. In: *H. Sauermann (Ed.) Contribution to experimental economics* 7, pp. 305–334.
- Young, H. P. (1993). “The Evolution of Conventions”. In: *Econometrica* 61.1, pp. 57–84.
- (2009). “Learning by trial and error”. In: *Games and Economic Behavior* 65, pp. 626–643.

Figure 4.1: Transition diagram for active, matched agent (period  $t + 1$ ).




Figure 4.2: Transition diagram for active, single agent (period  $t + 1$ ).

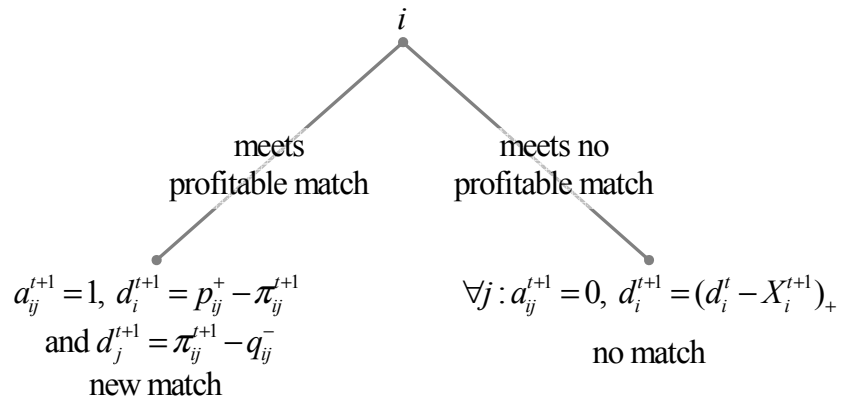


Table 4.1: Seller and buyer evaluations.

House	Sellers willingness to accept $q_{1j}^+ = q_{2j}^+ = q_{3j}^+$	Buyers' willingness to pay		
		$p_{i1}^+$	$p_{i2}^+$	$p_{i3}^+$
1	18,000	23,000	26,000	20,000
2	15,000	22,000	24,000	21,000
3	19,000	21,000	22,000	17,000

Figure 4.3: Imputation space for the sellers.





Figure 4.4: Core selection for the sellers.



## Chapter 5

### Complex cooperation: Agreements with multiple spheres of cooperation

## **Abstract**

A generalization of transferable utility cooperative games from the functional forms introduced by von Neumann and Morgenstern (1944) and Lucas and Thrall (1963) is proposed to allow for multiple membership. The definition of the core is adapted analogously and the possibilities for the cross-cutting of contractual arrangements are illustrated and discussed.

**Acknowledgements.** Nax thanks Peyton Young, Françoise Forges, Geoffroy de Clippel, Gérard Hamiache, Reza Saleh-Nejad, Oliver Walker, Chris Wallace, anonymous referees, as well as participants at the Gorman workshop and at the LSE Choice Group for helpful comments and discussions on earlier versions. This work was partly supported by the United States Office of Naval Research (Grant N00014-09-1-0751), by Instituto HPM Ascalone-Greco, and by the European Commission through the ERC Advanced Investigator Grant “Momentum” (Grant 324247).

## 5.1 Introduction

The coalitional game as defined by Von Neumann and Morgenstern, 1944 associates a unique worth with each coalition. Such a characterization is restrictive for many applications as it may be reasonable to allow the worth of one coalition to depend on the formation of other coalitions. Consequently, in Thrall and Lucas, 1963’s definition of a cooperative game, the worth of coalitions depends on the partitions of the rest of society, thus allowing different worth to be associated with each possible coalition depending on what coalitions are formed in the rest of society (“externalities”). This representation is still restrictive in the sense that it “presumes that coalitions are mutually exclusive, but in reality, a player might belong to multiple coalitions that interact with one another (e.g., a country might belong to both the United Nations and the European Union)” (Maskin, 2003). (See Haas, 1980; Charnovitz, 1998; Le Breton et al., 2013 for an international relations perspective on issue linkage through multilateral agreements.)

This note introduces a functional representation of a cooperative game where coalitions can form in multiple spheres of interaction simultaneously such that each coalition in each sphere is associated with a worth that depends on the overall coalition structure. Inherent to the model, therefore, is a new type of “cross externality”: the effect of forming coalitions across spheres. Such a formulation is relevant for many applications because, with multiple membership in the underlying application, a compartmentalized approach to the study of each sphere in isolation may lead to wrong conclusions concerning the stability of coalitional agreements. In a multiple membership setting, different layers may imbalance or balance each other depending on the structure of total spillovers (within and across spheres). Coalitions that seemed stable (or unstable) from the compartmentalized single-sphere



viewpoint may turn out to be destabilized (stabilized) by the multi-sphere game. To assess the stability of candidate agreements, we adapt definitions of the core (Gillies, 1959; Shapley, 1952) of the von Neumann-Morgenstern game (Von Neumann and Morgenstern, 1944) as done for the Lucas-Thrall game (Thrall and Lucas, 1963) in Hafalir, 2007, using an analogous “conjecture/ expectation formation approach” (Bloch and Nouweland, 2014) to recover the Bondareva-Shapley theorem (Bondareva, 1963; Shapley, 1967). To achieve this, the set of feasible deviations is restricted to a specific class. Further inspection of the resulting non-emptiness constraints reveals that different types of cross externalities create further opportunities for the cross-cutting of contractual arrangements. Our analysis builds on the work of Bloch and Clippel, 2010 who identify conditions for when non-emptiness of the core is facilitated through combining additively separable von Neumann-Morgenstern games (the single-sphere and no externalities case). (Not our lead example but some of our later examples are borrowed and generalized from theirs. See also Tijs and Brânzei, 2003 on the additivity of the core.) Our work also complements Diamantoudi et al., 2015’s generalization of Shapley, 1953 value in an environment like ours.

The rest of this note is structured as follows. Next, the model is motivated by means of a multimarket competition game. In Sections 3 and 4, we introduce the general game, define its core, and illustrate the core characteristics at hand of examples and observations. We conclude with some remarks.

## 5.2 A worked example

To motivate our model, we consider a multimarket Cournot economy with mergers and spillovers. (See, for example, Bloch, 1996; Ray and Vohra, 1999; Bloch and Nouweland, 2014 for single-market Cournot

competition games in this spirit.)

Example 1: A population of firms,  $N = \{f_1, \dots, f_n\}$ , competes in a multimarket industry,  $K = \{1, \dots, m\}$ , by setting production quantities. Each firm  $f$  is described by a vector of specializations,  $s_f = \{s_f^1, \dots, s_f^m\}$ , where each  $s_f^k$  is a real number representing firm  $f$ 's constant marginal costs in market  $k$  when no merger occurs.

In any market  $k$ , coalitions of firms  $S \subseteq N$  may merge and form a new firm. The resulting industry configuration,  $\mathcal{M}$ , describes the partitions in each market,  $\{\rho_1, \rho_m\}$ . Given  $\mathcal{M}$ , any firm  $S \in \rho^k$  produces quantity  $q_S^k$  in market  $k \in K$  at cost

$$C_S^k(q_S^k; \mathcal{M}) = c_S^k(\mathcal{M}) \times q_S^k + x_S^k(\mathcal{M}).$$

**Fixed costs of merger.**  $x_S^k(\mathcal{M})$ , the fixed cost of merging  $S$  in market  $k$ , is a real-valued function that depends on  $\mathcal{M}$  in the following way:

$$x_S^k(\mathcal{M}) = \begin{cases} 0 & \text{if } |S| = 1 \\ \kappa & \text{if } |S| > 1 \text{ and there exists } k' \neq k: S \in \rho_{k'} \\ \lambda & \text{if } |S| > 1 \text{ and there does not exist } k' \neq k: S \in \rho_{k'} \end{cases}$$

**Marginal costs of production.** Given any merger  $S \subseteq N$  in market  $k$ , the firms in  $S$  select the lowest marginal cost firm to be the only active firm amongst them in market  $k$ . Hence, the marginal cost of production of  $S$ ,  $c_S^k$ , as a result of the merger is given by

$$c_S^k = \min\{s_f^k\}_{f \in S}.$$

Given merger  $S \subseteq N$  in market  $k$ , the marginal cost of production of any coalition  $C$ ,  $c_C^k$ , in any other market  $k' \neq k$  is affected in the following way. For any  $C \in \rho_{k'}$ , we write  $\underline{c}_C^{k'}$  for  $\min\{s_f^{k'}\}_{f \in C}$ , i.e., for the marginal cost of the lowest marginal cost firm amongst  $C$  in market  $k'$ . For all  $C$  such that  $C \cap S = \emptyset$ ,  $c_C^{k'} = \underline{c}_C^{k'}$ . For all  $C$  such that

$C \cap S \neq \emptyset$ , given some  $\alpha \in (0, 1)$ ,

$$c_C^{k'} = \min\{\underline{c}_C^{k'}; \alpha \times c_S^{k'} + (1 - \alpha) \times \underline{c}_C^{k'}\}.$$

*The motivation for this marginal cost effect across markets is that firms connected by merger in one market may learn something about each others' production technologies and thus also improve (to some extent) their respective production technologies even in markets where they remain unmerged.*

**Demands.** *The demand of any product is the same in all markets (normalized to be equal-sized). Products are neither substitutes nor complements, meaning that all markets can be described by identical and independent linear demands. (These markets could be countries for example.) For any market  $k$ , therefore,*

$$p^k = 1 - Q^k \text{ where } Q^k \equiv \sum_{f \in N} q_f^k.$$

### 5.2.1 Oligopoly Externalities

A merger in a multimarket Cournot situation as introduced here has three different externality effects on the other firms in the same market and across markets. First, due to market consolidation, if merger occurs, the resulting quantity and price competition will change in that market, since the merged firms will be represented by the firm with the lowest marginal cost amongst them. Second, due to technology/knowledge spillovers across markets, the resulting quantity and price competition will also change in the markets where the merger did not occur because of the potential reduction in marginal costs (by how much is described by parameter  $\alpha$ ). Third, due to sharing of fixed costs merger, if the

same merger were to occur in more than one market, the fixed costs of merger per market would decline.

Due to the independence of the demand markets, the firms' optimization problems, given any industry configuration, can be solved for each market separately. The adjustments of equilibrium quantities and prices following mergers in any given market, however, have an effect in not only that same market because both the technology spillovers and the fixed cost effects may additionally influence the optimization problems in the other markets. A traditional representation of a cooperative game could not make these effects explicit. We shall illustrate these effects in more detail with a numerical illustration.

### 5.2.2 2-firm, 2-market Numerical Illustration

Take a symmetric two-firm, two-market case with  $s_{f_1}^1 = 4/9$ ,  $s_{f_1}^2 = 5/9$ ,  $s_{f_2}^1 = 5/9$ ,  $s_{f_2}^2 = 4/9$  (firm  $f_1$  is specialized in market 1 and firm  $f_2$  is specialized in market 2). Merger costs are  $\lambda = 2/108 > \kappa = 1/81$ . Choosing equilibrium outputs given the decision to merge in none, one, or both of the markets yields four cases: the no merger case, two one merger cases, and the full merger case. The equilibrium profits of these four cases are obtained by solving for the firms' profit-maximization problems. Table 1 summarizes the competition. (Writing (1),(2) means "no merger" in the underlying market and writing (1,2) means "merger".)

The direct effect of merger in market 1 is negative: profits fall from  $1 + 4 = 5$  to 4.75 if market 2 is not merged and from  $2.\bar{7} + 2.\bar{7} = 5.\bar{5}$  to 5.25 if market 2 is merged. If market 2 is merged, the overall cross effect on market 2 is positive: total profits in market 2 rise from 4.75 to 5.25. If market 2 is not merged, the profits in market 2 change from  $1 + 4 = 5$  to  $5 + \alpha \times 0.\bar{5}$ . If market 2 is not merged, the individual cross effect is negative on the strong firm in market 2 (profits fall from

4 to  $4 - \alpha \times 1.\bar{2}$ ), and positive on the weak firm (profits increase from 1 to  $1 + \alpha \times 1.\bar{7}$ ). The net total of the merger effects is therefore always positive if  $\alpha \times 0.\bar{5} > 0.25$ , *i.e.*, when  $\alpha > 0.45$ .

Since, *ceteris paribus*, mergers always decrease the worth of the merging market due to the high direct costs of merger, a partial view of one market suggests that merger is not in the firms' interests. When both markets are analyzed simultaneously, however, the cross-market effects of mergers are internalized. Since the cross effects are net positive if  $\alpha > 0.45$ , these effects would already render a single merger worthwhile overall.

When no merger takes place, each firm's profits from both markets are  $4 + 1 = 5$  and the total profits are 10. When one merger takes place, the firms can agree on sharing the total payoffs of  $(5 + \alpha \times 0.\bar{5}) + 4.75 = 9.75 + \alpha \times 0.\bar{5}$ . Under full merger, contracts can share the total profits of  $5.25 + 5.25 = 10.5$ . Therefore, no contracts can be written that Pareto-improve on contracts that result in full merger and share the total profits efficiently, paying each player at least 5 (which are the profits that each firm can guarantee itself from no merger). Whether a single merger already has a net-positive effect depends on whether  $\alpha > 0.45$  or not.

### 5.3 The model

This section generalizes the example to a representation of a cooperative game. Let  $N = \{f_1, f_2, \dots, f_n\}$  be the fixed population of agents. Write  $\rho$  for a partition of  $N$  and  $\rho(S)$  for the partition of some  $S \subset N$ . Let  $\mathbf{P}(N)$  be the set of partitions of  $N$  and  $\mathbf{P}(S)$  the set of partitions of  $S \subset N$ . Let  $K = \{1, \dots, m\}$  be the set of cooperative layers, that is, different spheres over which cooperation amongst  $S \subseteq N$  may ensue. Write  $\mathcal{M}$  for a society consisting of a partition of each layer,  $\mathcal{M} =$

$\{\rho_1, \dots, \rho_m\}$ , and  $\mathcal{M}_S = \{\rho_1(S), \dots, \rho_m(S)\}$  for a subsociety consisting of a partition of each layer of some  $S \subset N$  (in which case  $\mathcal{M}_S$  and  $\mathcal{M}_{N \setminus S}$  are “separable” subsocieties; *i.e.*, there is no coalition that includes members from both subsocieties).

Now,  $G(v, K, N)$  is a multiple membership game (MMG), defined by  $N$ ,  $K$  and  $v$ .  $v$  is the characteristic multiple membership function that assigns, for every layer  $k \in K$ , a worth in terms of transferable utility of  $v_k$  to each  $C \in \rho_k$  given  $\mathcal{M}$ : for any  $k \in K$ ,  $v_k(\cdot; \mathcal{M}) : \rho_k \rightarrow \mathbf{R}$  for all  $\rho_k \in \mathbf{P}(N)$ . Naturally, an MMG is a partition function game (PFG) as in Thrall and Lucas, 1963 if  $K$  consists of only one layer (when no multiple membership exists). With only one layer, the MMG/PFG further reduces to a characteristic function game (CFG) as in Von Neumann and Morgenstern, 1944 if, for any  $C \subseteq N$ ,  $v(C; \rho)$  is constant for all  $\rho \in \mathbf{P}(N)$  with  $C \in \rho$ .

### 5.3.1 Externalities

When multiple membership exists, externalities come in various kinds. In these notes, an externality is said to be present if one instance of its effect is present so that a game may exhibit different kinds of externalities over different parts of the game. This allows to model interesting situations like the above Cournot model: merger in one market has both positive and negative effects on the other firms and on the other markets.

The externalities will be defined using the notion of embedded coalitions. Given partition  $\rho$  of  $N$ ,  $C$  is an embedded coalition if  $C \in \rho$ . Partition  $\rho$  embeds  $\rho'$  if, for all  $C' \in \rho'$ , there is some  $C \in \rho$  such that  $C' \subseteq C$ . One externality is the “partition” externality, which is the externality known from PFGs: the intra-layer externality of an  $n(\geq 3)$ -player Cournot game, for example, where one firm’s payoff varies with the remaining firms’ decisions on whether to merge or not, is such an

externality.

**Partition externality.**  $G(v, k, N)$  exhibits a positive (or negative) partition externality if there exist  $\mathcal{M}, \mathcal{M}'$  such that  $\mathcal{M} \setminus \rho_k = \mathcal{M}' \setminus \rho'_k$ ,  $\rho_k$  embeds  $\rho'_k$  with  $C \in \rho_k, C \in \rho'_k$ , and

$$v_k(C; \mathcal{M}) > (\text{or } <) v_k(C; \mathcal{M}').$$

The other “cross” externality stems from the effects of the formation of coalitions in one layer on the payoffs of some coalition in another. This inter-layer effect is new and peculiar to multiple membership and cannot be expressed through existing cooperative game representations. In the multimarket Cournot example, the cross externality was the effect of merger in one product market on the firms’ profits in the other.

**Cross externality.**  $G(v, k, N)$  exhibits a positive (or negative) cross externality if there exist  $\mathcal{M}, \mathcal{M}'$  such that  $\mathcal{M} \setminus \rho_k = \mathcal{M}' \setminus \rho'_k$  with  $C \in \rho_{k'}, C \in \rho'_k$  for some  $k' \neq k$ ,  $\rho_k$  embeds  $\rho'_k$ , and

$$v_{k'}(C; \mathcal{M}) > (\text{or } <) v_{k'}(C; \mathcal{M}').$$

A subclass of cross externalities are “partition-cross” (“partition externalities across layers”). They have elements of cross and of partition externalities: coalition formation of one set of players  $S_1 \subseteq N$  in one layer affects the worth of coalitions of another  $S_2 \subseteq N$  in another layer with  $S_1 \cap S_2 = \emptyset$ . This occurs when, for example, a merger of firms one and two in one market affects the profits of firm three in another.

**Partition-cross externality.**  $G(v, k, N)$  exhibits a positive (or negative) partition-cross externality if there exist  $\mathcal{M}, \mathcal{M}'$  such that  $\mathcal{M} \setminus \rho_k = \mathcal{M}' \setminus \rho'_k$  with  $C \in \rho_{k'}, C \in \rho'_k$  for some  $k' \neq k$ ,  $\rho_k$  embeds  $\rho'_k$  while being identical w.r.t. the coalitions that all members of  $C$  join (*i.e.*, for all  $f$  such that  $f \in C \in \rho_{k'}, (f \in S \in \rho_k) \Leftrightarrow (f \in$

$S \in \rho'_k$ ) with the same  $S$  in both), and

$$v_{k'}(C; \mathcal{M}) > (\text{or } <) v_{k'}(C; \mathcal{M}').$$

A partition-cross externality is a partition externality where partitions  $\rho_{k_i}$  and  $\rho'_{k_i}$  are identical w.r.t. the coalitions that all members of  $C$  join: for all  $f$  such that  $f \in C \in \rho_{k_j}, (f \in S \in \rho_{k_i}) \Leftrightarrow (f \in S \in \rho'_{k_i})$  with the same  $S$  in both.

### 5.3.2 Feasible Deviations

In the absence of externalities and multiple membership (*i.e.*, in characteristic function games, CFGs), a deviation by some  $S \subset N$  when forming a coalition has a one-to-one association with a unique worth of  $S$  (Von Neumann and Morgenstern, 1944). In the presence of externalities, however, further *expectation conjectures* (assumptions about how the rest of society,  $N \setminus S$ , reacts to a coalitional deviation by  $S$ ) are needed (Von Neumann and Morgenstern, 1944; Aumann, 1967). For partition function games (PFGs), that is, in the presence of externalities within a single sphere (no multiple membership), Shenoy, 1979; Chander and Tulkens, 1997; Hart and Kurz, 1983; De Clippel and Serrano, 2008 propose definitions of the core dependent on different conjectures to evaluate the profitability of coalitional deviations. Bloch and Nouweland, 2014 provides an excellent discussion of these, also analyzing their axiomatic foundations (Hafalir, 2007 provides additional results on the externality structure relevant for the corresponding non-emptiness results for several of these cores.). Suppose the partition was  $\rho$  before  $S \subset N$  deviated and reorganized itself to form  $\rho(S)$ , then these are the existing conjecture rules that have been proposed in PFG environments (see Bloch and Nouweland, 2014 for a detailed classification and an axiomatic analysis):



1. *Max rule* (Bloch and Nouweland, 2014):  $(N \setminus S)$ , taking  $\rho(S)$  as given, organizes itself to  $\rho(N \setminus S)$  in order to maximize  $(N \setminus S)$ 's total worth
2. *Pessimistic* (Aumann, 1967; Hart and Kurz, 1983):  $(N \setminus S)$  organizes and forms  $\rho(N \setminus S)$  in order to minimize  $S$ 's total worth
3. *Optimistic* (Shenoy, 1979):  $(N \setminus S)$  organizes and forms  $\rho(N \setminus S)$  in order to maximize  $S$ 's total worth
4. *Singleton* (Chander and Tulkens, 1997; De Clippel and Serrano, 2008):  $(N \setminus S)$  breaks down into singletons
5. *Collective* (Bloch and Nouweland, 2014):  $(N \setminus S)$  forms one joint coalition (Bloch and Nouweland, 2014 call this rule  $\bar{N}$ -exogenous)
6. *Disintegrative* (Von Neumann and Morgenstern, 1944; Hart and Kurz, 1983): all  $C \in \rho$  such that  $C \cap S = \emptyset$  remain organized in the same way, all other coalitions  $C'$  from which members in  $S$  deviated break up into singletons
7. *Projective* (Hart and Kurz, 1983): all  $C \in \rho$  such that  $C \cap S = \emptyset$  remain organized in the same way, all other coalitions  $C'$  from which members in  $S$  deviated form coalitions amongst the remaining  $(C' \setminus S)$

Note that conjecture rules 1–3 depend on  $\rho(S)$  and on the underlying PFG, but not on the original partition  $\rho$ . Rules 4–5 depend only on  $S$ . Rules 6–7 depend on  $S$  and on the original partition  $\rho$ .

With multiple membership, in addition to the need of specifying a conjecture, we must specify what kinds of deviations are deemed feasible. The feasibility of deviations needs to be interpreted here because, for example, starting with the grand coalition in some layer, each  $S \subset N$  may deviate in many ways: in some or all of the layers, forming different coalitions in each layer or the same coalition in all layers. If cooperation is compartmentalized without cross externalities in between the layers, players may deviate in one layer but continue to form the grand coali-

tion in another layer. When cross externalities are present, however, the worth of coalitions vary with the coalition constellations across layers and deviators need to endogenize the cross external effects of their deviations. If  $S \subset N$  deviates and forms  $S$  in layer 1, for example, it cannot expect to form  $N$  in another layer because  $S$ 's members need to cooperate with  $(N \setminus S)$  to form this constellation. Therefore, this note only considers the following deviations:

**Feasible deviations.** Any  $S \subseteq N$  can form any subsociety  $\mathcal{M}_S \in \mathbf{P}(S)^m$  (a partition of  $S$  in every layer).  $\mathcal{M}_S$  and  $\mathcal{M}_{N \setminus S}$  must be separable.

From the feasible set of subsocieties available to  $S$ , it aims to form subsociety  $\widehat{\mathcal{M}}_S \in \mathbf{P}(S)^m$  that maximizes its total payoffs. For that, each  $S \subset N$  needs to conjecture how the rest of the population responds to its deviation. The reason for restricting deviations in this way is to guarantee that society  $\mathcal{M}$  after deviation by some  $S \subset N$  occurs is separable into subsocieties  $\mathcal{M}_S$  and  $\mathcal{M}_{N \setminus S}$ . If this is the case, then the above list of conjectures can be adapted directly.

Suppose  $Z$  represents any of the above conjectures so that  $Z$ , for every  $\mathcal{M}_S$  deviating from  $\mathcal{M}$ , specifies a resulting subsociety  $\mathcal{Z}((N \setminus S); \mathcal{M}_S) \in \mathbf{P}(N \setminus S)^m$  of  $(N \setminus S)$  (a partition of  $(N \setminus S)$  in every layer, but not necessarily the same one in all layers). Write  $\widehat{\mathcal{M}}_S(\widehat{\mathcal{N}})$  for the resulting overall society  $\{\widehat{\mathcal{M}}_S, \mathcal{Z}((N \setminus S); \widehat{\mathcal{M}}_S)\}$ . Hence,  $S$  forms the optimal subsociety  $\widehat{\mathcal{M}}_S$  such that, given conjecture  $Z$ ,

$$\sum_{k \in K} \sum_{C \in \rho_k(S)} v_k(C; \widehat{\mathcal{M}}_S(\widehat{\mathcal{N}})) = \max_{\mathcal{M}_S \in \mathbf{P}(S)^m} \sum_{k \in K} \sum_{C \in \rho_k(S)} v_k(C; \{\mathcal{M}_S, \mathcal{Z}((N \setminus S); \mathcal{M}_S)\}).$$

The finiteness of possible coalition structures guarantees the existence of such a (not necessarily unique) subsociety for any  $S \subseteq N$ . We will now define a function summarizing their worth.

**Conjectured worth function.** The conjectured worth function (CWF),

$z$ , summarizes the conjectured worth for all coalitions: given  $Z$ , for all  $C \subseteq N$ ,  $z(\cdot) : C \rightarrow \mathbf{R}$ . For any  $S \subseteq N$ ,  $z(S)$  is the largest feasible sum of payoffs for  $S$  under conjecture  $Z$ :

$$z(S) = \sum_{k \in K} \sum_{C \in \widehat{\rho_k(S)}} v_k(C; \widehat{\mathcal{M}_S(\mathcal{N})})$$

Note that  $z$  filters the information in the MMG to obtain a CFG view of deviating demands.

### 5.3.3 Superadditivity

When externalities exist, a detailed analysis of the effects of coalition formation may be needed to evaluate the global benefits of cooperation and a superadditivity assumption may be difficult to uphold. When one agent is able to take free ride on the coalition formed by others, for example, the grand coalition may no longer be the efficient coalition structure and it may indeed be insightful to work with a given coalition structure to analyze the effects of free ride.

In the presence of multiple membership and externalities, coalition formation may be mutually beneficial in some layer but not necessarily globally as negative cross externalities may exist. Suitably defined, MMGs may conversely be globally superadditive if the overall effect of coalition formation, which takes into account all external and direct effects, is positive for those that come together to cooperate even if coalition formation itself is not mutually beneficial in the separate layers.

The numerical illustration of the multimarket Cournot game for the cases when  $\alpha > 0.45$ , for instance, is superadditive because the total profits of the firms rise with every further merger: the no merger case has total payoffs of 10, compared with the  $10 + (\alpha - 0.45) \times 0.5$  of both one merger cases, and compared with the 10.5 of full merger. The below

definition of MMG superadditivity embeds definitions of superadditivity for CFGs and PFGs and implies efficiency of forming the grand coalition in all layers.

**Superadditivity:** An MMG is superadditive if, for all  $\mathcal{M}, \mathcal{M}'$  such

that  $\mathcal{M} \setminus \rho_k = \mathcal{M}' \setminus \rho'_k$ , and  $\rho_k$  embeds  $\rho'_k$  in layer  $k$ ,

$$\sum_{k \in K} \sum_{C \in \rho_k} v_k(C; \mathcal{M}) \geq \sum_{k \in K} \sum_{C \in \rho'_k} v_k(C; \mathcal{M}').$$

Superadditivity implies the efficiency of the “grand coalition” by which we mean society  $\{N\}$  (the grand coalition) forms in all layers.

When the game consists of a single layer without externalities (described by a CFG), the above definition implies the simple pairwise superadditivity that  $v(C \cup C') \geq v(C) + v(C')$  is to be satisfied for all  $(C, C') \subset N : C \cap C' = \emptyset$ . (Note that the implied sense of superadditivity when there is only one layer has also been defined as *full cohesiveness* (Hafalir, 2007, section 2.2 “Convexity”) in the contexts of PFGs, as opposed to a pairwise view of superadditivity (Hafalir, 2007, section 2.1 “Superadditivity”). Hafalir, 2007’s pairwise view of superadditivity does not imply the efficiency of the grand coalition.) Note that the optimization problem underlying  $z$ , which is a CFG, entails that  $z$  is superadditive by definition, even if the MMG is not superadditive: for any  $S, S' \subseteq N$  with  $S \cap S' = \emptyset$ ,  $z(S) + z(S') \leq z(S \cup S')$ .

## 5.4 Coalitional stability and the core

We now turn to the stability of an outcome. By outcome we mean  $(\mathcal{M}, x)$ ; a coalition structure together with an allocation of the common gains. To assess its stability, we will use the conjectured worth function. For allocation  $x$ , we write  $x = \{x_{f_1}, \dots, x_{f_n}\}$  such that each allocated player payoff  $x_f = \sum_{k \in K} x_f^k$  summarizes the payoffs to each  $f \in N$

obtained in all layers. Consequently, for some  $S \subseteq N$ ,  $x(S)$  is a vector of all-layer payoffs for the players in  $S$ . Naturally, an allocation must be feasible: given any  $\mathcal{M}$ ,  $\sum_{f \in N} x_f \leq \sum_{k \in K} \sum_{C \in \rho_k} v(C; \mathcal{M})$ .

Recall our numerical illustration of the multimarket Cournot game. Independent of  $\alpha = (0, 1)$ , one unique conjectured worth function is derived, *i.e.*,  $z$  is such that  $z(f_1) = z(f_2) = 1 + 4 = 5$  and  $z(f_1, f_2) = \max\{(5.25 + 5.25); (9.75 + \alpha \times 0.5)\} = 10.5$ . Note that no conjecture is needed for this assessment. It is easy to verify in this particular example that  $G(v, K, N)$  has a nonempty core: for an example of a core outcome, consider full merger with contract  $x = (5.25, 5.25)$ , paying both firms 5.25. This outcome is in the core because no firm can do better by deviating. In fact, any split of full merger paying each firm at least his individually rational payoff of 5 (what he gets from no merger) and the other the residual to 10.5 is a stable core allocation.

### 5.4.1 Core Stability

Assume  $G(v, K, N)$  is superadditive such that the grand coalition is efficient. Whether there exists a core-stable allocation supported by the grand coalition depends on  $v$  and on the conjecture. We now provide definitions for any given conjecture. The  $Z$ -core (based on conjecture  $Z$ ) can be defined using the conjectured worth function.

**Z-core:** Given  $Z$ , the  $Z$ -core of forming the efficient society of  $G(v, K, N)$  with total payoff allocation  $x$  is

$$\zeta(G(v, K, N); Z) = \{x \in \mathbf{R}^n; \sum_{f \in N} x_f \leq z(N) \text{ and } \sum_{f \in S} x_f \geq z(S) \forall (S \subseteq N)\}.$$

**Theorem.** *The  $Z$ -core of  $G(v, K, N)$  is nonempty if, and only if, its conjectured worth function  $z$  is balanced.*

The theorem is a (straightforward) recovery of the Bondareva-Shapley theorem via the conjectured worth function in our setup (see Bondareva, 1963 and Shapley, 1967 for independent proofs). What is interesting is that several characteristics can be identified to describe the core structure, which turns out to be very complex.

**Characteristic 1:** *If the cores of a superadditive MMG layer-by-layer separately are nonempty, the  $Z$ -core of the whole MMGs is also nonempty.*

While  $z$  is always additive over coalitions and layers,  $v$  does not need to be additive when externalities are present. In every layer, superadditivity implies that it is beneficial for members of any  $S \subseteq N$  to form the largest possible coalition  $\{S\}$ . Hence, whenever  $x$  is in a  $Z$ -core,  $\sum_{f \in N} x_f = z(N)$ . Now,  $z_k$  describes the game described by the conjectured worth function of layer  $k$ , *i.e.*, the conjectured CFG view of layer  $k$ . Given any  $z_k$ , a core stable allocation of forming the grand coalition in that layer exists if, and only if, every  $z_k$  is balanced. Since the sum of balanced games is

balanced, the  $Z$ -core of  $G(v, k, N)$  is, therefore forcedly, nonempty when all  $z_k$ s are balanced.

**Characteristic 2:** *In the presence of cross externalities but without partition and partition-cross externalities, the core is unambiguously defined (independent of conjecture).*

In the absence of partition and partition-cross externalities, in a society  $\mathcal{M}$  that is separable into  $\mathcal{M}_S$  and  $\mathcal{M}_{N \setminus S}$ , the worth of any coalition  $C \subseteq S$  is independent of  $\mathcal{M}_{N \setminus S}$  in all layers:  $v_k(C; \mathcal{M}) = v_k(C; \mathcal{M}')$  for all coalitions, layers and societies provided that it holds that  $\mathcal{M}_S = \mathcal{M}'_S$ ,  $(C \in \rho_k \in \mathcal{M})$  and  $(C \in \rho'_k \in \mathcal{M}')$ . Therefore, one unique game described by a characteristic worth function is derived, which implies one unambiguous definition of the core. This unambiguity is independent of the existence of cross externalities that are not partition-cross because deviators endogenize all other cross external variations that may still exist and affect them. The need to conjecture is therefore inherent to the presence of PFG-type (partition and partition-cross) externalities. The core of example 1, for instance, is unambiguously defined.

**Characteristic 3:** *In the presence of positive cross externalities, the core of the MMG may be nonempty even if coalition formation in any of the layers is, ceteris paribus, never beneficial.*

Example 1 as described by Table 1 illustrates this.

**Characteristic 4:** *In the presence of negative cross externalities, the core of forming the grand coalition in any layer of the MMG may be empty even if coalition formation in all layers is, ceteris paribus, always beneficial.*

Example 2: *Let  $n = k = 2$  and  $v$  be described by Table 2.*

Holding the coalition structure of one layer fixed, any coalition formation in the other layer is beneficial. However, due to the

negative cross externality of coalition formation in one layer on the other, the total worth of all coalitions is reduced as coalitions form. The core of forming the grand coalition in one or both of the layers of example 2 is empty:  $z(1) + z(2) = z(N) = (v_1(1) + v_2(1)) + (v_1(2) + v_2(2)) = 4 \times 1 = 4 > 3 = 0 + 0 + 3 = (v_1(1) + v_1(2)) + v_2(N) > 2 = 1 + 1 = v_1(N) + v_2(N)$ .

**Characteristic 5:** *Multiple membership may facilitate cooperation not because of cross external effects but because the layers “balance each other”: Even in the complete absence of externalities when all layers have empty cores, the core of an MMG may be nonempty (See Bloch and Clippel, 2010 “Examples 1 and 2” for a 4- and related 5-player examples.).*

Example 3: Let  $n = 5$ ,  $k = 2$  and let there be no externalities so that the MMG is described by two 5-player CFGs,  $v_1$  and  $v_2$ . Let  $v_1(N) = 1$ ,  $v_1(C) = 4/5 + \varepsilon$  (where  $\varepsilon$  is small) if  $|C| = 4$  and  $v_1(C) = 0$  otherwise. Let  $v_2(N) = 1$ ,  $v_2(C) = 3/5 + \varepsilon$  if  $|C| = 3$ , 4 and  $v_2(C) = 0$  otherwise.

$v_1$  is unbalanced: for the balanced collection of the 5 coalitions of size 4,  $\zeta_{|4|} = \{(1, 2, 3, 4), \dots, (2, 3, 4, 5)\}$ , with balancing weights  $\lambda_{|4|} = (1/4, \dots, 1/4)$ ,  $5 \times 1/4 \times v_1(i, j, k, l) = 5 \times 1/4 \times (4/5 + \varepsilon) = 1 + 5/4 \times \varepsilon > 1 = v_1(N)$ .  $v_2$  is unbalanced: for the balanced collection of the 10 coalitions of size 3,  $\zeta'_{|3|} = \{(1, 2, 3), \dots, (3, 4, 5)\}$ , with balancing weights  $\lambda'_{|3|} = (1/6, \dots, 1/6)$ ,  $10 \times 1/6 \times v_2(i, j, k) = 10 \times 1/6 \times (3/5 + \varepsilon) = 1 + 5/3 \times \varepsilon > 1 = v_2(N)$ . However, it is easy to verify that  $x = (2/5, 2/5, 2/5, 2/5, 2/5)$  is a core allocation of  $v$ :  $z$  associates  $z(N) = 2$ ,  $z(C) = 7/5 + 2 \times \varepsilon$  if  $|C| = 4$ ,  $z(C) = 3/5 + \varepsilon$  if  $|C| = 3$  and  $z(C) = 0$  otherwise.

**Characteristic 6:** *The presence of positive (or negative) cross and/or partition externalities may lead to inefficient herding.*

Example 4: Let  $n = 3$ ,  $k = 2$  and  $v_k(N; \{\{N\}, \{N\}\}) = 1$  for all  $k$ ,



$v_k(1; \{\rho_1, \rho_2\}) = 2 \forall i$  if  $\rho_1 = \rho_2 = \{(1), (2, 3)\}$  and  $v_k(C; M) = 0$  otherwise.

The *Pessimistic*-core of forming the inefficient grand coalition in both layers is nonempty because player 1 expects to receive 0 from being the singleton in both layers, e.g.,  $x = (2/3, 2/3, 2/3)$  is such a *Pessimistic*-core allocation. Inefficient herding results from the positive externality: the formation of the coalition of (2,3) in both layers creates worth for player 1, but player 1 is too pessimistic to agree to stay separate. The same effect may be due to negative externalities as a simple variation of  $v$  illustrates: consider, for example,  $v'$  with  $v'_k(N; \{\{N\}, \{N\}\}) = 1$  for all  $k$ ,  $v'_k(1; \{\rho_1, \rho_2\}) = 2 \forall k$  if  $\rho_1 = \rho_2 = \{(1), (2), (3)\}$  and  $v'_k(C; M) = 0$  otherwise.

## 5.5 Concluding remarks

This paper sets out to define the core of coalitional games with multiple membership externalities. The point of departure is the representation in partition function form as introduced by Thrall and Lucas, 1963. Inherent to our multiple membership game are two types of externalities; those from within a given layer of cooperation where a coalitional decision of one set of agents has payoff consequences for another set of agents, and those from across different layers of cooperation where coalitional decisions in one sphere of cooperation influence payoffs in another sphere. Recent contributions explore the consequences for core existence of the first externality type (Hafalir, 2007) and of the second Bloch and Clippel, 2010. Work that is complementary to ours concerns extensions of the Shapley value to multiple membership externality environments (Diamantoudi et al., 2015). Our work illustrates how the two externality types may interact with coalitional incentives to deviate. Moreover, our model highlights one crucial issue with defining

the core in the presence of multiple membership externalities, namely that of feasibility of deviations. In this note, we take a somewhat extreme stance and allow deviations by some subsociety only if they do not expect to form coalitions in any of the layers with any of the players outside of their subsociety. This assumption drives the analysis in this note, and we aim to relax this assumption in future work, likely in conjunction with an axiomatic approach.

## References

- Aumann, Robert J (1967). “A survey of cooperative games without side payments”. In: *Essays in mathematical economics*. Ed. by Martin Shubik. Princeton, NJ: Princeton University Press.
- Bloch, Francis (1996). “Sequential Formation of Coalitions in Games with Externalities and Fixed Payoff Division”. In: *Games and Economic Behavior* 14.1, pp. 90–123.
- Bloch, Francis and Geoffroy de Clippel (2010). “Cores of combined games”. In: *Journal of Economic Theory* 145.6, pp. 2424–2434.
- Bloch, Francis and Anne van den Nouweland (2014). “Expectation formation rules and the core of partition function games”. In: *Games and Economic Behavior* 88, pp. 339–353.
- Bondareva, Olga N (1963). “Some applications of linear programming methods to the theory of cooperative games”. In: *Problemy kibernetiki* 10, pp. 119–139.
- Chander, Parkash and Henry Tulkens (1997). “The core of an economy with multilateral environmental externalities”. In: *International Journal of Game Theory* 26.3, pp. 379–401.
- Charnovitz, Steve (1998). “Linking Topics in Treaties”. In: *University of Pennsylvania Journal of International Economic Law* 19.2, pp. 329–345.
- De Clippel, Geoffroy and Roberto Serrano (2008). “Marginal Contributions and Externalities in the Value”. In: *Econometrica* 76.6, pp. 1413–1436.
- Diamantoudi, Effrosyni et al. (2015). “Sharing the surplus in games with externalities within and across issues”. In: *Economic Theory*, pp. 1–29.

- Gillies, Donald B (1959). “Solutions to general non-zero-sum games”.  
 In: *Contributions to the Theory of Games*. Ed. by Albert William Tucker and Robert Duncan Luce. Vol. 4. 40. Princeton, NJ: Princeton University Press, pp. 47–85.
- Haas, Ernst B. (1980). “Why Collaborate? Issue-Linkage and International Regimes”. In: *World Politics* 32.3, pp. 357–405.
- Hafalir, Isa E (2007). “Efficiency in coalition games with externalities”. In: *Games and Economic Behavior* 61.2, pp. 242–258.
- Hart, Sergiu and Mordecai Kurz (1983). “Endogenous Formation of Coalitions”. In: *Econometrica* 51.4, pp. 1047–1064.
- Le Breton, Michel et al. (2013). “Stability and fairness in models with a multiple membership”. In: *International Journal of Game Theory* 42.3, pp. 673–694.
- Maskin, Eric (2003). “Bargaining, coalitions and externalities”. In: *Presidential Address to the Econometric Society*. Princeton, NJ: Institute for Advanced Study.
- Ray, Debraj and Rajiv Vohra (1999). “A Theory of Endogenous Coalition Structures”. In: *Games and Economic Behavior* 26.2, pp. 286–336.
- Shapley, Lloyd S (1952). “Notes on the n-person game, III: Some variants of the von Neumann-Morgenstern definition of solution”. In: *Rand Memorandum*. Santa Monica, CA: Rand Corporation.
- (1953). “A value for n-person games”. In: *Contributions to the Theory of Games*. Ed. by Harold W Kuhn and Albert W Tucker. Vol. 2. Princeton, NJ: Princeton University Press, pp. 307–317.
- (1967). “On balanced sets and cores”. In: *Naval Research Logistics* 14.4, pp. 453–460.

- Shenoy, Prakash P (1979). “On coalition formation: a game-theoretical approach”. In: *International Journal of Game Theory* 8.3, pp. 133–164.
- Thrall, Robert M and William F Lucas (1963). “N-person games in partition function form”. In: *Naval Research Logistics* 10.1, pp. 281–298.
- Tijs, Stef and Rodica Brânzei (2003). “Additive stable solutions on perfect cones of cooperative games”. In: *International Journal of Game Theory* 31.3, pp. 469–474.
- Von Neumann, John and Oskar Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.

Table 5.1: *Numerical illustration.*

	Industry configuration		Profits (scaled $\times 81$ )	
	market 1	market 2	market 1	market 2
<i>mergers:</i>				
<i>none</i>	(1), (2)	(1), (2)	4, 1	1, 4
<i>market 1</i>	(1,2)	(1), (2)	4.75	$(1 + \alpha 1.\bar{7}), (4 - \alpha 1.\bar{2})$
<i>market 2</i>	(1), (2)	(1,2)	$(4 - \alpha 1.\bar{2}), (1 + \alpha 1.\bar{7})$	4.75
<i>full merger</i>	(1,2)	(1,2)	5.25	5.25

Table 5.2: *Example 2.*

Society		Coalition worth	
layer 1	layer 2	layer 1	layer 2
(1), (2)	(1), (2)	1, 1	1, 1
(1,2)	(1), (2)	3	0, 0
(1), (2)	(1,2)	0, 0	3
(1,2)	(1,2)	1	1

## Chapter 6

Dynamics of financial  
expectations:

Super-exponential growth  
expectations and crises

## Abstract

We construct risk-neutral return probability distributions from S&P 500 options data over the decade 2003 to 2013, separable into pre-crisis, crisis and post-crisis regimes. The pre-crisis period is characterized by increasing realized and, especially, option-implied returns. This translates into transient unsustainable price growth that may be identified as a bubble. Granger tests detect causality running from option-implied returns to Treasury Bill yields in the pre-crisis regime with a lag of a few days, and the other way round during the post-crisis regime with much longer lags (50 to 200 days). This suggests a transition from an abnormal regime preceding the crisis to a “new normal” post-crisis. The difference between realized and option-implied returns remains roughly constant prior to the crisis but diverges in the post-crisis phase, which may be interpreted as an increase of the representative investor’s risk aversion.



## 6.1 Introduction

The Global Financial Crisis of 2008 brought a sudden end to a widespread market exuberance in investors' expectations. A number of scholars and pundits had warned *ex ante* of the non-sustainability of certain pre-crisis economic developments, as documented by Bezemer, 2011. Those who warned of the crisis identified as the common elements in their thinking the destabilizing role of uncontrolled expansion of financial assets and debt, the flow of funds, and the impact of behaviors resulting from uncertainty and bounded rationality. However, these analyses were strongly at variance with the widespread belief in the "Great Moderation" (Stock and Watson, 2003) and in the beneficial and stabilizing properties of financial derivatives markets by their supposed virtue of dispersing risk globally (Summers et al., 1999; Greenspan, 2005). In hindsight, it became clear to everyone that it was a grave mistake to ignore issues related to systemic coupling and resulting cascade risks (Bartram, Brown, and Hund, 2009; Hellwig, 2009). But could we do better in the future and identify unsustainable market exuberance *ex ante*, to diagnose stress in the system in real time before a crisis starts?

The present article offers a new perspective on identifying growing risk by focussing on growth expectations embodied in financial option markets. We analyze data from the decade around the Global Financial Crisis of 2008 over the period from 2003 to 2013.<sup>1</sup> We retrieve the full risk-neutral probability measure of implied returns and analyze its characteristics over the course of the last decade. Applying a change point detection method (Killick, Fearnhead, and Eckley, 2012), we endogenously identify the beginning and end of the Global Financial Crisis as indicated by the options data. We consistently identify the beginning and end of the Crisis to be June 2007 and May 2009,

---

<sup>1</sup>Related existing work has considered data from pre-crisis (Figlewski, 2010) and crisis (Birru and Figlewski, 2012).

which is in agreement with the timeline given by the Federal Reserve Bank of St. Louis, [2009](#).<sup>2</sup>

The resulting pre-crisis, crisis and post-crisis regimes differ from each other in several important aspects. First, during the pre-crisis period, but not in the crisis and post-crisis periods, we identify a continuing increase of S&P 500 expected returns. This corresponds to super-exponential growth expectations of the price. By contrast, regular expectation regimes prevail in the crisis and post-crisis periods. Second, the difference between realized and option-implied returns remains roughly constant prior to the crisis but diverges in the post-crisis phase. This phenomenon may be interpreted as an increase of the representative investor's risk aversion. Third, Granger-causality tests show that changes of option-implied returns Granger-cause changes of Treasury Bill yields with a lag of few days in the pre-crisis period, while the reverse is true at lags of 50 to 200 days in the post-crisis period. This role reversal suggests that Fed policy was responding to, rather than leading, the financial market development during the pre-crisis period, but that the economy returned to a “new normal” regime post-crisis.

The majority of related option market studies have used option data for the evaluation of risk. An early contribution to this strand of work is Aït-Sahalia and Lo, [2000](#) who proposed a nonparametric risk management approach based on a value at risk computation with option-implied state-price densities. Another popular measure of option-implied volatility is the Volatility Index (VIX), which is constructed out of options on the S&P 500 stock index and is meant to represent the market's expectation of stock market volatility over the next 30 days (Exchange, [2009](#)). Bollerslev and Todorov, [2011](#) extended the VIX framework to an “investor fears index” by estimating

---

<sup>2</sup>See section [6.3.2](#) for more details on market and policy events marking the Global Financial Crisis of 2008.

jump tail risk for the left and right tail separately. Bali, Cakici, and Chabi-Yo, [2011](#) define a general option-implied measure of riskiness taking into account an investor's utility and wealth leading to asset allocation implications. What sets our work apart is the focus on identifying the long and often slow build-up of risk during an irrationally exuberant market that typically precedes a crisis.

Inverting the same logic, scholars have used option price data to estimate the risk attitude of the representative investor as well as its changes. These studies, however, typically impose stationarity in one way or another. Jackwerth, [2000](#), for example, empirically derives risk aversion functions from option prices and realized returns on the S&P 500 index around the crash of 1987 by assuming a constant return probability distribution. In a similar way, Rosenberg and Engle, [2002](#) analyze the S&P 500 over four years in the early 1990s by fitting a stochastic volatility model with constant parameters. Bliss and Panigirtzoglou, [2004](#), working with data for the FTSE 100 and S&P 500, propose another approach that assumes stationarity in the risk aversion functions. Whereas imposing stationarity is already questionable in "normal" times, it is certainly hard to justify for a time period covering markedly different regimes as around the Global Financial Crisis of 2008. We therefore proceed differently and merely relate return expectations implicit in option prices to market developments, in particular to the S&P 500 stock index and yields on Treasury Bills. We use the resulting data trends explicitly to identify the pre-crisis exuberance in the trends of market expectations and to make comparative statements about changing risk attitudes in the market.

The importance of market expectation trends has not escaped the attention of many researchers who focus on 'bubbles' (Galbraith, [2009](#); Sornette, [2003](#); Shiller, [2005](#); Soros, [2009](#); Kindleberger and Aliber, [2011](#)). One of us summa-

rizes their role as follows: “In a given financial bubble, it is the expectation of future earnings rather than present economic reality that motivates the average investor. History provides many examples of bubbles driven by unrealistic expectations of future earnings followed by crashes” (Sornette, 2014). While there is an enormous econometric literature on attempts to test whether a market is in a bubble or not, to our knowledge our approach is the first trying to do so by measuring and evaluating the market’s expectations directly.<sup>3</sup>

This paper is structured as follows. Section 2 details the estimation of the risk-neutral return probability distributions, the identification of regime change points, and the causality tests regarding market returns and expectations. Section 3 summarizes our findings, in particular the evidence concerning pre-crisis growth of expected returns resulting in super-exponential price growth. Section 4 concludes with a discussion of our findings.

## 6.2 Materials and methods

### 6.2.1 Estimating risk-neutral densities

Inferring information from option exchanges is guided by the fundamental theorem of asset pricing stating that, in a complete market, an asset price is the discounted expected value of future payoffs under the unique risk-neutral measure (see e.g. Delbaen and Schachermayer, 1994). Denoting that measure by  $\mathbb{Q}$  and the risk-neutral density by  $f$ , respectively, the current price  $C_0$  of a standard European call option on a stock with price at maturity  $S_T$  and strike

---

<sup>3</sup>For the econometric literature regarding assessments as to whether a market is in a bubble or not see Stiglitz (1990) (and the corresponding special issue of the Journal of Economic Perspectives), Bhattacharya and Yu (2008) (and the corresponding special issue of the Review of Financial Studies), as well as Camerer (1989), Scheinkman and Xing (2003), Jarrow et al. (2011), Evanoff et al. (2012), Lleo and Ziemba, 2012, Anderson et al. (2013), Phillips et al. (2013), Hüslér, Sornette, and Hommes, 2013.

$K$  can therefore be expressed as

$$C_0(K) = e^{-r_f T} \mathbb{E}_0^{\mathbb{Q}} [\max(S_T - K, 0)] = e^{-r_f T} \int_K^{\infty} (S_T - K) f(S_T) dS_T, \quad (6.1)$$

where  $r_f$  is the risk-free rate and  $T$  the time to maturity. From this equation, we would like to extract the density  $f(S_T)$ , as it reflects the representative investor's expectation of the future price under risk-neutrality. Since all quantities but the density are observable, inverting equation (6.1) for  $f(S_T)$  becomes a numerical task.

Several methods for inverting have been proposed, of which Jackwerth, 2004 provides an excellent review. In this study, we employ a method by Figlewski, 2010 that is essentially model-free and combines standard smoothing techniques in implied-volatility space and a new method of completing the density with appropriate tails. Tails are added using the theory of Generalized Extreme Value distributions, which are capable of characterizing very different behaviors of extreme events.<sup>4</sup> This method cleverly combines mid-prices of call and put options by only taking into account data from at-the-money and out-of-the-money regions, thus recovering non-standard features of risk-neutral densities such as bimodality, fat tails, and general asymmetry.

Our analysis covers fundamentally different market regimes around the Global Financial Crisis. A largely nonparametric approach, rather than a parametric one, seems therefore appropriate, because an important question that we shall ask is whether and how distributions actually changed from one regime to the next. We follow Figlewski's method in most steps, and additionally weight points by open interest when interpolating in implied-volatility space – a proxy

---

<sup>4</sup>As Birru and Figlewski, 2012 note, the theoretically correct extreme value distribution class is the Generalized Pareto Distribution (GPD) because estimating beyond the range of observable strikes corresponds to the peak-over-threshold method. For our purposes, both approaches are known to lead to equivalent results.

of the information content of individual sampling points permitted by our data. We give a more detailed review of the method in appendix 6.4.

## 6.2.2 Data

We use end-of-day data for standard European call and put options on the S&P 500 stock index provided by Stricknet<sup>5</sup> for a period from January 1st, 2003 to October 23rd, 2013. The raw data includes bid and ask quotes as well as open interest across various maturities. For this study, we focus on option contracts with quarterly expiration dates, which usually fall on the Saturday following the third Friday in March, June, September and December, respectively. Closing prices of the index, dividend yields and interest rates of the 3-month Treasury Bill as a proxy of the risk-free rate are extracted from Thomson Reuters Datastream.

We apply the following filter criteria as in Figlewski, 2010. We ignore quotes with bids below \$0.50 and those that are larger than \$20.00 in the money, as such bids exhibit very large spreads. Data points for which the midprice violates no-arbitrage conditions are also excluded. Options with time to maturity of less than 14 calendar days are discarded, as the relevant strike ranges shrink to smaller and smaller lengths resulting in a strong peaking of the density.<sup>6</sup> We are thus left with data for 2,311 observations over the whole time period and estimate risk-neutral densities and implied quantities for each of these days.

---

<sup>5</sup>The data is accessible via stricknet.com, where it can be purchased retrospectively.

<sup>6</sup>Figlewski, 2010 points out that rollovers of hedge positions into later maturities around contract expirations may lead to badly behaved risk-neutral density estimates.

### 6.2.3 Subperiod classification

As the Global Financial Crisis had a profound and lasting impact on option-implied quantities, it is informative for the sake of comparison to perform analyses to subperiods associated with regimes classifiable as pre-crisis, crisis and post-crisis. Rather than defining the relevant subperiods with historical dates, we follow an endogenous segmentation approach for identifying changes in the statistical properties of the risk-neutral densities. Let us assume we have an ordered sequence of data  $x_{1:n} = (x_1, x_2, \dots, x_n)$  of length  $n$ , e.g. daily values of a moment or tail shape parameter of the risk-neutral densities over  $n$  days. A change point occurs if there exists a time  $1 \leq k < n$  such that the mean of set  $\{x_1, \dots, x_k\}$  is statistically different from the mean of set  $\{x_{k+1}, \dots, x_n\}$  (Killick, Fearnhead, and Eckley, 2012). As a sequence of data may also have multiple change points, various frameworks to search for them have been developed. The binary segmentation algorithm by Scott and Knott, 1974 is arguably the most established detection method of this kind. It starts by identifying a single change point in a data sequence, proceeds iteratively on the two segments before and after the detected change and stops if no further change point is found.

As in the case of estimating risk-neutral densities, we refrain from making assumptions regarding the underlying process that generates the densities and choose a nonparametric approach. We employ the numerical implementation of the binary segmentation algorithm by Killick, Fearnhead, and Eckley, 2012 with the cumulative sum test statistic (CUSUM) proposed by Page, 1954 to search for at most two change points. The idea is that the cumulative sum,  $S(t) := \sum_{i=1}^t x_i$ ,  $1 \leq t < n$ , will have different slopes before and after the change point. As opposed to moving averages, using cumulative sums allows rapid detection of both small and large changes. We state the mathematical

formulation of the test statistic in appendix 6.4.<sup>7</sup>

## 6.2.4 Determining lag-lead structures

Option-implied quantities may be seen as expectations of the (representative) investor under  $\mathbb{Q}$ . A popular question in the context of self-referential financial markets is whether expectations drive prices or vice versa. To get a feeling of the causality, we analyze the lag-lead structure between the time series based on the classical method due to Granger, 1969. Informally, ‘Granger causality’ means that the knowledge of one quantity is useful in forecasting another. Formally, given two time series  $X_t$  and  $Y_t$ , we test whether  $Y_t$  Granger-causes  $X_t$  at lag  $m$  as follows. We first estimate the univariate autoregression

$$X_t = \sum_{j=1}^m a_j X_{t-j} + \varepsilon_t, \quad (6.2)$$

where  $\varepsilon_t$  is an uncorrelated white-noise series. We then estimate the augmented model with lagged variables

$$X_t = \sum_{j=1}^m b_j X_{t-j} + \sum_{j=1}^m c_j Y_{t-j} + \nu_t, \quad (6.3)$$

where  $\nu_t$  is another uncorrelated white-noise series. An  $F$ -test shows if the lagged variables collectively add explanatory power. The null hypothesis “ $Y_t$  does not Granger cause  $X_t$ ” is that the unrestricted model (6.3) does not provide a significantly better fit than the restricted model (6.2). It is rejected if the coefficients  $\{c_j, j = 1, \dots, m\}$  are statistically different from zero as a group. Since the model is only defined for stationary time series, we will test for Granger causality with standardized incremental time series in identified

---

<sup>7</sup>Interested readers may consult Brodsky and Darkhovsky, 1993 as well as Csörgö and Horváth, 1997 for a deeper discussion of theory, applications, and potential pitfalls of these methods.



subperiods as described in section 6.2.3.

## 6.3 Results

### 6.3.1 First-to-fourth return moment analyses

We start by analyzing the moments and tail shape parameters of the option-implied risk-neutral densities over the whole period (see Figure 6.1). For comparability, we rescale the price densities by the S&P 500 index level  $S_t$ , i.e. assess  $f(S_T/S_t)$  instead of  $f(S_T)$ .<sup>8</sup> In general, we recover similar values to the ones found by Figlewski, 2010 over the period 1996 to 2008. The annualized option-implied log-returns of the S&P 500 stock index excluding dividends are defined as

$$r_t = \frac{1}{T-t} \int_0^\infty \log\left(\frac{S_T}{S_t}\right) f(S_T) dS_T. \quad (6.4)$$

They are on average negative with a mean value of  $-3\%$ , and exhibit strong fluctuations with a standard deviation of  $4\%$ . This surprising finding may be explained by the impact of the Global Financial Crisis and by risk aversion of investors as explained below. The annualized second moment, also called risk-neutral volatility, is on average  $20\%$  (standard deviation of  $8\%$ ). During the crisis from June 22nd, 2007 to May 4th, 2009, we observe an increase in risk-neutral volatility to  $29 \pm 12\%$ .

A skewness of  $-1.5 \pm 0.9$  and excess kurtosis of  $10 \pm 12$  indicate strong deviations from log-normality, albeit subject to large fluctuations.<sup>9</sup> During the crisis, we measure a third ( $-0.9 \pm 0.3$ ) and fourth moment ( $4.4 \pm 1.6$ ) of the risk-neutral densities closer to those of a log-normal distribution than before or after the

---

<sup>8</sup>We do not go into the analysis of the first moment, which, in line with efficient markets, is equal to 1 by construction of  $f(S_T/S_t)$  (up to discounting).

<sup>9</sup>For the sake of comparison, note that a log-normal distribution with standard deviation  $20\%$  has skewness of  $0.6$  and excess kurtosis of  $0.7$ . In particular, skewness is always positive.

crisis. Birru and Figlewski, 2012 find a similar dynamic using intraday prices for S&P 500 Index options. For the period from September 2006 until October 2007, they report an average skewness of  $-1.9$  and excess kurtosis of  $11.9$ , whereas from September to November 2008 these quantities change to  $-0.7$  and  $3.5$ , respectively.

As the fourth moment is difficult to interpret for a strongly skewed density, one must be careful with the implication of these findings. One interpretation is that, during crisis, investors put less emphasis on rare extreme events or potential losses, that is, on fat tails or leptokurtosis, while immediate exposure through a high standard deviation (realized risk) gains importance.<sup>10</sup> Another interpretation of the low kurtosis and large volatility observed during the crisis regime would be in terms of the mechanical consequences of conditional estimations. The following simple example illustrates this. Suppose that the distribution of daily returns is the sum of two Normal laws with standard deviations  $3\%$  and  $20\%$  and weights  $99\%$  and  $1\%$  respectively. This means that  $99\%$  of the returns are normally distributed with a standard deviation of  $3\%$ , and that  $1\%$  of the returns are drawn from a Gaussian distribution with a standard deviation of  $20\%$ . By construction, the unconditional excess kurtosis is non zero ( $27$  for the above numerical example). Suppose that one observes a rare spell of large negative returns in the range of  $-20\%$ . Conditional on these realizations, the estimated volatility is large, roughly  $20\%$ , while the excess kurtosis close to  $0$  a consequence of sampling the second Gaussian law (and Gaussian distributions have by construction zero excess kurtosis).

It is interesting to note that Jackwerth and Rubinstein, 1996 reported opposite behaviors in an early derivation of the risk-neutral probability distributions of

---

<sup>10</sup>In other words, this interpretation indicates that investors, during crisis, focus on the unfolding risk, while, during non-crisis regimes, investors worry more about possible/unlikely worst case scenarios. Related to this interpretation are hypothesis regarding human behavioral traits according to which risk-aversion versus risk-taking behaviors are modulated by levels of available attention (Gifford, 2013).

European options on the S&P 500 for the period before and after the crash of October 1987. They observed that the risk-neutral probability of a one-standard deviation loss is larger after the crash than before, while the reverse is true for higher-level standard deviation losses. The explanation is that, after the 1987 crash, option traders realized that large tail risks were incorrectly priced, and that the volatility smile was born as a result thereafter (Mackenzie, 2008).

The left tail shape parameter  $\xi$  with values of  $0.03 \pm 0.23$  is surprisingly small: a value around zero implies that losses are distributed according to a thin tail.<sup>11</sup> Moreover, with  $-0.19 \pm 0.07$ , the shape parameter  $\xi$  for the right tail is consistently negative indicating a distribution with compact support, that is, a finite tail for expected gains.

### 6.3.2 Regime change points

A striking feature of the time series of the moments and shape parameters is a change of regime related to the Global Financial Crisis, which is the basis of our subperiod classification. A change point analysis of the left tail shape parameter identifies the crisis period as starting from June 22nd, 2007 and ending in May 4th, 2009. As we obtain similar dates up to a few months for the change points in risk-neutral volatility, skewness and kurtosis, this identification is robust and reliable (see Table 6.1 for details). Indeed, the determination of the beginning of the crisis as June 2007 is in agreement with the timeline of the build-up of the financial crisis<sup>12</sup> (Federal Reserve Bank of

---

<sup>11</sup>When positive, the tail shape parameter  $\xi$  is related to the exponent  $\alpha$  of the asymptotic power law tail by  $\alpha = 1/\xi$ .

<sup>12</sup>(i) S&P's and Moody's Investor Services downgraded over 100 bonds backed by second-lien subprime mortgages on June 1, 2007, (ii) Bear Stearns suspended redemption of its credit strategy funds on June 7, 2007, (iii) S&P put 612 securities backed by subprime residential mortgages on credit watch, (iv) Countrywide Financial warned of "difficult conditions" on July 24, 2007, (v) American Home Mortgage Investment Corporation filed for Chapter 11 bankruptcy protection on July 31, 2007 and (vi) BNP Paribas, France's largest bank, halted

St. Louis, 2009), opening the gates of loss and bankruptcy announcements. Interestingly, when applying the analysis to option-implied returns instead, we detect the onset of the crisis only on September 5th, 2008, more than a year later. This reflects a time lag of the market to fully endogenize the consequences and implication of the crisis. This is in line with the fact that most authorities (Federal Reserve, US Treasury, etc.) were downplaying the nature and severity of the crisis, whose full blown amplitude became apparent to all only with the Lehmann Brother bankruptcy.

The identification of the end of the crisis in May 2009 is confirmed by the timing of the surge of actions from the Federal Reserve and the US Treasury Department to salvage the banks and boost the economy via “quantitative easing”, first implemented in the first quarter of 2009.<sup>13</sup> Another sign of a change of regime, which can be interpreted as the end of the crisis per se, is the strong rebound of the US stock market that started in March 2009, thus ending a strongly bearish regime characterized by a cumulative loss of more than 60% since its peak in October 2007.

Finally, note that the higher moments and tail shape parameters of the risk-neutral return densities in the post-crisis period from May 4th, 2009 to October 23, 2013 progressively recovered their pre-crisis levels.

### **6.3.3 Super-exponential return: bubble behavior before the crash**

Apart from the market free fall, which was at its worst in September 2008, the second most remarkable feature of the time series of option-implied stock

---

redemptions on three investment funds on Aug. 9, 2007 and so on.

<sup>13</sup>On March 18, 2009 the Federal Reserve announced to purchase \$750 billion of mortgage-backed securities and up to \$300 billion of longer-term Treasury securities within the subsequent year, with other central banks such as the Bank of England taking similar measures.

returns shown in kap6/Figures 6.1a and 6.2a is its regular *rise* in the years prior to the crisis. For the pre-crisis period from January 2003 to June 2007, a linear model estimates an average increase in the option-implied return of about 0.01% per trading day ( $p$ -value  $< 0.001$ ,  $R^2 = 0.82$ , more details can be found in Table 6.2). As a matter of fact, this increase is also present in the realized returns, from January 2003 until October 2007, i.e. over a slightly longer period, as shown in Figure 6.2a. Note, however, that realized returns have a less regular behavior than the ones implied by options since the former are realized whereas the latter are expected under  $\mathbb{Q}$ . An appropriate smoothing such as the exponentially weighted moving average is required to reveal the trend, see Figure 6.2a for more details.

In the post-crisis period, in contrast, the option-implied returns exhibit less regularity, with smaller upward trends punctuated by abrupt drops. We find that option-implied returns rise on average 0.003% per trading day from May 2009 to October 2013 ( $p$ -value  $< 0.001$ ). However, a coefficient of determination of  $R^2 = 0.20$  suggests that this period is in fact not well-described by a linear model.

To the best of our knowledge, super-exponential price growth expectations have not previously been identified as implied by options data. This finding has several important implications that we shall now detail.

The upward trends of both option-implied and realized returns pre-crisis signal a transient “super-exponential” behavior of the market price, here of the S&P500 index. To see this, if the average return  $r(t) := \ln[p(t)/p(t-1)]$  grows, say, linearly according to  $r(t) \approx r_0 + \gamma t$  as can be approximately observed in Figure 6.2a from 2003 to 2007, this implies  $p(t) = p(t-1)e^{r_0 + \gamma t}$ , whose solution is  $p(t) = p(0)e^{r_0 t + \gamma t^2}$ . In absence of the rise of return ( $\gamma = 0$ ), this recovers the standard exponential growth associated with the usual compounding of inter-

ests. However, as soon as  $\gamma > 0$ , the price is growing much faster, in this case as  $\sim e^{t^2}$ . Any price growth of the form  $\sim e^{t^\beta}$  with  $\beta > 1$  is faster than exponential and is thus referred to as “super-exponential.” Consequently, if the rise of returns is faster than linear, the super-exponential acceleration of the price is even more pronounced. For instance, Hüsler, Sornette, and Hommes, 2013 reported empirical evidence of the super-exponential behaviour  $p(t) \sim e^{e^t}$  in controlled lab experiments (which corresponds formally to the limit  $\beta \rightarrow \infty$ ). Corsi and Sornette, 2014 presented a simple model of positive feedback between the growth of the financial sector and that of the real economy, which predicts even faster super-exponential behaviour, termed transient finite-time singularity (FTS). This dynamics can be captured approximately by the novel FTS-GARCH, which is found to achieve good fit for bubble regimes (Corsi and Sornette, 2014). The phenomenon of super-exponential price growth during a bubble can be accommodated within the framework of a rational expectation bubble (Blanchard, 1979; Blanchard and Watson, 1982), using for instance the approach of Johansen, Sornette, and Ledoit, 1999; Johansen, Ledoit, and Sornette, 2000 (JLS model).<sup>14</sup> In a nutshell, these models represent crashes by jumps, whose expectations yield the crash hazard rate. Consequently, the condition of no-arbitrage translates into a proportionality between the crash hazard rate and the instantaneous conditional return: as the return increases, the crash hazard rate grows and a crash eventually breaks the price unsustainable ascension. See Sornette et al., 2013 for a recent review of many of these models.

Because super-exponential price growth constitutes a deviation from a long-term trend<sup>15</sup> that can only be transient, it provides a clear signature of a non-

---

<sup>14</sup>Alternative rational expectations frameworks include Sornette and Andersen, 2002; Lin and Sornette, 2011; Lin, Ren, and Sornette, 2014. Also related is the literature on mildly explosive bubbles (Phillips, Wu, and Yu, 2011; Phillips, Shi, and Yu, 2012).

<sup>15</sup>Long-term exponential growth is the norm in economics, finance and demographics. This simply reflects the Gibrat law of proportional growth (Gibrat, 1931), which has an

sustainable regime whose growing return at the same time embodies and feeds over-optimism and herding through various positive feedback loops. This feature is precisely what allows the association of these transient super-exponential regimes with what is usually called a “bubble” (Kaizoji and Sornette, 2009), an approach that has allowed bubble diagnostics ex-post and ex-ante (see e.g. Johansen, Sornette, and Ledoit, 1999; Sornette, 2003; Lin and Sornette, 2011; Sornette and Cauwels, 2014a; Sornette and Cauwels, 2014b).

### 6.3.4 Dynamics of realized and option-implied returns

Realized S&P 500 and option-implied S&P 500 returns exhibit different behaviors over time (Figure 6.2a). Note that this difference persists even after filtering out short-term fluctuations in the realized returns.<sup>16</sup> During the pre-crisis period (from January 2003 to June 2007), the two grow at roughly the same rate, but the realized returns grow are approximately 8% larger than the option-implied returns. This difference can be ascribed to the “risk premium” that investors require to invest in the stock market, given their aggregate risk aversion.<sup>17</sup> This interpretation of the difference between the two return quantities as a risk premium, which one may literally term “realized-minus-implied risk premium”, is based on the fact that the option-implied return is determined under the risk-neutral probability measure while the realized return

---

extremely broad domain of application (Yule, 1925; Simon, 1955; Saichev, Malevergne, and Sornette, 2009).

<sup>16</sup>Realized S&P 500 returns show more rapid fluctuations than option-implied ones, which is not surprising given that the former are realized whereas the latter are expected (under  $\mathbb{Q}$ ). In this section we only focus on dynamics on a longer timescale, thus Figure 6.2a presents realized returns smoothed by an exponential weighted moving average (EWMA) of daily returns over 750 trading days. Different values or smoothing methods lead to similar outcomes.

<sup>17</sup>To understand variations in the risk premium in relation to the identification of different price regimes, we cannot rely on many of the important more sophisticated quantitative methods for derivation of the the risk premium, but refer to the literature discussed in the introduction. There are many avenues for promising future research to develop hybrid approaches between these more sophisticated approaches and ours which a priori allows the premium to vary freely over time.

is, by construction, unfolding under the real-world probability measure.<sup>18</sup> In other words, the risk-neutral world is characterized by the assumption that all investors agree on asset prices just on the basis of fair valuation. In contrast, real-world investors are in general risk-averse and require an additional premium to accept the risks associated with their investments. During the crisis, realized returns plunged faster and deeper in negative territory than the option-implied returns, then recovered faster into positive and *growing* regimes post-crisis. Indeed, during the crisis, the realized-minus-implied risk premium surprisingly became negative.

While the option-implied returns exhibit a stable behavior punctuated by two sharp drops in 2010 and 2011 (associated with two episodes of the European sovereign debt crisis), one can observe that the realized returns have been increasing since 2009, with sharp drop interruptions, suggesting bubbly regimes diagnosed by transient super-exponential dynamics (Sornette and Cauwels, 2014b). Furthermore, the realized-minus-implied risk premium has steadily grown since 2009, reaching approximately 16% at the end of the analyzed period (October 2013), i.e. twice its pre-crisis value. This is qualitatively in agreement with other analyses (Graham and Harvey, 2013) and can be rationalized by the need for investors to be remunerated against growing uncertainties of novel kinds, such as created by unconventional policies and sluggish economic recovery.<sup>19</sup>

---

<sup>18</sup>The standard definition, which usually takes the expected 10-year S&P 500 return relative to a 10-year U.S. Treasury bond yield (Fernandez, 2013; Duarte and Rosa, 2013) captures different information.

<sup>19</sup>An incomplete list of growing uncertainties at that time is: instabilities in the middle-East, concerns about sustainability of China's growth and issues of its on-going transitions, and many other uncertainties involving other major economic players, such as Japan, India and Brazil, quantitative easing operations in the US, political will from European leaders and actions of the ECB to hold the eurozone together.



### 6.3.5 Granger causality between option-implied returns and the 3-month Treasury Bill

We now examine possible Granger-causality relationships between option-implied returns and 3-month Treasury Bill yields. First note that option-implied returns and the 3-month Treasury Bill yields reveal a much weaker correlation than between realized returns and option-implied returns. A casual glance at Figure 6.2b suggests that their pre-crisis behaviors are similar, up to a vertical translation of approximately 3%. To see if the Fed rate policy might have been one of the drivers of the pre-crisis stock market dynamics, we perform a Granger causality test in both directions. Since a Granger test is only defined for stationary time series, we consider first differences in option-implied S&P 500 returns and 3-month Treasury Bill yields, respectively. Precisely, we define

$$SP_t = r_t - r_{t-1}, \quad TB_t = y_t - y_{t-1}. \quad (6.5)$$

where  $r_t$  is the option-implied return (6.4) and  $y_t$  is the Bill yield at trading day  $t$ . Before testing, we standardize both  $SP_t$  and  $TB_t$ , i.e. we subtract the mean and divide by the standard deviation, respectively.

There is no evidence that Federal Reserve policy has influenced risk-neutral option-implied returns over this period, as a Granger causality test fails to reject the relevant null at any lag (see Table 6.3 and Figure 6.3a). The other direction of Granger causality is more interesting, revealing Granger-causal influence of the option-implied returns on the 3-month Treasury Bill. A Granger causality test for  $SP_t$  on  $TB_t$  rejects the null for a lag of  $m = 5$  trading days. This suggests that the Fed policy has been responding to, rather than leading, the development of the market expectations during the pre-crisis period. Previous works using a time-adaptive lead-lag technique had only documented

that stock markets led Treasury Bills yields as well as longer term bonds yields during bubble periods (Zhou and Sornette, 2004; Guo et al., 2011). It is particularly interesting to find a Granger causality of the forward-looking expected returns, as extracted from option data, onto a backward-looking Treasury Bill yield in the pre-crisis period and the reverse thereafter. Thus, expectations were dominant in the pre-crisis period as is usually the case in efficient markets, while realized monetary policy was (and still is in significant parts) shaping expectations post-crisis (as shown in Table 6.3 and Figure 6.3b). The null of no influence is rejected for Treasury Bill yields Granger causing option-implied returns lagged by 50 to 200 days. This is coherent with the view that the Fed monetary policy, developed to catalyze economic recovery via monetary interventionism, has been the key variable influencing investors and thus options/stock markets.

Analyses of Granger causality with respect to realized returns yield no comparable results. Indeed, mutual influences with respect to Bivariate Granger tests involving the first difference time series of realized returns (with both option-implied returns and Treasury Bill yields) confirm the results that would have been expected. Both prior to and after the crisis, Treasury Bill yields Granger-cause realized returns over long time periods ( $p < 0.1$  for lags of 150 and 200 trading days, respectively), whereas option-implied returns Granger-cause realized ones over short time periods ( $p < 0.01$  for a lag of 5 trading days).

## 6.4 Conclusion

We have extracted risk-neutral return probability distributions from S&P 500 stock index options from 2003 to 2013. Change point analysis identifies the

crisis as taking place from mid-2007 to mid-2009. The evolution of risk-neutral return probability distributions characterizing the pre-crisis, crisis and post-crisis regimes reveal a number of remarkable properties. Indeed paradoxically at first sight, the distributions of expected returns became very close to a normal distribution during the crisis period, while exhibiting strongly negative skewness and especially large kurtosis in the two other periods. This reflects that investors may care more about the risks being realized (volatility) during the crisis, while they focus on potential losses (fat left tails, negative skewness and large kurtosis) in quieter periods.

Our most noteworthy finding is the continuing increase of the option-implied average returns during the pre-crisis (from January 2003 to mid-2007), which more than parallels a corresponding increase in realized returns. While a constant average return implies standard exponential price growth, an increase of average returns translates into super-exponential price growth, which is unsustainable and therefore transient. This finding corroborates previous reports on increasing realized returns and accelerated super-exponential price trajectories, which previously have been found to be hallmarks of exuberance and bubbles preceding crashes.

Moreover, the comparison between realized and option-implied expected returns sheds new light on the development of the pre-crisis, crisis and post-crisis periods. A general feature is that realized returns adapt much faster to changes of regimes, indeed often overshooting. Interpreted as a risk premium, literally the “realized-minus-implied risk premium”, these overshoots can be interpreted as transient changes in the risk perceptions of investors. We find that the realized-minus-implied risk premium was approximately 8% in the pre-crisis, and has doubled to 16% in the post-crisis period (from mid-2009 to October 2013). This increase is likely to be associated with growing uncertain-

ties and concern with uncertainties, fostered possibly by unconventional financial and monetary policy and unexpectedly sluggish economic recovery.

Finally, our Granger causality tests demonstrate that, in the pre-crisis period, changes of option-implied returns lead changes of Treasury Bill yields with a short lag, while the reverse is true with longer lags post-crisis. In a way, the post-crisis period can thus be seen as a return to a “normal” regime in the sense of standard economic theory, according to which interest rate policy determines the price of money/borrowing, which then spills over to the real economy and the stock market. What makes it a “new normal” (El-Erian, 2011) is that zero-interest rate policies in combination with other unconventional policy actions actually dominate and bias investment opportunities. The pre-crisis reveals the opposite phenomenon in the sense that expected (and realized returns) lead the interest rate, thus in a sense “slaving” the Fed policy to the markets. It is therefore less surprising that such an abnormal period, previously referred to as the “Great Moderation” and hailed as the successful taming of recessions, was bound to end in disappointments as a bubble was built up (Sornette and Woodard, 2010; Sornette and Cauwels, 2014a).

These results make clear the existence of important time-varying dynamics in both equity and variance risk premia, as exemplified by the difference between the pre- and post-crisis periods in terms of the Granger causalities. The option-implied returns show that expectations have been changed by the 2008 crisis, and this confirms another massive change of expectations following the crash of October 1987, embodied in the appearance of the volatility smile (Mackenzie, 2008). We believe that extending our analysis to more crises will confirm the importance of accounting for changes of expectations and time-varying premia, and we will address these issues in future research.

## Returns and distributional moments implied by S&P 500 options

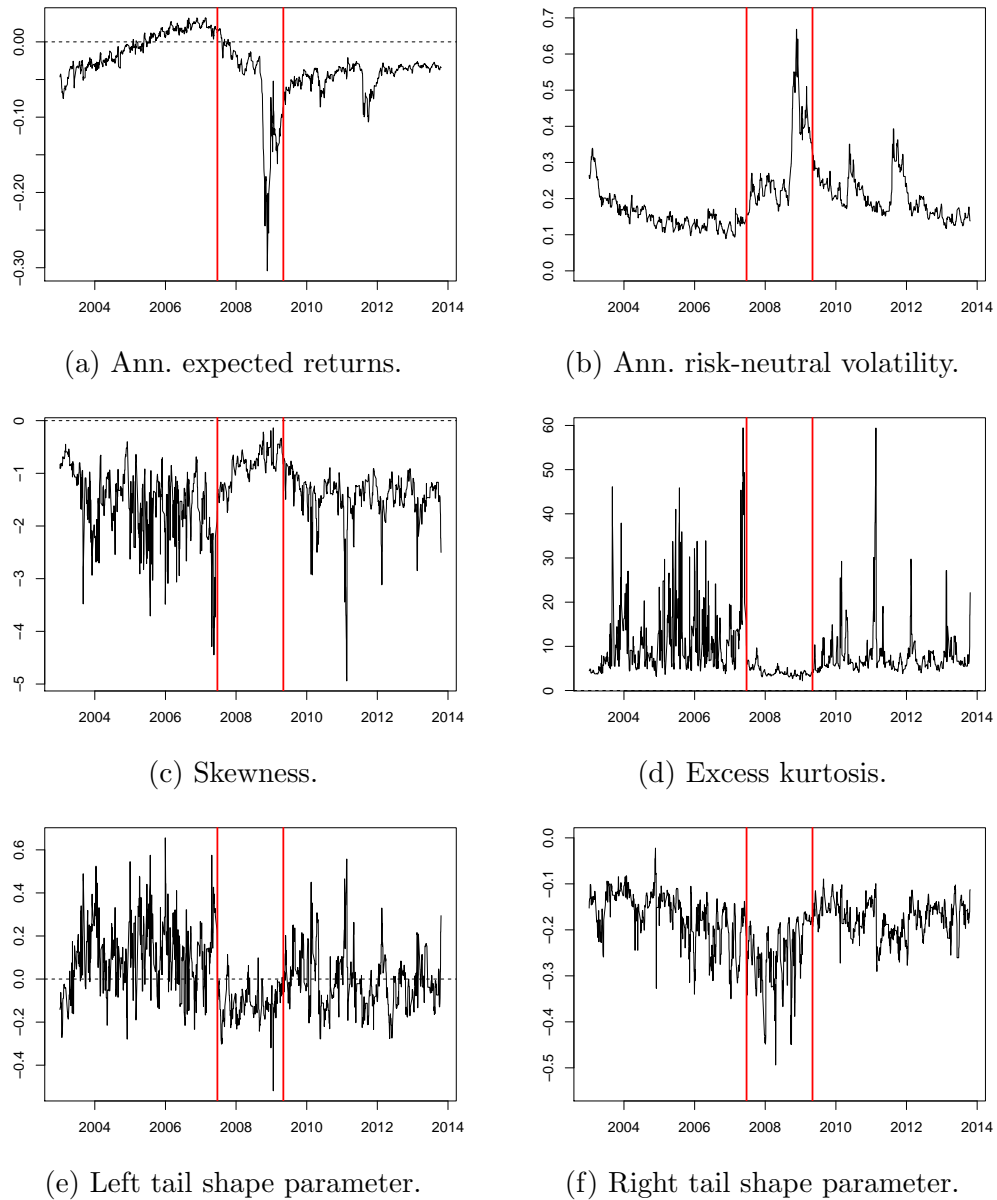




Figure 6.1: This figure presents returns and distributional moments implied by S&P 500 options. Structural changes around the financial crisis are identified consistently with a change point analysis of the means of the higher moments and tail shape parameters (vertical lines).

## Option-implied returns vs realized returns and Treasury Bill yields




(a) Annualized realized returns and option-implied S&P 500 returns. Realized returns are calculated by exponential weighted moving average (EWMA) smoothing of daily returns over 750 trading days.




(b) 3-month Treasury Bill yields and annualized option-implied S&P 500 returns (5-day moving averages).

Figure 6.2: This figure presents time series of option-implied S&P 500 returns, realized returns and Treasury Bill yields over the time period 2003–2013.

### Subperiod Granger causality tests



(a) Pre-crisis: January 1st, 2003 to June 22nd, 2007. No evidence for Treasury Bill yields Granger causing option-implied S&P 500 returns at any lag, but rather that option-implied S&P 500 returns Granger cause Treasury Bill yields at lags of a few trading days.



(b) Post-crisis: May 4th, 2009 to October 23rd, 2013. Treasury Bill yields Granger cause option-implied S&P 500 returns over a large range of lags.

Figure 6.3: Subperiod Granger causality tests on incremental changes in annualized option-implied S&P 500 returns and 3-month Treasury Bill yields. The  $p = 0.05$  line is plotted as dashed black.

### Risk-neutral density implied by options




Figure 6.4: Risk-neutral density implied by S&P 500 options from 2010-10-06 for index levels on 2010-12-18. The empirical part is directly inferred from option quotes, whereas tails must be estimated to account for the range beyond observable strike prices. Together, they give the full risk-neutral density. The method is reviewed in section 6.2.1 and appendix 6.4.



Table 6.1: Start and end dates of the Global Financial Crisis as identified by a change point analysis of statistical properties of option-implied risk-neutral densities. The dates found in the left tail shape parameter and higher moments identify consistently the crisis period as ca. June 2007 to ca. October 2009. Interestingly, the return time series signals the beginning only more than a year later, as September 2008. See section 6.2.3 for a review of the method, and 6.3.2 for a more detailed discussion of the results.

<i>Variable</i>	<i>Crisis start date</i>	<i>Crisis end date</i>
Left tail shape parameter	2007-06-22***	2009-05-04***
Right tail shape parameter	2005-08-08***	2009-01-22***
Risk-neutral volatility	2007-07-30***	2009-11-12***
Skewness	2007-06-22***	2009-10-19***
Kurtosis	2007-06-19***	NA <sup>a</sup>
Option-implied returns	2008-09-05***	2009-07-17***

*Note:*

\*p<0.1; \*\*p<0.01; \*\*\*p<0.001

<sup>a</sup> No change point indicating a crisis end date found.

Table 6.2: Results of a linear regression of option-implied returns of the S&P 500 index on time (trading days) by sub-period. In particular, a linear model fits well the pre-crisis, indicating the regular rise of expected returns, but not the post-crisis. This translates into super-exponential price growth expectations in the pre-crisis period. Standard deviations are in parentheses.

	<i>Option-implied returns (in percent):</i>		
	Pre-crisis	Crisis	Post-crisis
linear coefficient per trading day	0.009*** (0.0001)	-0.043*** (0.002)	0.003*** (0.0002)
Constant	-4.747*** (0.072)	2.836*** (0.485)	-5.775*** (0.097)
Observations	942	411	958
R <sup>2</sup>	0.820	0.520	0.196

*Note:*

\*p<0.1; \*\*p<0.01; \*\*\*p<0.001

Table 6.3: This table reports the results of a Granger-causality test of option-implied S&P 500 returns and Treasury Bill yields by sub-period. While we do not find evidence that Treasury Bill yields may have Granger-caused implied returns pre-crisis, there is Granger-influence in the other direction at a lag of 5 trading days both pre- and especially post-crisis. Notably, our test strongly suggests that post-crisis Treasury Bill yields have Granger-causal influence on option-implied returns at lags of 50 to 200 trading days.

Pre-crisis				
Lag	<i>S&amp;P Granger-causes T-Bill</i>		<i>T-Bill Granger-causes S&amp;P</i>	
	F-ratio <sup>a</sup>	Degrees of freedom	F-ratio <sup>a</sup>	Degrees of freedom
5	2.72*	5, 926	0.23	5, 926
50	0.84	50, 791	0.82	50, 791
100	0.82	100, 641	0.92	100, 641
150	0.92	150, 491	0.77	150, 491
200	0.86	200, 341	0.87	200, 341
250	0.95	250, 191	0.83	250, 191
Post-crisis				
Lag	<i>S&amp;P Granger-causes T-Bill</i>		<i>T-Bill Granger-causes S&amp;P</i>	
	F-ratio <sup>a</sup>	Degrees of freedom	F-ratio <sup>a</sup>	Degrees of freedom
5	1.95*	5, 942	0.56	5, 942
50	0.69	50, 807	1.37*	50, 807
100	0.79	100, 657	1.55**	100, 657
150	1.07	150, 507	1.32*	150, 507
200	1.16	200, 357	1.23*	200, 357
250	1.06	250, 207	1.18	250, 207

Note:

\*p<0.1; \*\*p<0.01; \*\*\*p<0.001

<sup>a</sup> Refers to the *F*-test for joint significance of the lagged variables.

# Estimating the risk-neutral density from option quotes

In this study, we estimate the option-implied risk-neutral density with a method developed by Figlewski, 2010, which is based on equation (6.1). For completeness, we shall briefly review the method as employed in this paper, but refer the interested reader to the original document for more detail. The raw data are end-of-day bid and ask quotes of European call and put options on the S&P 500 stock market index with a chosen maturity. Very deep out of the money options exhibit spreads that are large relative to the bid, i.e. carry large noise. Due to the redundancy of calls and puts, we may discard quotes with bid prices smaller than \$0.50. In this paper, we perform the calculation with mid-prices, which by inverting the Black-Scholes model translate into implied volatilities.

In a window of  $\pm\$20.00$  around the at-the-money level, the implied volatilities of put and call options are combined as weighted averages. The weights are chosen in order to ensure a smooth transition from puts to calls by gradually blending calls into puts when going to higher strikes. Below and above that window, we only use call and put data, respectively. We then fit a fourth order polynomial in implied volatility space. Here, we deviate slightly from Figlewski, 2010 because we use open interest as fitting weights. By doing so, we give more weight to data points carrying more market information. The Black-Scholes model transforms the fit in implied volatility space back to price space. The resulting density bulk is called “empirical density”.

To obtain a density estimate beyond the range of observable strike prices, we must append tails to the empirical part. Figlewski, 2010 proposes to add

tails of the family of generalized extreme value (GEV<sup>20</sup>) distributions with connection conditions: a) matching value at the 2%, 5%, 92% and 95% quantile points, and b) matching probability mass in the estimated tail and empirical density. An example can be seen in Figure 6.4. The empirical density together with the tails give the complete risk-neutral density.

## Change point detection

The following framework is used for significance testing in section 6.3.2 and Table 6.1. For more details, see Csörgö and Horváth, 1997. Let  $x_1, x_2, \dots, x_n$  be independent, real-valued observations. We test the “no change point” null hypothesis,

$$H_0 : \mathbb{E}(x_1) = \mathbb{E}(x_2) = \dots = \mathbb{E}(x_n), \quad (6)$$

against the “one change in mean” hypothesis,

$$H_1 : \text{there is a } k, 1 \leq k < n, \text{ such that } \mathbb{E}(x_1) = \dots = \mathbb{E}(x_k) \neq \mathbb{E}(x_{k+1}) = \dots = \mathbb{E}(x_n), \quad (7)$$

using the auxiliary functions

$$A(x) := \sqrt{2 \log \log x}, \quad D(x) := 2 \log \log x + \frac{1}{2} \log \log \log x - \frac{1}{2} \log \pi. \quad (8)$$

Then, following corollary 2.1.2 and in light of remark 2.1.2. Csörgö and Horváth, 1997, pp. 67-68, under mild regularity conditions,  $H_0$  and for large sample sizes, one has

$$P \left( A(n) \max_k \frac{1}{\hat{\sigma}_n} \left( \frac{n}{k(n-k)} \right)^{1/2} \left| S(k) - \frac{k}{n} S(n) \right| - D(n) \leq t \right) = \exp(-2e^{-t}), \quad (9)$$

---

<sup>20</sup>See Embrechts, Klüppelberg, and Mikosch, 1997 for a detailed theoretical discussion of GEV distributions and modeling extreme events.

where  $\hat{\sigma}_n$  is the sample standard deviation and  $S(t) := \sum_{i=1}^t x_i$  the cumulative sum of observations.

## References

- Aït-Sahalia, Yacine and Andrew W Lo (2000). “Nonparametric risk management and implied risk aversion”. In: *Journal of Econometrics* 94.1-2, pp. 9–51.
- Bali, Turan G, Nusret Cakici, and Fousseni Chabi-Yo (2011). “A generalized measure of riskiness”. In: *Management Science* 57.8, pp. 1406–1423.
- Bartram, S.M., G.W. Brown, and J.E. Hund (2009). “Estimating systemic risk in the international financial system”. In: *Journal of Financial Economics* 86.3, pp. 835–869.
- Bezemer, Dirk J. (2011). “The credit crisis and recession as a paradigm test”. In: *Journal of Economics Issues* 45.1, pp. 1–18.
- Birru, Justin and Stephen Figlewski (2012). “Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008”. In: *Journal of Financial Markets* 15.2, pp. 151–180.
- Blanchard, O.J. (1979). “Speculative bubbles, crashes and rational expectations”. In: *Economics Letters* 3.387-389.
- Blanchard, O.J. and M.W. Watson (1982). “Bubbles, rational expectations and speculative markets”. In: *Crisis in Economic and Financial Structure: Bubbles, Bursts and Shocks*. Ed. by P. Wachtel.
- Bliss, Robert R and Nikolaos Panigirtzoglou (2004). “Option-implied risk aversion estimates”. In: *The Journal of Finance* 59.1, pp. 407–446.
- Bollerslev, Tim and Viktor Todorov (2011). “Tails, fears, and risk premia”. In: *The Journal of Finance* 66.6, pp. 2165–2211.

- Brodsky, E. and B.S. Darkhovsky (1993). *Nonparametric Methods in Change Point Problems*. Mathematics and Its Applications. Springer.
- Corsi, Fulvio and Didier Sornette (2014). “Follow the money: The monetary roots of bubbles and crashes”. In: *International Review of Financial Analysis* 32, pp. 47–59.
- Csörgö, Miklós and Lajos Horváth (1997). *Limit Theorems in Change-point Analysis*. Wiley series in probability and statistics. Wiley.
- Delbaen, Freddy and Walter Schachermayer (1994). “A general version of the fundamental theorem of asset pricing”. In: *Mathematische Annalen* 300.1, pp. 463–520.
- Duarte, Fernando and Carlo Rosa (2013). “The Equity Risk Premium: A Consensus of Models”. In: *SSRN WP 2377504*.
- El-Erian, Mohamed (2011). “Spain is not Greece and need not be Ireland”. In: *Financial Times* Retrieved 2011-08-18, February 3.
- Embrechts, Paul, Claudia Klüppelberg, and Thomas Mikosch (1997). *Modelling extremal events: for insurance and finance*. Vol. 33. Springer.
- Exchange, Chicago Board Options (2009). “The CBOE volatility index – VIX”. In: *White Paper*.
- Federal Reserve Bank of St. Louis (2009). “The Financial Crisis: A Timeline of Events and Policy Actions”. In: *timeline.stlouisfed.org*.
- Fernandez, Pablo (2013). “The Equity Premium in 150 Textbooks”. In: *SSRN WP 1473225*.
- Figlewski, Stephen (2010). “Estimating the Implied Risk Neutral Density”. In: *Volatility and Time Series Econometrics*. Ed. by Tim Bollerslev, Jeffrey Russell, and Mark Watson. Oxford: Oxford University Press.
- Galbraith, J.K. (2009). *The Great Crash 1929*. Mariner Books (Reprint Edition).

- Gibrat, R. (1931). “Les Inégalités économiques: Aux Inégalités des Richesses, à la Concentration des Entreprises, Aux Populations des Villes, Aux Statistiques des Familles, etc., d’une Loi Nouvelle: La Loi de l’Effect Proportionnel”. In:
- Gifford, S. (2013). “Risk and uncertainty”. In: *Z.J. Acs, D.B. Audretsch (eds.), Handbook of Entrepreneurship Research* International Handbook Series on Entrepreneurship 5, pp. 303–318.
- Graham, John R. and Campbell R. Harvey (2013). “The Equity Risk Premium in 2013”. In: *SSRN WP 2206538*.
- Granger, C. W. J. (1969). “Investigating Causal Relations by Econometric Models and Cross-spectral Methods”. In: *Econometrica* 37.3, pp. 424–438.
- Greenspan, Alan (2005). *Consumer Finance*. Federal Reserve System’s Fourth Annual Community Affairs Research Conference. Federal Reserve Board.
- Guo, Kun et al. (2011). “The US stock market leads the Federal funds rate and Treasury bond yields”. In: *PloS One* 6.8, e22794.
- Hellwig, Martin (2009). “Systemic Risk in the Financial Sector: An Analysis of the Subprime-Mortgage Financial Crisis”. In: *De Economist* 157.2, pp. 129–207.
- Hüsler, Andreas, Didier Sornette, and Cars H Hommes (2013). “Super-exponential bubbles in lab experiments: evidence for anchoring over-optimistic expectations on price”. In: *Journal of Economic Behavior & Organization* 92, pp. 304–316.
- Jackwerth, Jens Carsten (2000). “Recovering risk aversion from option prices and realized returns”. In: *Review of Financial Studies* 13.2, pp. 433–451.
- (2004). *Option-implied risk-neutral distributions and risk aversion*. Research Foundation of AIMR Charlottesville.

- Jackwerth, Jens Carsten and Mark Rubinstein (1996). “Recovering probability distributions from option prices”. In: *The Journal of Finance* 51.5, pp. 1611–1631.
- Johansen, Anders, Olivier Ledoit, and Didier Sornette (2000). “Crashes as critical points”. In: *International Journal of Theoretical and Applied Finance* 3.2, pp. 219–255.
- Johansen, Anders, Didier Sornette, and Olivier Ledoit (1999). “Predicting financial crashes using discrete scale invariance”. In: *Journal of Risk* 1.4, pp. 5–32.
- Kaizoji, T. and D. Sornette (2009). “Market Bubbles and Crashes”. In: *The Encyclopedia of Quantitative Finance*.
- Killick, Rebecca, Paul Fearnhead, and IA Eckley (2012). “Optimal detection of changepoints with a linear computational cost”. In: *Journal of the American Statistical Association* 107.500, pp. 1590–1598.
- Kindleberger, Charles P. and Robert Z. Aliber (2011). *Manias, Panics and Crashes: A History of Financial Crises*. Palgrave Macmillan; Sixth Edition.
- Lin, L. and D. Sornette (2011). “Diagnostics of Rational Expectation Financial Bubbles with Stochastic Mean-Reverting Termination Times”. In: *The European Journal of Finance*.
- Lin, Li, R.E. Ren, and Didier Sornette (2014). “The Volatility-Confined LPPL Model: A Consistent Model of ‘Explosive’ Financial Bubbles With Mean-Reversing Residuals”. In: *International Review of Financial Analysis* 33, pp. 210–225.
- Lleo, Bastien and William T. Ziemba (2012). “Stock market crashes in 2007 – 2009: were we able to predict them?” In: *Quantitative Finance* 12.8, pp. 1161–1187.
- Mackenzie, Donald (2008). *An Engine, Not a Camera: How Financial Models Shape Markets*. The MIT Press.



- Page, ES (1954). “Continuous inspection schemes”. In: *Biometrika* 41.1/2, pp. 100–115.
- Phillips, P. C. B., S.-P. Shi, and J. Yu (2012). “Testing for multiple bubbles 1: Historical episodes of exuberance and collapse in the S&P 500”. In: *Cowles Foundation Discussion Paper No. 1843*.
- Phillips, P. C. B., Yangru Wu, and J. Yu (2011). “Explosive behavior in the 1990s Nasdaq: when did exuberance escalate asset values?” In: *International Economic Review* 52.1, pp. 201–226.
- Rosenberg, Joshua V and Robert F Engle (2002). “Empirical pricing kernels”. In: *Journal of Financial Economics* 64.3, pp. 341–372.
- Saichev, A., Y. Malevergne, and D. Sornette (2009). “A Theory of Zipf’s Law and beyond”. In: *Lecture Notes in Economics and Mathematical Systems* 632, pp. 1–171.
- Scott, AJ and M Knott (1974). “A cluster analysis method for grouping means in the analysis of variance”. In: *Biometrics* 30.3, pp. 507–512.
- Shiller, Robert J (2005). *Irrational exuberance*. Random House LLC.
- Simon, H.A. (1955). “On a class of skew distribution functions”. In: *Biometrika* 52, pp. 425–440.
- Sornette, D. (2014). “Physics and Financial Economics (1776-2013): Puzzles, Ising and agent-based models”. In: *Rep. Prog. Phys.* 77, 062001 (28 pp.)
- Sornette, D. and J. V. Andersen (2002). “A Nonlinear Super-Exponential Rational Model of Speculative Financial Bubbles”. In: *International Journal of Modern Physics C* 13.2, pp. 171–188.
- Sornette, D. and P. Cauwels (2014a). “1980-2008: The Illusion of the Perpetual Money Machine and what it bodes for the future”. In: *Risks* 2, pp. 103–131.
- (2014b). “Financial Bubbles: Mechanism, diagnostic and state of the world (Feb. 2014)”. In: *Review of Behavioral Economics* (in press).

- Sornette, Didier (2003). *Why stock markets crash: critical events in complex financial systems*. Princeton University Press.
- Sornette, Didier and Ryan Woodard (2010). “Financial bubbles, real estate bubbles, derivative bubbles, and the financial and economic crisis”. In: *Economics Approaches to Large-Scale Business Data and Financial Crisis*. Springer, pp. 101–148.
- Sornette, Didier et al. (2013). “Clarifications to Questions and Criticisms on the Johansen-Ledoit-Sornette bubble Model”. In: *Physica A: Statistical Mechanics and its Applications* 392.19, pp. 4417–4428.
- Soros, G. (2009). *The Crash of 2008 and What it Means: The New Paradigm for Financial Markets*. Public Affairs; Revised edition.
- Stock, James H. and Mark W. Watson (2003). “Has the Business Cycle Changed and Why?” In: *NBER Macroeconomics Annual 2002, MIT Press* 17, pp. 159–230.
- Summers, Lawrence et al. (1999). *Over-the-Counter Derivatives Markets and the Commodity Exchange Act*. Report of The President’s Working Group on Financial Markets.
- Yule, G. U. (1925). “A Mathematical Theory of Evolution, based on the Conclusions of Dr. J. C. Willis, F.R.S.” In: *Philosophical Transactions of the Royal Society B* 213.402-410, pp. 21–87.
- Zhou, Wei-Xing and Didier Sornette (2004). “Causal slaving of the US treasury bond yield antibubble by the stock market antibubble of August 2000”. In: *Physica A: Statistical Mechanics and its Applications* 337.3, pp. 586–608.

## Chapter 7

Meritocratic mechanism  
design:

Theory and experiments

## Abstract

One of the fundamental tradeoffs underlying society is that between efficiency and equality. The challenge for the right design of many institutions and social mechanisms is to strike the right balance between the two –often conflicting– goals. Game-theoretic models of public-goods provision under ‘meritocratic matching’ succinctly capture this tradeoff: under zero meritocracy (societal order is random), theory predicts maximal inefficiency but perfect equality; higher levels of meritocracy (society matches contributors with contributors) are predicted to improve efficiency but come at the cost of growing inequality. This chapter is split into a theory part and an experimental part. In the theory part, we study the model’s stability properties and the predictions concerning the efficiency-equality tradeoffs in the context of voluntary contribution games. In the experimental part, we analyze behavior from an experiment that we conducted to test this tradeoff behaviorally. We make the astonishing finding that, notwithstanding theoretical predictions, higher levels of meritocracy increase both efficiency and equality, that is, meritocratic matching dissolves the tradeoff. Fairness considerations can explain the departures from theoretical predictions including the behavioral phenomena that lead to dissolution of the efficiency-equality tradeoff.

# Part 1: Theory

## Abstract

We study stability properties and the efficiency-equality tradeoff of a class of *meritocratic* group-matching mechanisms in the context of voluntary contribution games. The mechanisms assort players by their contributions, resulting equilibria critically depending on matching fidelity. Efficiency and stability summarize as follows. For low levels of meritocracy, the only equilibrium state is inefficient. Above a first threshold, several more efficient equilibria emerge, but only the inefficient equilibrium is stable. Above a second threshold, near-efficient equilibria become stable. This operationalization sheds light on critical transitions, enabled by meritocratic matching, between low-efficiency and high-efficiency equilibria. Transitions to more efficient equilibria come at inequality costs, implying a hard efficiency-equality tradeoff. Our analysis reveals that welfare is generally maximized at the second threshold level, at no meritocracy only under extreme inequality aversion, with low rates of return or in small populations.

**Acknowledgements.** The authors would like to thank Lukas Bischofberger for help with simulations, Bary Pradelski, Anna Gunnthorsdottir, Matthias Leiss, Michael Mäs, Francis Dennig and Stefan Seifert for helpful comments on earlier drafts, Peyton Young for help with framing of the questions, Luis Cabral for help with proposition (16), Ingela Alger and Jörgen Weibull for a helpful discussion, and finally members of GESS at ETH Zurich, the participants at *Norms Actions Games 2014* and at the *25<sup>th</sup> International Conference on Game Theory 2014* at Stony Brook, as well as anonymous referees for helpful feedback. All remaining errors are ours.

## 7.1 Motivation

The argument in favor of ‘meritocracy’ (Young, 1958a) is that meritocratic regime incentives, such as rewarding effort or performance, promise efficiency gains. In environments such as education, job matching, or marriage markets, however, the downside of meritocracy is that, through these incentives, inequalities may be exacerbated (Young 1958a, Greenwood et al. 2014). A meritocratic regime may therefore turn out to be ultimately undesirable from the perspective of a social planner if he is averse to inequality (Arrow, Bowles, and Durlauf 2000a).

In this paper, we study a class of meritocratic mechanisms in the context of voluntary contributions games. Players first simultaneously make their contributions and are then matched based on these decisions. We shall consider regimes in which players are assortatively but imperfectly grouped by their contributions, and shall refer to this as *meritocratic matching*. If the regime is not meritocratic enough (groups form too randomly), the only equilibrium outcome is universal non-contribution, resulting in complete inefficiency. For sufficiently meritocratic regimes, universal non-contribution remains an equilibrium, but more efficient equilibria may also be enabled. Positive contributions may become best replies under sufficient meritocracy because it promises to be matched in better groups with other players who also contribute positive amounts. Equilibria with higher contributions improve efficiency, but they come with payoff heterogeneity and inequality, and their stability is not guaranteed.

From the perspective of social planning, one must therefore address two orthogonal questions. First, given a fixed level of meritocracy, which equilibrium is stable? Second, what level of meritocracy maximizes welfare? Our analysis suggests that, other than in the aforementioned contexts such as education,

job matching, or marriage markets, an intermediate but substantial level of meritocracy generally maximizes welfare in our setting. This result surprises as one would expect that welfare comparisons depend more subtly on the social planner’s degree of inequality aversion. We obtain our result from focus on states that are *stochastically stable* (Foster and Young 1990, Young 1993), requiring the social planner to choose amongst stable states.

The rest of this paper is structured as follows. Next, we discuss related literature, including voluntary contributions mechanisms and the broad conceptual approach. In section 3, we develop a formal model of meritocratic matching, calculate its equilibria, and detail the stability and welfare properties. We conclude in section 4.

## 7.2 Related literature

Specifically, this paper provides theoretical underpinnings for *meritocratic matching* mechanisms in the context of voluntary contributions games. Our main contribution is to generalize such mechanisms to a full range of regimes, ranging from no meritocracy to full meritocracy, and to explore their stability and welfare properties. The first meritocratic matching mechanism, corresponding to full meritocracy in our model, was recently theoretically and experimentally analyzed in a seminal paper by Gunnthorsdottir et al., 2010. Several experimental studies have shown that such a mechanism results in near-efficient contribution levels (see Gunnthorsdottir et al. 2010; Gunnthorsdottir and Thorsteinsson 2010; Gunnthorsdottir, Vragov, and Shen 2010).<sup>1</sup> The present work extends existing the existing investigation in three directions. First, the fidelity of meritocratic matching is allowed to be imperfect

---

<sup>1</sup>See also a related experimental investigation of pairwise tournaments à la Becker, 1973; Cole, Mailath, and Postlewaite, 1992 played under a similar mechanism by “[Efficient Investment via Assortative Matching: A laboratory experiment](#)”.



allowing to close the gap between full meritocracy and random (re-)matching. Second, given any intermediate degree of meritocracy, the stability of alternative equilibria is assessed using evolutionary refinement concepts. Third, we compare the welfare of different meritocracies.

More generally, our work is intended, on the one hand, as a contribution to the game-theoretic underpinnings of meritocracy in relation to the efficiency-equality tradeoff, and, on the other hand, to the study of voluntary contributions mechanisms in the context of local public goods provision games. Voluntary contributions games (together with the ultimatum games) are the principal fruitflies of experimental economics; see Ledyard 1995 and Chaudhuri 2011a for reviews of that literature. The first formal model of a voluntary contributions game was introduced in Isaac, McCue, and Plott, 1985a, Isaac and Walker, 1988. A repeated implementation of this game with random group re-matching is due to Andreoni, 1988, and this implementation represents the “no meritocracy” regime in our model.<sup>2</sup> The important feature of “no meritocracy” is that group matching is essentially exogenous, and contribution decisions play no role. An important avenue has been to study alternative, non-random, mechanisms, as we consider here by matching contributions assortatively, an approach pioneered by Gunnthorsdottir et al., 2010.<sup>3</sup> A common finding of the emerging literature on endogenous group formation is that a variety of suitable mechanisms/dynamics of non-random group formation can stabilize higher contribution levels. In our case, this is achieved

---

<sup>2</sup>Many experiments use variants of Andreoni’s implementation (e.g. Andreoni, 1988; Andreoni, 1993; Andreoni, 1995; Palfrey and Prisbrey, 1996; Palfrey and Prisbrey, 1997; Goeree, Holt, and Laury, 2002; Ferraro and Vossler, 2010; Fischbacher and Gächter, 2010; Bayer, Renner, and Sausgruber, 2013; Nax et al., 2013).

<sup>3</sup>There are other mechanisms. For example, Cinyabuguma, Page, and Putterman, 2005, Charness and Yang, 2008 consider endogenous group formation via voting; Ehrhart and Keser, 1999, Ahn, Isaac, and Salmon, 2008 study the effects of free group entry and exit; Coricelli, Fehr, and Fellner, 2004 analyze roommate-problem stable matching in pairwise-generated public goods; Page, Putterman, and Unel, 2005 study rematching based on reputation; Brekke, Nyborg, and Rege, 2007, Brekke et al., 2011 consider the effects of signaling.

via meritocratic matching.

Although the term “meritocracy” was only introduced in 1958 by Michael Young (Young, 1958a), the meritocratic principle underlying institutional mechanisms can be traced back in the history of many independent cultures. Indeed, several institutions of early modern civilizations (e.g. ancient China and ancient Greece) were designed explicitly to be meritocratic, and such practice was advocated by their thinkers (e.g. Confucius, Aristotle, and Plato). Historically, these institutions included the selection of officials and councilmen, reward and promotion schemes, and access to education.<sup>4</sup> Until today, meritocratic institutions like the Chinese civil service examination are in place. Other examples include honorary circles, bonus wage schemes, etc. Our paper is a first stab at studying the stability and welfare properties of such mechanisms, here in the context of voluntary contributions games.<sup>5</sup>

Meritocracy’s incentive given to contributory behaviors by our mechanisms is the promise of being matched into a better group. Importantly, meritocracy ‘works’ in our model –if it does– despite agents maximizing only in their own material payoff. Our paper therefore complements research on cooperative phenomena that arise from non-selfish preferences and altruism (Simon 1990, Bowles and Gintis 2011), in particular in public goods games (e.g. Fehr and Camerer, 2007). In the terminology of Allchin, 2009, our paper therefore studies a ‘system’ rather than moral ‘acts’ or ‘intentions’. In our mechanism, the system assortments contributions, i.e. actions. Other assortative systems are also known to lead to cooperation in social dilemma situations. In evolutionary biology, for example, kin selection does (e.g. Hamilton 1964a, Hamilton 1964b, Nowak 2006), so does local interaction and/or assortative matching of prefer-

---

<sup>4</sup>See, for example, Lane, 2004 for a description of the reward and promotion scheme in Genghis Khan’s army. Another famous example is China’s civil service examination (Miyazaki, 1976).

<sup>5</sup>There is a related literature on team-based reward schemes in labour market applications (e.g. Dickinson and Isaac 1998; Irlenbusch and Ruchala 2008).

ences (Alger and Weibull 2013, Grund, Waloszek, and Helbing 2013).

What distinguishes our meritocratic matching mechanism is that the social planner that we have in mind can observe actions and is able to group people, but cannot transfer payoffs or change group size. Meritocratic matching is therefore a mechanism that leaves the basic payoffs of the game unchanged. In principle, we are agnostic as to the origins of such an institution, even though there is evidence that players endogenously may implement such a system over time (Ones and Putterman, 2007). Players contribute based on purely egoistic and fully rational motivations, without reputation-sensitive concerns, and without hope to be “recognized” for contributing or fear to be “stigmatized” for free-riding (e.g. Andreoni and Petrie 2004, Samek and Sheremeta 2014).

### 7.3 Meritocratic matching

Before we proceed to formalize the set-up of our model, we would like to provide more intuition for the basic flavour of *meritocratic matching*. While none of the following real-world examples of institutions coincides one-to-one with meritocratic matching as it will be instantiated in our simple model of a linear and symmetric public goods game, they do mirror meritocratic matching’s key features. Importantly, all of these real-world examples are typically both noisy and not always fair. The first example is school/university admission. Entrance examinations to schools or universities assort individuals based on an imperfect measure of applicants’ adequacies for different streams of education and to enter different schools. An important feature of this sorting mechanism is that the resulting differences in educational quality amongst the different schools are not only determined by the institutional design, but also by the

different quality levels of students present in them. Better students tend to study with better students, and worse students with worse students. The incentive to work hard for the examinations is getting into a good school. The second example is team-based payment. Imagine an assortative employment regime with team-based payments that rewards employees for performance by matching them with similarly performant other employees. Real-world situations with this structure include trading desks in large investment banks, and again this type of competitive grouping incentivizes hard work through promise of being matched into better teams. Team formation in professional sports also has features that are similar to this: performant athletes tend to be rewarded by joining successful teams with better contracts.

### 7.3.1 The model

Suppose population  $N = \{1, 2, \dots, n\}$  plays the following game, of which all aspects are common knowledge. The game is divisible into three steps. First, players make simultaneous voluntary contributions. Second, players receive ranks that imperfectly represent their contributions. Third, groups and payoffs realize based on the ranking.

#### *Step 1. Voluntary contributions*

Player  $i \in N$  decide simultaneously whether to *contribute* or to *free-ride*; we shall write  $c_i = 0$  for free-riding and  $c_i = 1$  for contributing, yielding the contribution vector  $c = \{c_i\}_{i \in N}$ . Given some player  $i$ , denote by  $c_{-i}$  the contribution vector excluding him.

*Note:* The restriction to a binary action (all-or-nothing) action set comes without loss of generality in the stage game as it is a general feature of equilibria to polarize in this way (see Gunnthorsdottir et al. 2010). We commit to this

binary action structure in order to facilitate our evolutionary analysis.

***Step 2. Ordering as a function of contributions***

An authority imperfectly observes the contribution vector  $c$  and/ or imperfectly ranks players according to observed contributions. The measure of ranking precision is given by parameter  $\beta \in [0, 1]$ . The characteristics of the regimes summarize as follows: (i) *no meritocracy* ( $\beta = 0$ ), all rankings are equally likely, and all players have the same expected rank; (ii) in *full meritocracy* ( $\beta = 1$ ), only “perfect” rankings are possible so that all contributors will have a higher rank than all free-riders; (iii) in the *intermediate meritocracy* range, when  $\beta \in (0, 1)$ , all rankings have positive probability, but enough contributors have a higher expected rank than free-riders.

Formally, let  $\Pi = \{\pi_1, \pi_2, \dots, \pi_{n!}\}$  be the set of *orderings* (permutations) of  $N$ . Given any  $\pi \in \Pi$ , denote by  $k_i$  the case when rank  $k \in \{1, 2, \dots, n\}$  is taken by player  $i \in \{1, 2, \dots, n\}$ . Write  $\hat{\pi}$  for a *perfect* ordering if, for all pairs of players  $i, j$ ,  $k_i < k_j \Rightarrow c_i \geq c_j$ , that is, all free-riders are ranked below contributors. Any other ordering is called a *mixed* ordering, and is denoted by  $\tilde{\pi}$  (i.e. at least one free-rider is ranked above a contributor). Given regime  $\beta \in [0, 1]$ , the probability distribution over orderings,  $P(\Pi)$ , is a function of  $\beta$  and  $c$ ,  $P(\Pi) = F(c, \beta)$ . Write  $f_\pi^\beta$  for the probability of a particular ordering,  $\pi \in \Pi$ , under  $\beta$ . Similarly, write  $f_{ik}^\beta$  for the probability that agent  $i$  takes rank  $k$  given  $\beta$ , and  $\bar{k}_i^\beta$  for  $i$ 's expected rank. We shall write  $\bar{k}_i^\beta(c_i)$  to indicate that  $i$ 's expected rank is a function of his contribution. Finally, define  $\mathbf{E} \left[ \bar{k}_i^\beta(c_i = 0) - \bar{k}_i^\beta(c_i = 1) \right]$  as the *expected rank difference* from contributing versus free-riding.

We shall assume that all functions  $f$  are continuous in  $\beta$ , and that the following properties are the key ingredients to constitute a ‘meritocratic matching’ mechanism:

(i) *no meritocracy.* if  $\beta = 0$ , then, for any  $c$ ,  $f_\pi^0 = 1/n!$  for all  $\pi \in \Pi$ ; hence

$$\bar{k}_i^\beta = \frac{(n+1)}{2} \quad \forall i$$

(ii) *full meritocracy.* if  $\beta = 1$ , then, for any  $c$  with  $\sum_{i \in N} c_i = m$ ,  $f_\pi^1 = 0$  for

all mixed orderings  $\tilde{\pi}$ , and  $f_{\hat{\pi}}^1 = \frac{1}{m!(n-m)!}$  for all perfect orderings  $\hat{\pi}$ ; hence

$$k_i^\beta(c_i = 1) = \frac{m+1}{2} \text{ for all } i \text{ with } c_i = 1, \text{ and } k_j^\beta(c_j = 0) = \frac{n+m+1}{2} \text{ for all } j$$

with  $c_j = 0$

(iii) *imperfect meritocracy.* if  $0 < \beta < 1$ , then, for all players  $i$  and for any

$c_{-i}$ ,

$$\mathbf{E} \left[ \bar{k}_i^\beta(c_i = 0) - \bar{k}_i^\beta(c_i = 1) \right] > 0, \quad (7.1)$$

$$\partial \mathbf{E} \left[ \bar{k}_i^\beta(c_i = 0) - \bar{k}_i^\beta(c_i = 1) \right] / \partial \beta > 0. \quad (7.2)$$

### Step 3. Grouping as a function of orderings

**Groupings.** Finally, groups form based on the ranking and payoffs realize based on the contributions made in each group. Given  $\pi$ , we assume that  $m$  groups  $\{S_1, S_2, \dots, S_m\}$  of a fixed size  $s < n$  form the partition  $\rho$  of  $N$  (where  $s = n/m > 1$  for some  $s, m \in \mathbf{N}^+$ ): every group  $S_p \in \rho$  (s.t.  $p = 1, 2, \dots, m$ ) consists of all players  $i$  for whom  $k_i \in ((p-1)s + 1, ps]$ .

**Payoffs.** Given contributions  $c$  and partition  $\rho$ , the contributions in each group are multiplied by a *rate of return*  $r > 1$ . Each  $i \in N$  receives a payoff  $\phi_i(c_i | c_{-i}, \rho)$ . Let  $\phi = \{\phi_i\}_{i \in N}$  be the payoff vector. Formally, when  $i \in S$ , given the *marginal per capita rate of return*  $R := r/s$ ,  $i$  receives

$$\phi_i(c_i | c_{-i}, \rho) = \underbrace{(1 - c_i)}_{\text{remainder from budget}} + \underbrace{(R) * \sum_{j \in S} c_j}_{\text{return from the public good}}. \quad (7.3)$$

It is standard to assume that  $R \in (1/s, 1)$ , in which case contributing is socially beneficial under all mechanisms, but a strictly dominated strategy under “no

meritocracy” (details are provided in the analysis of the Nash equilibria in the next section).

## Examples

Under meritocratic matching, a player’s expected rank difference (expression (7.1)) is always positive and increasing in  $\beta$  (expression (7.2)). There are many functional assumptions that satisfy these requirements, one of which is the following:

**Meritocratic matching via logit.** Given  $\beta$  and  $c$ , let  $l_i := \frac{\beta c_i}{1-\beta}$ . Suppose ranks are assigned according to the following logit-response ordering: if any arbitrary number of  $(k-1)$  ranks from 1 to  $(k-1) < n$  have been taken by some set of players  $S \subset N$  (with  $|S| = k-1$ ), then any player’s  $i \in \{N \setminus S\}$  probability to take rank  $k$  is

$$p_i(k) = \frac{e^{l_i}}{\sum_{j \in N \setminus S} e^{l_j}}. \quad (7.4)$$

**Other interpretations.** Other interpretations of  $\beta \in [0, 1]$  are (i)  $\beta$  represents the probability to enter the group-based mechanism and  $1 - \beta$  the probability to enter the voluntary contributions mechanism, or (ii)  $(1 - \beta)/\beta$  represents normally distributed noise  $\delta^2$  added to the contribution vector  $c$  so that contributions are only imperfectly observable, after which the group-based mechanism (with  $\beta = 1$ ) is applied to  $x \sim N(c, \delta^2)$ .

### 7.3.2 Nash equilibria

From expression (7.3), the expected payoff of contributing  $c_i$  given  $c_{-i}$  for any  $i$  is

$$\underbrace{\mathbf{E}[\phi_i(c_i|c_{-i})]}_{\text{expected return from } c_i} = \underbrace{1}_{(i) \text{ budget}} - \underbrace{(1-R)*c_i}_{(ii) \text{ sure loss on own contribution}} + \underbrace{R * \mathbf{E} \left[ \sum_{j \neq i: j \in S_i^\pi} c_j | c_i \right]}_{(iii) \text{ expected return from others' contributions}}, \quad (7.5)$$

where  $S_i^\pi \in \rho$  is the subgroup into which player  $i$  is grouped. Note that term (iii), the expected return from others' contributions, is a function of one's own contribution due to meritocratic matching, which, if  $c_i = 1$ , is increasing in both  $c_{-i}$  and  $\beta$ .

First, let us consider candidates for Nash equilibria in pure strategies. Write  $1^m$  for “ $m$  players contribute, all others free-ride”, and  $1_{-i}^m$  for the same statement excluding player  $i$ . The following two conditions must hold for  $1^m$  to constitute a Nash equilibrium:

$$\mathbf{E}[\phi_i(1|1_{-i}^m)] \geq \mathbf{E}[\phi_i(0|1_{-i}^{m-1})] \quad (7.6)$$

$$\mathbf{E}[\phi_i(0|1_{-i}^m)] \geq \mathbf{E}[\phi_i(1|1_{-i}^{m+1})] \quad (7.7)$$

A special case is  $1^0$  when all players free-ride, and we shall reserve the expression  $1^m$  to refer to cases with  $m > 0$ . It is easy to verify that  $1^0$  is always a Nash equilibrium (see Appendix A, proposition 10). Gunthorsdottir et al., 2010 show that, when  $\beta = 1$ , there exists a Nash equilibrium of the form  $1^m$  with  $m > 0$  provided  $R \geq \frac{n-s+1}{ns-s^2+1} =: \underline{mpcr}$ . We shall extend this analysis to show that, given any  $R > \underline{mpcr}$ , there exists a  $\beta < 1$  such that there exists a Nash equilibrium of the form  $1^m$  with  $m > 0$  (see Appendix A, proposition 11). The minimum level of  $\beta$ , denoted by  $\underline{\beta}$ , for which such a Nash equilibrium



exists, is an implicit function that is decreasing in  $R$  provided  $R > \underline{mpcr}$ .

Second, we consider Nash equilibria in symmetric mixed strategies. Write  $1^p$  for “all players contribute with probability  $p > 0$ ”, and  $1_{-i}^p$  for the same statement excluding some player  $i$ . Again we require  $p > 0$  to distinguish from the universal free-riding state. The following condition must hold for  $1^p$  to constitute a Nash equilibrium:

$$\mathbf{E} [\phi_i(0|1_{-i}^p)] = \mathbf{E} [\phi_i(1|1_{-i}^p)]. \quad (7.8)$$

We shall prove that, for every  $\beta$ , there exists a  $R \in (\underline{mpcr}, 1)$  such that there exist two Nash equilibria of the form  $1^p$  with  $p > 0$ , one with a high  $\bar{p}$  and one with a low  $\underline{p}$  (see Appendix A, proposition 13). Write  $\overline{mpcr}$  for the necessary marginal per capita rate of return when  $\beta = 1$ . Expressed differently, given any  $R > \overline{mpcr}$ , there exists a  $\beta < 1$  such that there exist two Nash equilibria of the form  $1^p$  with  $\underline{p}, \bar{p}$  such that  $1 > \bar{p} > \underline{p} > 0$ .

It should be noted that the particular interest of this paper is the analysis of the evolutionary stability and welfare analysis of the system’s equilibria as a function of the meritocratic matching parameter  $\beta$ . We shall therefore assume that our implicit bound  $\overline{mpcr}$  is satisfied, meaning that all equilibria are at least guaranteed to exist. Thus, our work complements the analysis of Gunthorsdottir et al., 2010, where the focus of analysis is the dependence of equilibria existence on the model parameters including the rate of return for the case when  $\beta = 1$ . Note that this bound becomes generally satisfied for large  $n$  (see Appendix A, remark 8).

For the case when  $R > \overline{mpcr}$ , the following observations summarize the equilibrium analysis:

*A. Free-ride trumps contribute.*

$$\mathbf{E} [\phi_i(0|1_{-i}^m)] > \mathbf{E} [\phi_i(1|1_{-i}^m)] \text{ for } \beta < \underline{\beta} \text{ and for any } 1_m \geq 0$$

Observation A states that, when there is not enough meritocracy, then free-riding is a better reply given any set of actions by the other players.

*B1. Free-ride trumps contribute.*

$$\mathbf{E} [\phi_i(0|1_{-i}^p)] > \mathbf{E} [\phi_i(1|1_{-i}^p)] \text{ for } \beta \geq \underline{\beta} \text{ and for any } p < \underline{p} \text{ or } p > \bar{p}$$

*B2. Contribute trumps free-ride.*

$$\mathbf{E} [\phi_i(0|1_{-i}^p)] < \mathbf{E} [\phi_i(1|1_{-i}^p)] \text{ for } \beta \geq \underline{\beta} \text{ and for any } p \in (\underline{p}, \bar{p})$$

Observation B states that, when meritocracy is above a “necessary meritocracy” level ( $\beta \geq \underline{\beta}$ ), then contributing is a better reply for intermediate proportions of contributing of contributions (the range of which is given by the symmetric mixed strategy Nash equilibria probabilities), and free-riding is a better reply outside that range.

*C. Contribute-free-ride indifference.*

$$\mathbf{E} [\phi_i(0|1_{-i}^p)] = \mathbf{E} [\phi_i(1|1_{-i}^p)] \text{ for } \beta \geq \underline{\beta} \text{ and for } p = \underline{p} \text{ or } \bar{p}$$

Observation C is the condition for a symmetric mixed strategy Nash equilibria to exist. We shall refer to  $\bar{p}$  as the “near-efficient” symmetric mixed strategy Nash equilibrium, and to  $\underline{p}$  as the “less efficient” symmetric mixed strategy Nash equilibrium.

### 7.3.3 Stability

In this section, we shall analyze the stability properties of states in terms of *evolutionary stability* (Maynard Smith and Price, 1973) under replicator dynamics (Taylor and Jonker, 1978; Weibull, 1995; Helbing, 1996) and in terms of *stochastic stability* (Foster and Young, 1990) under constant error rates (Kandori, Mailath, and Rob, 1993; Young, 1993). The motivation for

this analysis is that we view  $\beta$  as a policy choice. We want to understand how the stability of different equilibria depends on the level of meritocracy in matching. The analysis of evolutionary stability will provide us with the candidates for stability, and stochastic stability with a unique prediction for every level of meritocracy.

We shall begin by defining the following dynamic game played by agents that we shall assume act myopically. A large population  $N = \{1, 2, \dots, n\}$  plays our game in continuous time. Let a state of the process be described by  $p$ , which is a proportion of players contributing, while the remaining  $(1 - p)$  free-ride. Let  $\Omega = [0, 1]$  be the state space.

### Evolutionary (bi-)stability

Suppose the two respective population proportions grow according to the following *replicator equation* (Maynard Smith and Price 1973, Taylor and Jonker 1978, Helbing 1996):

$$\partial p / \partial t = (1 - p)p (\mathbf{E} [\phi_i(1|1^p)] - \mathbf{E} [\phi_i(0|1^p)]) \quad (7.9)$$

**Evolutionarily stable states.** A state where a proportion  $\bar{p}$  of players plays  $c_i = 1$  is *evolutionarily stable* (ESS) if, for all  $p \in [0, 1]$  in some arbitrarily small  $\epsilon$ -neighbourhood around  $\bar{p}$ ,  $\partial p / \partial t > 0$  at  $p < \bar{p}$ ,  $\partial p / \partial t = 0$  at  $p = \bar{p}$ , and  $\partial p / \partial t < 0$  at  $p > \bar{p}$ .<sup>6</sup>

It turns out that the free-riding equilibrium is always evolutionarily stable. In addition, the high-efficiency symmetric mixed-strategy equilibrium, when it exists, is also evolutionarily stable.

---

<sup>6</sup>We shall speak of evolutionarily stable ‘states’ here instead of evolutionarily stable ‘strategies’ because of the asymmetry of the state.


**Lemma 5.** *Given population size  $n$ , group size  $s$  such that  $n > s > 1$  and rate of return  $r$  such that  $R \in (\overline{mPCR}, 1)$ , there exists a  $\underline{\beta} > 0$  below which the only ESS is the free-riding Nash equilibrium. When  $\beta > \underline{\beta}$ , the free-riding Nash equilibrium remains ESS, and, in addition, the population proportions given by the near-efficient symmetric mixed-strategy Nash equilibrium is also an ESS.*

*Proof.* The proof of Lemma 5 and the cut-off structure of the ESS as given by the analysis of symmetric mixed strategy Nash equilibria in Proposition 11 (see Appendix A for both) led to the summary of best replies as given by Observations A-C. Denote by  $\underline{\beta}$  the **necessary meritocracy** level in Proposition 11. Observation A implies that the the only ESS when  $\beta < \underline{\beta}$  is given by the free-riding Nash equilibrium because there is only one Nash equilibrium. Observations B1 implies that the free-riding Nash equilibrium is also ESS when  $\beta \geq \underline{\beta}$ . Observation B1, B2 and C, jointly, imply that population proportions given by the near-efficient symmetric mixed-strategy Nash equilibrium also describe an ESS since it is a local attractor.  $\square$

**Remark 6.** *As replicator dynamics increase the population size (population becomes ‘large’), the possible interval for  $R$  converges to  $(1/s, 1)$  (as is proven in Proposition 16 in Appendix A). In that sense, Lemma 5 is a general observation about the near-efficient symmetric mixed-strategy Nash equilibrium for any rate of return.*

Figure 7.1 illustrates the implied replicator phase transitions for proportions of players contributing as a function of  $\beta$  under meritocratic matching via logit (Equation 7.4) for  $s = 4$  and  $r = 1.6$  starting with  $n = 16$  (note that the phase transitions assume the long-run behavior as the population becomes large). In particular, the figure shows how, for large enough values of  $\beta$ , a relatively

Figure 7.1: Evolutionary stability of population strategies for an economy initialized with  $s = 4$ ,  $r = 1.6$  and  $n = 16$ .



*In any case when  $\beta < \underline{\beta}$ , and when  $\beta > \underline{\beta}$  then if  $p$  is either in excess of the near-efficient symmetric mixed-strategy Nash equilibrium ( $p > \bar{p}$ ) or short of the less-efficient symmetric mixed-strategy Nash equilibrium ( $p < \bar{p}$ ),  $\partial p / \partial t < 0$  (replicator tendency is down). When  $\beta > \underline{\beta}$  and  $\bar{p} > p > \underline{p}$ , then  $\partial p / \partial t > 0$  (replicator tendency is up). Depending on the location along the bifurcation, the evolutionarily stable states are therefore when either  $p = 0$  (free-riding Nash equilibrium) and when  $p$  is set according to the near-efficient symmetric mixed-strategy Nash equilibrium ( $p = \bar{p}$ ). Solid lines in the figure indicate stable equilibria, dashed lines indicate unstable equilibria.*

small ‘jump up’ is needed starting at the free-riding equilibrium to reach the basin of attraction of the high-contribution equilibrium. By contrast, for low values of  $\beta$ , a small ‘draw down’ is sufficient to fall out of the high equilibrium into the free-riding equilibrium.

## Stochastic stability

Given the possible bi-stability, which equilibrium is more stable? To answer this question, we suppose, instead of replicator dynamics, that population  $N$  remains fixed, but that individual best-reply dynamics are perturbed by individual errors. Suppose further that individuals are activated by independent Poisson clocks. The distinct times at which one agent becomes active will be called *time steps*  $t = 1, 2, \dots$ . When individual  $i$  at time  $t$  is activated (the uniqueness of only one agent's activation is guaranteed by the independence of the Poisson clocks), all agents  $j \neq i$  continue playing their previous strategy ( $c_j^t = c_j^{t-1}$ ), while  $i$  plays a best reply with probability  $1 - \epsilon$ , but takes the opposite action with probability  $\epsilon$ . When both actions are best replies,  $i$  replies by playing  $c_j^t = c_j^{t-1}$  with probability  $1 - \epsilon$  and  $c_j^t = 1 - c_j^{t-1}$  with probability  $\epsilon$ .

**State.** Let a state of the process be defined by  $p^t = \frac{1}{n} \sum_{i \in N} c_i^t$ .

Let us begin with a couple of observations. First, the perturbed process (when  $\epsilon > 0$ ) is ergodic, that is, it reaches every state from any state with positive probability in finitely many steps (at most  $n$ ). The process, therefore, has a unique stationary distribution over  $\Omega$ . Second, for any given level of  $\beta$ , the absorbing states of the unperturbed process (when  $\epsilon = 0$ ) are the various Nash equilibria in pure strategies of the game as identified in section 3.2 (and in particular the free-riding Nash equilibrium and the near-efficient pure-strategy Nash equilibrium).

**Stochastically stability.** A state  $p$  is *stochastically stable* (Foster and Young, 1990) if the stationary distribution as  $\epsilon \rightarrow 0$  places positive weight on  $p$ .

It will be useful to define the “critical mass” necessary to destabilize a given state  $p$ .

**Critical mass.** Let the *critical mass*,  $\mathcal{M}_p^\beta \in [0, n-1]$ , necessary to destabilize state  $p$  given  $\beta$  be the minimum number of players  $|S|$  needed to switch strategy simultaneously corresponding to an arbitrary set of players,  $S \subset N$ , such that as a result of their switch playing current strategy for at least one player in  $N \setminus S$  ceases to be a best reply.

**Lemma 7.** *The stochastically stable state is the near-efficient pure-strategy Nash equilibria if  $\mathcal{M}_0^\beta < \mathcal{M}_{\underline{p}}^\beta$ , the free-riding Nash equilibrium when  $\mathcal{M}_0^\beta > \mathcal{M}_{\underline{p}}^\beta$ , and both when  $\mathcal{M}_0^\beta = \mathcal{M}_{\underline{p}}^\beta$ .*

*Proof.* When pure strategy Nash equilibria exist, stochastically stable states must be pure strategy Nash equilibria of the unperturbed process. Candidates are the free-riding Nash equilibrium and the  $\binom{n}{m}$  near-efficient pure-strategy Nash equilibria.

Obviously, the critical mass for any non-equilibrium state  $p$  is  $\mathcal{M}_p^\beta = 0$  for all values of  $\beta$ . When  $\beta < \underline{\beta}$ , there exists no critical mass to destabilize the unique equilibrium which is the free-riding Nash equilibrium;  $\mathcal{M}_0^\beta = \emptyset$ . In other words, the free-riding Nash equilibrium is the only absorbing state and therefore the unique stochastically stable state. When  $\beta = \underline{\beta}$ , the near-efficient pure-strategy Nash equilibrium has a critical mass of  $\mathcal{M}_{\underline{p}}^\beta = 1$ . When  $\beta > \underline{\beta}$ , for all less-efficient  $p \geq \underline{p}$ , the critical mass is  $\mathcal{M}_{\underline{p}}^\beta = 1$  because one more contribution of some player incentivizes other non-contributors to contribute (see Observations A, B1, B2), or one contribution fewer incentivizes all to not contribute. Moreover, for  $\beta > \underline{\beta}$ ,  $\Delta\mathcal{M}_0^\beta/\Delta\beta < 0$  and  $\Delta\mathcal{M}_{\underline{p}}^\beta/\Delta\beta > 0$  provided  $\Delta\beta$  is large enough. If  $\mathcal{M}_{\underline{p}}^1 > \mathcal{M}_0^1$  at  $\beta = 1$ , then, since  $\mathcal{M}_{\underline{p}}^\beta < \mathcal{M}_0^\beta$ , it must be that there exists a  $\bar{\beta} \in (\underline{\beta}, 1)$  above which the near-efficient pure-strategy Nash equilibrium has a larger critical mass than the free-riding Nash equilibrium.

The proof of the lemma is now a direct application of Theorem 3.1 in Young,

1998, and follows from the fact that the resistances of transitions between  $p = \bar{p}$  and  $p = 0$  are given by the critical masses, thus yielding the stochastic potential for each candidate state.  $\square$

### 7.3.4 Welfare

Finally, we turn to our welfare analysis. We shall compare the efficiency and equality properties of equilibria induced by stochastically stable outcomes under varying meritocracy levels. We use this comparison to assess, given a general class of social welfare functions, which meritocracy level is welfare-optimal for a given social planner.

First, we shall introduce some notation.

**Outcome.** Let  $(\rho, \phi)$  describe an *outcome*, that is, realized groups and payoffs.

We commit to a class of social welfare functions based on Atkinson, 1970. This representation has the advantage that –just as our parameter  $\beta$  governed meritocracy– we can characterize a continuous range of social planner preferences.

**Social welfare.** Given outcome  $(\rho, \phi)$ , let  $W_e(\phi)$  be the *social welfare function* measuring its welfare given the *inequality aversion* parameter  $e \in [0, \infty)$ :

$$W_e(\phi) = \frac{1}{n(1-e)} \sum_{i \in N} \phi_i^{1-e} \quad (7.10)$$

When  $e = 1$ , it is standard that  $W_1(\phi) = \frac{1}{n} \prod_{i \in N} \phi_i$ , i.e. be the Nash product.

Expression (7.10) is a variant of the social welfare function introduced by Atkinson, 1970. It nests both the **Utilitarian** (Bentham) and **Rawlsian** so-



cial welfare functions.<sup>7</sup> When  $e = 0$ , expression (7.10) reduces to  $W_0(\phi) = \frac{1}{n} \sum_{i \in N} \phi_i$ , i.e. a Utilitarian social welfare function measuring the state's efficiency. When  $e \rightarrow \infty$ , expression (7.10) approaches  $W_\infty(\phi) = \min(\phi_i)$ , i.e. a Rawlsian social welfare function measuring the state's worst-off utility. Obviously, a Utilitarian social planner prefers the near-efficient pure-strategy Nash equilibrium to the free-riding equilibrium. The *ex-post* Rawlsian-optimal equilibrium, however, could be the free-riding Nash equilibrium with perfect equality of payoffs (equal to one for every player) if any player in the near-efficient pure strategy Nash equilibrium receives a payoff of less than one. **Harsanyi's** social welfare approach (Harsanyi, 1953), on the other hand, would always prefer the near-efficient pure-strategy Nash equilibrium if every contributor and every free-rider is in expectation (i.e. *ex ante*) better-off.<sup>8</sup>

Which equilibrium is preferable in terms of social welfare for any given social welfare function depends on the social planner's relative weights on efficiency and equality and is related to whether an *ex ante* or an *ex post* view is taken with regards to payoff dominance (Harsanyi and Selten, 1988a).<sup>9</sup> Critical for this assessment is the inequality aversion  $e$ . For the economy illustrated in Table 7.1 (with  $n = 16$ ,  $s = 4$  and  $r = 1.6$ ), suppose a social planner considers moving from  $\beta = 0$  to  $\beta = 1$ . To assess this, he makes an *ex-post*  $W_e$ -comparison. It turns out that for any  $W_e$  with  $e < 10.3$  he prefers the near-efficient pure-strategy Nash equilibrium, while for a  $W_e$  with  $e \geq 10.3$  he prefers the free-riding Nash equilibrium.<sup>10</sup>

**Welfare assessment assumption.** Suppose the social planner sets  $\beta \in [0, 1]$

so as to maximize  $\mathbf{E}[W_e(\phi)]$ , where  $\phi$  are expected to be realizations of

<sup>7</sup>See, for example, Jones-Lee and Loomes, 1995 for a discussion of this generalization.

<sup>8</sup>Harsanyi's social welfare function is  $W_H(\phi) = \frac{1}{n} \sum_{i \in N} \mathbf{E}[\phi_i]$ . See, for example, Binmore, 2005 for a discussion of the Rawlsian and Harsanyi's 'original position' approach.

<sup>9</sup> $\phi$  *payoff-dominates*  $\phi'$  if  $\phi_i \geq \phi'_i$  for all  $i$ , and there exists a  $j$  such that  $\phi_j > \phi'_j$ .

<sup>10</sup>With  $e = 10.3$ ,  $W_e$  requires efficiency gains of more than twice the amount lost by any player to compensate for the additional inequality.

stochastically stable states. Moreover, assume the social planner expects the near-efficient pure strategy Nash equilibrium (here denoted by  $\bar{p}$ ) to be played when  $\mathcal{M}_0^\beta = \mathcal{M}_{\bar{p}}^\beta$  (both are stochastically stable).

**Proposition 8.** *For any  $R > \max\{\overline{mpcr}, 1/(s-1)\}$ , there exists a population size  $n < \infty$  such that  $\mathbf{E}[W_e(\phi); \beta = \bar{\beta}] > \mathbf{E}[W_e(\phi); \beta]$  at “sufficient meritocracy” ( $\beta = \bar{\beta}$ ) for all  $\beta \neq \bar{\beta}$  given any parameter of inequality aversion  $e \in [0, \infty)$ .*

*Proof.* Suppose there exists a  $\bar{\beta} \in (\underline{\beta}, 1)$  above which the near-efficient pure-strategy Nash equilibrium is stochastically stable. Write  $q_1^n$  for the probability of having more than one free-rider in any group for a realized outcome  $(\rho, \phi)$  given  $n < \infty$ . Since the number of free-riders does not increase as  $n$  increases,  $\partial q_1^n / \partial n < 0$ . Since contributors in groups with at most one free-rider receive a payoff strictly greater than one ( $(s-1)R > 1$ ), we have  $\mathbf{E}[W_e(\phi); \beta] > (1 - q_1^n) \times W_e(\phi_i = (s-1)R \forall i)$ . Because, given any  $\beta < 1$ ,  $\partial q_1^n / \partial n < 0$ , there therefore exists  $n < \infty$  above which  $\mathbf{E}[W_e(\phi)] > W_e(\phi_i = 1 \forall i)$ .  $\square$

**Remark 9.**  $\mathbf{E}[W_e(\phi); \beta = \bar{\beta}] > \mathbf{E}[W_e(\phi); \beta]$  at “sufficient meritocracy” ( $\beta = \bar{\beta}$ ) is also the case for  $n$  smaller than implied by the proposition when (a)  $e$  is set below some bound  $e < \infty$  and/ or (b) set above some bound  $R > 1/(s-1)$ .

### 7.3.5 Summary

In our analysis, we have addressed three issues. First, we assessed the robustness of equilibrium predictions for meritocracy levels everywhere in between “no meritocracy” and “full meritocracy”. We found that the minimum meritocracy threshold (“necessary meritocracy”) that may enable equilibria with high contributions decreases with the population size, the number of groups and with the rate of return. Second, we analyzed the stability properties

Table 7.1: Stem-and-leaf plot of individual payoffs for the free-riding Nash equilibrium when  $\beta = 0$  and for the near-efficient pure-strategy Nash equilibrium when  $\beta = 1$  with  $n = 16$ ,  $s = 4$ ,  $r = 1.6$  and  $\beta = 1$ .

near-efficient pure-strategy NE when $\beta = 1$	payoff	free-riding NE when $\beta = 0$
0	0.0	0
0	0.2	0
0	0.4	0
0	0.6	0
13 14 ( $c_i = 1$ ) 2	0.8	0
0	1.0	16 ( $c_i = 0$ )
0	1.2	0
0	1.4	0
1 2 3 4 5 6 7 8 9 10 11 12 ( $c_i = 1$ ) 12	1.6	0
15 16 ( $c_i = 0$ ) 2	1.8	0
24.4	efficiency	16

*The stem of the table are payoffs. The leafs are the number of players receiving that payoff (with their contribution decision), and the individual ranks of players corresponding to payoffs in the two equilibria. At the bottom, the efficiencies of the two outcomes are calculated. Note that the near-efficient pure-strategy Nash equilibrium is more efficient, whereas the free-riding Nash equilibrium is more equitable.*

of the equilibria. It turned out that there exists a second threshold (“sufficient meritocracy”) between “necessary meritocracy” and “full meritocracy”, above (below) which the high contributions equilibria (zero contributions equilibrium) is stable. Qualitatively, the same comparative statics apply to this second threshold as with respect to the first. Third, we assessed the relative welfare properties of the candidate stable equilibria to identify, given varying degrees of inequality aversion, the uniquely welfare-maximizing regime. We found that setting meritocracy at “sufficient meritocracy” maximizes welfare for any inequality-averse social welfare functions when the population is large enough. Group size does not matter. For smaller populations, the same result holds if (a) the inequality aversion is not extreme and (b) the rate of return is high. Only for extremely inequality-averse social planners should efficiency be sacrificed and meritocracy be set to zero.

## 7.4 Theoretical predictions

There are two reasons why a social planner in our model should generally go for an intermediate level of meritocracy. First, compared with no meritocracy, levels of meritocracy above a first threshold we termed “necessary meritocracy” gain a lot of efficiency. Second, compared with even higher levels of meritocracy, marginally less meritocracy gains (a lot of) equality without losing (much) efficiency. Ideally, the social planner would therefore like to reduce meritocracy down to the level of necessary meritocracy, but obeying stability forces him to settle at one and the same “sufficient meritocracy” level in general.

Our findings seem to contradict the general social choice theory wisdom that meritocracy leads to inequality. The reason for this contradiction is that, on the one hand, we focus on situations that are strictly non-constant sum, and, on the other hand, that we do neither consider repeated game effects such as inheritance, wealth, or reputation, nor do we allow for heterogeneity in the population. The former is a crucial feature of our model and a fundamental difference in environments compared to what is usually considered. It has important implications regarding the role of meritocracy. The latter restrictions come with serious loss of generality. It is an avenue left for future research to enrich our model to allow for such features and to evaluate their welfare consequences. In a way, the purpose of this paper was to “resurrect meritocracy” in a specific interactive setting where it represents an almost unambiguously beneficial mechanism. We view this as a first step toward a much larger research agenda that aims at a more subtle assessment of meritocracy than recently voiced perceptions culminating in statements like “the meritocracy of capitalism is a big, fat lie”.<sup>11</sup>

---

<sup>11</sup>This is how The Guardian’s Heidi Moore summarizes Thomas Piketty’s bestseller book on inequality in her article “Thomas Piketty is a rock-star economist – can he re-write the

# Appendices

## Appendix A: Nash equilibria

**Proposition 10.** *For any population size  $n > s$ , group size  $s > 1$ , rate of return  $r \in (1, s)$ , and meritocratic matching factor  $\beta \in [0, 1]$ , there always exists a **free-riding Nash equilibrium** such that all players free-ride.*

The proof of Proposition 10 follows from the fact that, given any  $\beta$  and for  $c_{-i}$  such that  $\sum_{j \neq i} c_j = 0$ , we have:

$$1 = \mathbf{E}[\phi_i(0|c_{-i})] > \mathbf{E}[\phi_i(1|c_{-i})] = R. \quad (7.11)$$

Equation 7.11, in words, means that it is never a best response to be the only contributor for any level of  $\beta$ . If, for any level of  $\beta$ , given any  $c_{-i}$ ,  $\mathbf{E}[\phi_i(0|c_{-i})] > \mathbf{E}[\phi_i(1|c_{-i})]$  holds for all  $i$ , then we have a situation where free-riding is the strictly dominant strategy. In that case, for any level of meritocracy ( $\beta$ ), universal free-riding is the unique Nash equilibrium. We shall proceed to show that this is not the case if the marginal per capita rate of return ( $R$ ) and the meritocratic matching fidelity ( $\beta$ ) are high enough.

Recall that  $1^m$  stands for “ $m$  players contribute, all others free-ride”,  $1_{-i}^m$  for the same statement excluding player  $i$ , and that  $\underline{mpcr} = \frac{n-s+1}{ns-s^2+1}$ .

**Proposition 11.** *Given population size  $n > s$ , group size  $s > 1$  and rate of return  $r$  such that  $R \in (\underline{mpcr}, 1)$ , there exists a **necessary meritocracy level**,  $\underline{\beta} \in (0, 1)$ , above which there is a **pure-strategy Nash equilibrium**, where  $m > 0$  agents contribute and the remaining  $n - m$  agents free-ride.*

---

American dream?” (April 27, 2014).

*Proof.* The following two conditions must hold for Proposition 11 to be true:

$$\mathbf{E} [\phi_i(1|1_{-i}^m)] \geq \mathbf{E} [\phi_i(0|1_{-i}^{m-1})] \quad (7.12)$$

$$\mathbf{E} [\phi_i(0|1_{-i}^m)] \geq \mathbf{E} [\phi_i(1|1_{-i}^{m+1})] \quad (7.13)$$

The proof for the existence of an equilibrium in which some appropriate (positive) number of contributors  $m$  exists for the case when  $\beta = 1$  and  $R \geq \underline{mpcr}$  follows from Theorem 1 in Gunnthorsdottir et al., 2010, in which case both equations (7.12) and (7.13) are strictly satisfied.

The fixed point argument behind that result becomes clear by inspection of terms (ii) and (iii) in expression (7.5): namely, the decision to contribute rather than to free-ride is a trade-off between (ii), ‘the sure loss on own contribution’, which is zero for free-riding, versus (iii), ‘the expected return on others’ contributions’, which may be larger by contributing rather than by free-riding depending on how many others also contribute. Obviously, when  $c_{-i}$  is such that  $\sum_{j \neq i} c_j = 0$  or  $\sum_{j \neq i} c_j = (n - 1)$  (i.e. if either all others free-ride or all others contribute), it is the case that  $\phi_i(0|c_{-i}) > \phi_i(1|c_{-i})$ . Hence, in equilibrium,  $0 < m < n$ .

Now suppose  $1^m$  describes a pure-strategy Nash equilibrium for  $\beta = 1$  with  $0 < m < n$  and  $R \in (\underline{mpcr}, 1)$  in which case equations (7.12) and (7.13) are strictly satisfied. Note that  $\beta$  has a positive effect on the expected payoff of contributing and a negative effect on the expected payoff of free-riding:

$$\partial \mathbf{E} [\phi_i(1|1_{-i}^m)] / \partial \beta > 0 \quad (7.14)$$

$$\partial \mathbf{E} [\phi_i(0|1_{-i}^m)] / \partial \beta < 0 \quad (7.15)$$

When  $\beta = 0$ , we know that  $\phi_i(1|1_{-i}^m) = R < \phi_i(0|1_{-i}^m) = 1$  for any  $m$ . However,

by existence of the equilibrium with  $m > 0$  contributors when  $\beta = 1$ , provided that  $R > \underline{mpcr}$  is satisfied, there must exist some maximum value of  $\underline{\beta} \in (0, 1)$ , at which either equation (7.12) or equation (7.13) first binds due to continuity of expressions (7.14) and (7.15) in  $\beta$ . That level is the bound on  $\beta$  above which the pure-strategy Nash equilibrium with  $m > 0$  exists.  $\square$

**Remark 12.** *Note that, for a finite population of size  $n$ , a group size  $s$  larger than one implies that  $\underline{mpcr} > 1/s$  for Proposition 11 to be true, but as  $n \rightarrow \infty$ ,  $\underline{mpcr}$  converges to  $1/s$ .<sup>12</sup>*

A special case of a pure-strategy Nash equilibrium is the **near-efficient pure-strategy Nash equilibrium** (see Gunthorsdottir et al., 2010): in our set-up, the near-efficient pure-strategy Nash equilibrium generalizes to the the pure-strategy Nash equilibrium in which  $m$  is chosen to be the largest value given  $n, s, r$  for which equations (7.12) and (7.13) hold. For that  $m$  to be larger than zero  $\beta$  needs to be larger than  $\underline{\beta}$  (Proposition (11)).

Now we shall compare the asymmetric equilibria in pure strategies (in particular the near-efficient pure-strategy Nash equilibrium) with symmetric mixed-strategy Nash equilibria. For this, we define  $p_i \in [0, 1]$  as a mixed strategy with which player  $i$  plays ‘contributing’ ( $c_i = 1$ ) while playing ‘free-riding’ ( $c_i = 0$ ) with  $(1 - p_i)$ . Write  $p = \{p_i\}_{i \in N}$  for a vector of mixed strategies. Write  $1^p$  for “all players play  $p$ ”, and  $1^p_{-i}$  for the same statement excluding some player  $i$ .

**Proposition 13.** *Given population size  $n > s$  and group size  $s > 1$ , there exists a rate of return  $r$  such that  $R \in [\underline{mpcr}, 1)$  beyond which there exists a **necessary meritocracy** level,  $\underline{\beta} \in (0, 1)$ , such that there always are two mixed strategy profiles, where every agent places weight  $p > 0$  on contribut-*

---

<sup>12</sup>It is easy to check that  $\lim_{n \rightarrow \infty} \underline{mpcr} = 1/s$ .

ing and  $1 - p$  on free-riding, that constitute **symmetric mixed-strategy Nash equilibrium**. One will have a high  $\bar{p}$  (the near-efficient symmetric mixed-strategy Nash equilibrium) and one will have a low  $\underline{p}$  (the less-efficient symmetric mixed-strategy Nash equilibrium).

*Proof.* The symmetric mixed-strategy Nash equilibrium exists if there exists a  $p \in (0, 1)$  such that, for any  $i$ ,

$$\mathbf{E} [\phi_i(0|1_{-i}^p)] = \mathbf{E} [\phi_i(1|1_{-i}^p)], \quad (7.16)$$

because, in that case, player  $i$  has a best response also playing  $p_i = p$ , guaranteeing that  $1^p$  is a Nash equilibrium. Proposition 11 implies that, if  $R > \underline{mpcr}$ , equations (7.12) and (7.13) are strictly satisfied when  $\beta = 1$  for  $m$  contributors corresponding to the near-efficient pure-strategy Nash equilibrium. Indeed, expressions (7.12) and (7.13) imply lower and upper bounds (see Gunnthorsdottir et al. 2010) on the number of free-riders given by

$$l = \frac{n - nR}{1 - R + nR - r}, \quad u = 1 + \frac{n - nR}{1 - R + nR - r}. \quad (7.17)$$

*Part 1.* First, we will show, given any game with population size  $n$  and group size  $s$ , for the case when  $\beta = 1$ , that there is (i) at least one symmetric mixed-strategy Nash equilibrium when  $R \rightarrow 1$ ; (ii) possibly none when  $R = \underline{mpcr}$ ; and (iii) a continuity in  $R$  such that there is some intermediate value of  $R \in [\underline{mpcr}, 1)$  above which at least one symmetric mixed-strategy Nash equilibrium exists but not below.

(i) Because  $\partial \mathbf{E} [\phi_i(c_i|1_{-i}^p)] / \partial p > 0$  for all  $c_i$ , there exists a  $p \in (\frac{m-1}{n}, \frac{m+1}{n})$  such that expression (7.16) holds if  $R \rightarrow 1$ . This is the standard symmetric mixed-strategy Nash equilibrium, which always exists in a symmetric two-



action  $n$ -person game where the only pure-strategy equilibria are asymmetric and of the same kind as the near-efficient pure-strategy Nash equilibrium (see the proof of Theorem 1 in Cabral 1988). In this case, the presence of the free-riding Nash equilibrium makes no difference because the incentive to free-ride vanishes as  $R \rightarrow 1$ .

(ii) If  $R = \underline{mPCR}$ , one or both of the equations, (7.12) or (7.13), bind. Hence, unless expression (7.16) holds exactly at  $p = m/n$  (which is a limiting case in  $n$  that we will address in proposition 16), there may not exist any  $p$  such that expression (7.16) holds. This is because the Binomially distributed proportions of contributors implied by  $p$ , relatively speaking, place more weight on the incentive to free-ride than to contribute because universal free-riding is consistent with the free-riding Nash equilibrium while universal contributing is not a Nash equilibrium. In this case, the incentive to free-ride is too large for a symmetric mixed-strategy Nash equilibrium to exist.

(iii)  $\partial \mathbf{E} [\phi_i(c_i|1_{-i}^p)] / \partial r$  is a different linear, positive constant for both  $c_i = 0$  and  $c_i = 1$ . At and above some intermediate value of  $R$ , therefore, there exists a  $p \in (0, 1)$  such that, if played in a symmetric mixed-strategy Nash equilibrium, the incentive to free-ride is mitigated sufficiently to establish equation (7.16). We shall refer to this implicit minimum value of  $R$  by  $\overline{mPCR}$ .

Finally, for any  $p > 0$  constituting a symmetric mixed-strategy Nash equilibrium when  $\beta = 1$ ,  $\mathbf{E} [\phi_i(0|1_{-i}^p)] = \mathbf{E} [\phi_i(1|1_{-i}^p)] > 1$ . Because of this, a similar argument as in Proposition 11 applies to ensure the existence of some  $\underline{\beta} \in (0, 1)$  above which the symmetric mixed-strategy Nash equilibrium continues to exist when  $R > \overline{mPCR}$ : because, at  $\beta = 1$ , equations (7.12) and (7.13) are strictly satisfied and  $\mathbf{E} [\phi_i(0|1_{-i}^p)] = \mathbf{E} [\phi_i(1|1_{-i}^p)] > 1$ , there therefore must exist some  $\beta < 1$  and  $p' < p$  satisfying equation (7.16) while still satisfying  $\mathbf{E} [\phi_i(0|1_{-i}^{p'})] = \mathbf{E} [\phi_i(1|1_{-i}^{p'})] > 1$ . Note that this implicit bounds here may be

different from that in Proposition 11.

*Part 2.* If  $R > \overline{m\overline{p}c\overline{r}}$  and  $\beta > \underline{\beta}$ , existence of two equilibria with  $\overline{p} > \underline{p} > 0$  is shown by analysis of the comparative statics of equation (7.16).

First note that, for any  $R > \overline{m\overline{p}c\overline{r}}$  and  $\beta > \underline{\beta}$ ,  $\partial \mathbf{E} [\phi_i(0|1_{-i}^p)] / \partial \beta < 0$  while  $\partial \mathbf{E} [\phi_i(1|1_{-i}^p)] / \partial \beta > 0$ .  $p$  therefore has to take different values for equation (7.16) to hold for two different values of  $\beta$  above  $\underline{\beta}$ . Unclear is whether it has to take a higher or lower value. Note also that both  $\partial \mathbf{E} [\phi_i(0|1_{-i}^p)] / \partial p > 0$  and  $\partial \mathbf{E} [\phi_i(1|1_{-i}^p)] / \partial p > 0$  for all  $\beta \in (0, 1)$ . We can rearrange the partial derivative with respect to  $\beta$  of Expression 7.16, and obtain

$$\partial p / \partial \beta = \frac{\partial \mathbf{E} [\phi_i(1|1_{-i}^p)] / \partial \beta - \partial \mathbf{E} [\phi_i(0|1_{-i}^p)] / \partial \beta}{\partial \mathbf{E} [\phi_i(0|1_{-i}^p)] / \partial p - \partial \mathbf{E} [\phi_i(1|1_{-i}^p)] / \partial p}. \quad (7.18)$$

Expression 7.18 is negative if the denominator is negative, because the numerator is always positive.


**Claim 14.** *The denominator of Equation 7.18 is negative when  $p$  is low, and positive when  $p$  is high.*

Write  $\overline{w}_{c_i}^i$  and  $\underline{w}_{c_i}^i$  respectively for the probabilities with which agent  $i$  is matched in an above- or below-average group when playing  $c_i$  where the average is taken over contributions excluding  $i$ . Write  $\mathbf{E} [\overline{\phi}_i(c_i|1_{-i}^p)]$  and  $\mathbf{E} [\underline{\phi}_i(c_i|1_{-i}^p)]$  for the corresponding expected payoffs.

Recall that, for  $\beta > 0$  and  $1_{-i}^p \in (0, 1)$ , Expression 7.2 holds, where  $\widehat{k}$  is compatible with a perfect ordering  $\widehat{\pi}$ , and  $\widetilde{k}$  is any rank compatible with a mixed ordering  $\widetilde{\pi}$ . When  $1_{-i}^p = 0$  or  $1_{-i}^p = 1$ , the probability of agent  $i$  to take rank  $j$ ,  $f_{ij}^\beta$ , depends on his choice of  $c_i$ , but  $\overline{w}_{c_i}^i = \underline{w}_{c_i}^i = 0$  for any choice of contribution  $c_i$ .

For  $p(0, 1)$ , we shall rewrite  $\partial \mathbf{E} [\phi_i(0|1_{-i}^p)] / \partial p$  in the denominator of Equation

Figure 7.2: Expected payoffs of contributing versus free-riding if all others play  $p$  and the meritocratic matching fidelity is  $\beta$  for the economy with  $n = 16$ ,  $s = 4$ ,  $r = 1.6$ .



Expected values of  $\phi_i(0|1_{-i}^p)$  and  $\phi_i(1|1_{-i}^p)$  are plotted as functions of probability  $p$  and meritocratic matching fidelity  $\beta$  under meritocratic matching via logit (equation 7.4). The two planes intersect at the bifurcating symmetric mixed-strategy Nash equilibrium-values of  $\bar{p}$  and  $\underline{p}$  (see Proposition 13). Notice that the expected values of both actions increase linearly in  $p$  when the meritocratic matching fidelity is zero but turn increasingly S-shaped for larger values, until they intersect at  $\bar{p}$  and  $\underline{p}$ . Note that Figure 1 is a birds-eye view of this figure.

7.18 as

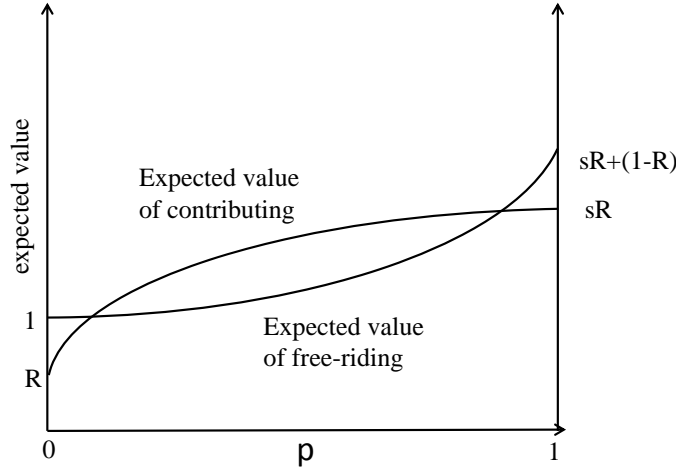
$$\frac{\partial}{\partial p} [\bar{w}_0^i * \mathbf{E} [\bar{\phi}_i(0|1_{-i}^p)]] + \frac{\partial}{\partial p} [\underline{w}_0^i * \mathbf{E} [\underline{\phi}_i(0|1_{-i}^p)]] \quad (7.19)$$

and  $\partial \mathbf{E} [\phi_i(1|1_{-i}^p)] / \partial p$  as

$$\frac{\partial}{\partial p} [\bar{w}_1^i * \mathbf{E} [\bar{\phi}_i(1|1_{-i}^p)]] + \frac{\partial}{\partial p} [\underline{w}_1^i * \mathbf{E} [\underline{\phi}_i(1|1_{-i}^p)]] . \quad (7.20)$$

Notice that, for large  $\beta$ ,  $\underline{w}_0^i \gg \bar{w}_0^i$  when  $p$  is close to zero, and  $\underline{w}_1^i \ll \bar{w}_1^i$  when  $p$  is close to one. Moreover, notice that the existence of the pure-strategy Nash equilibrium with high contribution for high levels of  $\beta$  ensures that  $\mathbf{E} [\phi_i(0|1_{-i}^p)]$  is not always larger than  $\mathbf{E} [\phi_i(1|1_{-i}^p)]$ . It therefore follows from continuity in  $\beta$  that Expression 7.20 exceeds Expression 7.19 when  $p$  is low,

Figure 7.3: Expected payoffs of contributing versus free-riding if all others play  $p$  and the meritocratic matching fidelity is  $\beta > \underline{\beta}$ .



Expected values of  $\phi_i(0|1_{-i}^p)$  and  $\phi_i(1|1_{-i}^p)$  are plotted as functions of probability  $p$  for some fixed  $\beta > \underline{\beta}$ . The two planes intersect at the bifurcating symmetric mixed-strategy Nash equilibrium-values of  $\bar{p}$  and  $\underline{p}$  (see Proposition 13). The relative slopes of the two curves illustrate the proposition. Note that this figure is a slice through Figure 2 along a value of  $\beta > \underline{\beta}$ .

and that Expression 7.19 exceeds Expression 7.20 when  $p$  is high, hence the denominator of Equation 7.18 is negative when  $p$  is low, and positive when  $p$  is high. Figure 3 illustrates.  $\square$

**Remark 15.** Note that the *necessary meritocracy* level  $\underline{\beta}$  in Propositions 11 and 13 need not be the same. We shall write  $\underline{\beta}$  for whichever level is larger.

**Proposition 16.** Given group size  $s > 1$ , then, if  $\beta = 1$ , as  $n \rightarrow \infty$  (i)  $1^m/n$  of the near-efficient pure-strategy Nash equilibrium and  $\bar{p}$  of the near-efficient symmetric mixed-strategy Nash equilibrium converge, and (ii) the range of  $R$  for which these equilibria exist converges to  $(1/s, 1)$ .

*Proof.* Suppose  $R > \overline{mPCR}$ , i.e. that both symmetric mixed-strategy Nash equilibrium and near-efficient pure-strategy Nash equilibrium exist. Let  $1^m$

describe the near-efficient pure-strategy Nash equilibrium and  $1^p$  describe the near-efficient symmetric mixed-strategy Nash equilibrium. Recall that expressions under (7.17) summarize the lower and upper bound on the number of free-riders,  $(n - m)$  in the near-efficient pure-strategy Nash equilibrium. Taking  $\lim_{n \rightarrow \infty}$  for those bounds implies a limit lower bound of  $\frac{1}{1+n \frac{R-r/n}{1-R}}$ , and a limit upper bound of the expected proportion of free-riders of  $\frac{1}{n} + \frac{1}{1+n \frac{R-r/n}{1-R}}$ , and thus bounds on the number of free-riders that contain at most two integers and at least one free-rider. (Notice that the limits imply that exactly one person free-rides as  $R \rightarrow 1$ .) We know that, if there is one more free-rider than given by the upper bound, then equation (7.13) is violated. Similarly, if there is one fewer free-rider than given by the lower bound, then equation (7.12) is violated.

With respect to the near-efficient symmetric mixed-strategy Nash equilibrium, recall that Expression 7.16 must hold; i.e.  $\mathbf{E}[\phi_i(0|1^p_{-i})] = \mathbf{E}[\phi_i(1|1^p_{-i})]$ . We can rewrite  $\mathbf{E}[\phi_i(c_i|1^p_{-i})]$  as  $\mathbf{E}[\phi_i(c_i|B)]$ , where  $B$  is the number of other players actually contributing (playing  $c_i = 1$ ), which is distributed according to a Binomial distribution  $Bin(p, n)$  with mean  $\mathbf{E}[B] = np$  and variance  $\mathbf{V}[B] = np(1-p)$ . As  $n \rightarrow \infty$ , by the law of large numbers, we can use the same bounds obtained for the near-efficient pure-strategy Nash equilibrium to bound  $(B/n) \in [(n-u)/n, (n-l)/n]$ , which converges to the unique  $p$  at which expression (7.16) actually holds.<sup>13</sup>

Suppose all players contribute with probability  $p$  corresponding to the near-efficient symmetric mixed-strategy Nash equilibrium limit value. Then,  $\lim_{n \rightarrow \infty} \mathbf{V}[(B/n)] = \lim_{n \rightarrow \infty} \frac{p(1-p)}{n} = 0$  for the actual proportion of contributors. Hence, the limit for the range over  $R$  necessary to ensure existence converges to that of the near-efficient pure-strategy Nash equilibrium, which by Remark 12 is  $(1/s, 1)$ .  $\square$

---

<sup>13</sup>Details concerning the use of the law of large numbers can be followed based on the proof in Cabral, 1988.

**Remark 17.** *In light of the limit behavior, it is easy to verify, ceteris paribus, that the value of the marginal per capita rate of return necessary to ensure existence of the symmetric mixed-strategy Nash equilibrium is decreasing in population size  $n$ , but increasing in group size  $s$ ; i.e. decreasing in relative group size  $s/n$ .*

## Part 2: Experiments

### Abstract

One of the fundamental tradeoffs underlying society is that between efficiency and equality. The challenge for institutional design is to strike the right balance between these two goals. Game-theoretic models of public-goods provision under ‘meritocratic matching’ succinctly capture this tradeoff: under zero meritocracy (society is randomly formed), theory predicts maximal inefficiency but perfect equality; higher levels of meritocracy (society matches contributors with contributors) are predicted to improve efficiency but come at the cost of growing inequality. We conduct an experiment to test this tradeoff behaviorally and make the astonishing finding that, notwithstanding theoretical predictions, higher levels of meritocracy increase both efficiency and equality, that is, meritocratic matching dissolves the tradeoff. Fairness considerations can explain the departures from theoretical predictions including the behavioral phenomena that lead to dissolution of the efficiency-equality tradeoff.

**Acknowledgements.** The authors acknowledge support by the European Commission through the ERC Advanced Investigator Grant ‘Momentum’ (Grant No. 324247). Further, we thank Bary Pradelski, Anna Gunnthorsdottir, Michael Mäs, Stefan Seifert, Jiabin Wu, Yoshi Saijo, Yuji Aruka, and Guillaume Hollar for helpful discussion and comments on earlier drafts, and finally members of GESS at ETH Zurich as well as seminar participants at the *Behavioral Studies Colloquium* at ETH Zurich, at the 25<sup>th</sup> *International Conference on Game Theory 2014* at Stony Brook, at the *Choice Group* at LSE and at the Kochi University of Technology for helpful feedback.

## 7.5 The efficiency-equality tradeoff

Making policy decisions often requires tradeoffs between different goals. One of the most fundamental tradeoffs is that between efficiency and equality. The basic idea of institutional *meritocracy* (Young, [1958b](#)) is to devise a system of rewards that “is intended to encourage effort and channel it into socially productive activity. To the extent that it succeeds, it generates efficient economy. But that pursuit of efficiency necessarily creates inequalities. And hence society faces a tradeoff between equality and efficiency.” (Arthur M. Okun, *Equality and efficiency, the big tradeoff*, The Brookings Institution, 1975, p. 1.)

One could argue that inherent to this statement is the view that societal activity can be modeled in the language of game theory as a public-goods provision/voluntary contributions game (Isaac, McCue, and Plott, [1985b](#); Ledyard, [1997](#); Chaudhuri, [2011b](#)). In the baseline model, voluntary contributions games create no incentives for contributors and universal free-riding is the only stable equilibrium (Nash, [1950](#)). In such a setting, the “tragedy of the commons” cannot be circumvented (Hardin, [1968](#)). However, even if this outcome is maximally inefficient, one positive thing about it is that it comes with a very high degree of equality (at the cost of low average payoffs). For this reason, the outcome of universal free-riding has been controversially associated with extreme forms of socialism (Mises, [1922](#); Hayek, [1935](#)). Fortunately, an array of mechanisms exists with the potential to foster contributions to public goods. One such mechanism that has been extensively studied in the literature is punishment (Fehr and Gächter, [2000](#); Ledyard, [1997](#); Chaudhuri, [2011b](#)). However,



mechanisms such as punishment tend to be “leaky buckets” (Okun, 1975), in the sense that some of the efficiency gains generated by the increase in contributions are spent in order to uphold them (e.g. on punishment costs).

An alternative mechanism, discussed here, is ‘meritocratic matching’ (Nax, Murphy, and Helbing, 2014) which is inspired by a recent, seminal paper introducing the “group-based mechanism” (Gunnthorsdottir et al., 2010). Meritocratic matching generalizes the group-based mechanism by introduction of an additional parameter that measures the degree of imprecision inherent to the mechanism and thus bridges the no-mechanism and group-based mechanism continuously. Matching is said to be “meritocratic” because cooperators are matched with cooperators, and defectors are matched with defectors (Gunnthorsdottir et al., 2010; Nax, Murphy, and Helbing, 2014), hence “merit” is associated with contribution decisions. Under meritocratic matching, near-efficient outcomes are supported by payoff-dominant equilibria (Nash, 1950; Harsanyi and Selten, 1988b) provided the rate of return (Gunnthorsdottir et al., 2010) and the level of meritocracy exceed certain thresholds (Gunnthorsdottir et al., 2010). The reason for this is that agents have incentives to contribute more in order to be grouped with other high-contributors. As a result, only a small fraction of free-riders continues to exist in these equilibria. Such equilibria are excellent predictors of the population’s distribution of play under ‘full meritocracy’ (Gunnthorsdottir et al., 2010; Gunnthorsdottir and Thorsteinsson, 2010; Gunnthorsdottir, Vragov, and Shen, 2010; Rabanal and Rabanal, 2010). Unfortunately, the new equilibria, however desirable in terms of efficiency vis-à-vis tragedy of the commons, typically feature a higher degree

of inequality.<sup>14</sup> The contrast between these two outcomes is well illustrated by the tensions that would exist between an ideal Bentham (utility-maximizing) social planner, on the one hand, and an ideal Rawlsian (inequality-minimizing) social planner on the other: in many games, the Bentham (Bentham, 1907) would strictly favor perfect action-assortativity, while the Rawlsian (Rawls, 1971) would rather prefer complete non-assortativity. In comparison, a real-world social planner typically exercises a certain degree of ‘inequality aversion’, aiming for an outcome between these two extremes (Atkinson, 1970).

Meritocratic matching differs from what is commonly associated with meritocracy in terms of reward/punishment mechanisms. Nevertheless, in the real world, many mechanisms and institutions exist that are based on the logic of meritocratic matching. Admissions to schools or types of education, for example, are often based on rewards of past school or exam performances which are a function of the work/effort applicants had invested. An important determinant of what makes places that are more competitive to enter ‘better’ is the promise of being matched with others who also performed well in the best. Similarly, in professional team sports, clubs aim to hire athletes with good track records, and athletes join teams in order to be matched with others. Basically, meritocratic matching mirrors the key features of any system that features team-based pay such as on trading desks.

Essentially, the efficiency-equality tradeoff in designing a meritocratic matching regime boils down to the choice of a systemic degree of assortativity, i.e.

---

<sup>14</sup>The new equilibria always have positive variance, while the free-riding equilibrium has variance zero. In what cases this translates into more inequality depends both on the particular structure of the equilibrium given a game and on the measure of inequality that is applied.

the selection of a certain degree of meritocracy. This tradeoff is at the heart of social choice theory (see e.g. (Arrow, 1951; Sen, 1970; Arrow, Bowles, and Durlauf, 2000b)) and welfare economics (see e.g. (Samuelson, 1980; Feldman, 1980; Atkinson, 2012)). Zero meritocracy represents maximal equality, but also minimal efficiency; full meritocracy represents the opposite. For any degree of inequality aversion away from the two extremes (given by (Bentham, 1907) and (Rawls, 1971)), there exist, at least in theory, an intermediate degree of meritocracy that maximizes social welfare (Nax, Murphy, and Helbing, 2014). Unfortunately, this is a difficult tradeoff as the buckets are leaky in both directions: reducing meritocracy increases equality at the expense of efficiency, and increasing meritocracy increases efficiency at the expense of equality.

In this paper we set out to test this tradeoff experimentally by analysis of intermediate regimes of meritocracy. We are thus the first to bridge the rich experimental literature on public-goods games under random interactions (zero meritocracy) (Ledyard, 1997; Chaudhuri, 2011b) with the more recent literature on full meritocracy (group-based mechanisms) (Gunnthorsdottir et al., 2010; Rabanal and Rabanal, 2010). The experiments reveal that the strict tradeoff implied by theory is dissolved in practice. Higher degrees of meritocracy turn out to increase welfare for any symmetric and additive objective function (Atkinson, 1970), including Bentham utility-maximization (Bentham, 1907) and Rawlsian inequality minimization (Rawls, 1971). In other words, meritocracy increases both efficiency *and* equality, leading to unambiguous welfare improvements as we illustrate for a variety of measures. We argue that the dissolution of the tradeoff is driven by the agents' distastes of 'meritocratic' unfairness, and by the corrections to their actions that these considerations im-

ply. The view of fairness that we adopt and test here generalizes the concept of distributive fairness/ inequity aversion (Fehr and Schmidt, 1999; Ockenfels and Bolton, 2000) to settings with positive levels of meritocracy. This fairness definition is a game-theoretic application of a notion related to systemic fairness (Adams, 1965; Greenberg, 1987), which has been long recognized in organizational theory, but not previously applied to game theory (and the problem of public-goods provision in particular). The patterns associated with reactions to between-group comparisons, however, have been noted as robust phenomena without being interpreted as driven by norms of fairness (Bohm and Rokenbach, 2013).

Among our results are the following key findings:

1. *Efficiency increases with meritocracy.* Perfect meritocracy is near-efficient and coincides with the theoretically predicted levels. The zero meritocracy regime lies above the efficiency levels implied by the theoretical equilibrium assuming self-regarding rational choice. For intermediate meritocracy levels, efficiency is above that of zero meritocracy, but below the theoretically expected equilibrium values.
2. *Equality, in contrast to theoretical predictions, also increases with meritocracy.* This finding is robust with regard to several inequality measures, including the payoff of the worst-off subject. In our settings, the often-cited tradeoff between equality and efficiency turns out to be a theoretical construct, rather than a behavioral regularity.
3. *Fairness considerations can explain the dissolution of the tradeoff between efficiency and equality.* According to our definition, agent  $A$  considers the

outcome of the game “unfair” if another agent  $B$  contributed less than  $A$ , but  $B$  was placed in a better group. As a consequence, agent  $A$  is assumed to respond by decreasing his/her contribution.

4. *Higher meritocracy levels increase agents’ sensitivity to unfair group matching in lower meritocracy levels.* Our experimental setup expose each participant to two distinct levels of meritocracy. When the second part of the experiment is restarted at a lower meritocratic regime, it turns out that agents’ distaste for unfair group matching is magnified.

## 7.6 The experiment

### 7.6.1 The underlying meritocracy game

A fixed population of  $n$  agents plays the following public-goods game repeatedly through periods  $T = \{1, 2, \dots, t\}$ . First, each agent  $i$  simultaneously decides to contribute any number of coins  $c_i$  between zero and his full budget  $B > 0$ . The amount not contributed goes straight to his/her private account. The ensemble of players’ decisions yields the contribution vector  $c$ . Second, Gaussian noise with mean zero and variance  $\sigma^2 \geq 0$ . Third,  $k$  groups of a fixed size  $s < n$  (such that  $s * k = n$ ) are formed according to the ranking of the values  $c'$  (with random tie-breaking). That is, the highest  $s$  contributors form group  $G_1$ , the next highest  $s$  contributors form  $G_2$ , etc. The resulting group partition is  $\rho = \{G_1, G_2, \dots, G_k\}$ . Finally, based on the grouping and the initial contributions vector  $c$ , payoffs  $\phi$  are computed. Each player  $i$  in a group  $G_i$

with other players  $j \neq i$  receives:

$$\underbrace{\phi_i(c)}_{\text{payoff}} = \underbrace{(B - (1 - m) * c_i)}_{\text{return from private account}} + \underbrace{\sum_{j \in G_{-i}} m * c_j}_{\text{return from group account}}, \quad (7.21)$$

where  $m$  represents the marginal per capita rate of return, and  $G_{-i}$  indicates the members of group  $G_i$  excluding  $i$ .

**NOTE** that the game is equivalent to play under the group-based mechanism (here, ‘perfect meritocracy’) (Gunnthorsdottir et al., 2010) if  $\sigma^2 = 0$ , and that the case of  $\sigma^2 \rightarrow \infty$  corresponds to random re-matching (here, ‘zero meritocracy’) (Andreoni, 1988).

### Equilibrium play

To highlight the structure of the Nash equilibria (Nash, 1950) for this class of games, it is useful to evaluate the value of the expected payoff  $\mathbf{E}[\phi_i(c)]$  during the decision stage, i.e. before groups are formed. In Eq. (7.21), the first term, i.e. the private-account return, is completely determined by the agent’s contribution choice. The second term, i.e. the group-account return, however, depends on the players’ contributions in a probabilistic way. In the case of zero meritocracy (i.e. random re-matching) ( $\sigma^2 = \infty$ ),  $\mathbf{E}[\phi_i(c)]$  is strictly decreasing in the player’s own contribution because the marginal per capita rate of return is less than one. Under zero meritocracy, the player’s own contribution has no effect on group matching, and, therefore, the only equilibrium is universal free-riding. Conversely, for positive levels of meritocracy, the player’s contribu-

tion choice influences the probability of being ranked in a high group. Hence, making a positive contribution is a tradeoff between the sure loss on the own contribution and the promise of a higher return from the group-account. However, the chances of being ranked in a better group are decreasing with growing variance. As a result, new Nash equilibria with positive contribution levels may emerge: indeed, Nax et al., 2013 generalizes the results by Gunnthorsdottir et al., 2010 showing that, if the level of meritocracy stays sufficiently large in addition to some bound on  $r$ , there exist a near-efficient pure-strategy Nash equilibria in which a large majority of players contributes the full budget  $B$  and a small minority of players contributes nothing.<sup>15</sup>

## 7.6.2 Choice of experimental parameters

In order to ensure comparability with the literature on voluntary contributions games under random re-matching (Andreoni, 1988) (as reviewed by Ledyard 1997; Chaudhuri 2011b) and particularly under the group-based mechanisms (Gunnthorsdottir et al., 2010), we set the group size  $s = 4$  and the marginal per capita rate of return  $m = 0.5$  (as in Gunnthorsdottir et al. 2010). Due to laboratory capacity restrictions we set  $n = 16$ . Finally, we need to set different meritocracy levels as represented by variance  $\sigma^2$  other than  $\sigma^2 = 0$  and  $\sigma^2 = \infty$ .

In order to determine the right and meaningful levels of variance levels, we conducted a series of 16 experimental sessions on Amazon’s Mechanical Turk

---

<sup>15</sup>Universal free-riding continues to be an equilibrium too. See Theorem 1 in Ref. (Gunnthorsdottir et al., 2010) and Propositions 6 and 7 in Ref. (Nax, Murphy, and Helbing, 2014) for detailed proof and game-theoretic characterization of these equilibria.

(MTurk) with a total of 256 participants using our new NodeGame software. Details about the experiment can be found in Appendix B. In each session, all participants played a game with different variance levels which were  $\sigma^2 = \{0, 2, 4, 5, 10, 20, 50, 100, 1000, \infty\}$ . For all variance levels below  $\sigma^2 = 100$ , the near-efficient Nash equilibria exist in the stage game. For higher variance levels, the free-riding Nash equilibrium is the unique Nash equilibrium of the stage game.

Each game was repeated for 25 (or 20) successive rounds. We evaluated the level of variance starting at which the mechanism started (i) to differ from the levels implied by the near-efficient Nash equilibria under  $\sigma^2 = 0$  and (ii) not to stabilize, and we found these variance levels to be (i)  $\sigma^2 = 3$  and (ii)  $\sigma^2 = 3$ . Appendix C contains details. Hence, we settled for the following four variances for our laboratory experiment:  $\sigma^2 = \{0, 3, 20, \infty\}$ . We use the following terminology. For  $\sigma^2 = 0$  we use PERFECT-MERIT, and for  $\sigma^2 = \infty$  we use NO-MERIT. For the intermediate values we use HIGH-MERIT ( $\sigma^2 = 3$ ) and LOW-MERIT ( $\sigma^2 = 20$ ).

**NOTE** that in the case of these four levels of variance tested in this study, the predicted stage-game Nash equilibria are as follows. For  $\sigma^2 = \infty$  (NO-MERIT), the unique stage-game Nash equilibrium is universal free-riding, which is also a Nash equilibrium for all the other variance levels. For  $\sigma^2 = \{0, 3, 20\}$ , moreover, there exist  $\binom{n}{2}$  alternative pure-strategy equilibria where exactly two players free-ride while all others contribute fully. Details on equilibria can be found in Appendix A.



### 7.6.3 The laboratory experiment

We ran 12 experimental sessions with a total of 192 participants at the De-Scil Laboratory (ETH Zurich) using the same NodeGame software as in the pre-tests. Details about the experiment can be found in Appendix B. In each session, all participants played two repeated games, one after the other, each one with one of the different variance level  $\sigma^2 = \{0, 3, 20, \infty\}$ . Each session, therefore, represented a unique order of two of the four possible variance levels (leading to 12 sessions to account for every possible ordered pair). Each repeated game was played for 40 successive rounds ( $T = \{1, 2, \dots, 40\}$ ), with population size  $n = 16$ , group size  $s = 4$ , and marginal per capita rate of return  $m = 0.5$ .<sup>16</sup>

Fig. 7.4 illustrates how the laboratory results fit in with the MTurk pre-tests.

## 7.7 Experimental results

Overall, we found a significant difference in the mean level of contributions among the four treatments (linear mixed model LMM:  $F_{3,8} = 36.8, P < 0.0001$ ), as Fig. 7.5 illustrates.

In the following, we first study efficiency, inequality and fairness, focusing on the first part of the experiment. Then, we use the second part of the experiment

---

<sup>16</sup>We would have liked to reproduce the 80 rounds of play by Gunthorsdottir et al., 2010, but due to time restrictions as in how long we could keep subjects in the laboratory, we decided to halve this amount in order to be able to run two variance levels per person. Each session lasted roughly one hour.




Figure 7.4: **Average contribution levels different variance levels.** Contribution levels increase as meritocracy increases. In perfect meritocracy, contribution levels are near efficient and approximately coincide with theoretical predictions. As noise is increased, that is, as meritocracy decreases, the contributions start dropping but not as starkly as predicted by theory.




Figure 7.5: **Average contribution levels for perfect-, high-, low-, and no-meritocracy, respectively associated with the values of  $\sigma^2 = \{0, 3, 20, \infty\}$ .** Contribution levels increase as meritocracy increases. In perfect meritocracy, contribution levels are near efficient and approximately coincide with theoretical predictions. Meritocratic treatments are mostly stable over the forty rounds of the game, and do not follow the contribution decay of the random treatment. Error bars represent the 95%-confidence intervals.

to determine the agents' sensitivity to changes in meritocracy levels.

### 7.7.1 Efficiency

In this section, we evaluate the effect of meritocracy on total payoffs generated, i.e. on efficiency. Theory predicts Gunnthorsdottir et al., 2010; Nax, Murphy, and Helbing, 2014 that equilibria supported by higher meritocracy levels are more efficient, and we shall show that this predictions holds true in the lab, confirming previous experimental results (Gunnthorsdottir et al., 2010; Gunnthorsdottir and Thorsteinsson, 2010; Gunnthorsdottir, Vragov, and Shen, 2010). Indeed, the levels of efficiency supported by the payoff-dominant equilibria under meritocracy regimes LOW-MERIT, HIGH-MERIT and PERFECT-MERIT represent relatively accurate predictions, while the complete inefficiency prediction of the unique, zero-contribution Nash equilibrium under no-meritocracy (NO-MERIT) understates the achieved efficiency levels in the order of standard magnitudes (Ledyard, 1997; Chaudhuri, 2011b).

We measure efficiency as the average payoff over players,  $\bar{\phi} = \frac{\sum_{i \in N} \phi_i}{n}$ , over the forty rounds. As shown in Fig. 7.6, when climbing up the meritocracy ladder we find increases in efficiency from  $\sigma^2 = \infty$  (NO-MERIT) through  $\sigma^2 = \{20, 3\}$  to  $\sigma^2 = 0$  (PERFECT-MERIT).

Overall, we observe significant differences in the mean of realized payoffs among the four treatments (linear mixed model LMM:  $F_{3,8} = 36.95$ ,  $P < 0.0001$ ). Taking NO-MERIT as a baseline, LOW-MERIT led to an increase in the average realized payoff of 7.1611 (Likelihood Ratio Test LRT:  $\chi_{(1)} = 12.7$ ,  $P = 0.0004$ ),

HIGH-MERIT to an increase of 8.1964 (LRT:  $\chi_{(1)} = 17.48, P < 0.0001$ ), and PERFECT-MERIT to an increase of 8.8287 (LRT:  $\chi_{(1)} = 16.22, P < 0.0001$ ). These levels correspond to roughly double those of NO-MERIT. Computing the most conservative (Bonferroni) adjusted  $p$ -values on all pair-wise differences reveals that the treatment with variance  $\infty$  is significantly different ( $P < 0.0001$ ) from the other three variance levels  $\sigma^2 = \{0, 3, 20\}$ , which are themselves not significantly different from each other.

For intermediate meritocracy regimes  $\sigma^2 = \{20, 3\}$ , efficiency is significantly below the level implied by the respective payoff-dominant equilibria (Harsanyi and Selten, 1988b), but only by less than five percent. Conversely, under full meritocracy  $\sigma^2 = 0$ , efficiency is above and within five percent of equilibrium. Note that contribution levels resemble the levels implied by the symmetric mixed-strategy Nash equilibrium identified in Ref. (Nax, Murphy, and Helbing, 2014), but do not perfectly coincide with them, as intermediate contribution levels continue to be selected under  $\sigma^2 = \{20, 3\}$ , which are dominated even in the mixed equilibrium.

The contribution patterns under  $\sigma^2 = 0$  confirm the qualitative patterns of contributions found in Gunthorsdottir et al., 2010, instead now we have  $n = 16$ . For  $\sigma^2 = \infty$ , we have the same pattern of contributions that, on average, roughly halve every 10-20 rounds as found in many related studies (Ledyard, 1997; Chaudhuri, 2011b).




Figure 7.6: Analysis of efficiency based on smoothed distributions of average payoffs over 40 rounds for perfect-, high-, low-, and no-meritocracy, respectively associated with the values of  $\sigma^2 = \{0, 3, 20, \infty\}$ . Efficiency, measured as average payoff, increases as meritocracy increases. Black solid lines indicate the mean payoff as implied by the respective payoff-dominant Nash equilibria, red solid lines indicate the mean payoff observed in the experiment, red-shaded areas indicate the 95%-confidence intervals of the mean. Blue dots indicate the payoff of the worst-off player (note that the worst-off player in every equilibrium receives twenty ‘coins’).

### 7.7.2 Equality

Recall the theory prediction from Nax, Murphy, and Helbing, 2014 that equilibria supported by higher meritocracy levels feature more inequality in the distribution of payoffs. In this section, we shall show that laboratory evidence yields diametrically opposite results; namely, higher meritocracy levels lead to outcomes that are more equal in terms of payoff distributions.

One can identify two measures of payoff inequality directly from the moments of the payoff distribution: (i) the payoff of the worst-off (Rawls, 1971),  $\underline{\phi} = \min\{\phi_i\}$ , and (ii) the variance of payoffs,  $\sigma^2 = \frac{\sum_{i \in N} (\phi_i - \bar{\phi})^2}{n}$ . A more sophisticated third alternative is (iii) the Gini coefficient. In terms of all measures, our analysis shows that equality increases with meritocracy. Note that the following results are also robust to other measures of inequality (Cowell, 2011) (see *appendix*).

Fig. 7.7 shows that, like efficiency, equality also increases from  $\sigma^2 = \infty$  (NO-MERIT) through  $\sigma^2 = \{20, 3\}$  to  $\sigma^2 = 0$  (PERFECT-MERIT). These increases are reflected by differences in the Gini coefficient, and by the order of the payoff of the worst-off – Rawlsian inequality. Under NO-MERIT, equality is significantly below the level implied by equilibrium. For all three positive levels of meritocracy, equality is above that achieved by NO-MERIT and above the theoretically implied levels. Details about the statistical tests can be found in the Statistical Analysis section of *Materials and Methods* section.

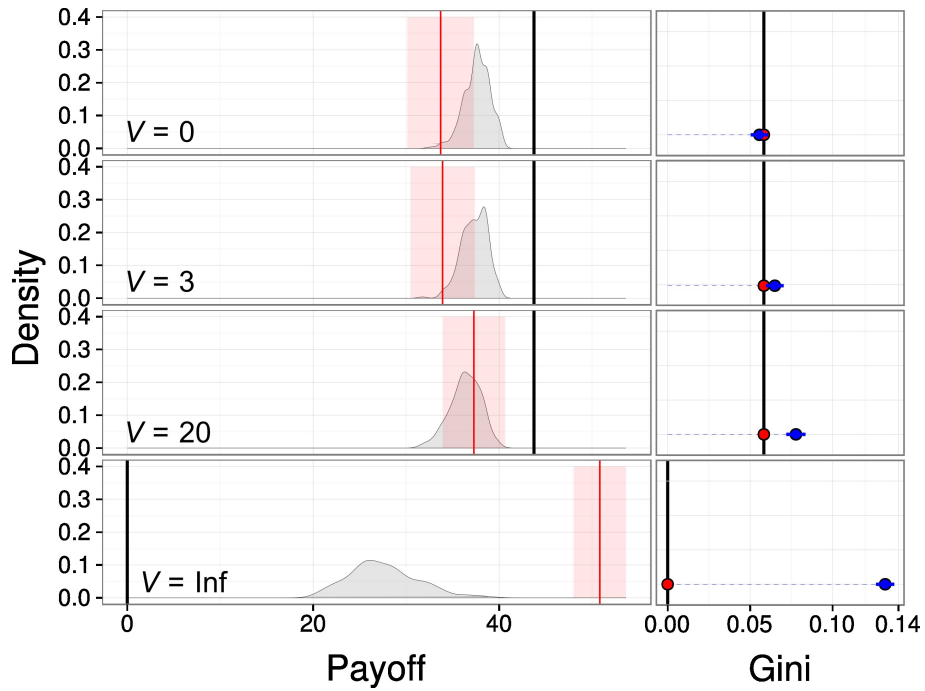


Figure 7.7: Level of payoff equality for perfect-, high-, low- and no-meritocracy, respectively associated with the values of  $\sigma^2 = \{0, 3, 20, \infty\}$ . Inequality, measured by the variance of payoff and by the Gini coefficient, decreases, as meritocracy increases. Left panel: Smoothed distributions of average payoffs over 40 rounds. Black solid lines indicate the variance of the payoffs as given by the respective payoff-dominant Nash equilibria, red solid lines indicate the mean variance observed in the experiment, red-shaded areas indicate the 95%-confidence intervals of the mean variance. Right panel: Average Gini coefficient of the distribution of payoffs with 95%-confidence intervals. Black solid lines and red dots indicate the Gini coefficient implied by the equilibrium (without fairness considerations).



### 7.7.3 Fairness

Theory predicted that, as meritocracy is increased, that play features higher efficiency at the cost of growing inequality. Results confirm the theory predictions in terms of efficiency at least qualitatively, but diametrically contradict those regarding equality. We found that Nash predictions work well on aggregate in the meritocracy regimes LOW-MERIT, HIGH-MERIT and PERFECT-MERIT, but not for NO-MERIT. In this section, we explore the role of individuals' fairness considerations in explaining these deviations from standard Nash predictions. We shall find evidence for meritocratic fairness concerns that could explain these phenomena and that generalize well-known fairness considerations (Fehr and Schmidt, 1999; Ockenfels and Bolton, 2000) in the meritocracy context, allowing for a systemic understanding of the payoff structure.

**Meritocratic fairness = distributional fairness + strategic concerns**

In public-goods games with completely random interactions, i.e. in environments with zero meritocracy, a payoff allocation is considered unfair if a player contributed more than the average of the other group members. The larger the variance in payoffs is, the larger the degree of unfairness. From the perspective of an individual player, unfairness can be *advantageous*, if he/she contributed less than the average, or *disadvantageous* in the opposite situation (Fehr and Schmidt, 1999). It has been shown that unfair allocations influence players' utilities negatively and that agents respond to unfairness by adjusting their contributions (Fehr and Schmidt, 1999; Ockenfels and Bolton, 2000). Disad-

vantageous unfairness has an accentuated negative effect on a player's utility, while advantageous unfairness has a negative but weaker effect. This gain-loss asymmetry is of course related to some of the most robust findings in experimental economics (Kahneman and Tversky, 1979; Tversky and Kahneman, 1991; Erev, Ert, and Yechiam, 2008). The consequences of the distaste for unfairness are such that, on average, a player responds by decreasing (increasing) his/her contribution after experiencing disadvantageous (advantageous) unfairness (Fehr and Schmidt, 1999; Ockenfels and Bolton, 2000). Importantly, the tendency to decrease is stronger than the tendency to increase due to the asymmetry in distastes. The typical contribution pattern found in repeated public goods experiments (intermediate contribution levels at the beginning, followed by a decay over time) can therefore be explained by heterogeneity in social preferences, and reactions to (un)fairness and reciprocity (Ledyard, 1997; Chaudhuri, 2011b).

It is reasonable to conjecture that fairness considerations continue to matter in the presence of meritocracy. It is not clear, however, exactly how they are likely to matter. The additional subtlety comes from the fact that contributions now play a double role. On the one hand, they determine a player's payoff within a given group. On the other hand, they also determine the group into which the player is matched. Expressed differently, a *(meritocratic fairness concern) = (regular fairness concern) + (strategic concern)*. Consequently, a player cannot be expected to respond to unfairness in the same way as outlined above for the case of zero meritocracy. For example, a player matched into a low-contribution group where he/she currently is the highest contributor may not respond by decreasing his/her contribution, but rather by strategi-

cally increasing it in order to enter a better group in the next round. In order to account for this more complex reasoning, we generalize the concept of distributional fairness of Refs. (Fehr and Schmidt, 1999; Ockenfels and Bolton, 2000) to a definition of ‘meritocratic’ fairness, and we shall use it to explain the deviations from equilibrium predictions in the intermediate meritocracy regimes (HIGH-MERIT and LOW-MERIT).

### Meritocratic fairness: definition

We define a payoff allocation as *fair* in terms of meritocracy if all players are matched into a group with average contribution levels that are compatible with their own contribution, in the sense that they are more similar relative to other players’ contributions in the other groups. In particular, a payoff allocation is considered unfair if there exist at least one player who contributed less (more) than others who are matched into groups with a lower (higher) average contribution level. The more players are matched into incompatible groups, and the larger the difference in average group payoffs, the higher the level of meritocratic unfairness. More formally, *meritocratic unfairness* of a given payoff allocation is measured by the following two quantities:

$$\begin{aligned} MU_{Dis} &= \frac{1}{n-s} * \sum_{j \in N} \max(\Delta_{ij}, 0) * \max(\Delta_{G_j G_i}, 0), \\ MU_{Adv} &= \frac{1}{n-s} * \sum_{j \in N} \max(\Delta_{ji}, 0) * \max(\Delta_{G_i G_j}, 0), \end{aligned} \tag{7.22}$$

where for any pair of players,  $i$  and  $j$  in groups  $G_i$  and  $G_j$  ( $i \neq j$ ),  $\Delta_{ij}$  represents

the difference in contributions  $c_i - c_j$ , and  $\Delta_{G_i G_j}$  is the difference in average group contributions  $\frac{1}{4} \sum_{k \in G_i} c_k - \frac{1}{4} \sum_{k \in G_j} c_k$ .

In line with previous behavioral findings in studies investigating distributional fairness (Fehr and Schmidt, 1999; Ockenfels and Bolton, 2000), we assume that disadvantageous unfairness has a more accentuated negative effect than advantageous unfairness. The consequences of the distaste for meritocratic unfairness in repeated random interactions are such that, on average, a player responds by decreasing (increasing) his/her contribution after experiencing disadvantageous (advantageous) meritocratic unfairness. Note that, under this definition, every outcome is meritocratic and fair with probability one under perfect meritocracy (when  $\sigma^2 = 0$ ).

Our assumptions regarding meritocratic fairness lead to the following predictions:

- In environments with zero meritocracy, our predictions coincide with those of Ref. (Fehr and Schmidt, 1999; Ockenfels and Bolton, 2000), that is, we expect the typical contribution pattern (intermediate contributions levels at the beginning, then decay over time). The decay is driven by the asymmetry in behavioral responses to disadvantageous versus advantageous unfairness.
- Under perfect meritocracy, starting at the near-efficient Nash equilibrium prediction, we do not expect significant departures from such a state as there is no inherent meritocratic unfairness (by definition).
- For the intermediate meritocracy levels (HIGH-MERIT and LOW-MERIT),

starting at the near-efficient Nash equilibrium prediction, we expect decreases as unfairness is expected to occur even in equilibrium. However, other than under zero meritocracy, downward corrections of contributions will not trigger an overall downward decay of contributions because higher amounts become better and fair replies again than contributing zero once substantial decreases of contributions occurred, which were themselves triggered by disadvantageous unfairness. This is due to the fact that there are then new strategic concerns.

### Meritocratic fairness: results

Fig. 7.8 shows the distributions of meritocratic unfairness across different treatments. Similarly to efficiency and inequality, we find increases in fairness from NO-MERIT through all meritocracy levels up to PERFECT-MERIT, and these increases are significant (LMM:  $F_{3,8} = 53.74$ ,  $P < 0.0001$ ).

Meritocratic unfairness translates directly into departures from the levels of contribution predicted by theory. In particular, we studied how the unfairness level experienced in the previous round impacts the decision to contribute in the following round. To do so, we performed a multilevel regression of between-rounds contribution adjustments with subject and session as random effects, and we tested several models for both distributional (Fehr and Schmidt, 1999) and meritocratic fairness (statistical details are given in the Statistical Analysis section in *Materials and Methods* section and regression tables are available in the Supplementary Information). As expected, applying the notion of distributional fairness *as it is* to a meritocratic environment is not straightforward:

the results of the regressions for distributional fairness are often inconsistent across treatments, and, even in many cases contrary to the predictions of the theory. On the other hand, meritocratic unfairness proved a good predictor of the contribution adjustments between rounds across all treatments. Therefore, meritocratic fairness can be seen as natural generalization of distributional fairness in games with positive levels of meritocracy.

#### 7.7.4 Sensitivity

So far, we have shown that (i) both efficiency *and* equality increase with meritocracy, and that (ii) considerations of ‘meritocratic’ fairness can explain deviations from the theoretically expected equilibrium. In this section, we show that changes in the level of experienced meritocracy have significant implications as well. In particular, we test whether participants coming from a higher (lower) meritocracy level in part 1 are more (less) sensitive to meritocratic unfairness in part 2.

For this analysis, we used the data pertaining of part 2 of the experiment, controlling for which meritocracy level was played in part 1. We divided the dataset in two subsets, depending on whether participants in part 2 experienced a higher or lower meritocracy level than in part 1. In order to obtain a balanced design with respect to the direction of meritocracy changes, we further sampled the data from part 2 to include only the intermediate regimes of meritocracy ( $\sigma^2 = \{3, 20\}$ ). In this way, both conditions could be tested against perfect meritocracy, zero meritocracy, and one intermediate regime. We created a dummy variable for “contribution goes down” (0;1) and performed a

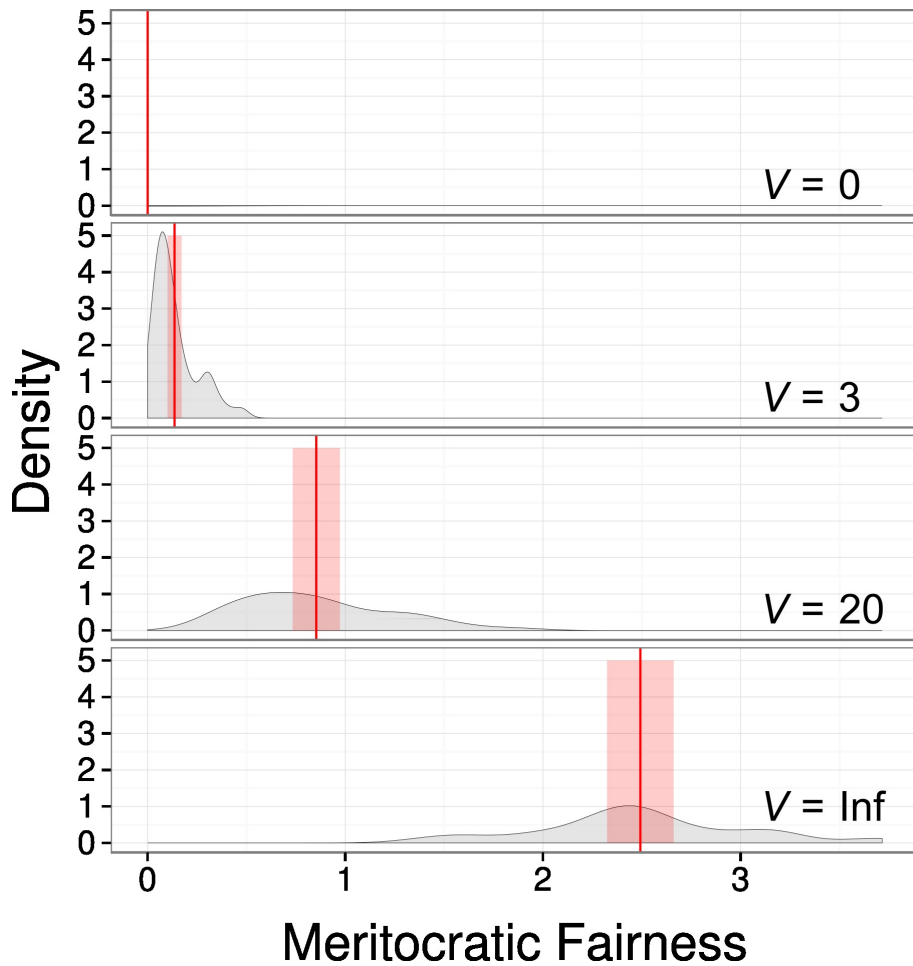


Figure 7.8: Meritocratic unfairness for perfect-, high-, low-, and no-meritocracy, respectively associated with the values of  $\sigma^2 = \{0, 3, 20, \infty\}$ . Smoothed distribution of average meritocratic unfairness per round. Unfairness decreases as meritocracy increases. Red solid lines indicate the mean level of meritocratic unfairness observed in the experiment, red-shaded areas indicate the 95%-confidence intervals of the mean.

multilevel logistic regression with subject and session as random effects. We used the level of disadvantageous meritocratic unfairness experienced in the previous round as a predictor of whether contribution is expected to go up or down in the next round.

Our main finding is that the distaste for meritocratic unfairness is exacerbated after having played a more meritocratic regimes in part 1. That is, if a participant experienced meritocratic unfairness in the previous round, he/she is more likely to reduce the own contribution in the current round if the level of meritocracy in part 2 is lower than in part 1 (Logistic Mixed Regression LMR:  $Z = 2.521, P = 0.0117$ ). The effect in the opposite direction – a lower meritocracy level in part 1 than in part 2 – is not significant (LMR:  $Z = 1.522, P = 0.128$ ).

The different sensitivity to meritocratic unfairness leads to different levels of efficiency and equality overall. Sessions in part 2 with higher sensitivity to meritocratic unfairness – i.e. descending the meritocracy ladder – have significantly lower average payoff (One-sided Kolmogorov-Smirnoff KS:  $D^+ = 0.1531, P < 0.0001$ ), and significantly higher inequality – measured by the average Gini coefficient per round ( $D^+ = 0.1583, P = 0.0494$ ). These results confirm once again that, in our settings, increases in efficiency are followed by inequality reduction, and that meritocratic fairness considerations can explain the dissolution of the classical efficiency-equality tradeoff.



## 7.8 Discussion

Economic theory has identified the efficiency-equality tradeoff as one of the most fundamental tradeoffs underlying society ([gauthier'morals'1986](#); Arrow, [1951](#); Sen, [1970](#); Okun, [1975](#); Arrow, Bowles, and Durlauf, [2000b](#)). In our study, we decided to analyze an environment that succinctly captures the essence of this tradeoff. The well-known public-goods (voluntary-contribution) game (Isaac, McCue, and Plott, [1985b](#)) perfectly suited our task, since it naturally relates to many important real-life issues such as climate change, collective action, common-pool resource problems, etc. (Ostrom, [1990](#); Ostrom, [1999](#)). For this, it has received tremendous attention in the theoretical and experimental literature in and outside of economics (Chaudhuri, [2011b](#)).

The standard case of random re-matching and a recently proposed and seminal group-based mechanism (Gunnthorsdottir et al., [2010](#)) were generalized to a class of mechanisms called “meritocratic matching” (Nax, Murphy, and Helbing, [2014](#)). Here, we test these mechanism, we made the astonishing finding that agents seem to be able to ‘make the better system work’. That is, meritocratic mechanisms that promise higher efficiency from a theoretic point of view, also turn out to benefit the worst-off and to improve overall distributional equality, despite theory predicting the opposite (Nash, [1951](#)). The reason for this unexpected finding lies in agents’ attempts to improve ‘fairness’ by adjustments of their actions in order to counter situations in which particular agents are better-off (worse-off) despite being associated with low (high) ‘merit’. This fairness concept not only explains our results in the new class of assortative games studied by us, but also remains a significant explanatory variable in

games with random interactions, and is consistent with previous results for this class of games. The criterion of ‘meritocratic’ fairness is formally different from the standard formulation of ‘distributional’ fairness (Fehr and Schmidt, 1999; Ockenfels and Bolton, 2000), but for random interaction environments their predictions agree qualitatively. In meritocratic environments, due to the double-role of contributions inherent in the matching mechanism (both as a group-sorting device and as a payoff determinant within groups), the concept of ‘meritocratic’ fairness is indeed a natural extension of classical fairness criteria when agents are aware of this double-nature.

The results of our study show that meritocracy can dissolve the fundamental tradeoff between efficiency and equality. Creating a public good does not necessarily generate inefficiencies, nor it requires the intervention of a central coercive power for their suppression. Fairness preferences and suitable institutional settings, such as well-working merit-based matching mechanisms, can align agents’ incentives, and shift the system towards more cooperative and near-efficient Nash equilibria. Overall, the results of our experiment lend credibility to agents’ sensitivity to the famous quote associated with Virgil that “The noblest motive is the public good.”

# Appendix: Materials and methods

## Appendix A: Equilibrium structure

Our stage games with  $n = 16$ ,  $s = 4$ ,  $B = 20$  and  $m = 0.5$  have the following equilibria dependent on which variance level of  $\sigma^2 = \{0, 3, 20, \infty\}$  is played. When  $\sigma^2 = \infty$  (NO-MERIT), the only equilibrium is  $c_i = 0$  for all  $i$ .  $c_i = 0$  for all  $i$  is also an equilibrium for all other variance levels. In that equilibrium, all players receive a payoff of  $\phi_i = 20$ . However, when  $\sigma^2 = \{0, 3, 20\}$ , there also exist exactly  $\binom{n}{k}$  unique pure-strategy equilibria such that  $c_i = 0$  for exactly two agents and  $c_j = 20$  for the remaining fourteen. In that equilibrium, for the case when  $\sigma^2 = 0$  (PERFECT-MERIT), payoffs are such that twelve of the fourteen players who contribute  $c_i = 20$  are matched in groups with each other and receive  $\phi_i = 40$ . The remaining four players are matched in the worst group. Of those, the two players who contribute  $c_i = 0$  receive a payoff of  $\phi_i = 40$ , while the two players who contribute  $c_i = 20$  receive a payoff of  $\phi_i = 20$ . For the cases when  $\sigma^2 = 3$  (HIGH-MERIT)/ $\sigma^2 = 20$  (LOW-MERIT), payoffs in the last group are as in the case when  $\sigma^2 = 0$  (PERFECT-MERIT) in over 99.9%/ 99% of all cases. In the remaining cases, payoffs are such that 6 out of fourteen players who contribute  $c_i = 20$  are matched in groups with each other and receive  $\phi_i = 40$ . The remaining 6 players who contribute  $c_i = 20$  are matched in a group with one player who contributes  $c_i = 0$  and receives a payoff of 30. The two players who contribute  $c_i = 0$  receive a payoff of  $\phi_i = 50$  each. The near-efficient Nash equilibrium collapses when the variance reaches a level of about  $\sigma^2 = 100$  (see propositions 6 and 7 in Ref. (Nax, Murphy, and

Helbing, 2014)).

## Appendix B: Experimental design

A total of 192 voluntary participants took part in one session consisting of two separate games each. Each session lasted roughly one hour. There were 16 participants in each session and 12 sessions in total. All sessions were conducted at the ETH Decision Science Laboratory (DeSciL) in Zürich, Switzerland, using the experimental software NodeGame (*nodegame.org*). DeSciL recruited the subjects using the Online Recruitment System for Economic Experiments (ORSEE). The experiment followed all standard behavioral economics procedures and meets the ethical committee guidelines. Decisions, earnings and payments were anonymous. Payments were administered by the DeSciL administrators. In addition to a 10 CHF show-up fee, each subject was paid according to a known exchange rate of 0.01 CHF per coin. Overall, monetary rewards ranged from 30 to 50 CHF, with a mean of 39 CHF.

Each session consisted of two games, each of which was a forty-round repetition of the same underlying stage game, namely a public-goods game. The same fixed budget was given to each subject every period. Each game had separate instructions that were distributed at the beginning of each game. After reading the instructions, all participants were quizzed to make sure they understood the task. The two games differ with respect to the variance level that is added to players' contributions. There were four variance levels ( $\sigma^2 = \{0, 3, 20, \infty\}$ ), and each game had equivalent instructions. Instructions contained full information about the structure of the game and about the payoff consequences to

themselves and to the other agents. We played every possible pair of variance levels in both orders to have an orthogonal balanced design, which yields a total of 12 sessions. As the game went on, players learnt about the other players' previous actions and about the groups that formed. Each of our 192 participants made forty contribution decisions in each of the two games in his session. This yields 80 choices per person per session, hence a total of 15,360 observations. More details, including a copy of a full instructions set and the quiz questions, are provided in the Supplementary Information Appendix.

### **Instructions of the lab experiment**

Each experimental session consisted of two separate games (part 1, part 2), each played with a different variance level. We exhausted all possible pair of variance levels in both orders, for a total of 12 different combinations. Consequently, we prepared 12 different instruction texts that took into account whether a variance level was played in the first or in the second part, and in the latter case also considered which variance level was played in part 1.

Together with the main instructions sheet, we provided an additional sheet containing tabulated numerical examples of fictitious game-rounds played at the current variance level. This aimed to let participants get an intuitive feeling of the consequences of noise on contributions and final payoffs.

All instructions texts can viewed at the address <http://nodegame.org/games/merit/>. Here we report the instruction text for variance level equal 20 played in the part 1.

## **Instructions for Variance Level = 20, Part 1**

Welcome to the experiment and thanks for your participation. You have been randomly assigned to an experimental condition with 16 people in total. In other words you and 15 others will be interacting via the computer network for this entire experimental session.

The experiment is divided into two parts and each part will last approximately 30-40 minutes long. Both parts of the experiment contribute to your final earnings. The instructions for the first part of the experiment follow directly below. The instructions for the second part of the experiment will be handed out to you only after all participants have completed the first part of the experiment. It is worth your effort to read and understand these instructions well. You will be paid based on your performance in this study; the better you perform, the higher your expected earnings will be for your participation today.

### **Your decision.**

In this part you will play 40 independent rounds. At the beginning of each round, you will receive 20 “coins”. For each round, you will have to decide how many of your 20 coins to transfer into your “personal” account, and how many coins to transfer into a “group” account. Your earnings for the round depend on how you and the other participants decide to divide the coins you have received between the two accounts.

### **Group matching with noise.**

For each round you will be assigned to a group of 4 people, that is, you and

three other participants. In general, groups are formed by ranking each individual transfer to the group account, from the highest to the lowest. Group 1 is generally composed of those participants who transferred the most to the group account; Group 4 is generally composed of those who transferred the least to the group account. The other groups (2 and 3) are between these two extremes.

However, the sorting process is noisy by design; contributing more will increase a participant's chances of being in a higher ranked group, but a high ranking is not guaranteed. Technical note- The noisy ranking and sorting is implemented with the following process:

1. **Step 1:** Preliminary ordering. A preliminary list is created in which transfers to the group account are ranked from highest to lowest. In case two or more individuals transfer the same amount, their relative position in the ranking will be decided randomly.
2. **Step 2:** Noisy ordering. From every participant's actual transfer to the group account, we obtain a unique noisy contribution by adding an i.i.d. (independent and identically distributed) normal variable with mean 0 and variance 20. The noisy contributions are then ranked from 1 to 16 from highest to lowest, and a final list is created.
3. **Step 3:** Group matching. Based on the final list created at Step 2 (the list with noise), the first 4 participants on that list form Group 1, the next 4 people in the list form Group 2, the third 4 people in the list form Group 3, and the last 4 people form Group 4.

**Return from personal account.**

Each coin that you put into your personal account results in a simple one-to-one payoff towards your total earnings.

**Return from group account.**

Each coin that you put into the group account will pay you back some positive amount of money, but it depends also on how much the other group members have transferred to the group account, as described below.

The total amount of coins in your group account is equal to the sum of the transfers to the group account by each of the group members. That amount is then multiplied by 2 and distributed equally among the 4 group members. In other words, you will get a return equal to half of the group account total.

**Final Earnings**

Your total earnings for the first part of the experiment are equal to the sum of all your rounds' earnings. One coin is equal to 0.01 CHF. This may not appear to be very much money, but remember there are 40 rounds in this part of the experiment so these earnings build up.

**Example**

Here is an example of one round to demonstrate this decision context, the noisy sorting into different groups, and the different resulting payoffs. In the table below, pay attention to the following facts:

- Groups are roughly formed by ranking how much participants transferred to the group account, but this is not a perfect ranking. For example,



participant #8 transferred less to the group account than participant #10, but the noisy sorting process placed him in a higher ranked group.

- Participant #7 transferred 14 of his coins to the group account. This means that he transferred 6 to his personal account. Due to noisy sorting he was ranked first, and assigned to Group 1. The other participants in Group 1 transferred a total of 64 coins to the group account. This amount is doubled and redistributed evenly back to the 4 members of the group this is 32 for each participant. So then participant #7 earned 38 coins for this round.
- Participant #12 transferred 7 coins to the group account and transferred the remaining 13 coins to his personal account. He was sorted (with noise) into Group 3 and this group transferred 46 coins in total. This resulted in 23 coins being returned to each of the group members, and thus his total payoff is 36 coins (23 returned from the group account and the 13 he kept in his personal account).

Player ID		Transfer Group to group account	Transfer to personal account	Total to group account	Amount returned to player	Total earnings for the round
7	1	14	6	64	32	38
6	1	13	7	64	32	39
14	1	16	4	64	32	36
4	1	8	12	64	32	44
1	2	14	6	51	25.5	31.5
3	2	20	0	51	25.5	25.5
8	2	11	9	51	25.5	34.5
11	2	19	1	51	25.5	26.5
10	3	17	3	46	23	26
12	3	7	13	46	23	36
16	3	6	14	46	23	37
5	3	16	4	46	23	27
9	4	10	10	18	9	19
2	4	1	19	18	9	28
13	4	5	15	18	9	24
15	4	2	18	18	9	27

Additional examples are provided in a separate sheet for your own reference.

## Quiz

Subjects were given a quiz after instructions to test their understanding of the game. Only after “passing” the quiz were subjects allowed to begin play. Details about the quiz can be found at <http://nodegame.org/games/merit/>.

## Graphical interface of the experiment

The experiment was implemented using the experimental software nodeGame **nodegame** website Besides, offering a textual response of the actions of the players, we also offer a visual summary with contributions bars ordered by group. More details about the interface, and the implementation are available at the url: <http://nodegame.org/games/merit/>

## Appendix C: Statistical analyses

### Equality analysis

Overall, we found a significant difference in the variance of realized payoffs in each round among the four treatments (LMM:  $F_{3,8} = 7.27, P < 0.0113$ ). When computing Bonferroni adjusted  $p$ -values, the treatment with variance  $\infty$  was found significantly different ( $P = 0.0003; P = 0.0004; P = 0.0086$ ) from the other three variance levels ( $\sigma^2 = \{0, 3, 20\}$ ), which are themselves not significantly different from each other. Taking NO-MERIT as a baseline, LOW-MERIT led to a decrease in the variance of realized payoffs in each round of -13.546 (LRT  $\chi_{(1)} = 8.13, P = 0.0043$ ), HIGH-MERIT to a decrease of -16.914 (LRT  $\chi_{(1)} = 9.89, P = 0.0016$ ), and PERFECT-MERIT to a decrease of -17.122 (LRT  $\chi_{(1)} = 6.78, P = 0.0091$ ).

Similarly, the Gini index differs significantly among the four treatments (LMM:  $F_{3,20} = 42.0, P < 0.0001$ ). Taking NO-MERIT as a baseline, LOW-MERIT led to a decrease in the variance of realized payoff in each round of -0.058901 (LRT  $\chi_{(1)} = 18.18, P < 0.0001$ ), HIGH-MERIT to a decrease of -0.071843 (LRT  $\chi_{(1)} = 22.28, P < 0.0001$ ), and PERFECT-MERIT to a decrease of -0.075453 (LRT  $\chi_{(1)} = 22.06, P < 0.0001$ ). Computing Bonferroni adjusted  $p$ -values for all pair-wise differences reveals that the treatment with variance  $\infty$  is significantly different ( $P < 0.0001$ ) from the other three variance levels ( $\sigma^2 = \{0, 3, 20\}$ ), which are themselves not significantly different from each other (see Fig. 7.7).

## Fairness analysis

We find a significant difference in the experienced levels of meritocratic unfairness in each round among the four treatments (LMM:  $F_{3,8} = 53.74, P < 0.0001$ ). When computing Bonferroni adjusted  $p$ -values we find that – excluding PERFECT-MERIT for which meritocratic unfairness is always zero by definition – all treatments are statistically significantly different from each other (HIGH-MERIT vs LOW-MERIT  $P = 0.0071$ , all the other pair-wise comparisons  $P < 0.0001$ ). Taking NO-MERIT as a baseline, LOW-MERIT led to a decrease in the experienced meritocratic unfairness in each round of -1.66 (LRT  $\chi_{(1)} = 11.76, P = 0.0006$ ), HIGH-MERIT to a decrease of -2.36 (LRT  $\chi_{(1)} = 18.92, P < 0.0001$ ).

We also analyzed the effect of meritocratic (dis)advantageous unfairness on contribution adjustments between rounds, by performing a multilevel regression with subject and session as random effects. Our findings reveal that disadvantageous unfairness leads to decreases in treatments LOW-MERIT  $-0.18^{***}(0.05)$ , and NO-MERIT  $-0.25^{***}(0.03)$ ). For HIGH-MERIT the decrease is consistent in sign and size, but not statistically significant  $-0.39(0.21)$ . However, if HIGH-MERIT and LOW-MERIT are pooled together the effect turns out to be significant  $-0.25^{***}(0.03)$ . Meritocratic disadvantageous fairness can, therefore, originate significant differences between the theoretical equilibrium predictions and experimentally observed behavior. Advantageous unfairness leads to increases under some but not under all regimes. Full regression tables are available in the remainder of this Appendix.

## Fairness regressions

Here we report the results of the mixed-effects regressions of meritocratic and distributional fairness on contributions adjustments between rounds in part 1 and part 2 of the experiment. As we argued in the main text, distributional fairness cannot easily be generalized to the case of assortative matching. Here we show that a naïve extension of the formula in Fehr and Schmidt, 1999 fails to reproduce the results predicted by theory. In fact, both within-group and across-groups distributional fairness under assortativity often lead to the contradictory result that disadvantageous fairness implies an increase in the contribution levels. However, by taking into account assortativity in the formula of distributional fairness, we developed an extension that is able to reproduce the results predicted by the theory for all treatments.

### Meritocratic fairness

In tables 7.2 and 7.3, meritocratic unfairness is used as a predictor. `lag.merit.fair.dis` and `lag.merit.fair.adv` are respectively the amount of *disadvantageous* and *advantageous* meritocratic unfairness experienced by a player in the previous round, measured according to the equations in 7.22 in appendix.

### Distributional fairness

The results of the regressions for distributional fairness are shown in tables 7.4, 7.5, 7.6 and 7.7. Based on the original formula in Ref. Fehr and Schmidt, 1999, we tried two different extensions of the notion of distributional fairness for

Table 7.2: **Meritocratic fairness predicts contribution differential.**  
**(Part 1)** The sign of the regression coefficient is always consistent with theory predictions. HIGH-MERIT is significant if pooled together with LOW-MERIT.

	HIGH-MERIT	LOW-MERIT	HIGH-MERIT&LOW-MERIT	NO-MERIT
(Intercept)	0.25 (0.16)	0.15 (0.16)	0.03 (0.19)	0.03 (0.19)
lag.merit.fair.dis	-0.39 (0.21)	-0.18*** (0.05)	-0.25*** (0.03)	-0.25*** (0.03)
lag.merit.fair.adv	-0.91** (0.30)	0.06 (0.06)	0.15*** (0.03)	0.15*** (0.03)
AIC	12314.36	12284.05	12359.50	12359.50
BIC	12347.56	12317.24	12392.70	12392.70
Log Likelihood	-6151.18	-6136.02	-6173.75	-6173.75
Num. obs.	1872	1870	1872	1872

\*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$

Table 7.3: **Meritocratic fairness predicts contribution differential.**  
**(Part 2)** The sign of the regression coefficient is always consistent with theory predictions. HIGH-MERIT is significant if pooled together with LOW-MERIT.

	HIGH-MERIT	LOW-MERIT	HIGH-MERIT&LOW-MERIT	NO-MERIT
(Intercept)	0.13 (0.16)	0.16 (0.17)	0.11 (0.11)	0.38* (0.18)
lag.merit.fair.dis	-0.45 (0.28)	-0.29*** (0.07)	-0.29*** (0.06)	-0.26*** (0.02)
lag.merit.fair.adv	-0.57 (0.32)	0.00 (0.07)	-0.02 (0.07)	0.04 (0.02)
AIC	12288.63	12419.05	24699.24	12123.03
BIC	12321.83	12452.25	24736.60	12156.23
Log Likelihood	-6138.31	-6203.53	-12343.62	-6055.51
Num. obs.	1872	1871	3743	1872

\*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$

meritocratic environments. First, we computed distributional fairness for each player only taking into account the other players within the group into which he/she was matched (Within-group distributional fairness). The regressors in this case are called: `lag.distr.fair.group.dis` and `lag.distr.fair.group.adv`.

Then, we also computed distributional fairness across all players, regardless of the group they belonged to (Across-group distributional fairness). The regressors for across-group distributional fairness are called: `lag.distr.fair.dis` and `lag.distr.fair.adv`.

Table 7.4: **Within-group distributional fairness predicts contribution differential. (Part 1)** The sign of the regression coefficient is often inconsistent with theory predictions.

	PERFECT-MERIT	HIGH-MERIT	LOW-MERIT	HIGH-MERIT & LOW-MERIT	NO-MERIT
(Intercept)	-0.79*** (0.23)	-1.39*** (0.22)	-1.32*** (0.21)	-1.39*** (0.15)	1.40** (0.45)
lag.distr.fair.group.dis	-0.03 (0.04)	0.13** (0.05)	0.01 (0.05)	0.06* (0.03)	-0.70*** (0.04)
lag.distr.fair.group.adv	0.76*** (0.04)	0.99*** (0.04)	0.77*** (0.04)	0.88*** (0.03)	0.28*** (0.04)
AIC	11682.40	11933.18	12025.27	23952.86	11968.23
BIC	11715.59	11966.38	12058.46	23990.22	12001.43
Log Likelihood	-5835.20	-5960.59	-6006.64	-11970.43	-5978.12
Num. obs.	1872	1872	1870	3742	1872

\*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$



Table 7.5: **Within-group distributional fairness predicts contribution differential. (Part 2)** The sign of the regression coefficient is often inconsistent with theory predictions.

	PERFECT-MERIT	HIGH-MERIT	LOW-MERIT	HIGH-MERIT & LOW-MERIT	NO-MERIT
(Intercept)	-0.93*** (0.25)	-1.54*** (0.40)	-1.25*** (0.23)	-1.43*** (0.22)	1.60*** (0.38)
lag.distr.fair.group.dis	-0.10* (0.04)	0.05 (0.04)	-0.06 (0.05)	0.00 (0.03)	-0.61*** (0.03)
lag.distr.fair.group.adv	0.88*** (0.04)	1.19*** (0.04)	0.86*** (0.04)	1.02*** (0.03)	0.15*** (0.03)
AIC	11856.01	11799.36	12109.33	23935.12	11827.92
BIC	11889.21	11832.55	12142.53	23972.48	11861.12
Log Likelihood	-5922.01	-5893.68	-6048.67	-11961.56	-5907.96
Num. obs.	1871	1872	1871	3743	1872

\*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$

Table 7.6: **Across-group distributional fairness predicts contribution differential. (Part 1)** The sign of the regression coefficient is often inconsistent with theory predictions.

	PERFECT-MERIT	HIGH-MERIT	LOW-MERIT	HIGH-MERIT & LOW-MERIT	NO-MERIT
(Intercept)	-1.42*** (0.26)	-2.40*** (0.34)	-2.20*** (0.34)	-2.23*** (0.24)	1.04* (0.40)
lag.distr.fair.dis	0.22*** (0.03)	0.39*** (0.04)	0.33*** (0.04)	0.35*** (0.03)	-0.44*** (0.05)
lag.distr.fair.adv	0.44*** (0.08)	0.59*** (0.10)	0.43*** (0.08)	0.48*** (0.06)	0.13* (0.05)
AIC	11934.03	12223.59	12225.86	24434.15	12277.90
BIC	11967.23	12256.79	12259.05	24471.51	12311.10
Log Likelihood	-5961.02	-6105.80	-6106.93	-12211.07	-6132.95
Num. obs.	1872	1872	1870	3742	1872

\*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$

Table 7.7: **Across-group distributional fairness predicts contribution differential. (Part 2)** The sign of the regression coefficient is often inconsistent with theory predictions.

	PERFECT-MERIT	HIGH-MERIT	LOW-MERIT	HIGH-MERIT & LOW-MERIT	NO-MERIT
(Intercept)	-2.15*** (0.30)	-1.98*** (0.30)	-2.19*** (0.35)	-2.01*** (0.23)	1.96*** (0.48)
lag.distr.fair.dis	0.21*** (0.03)	0.29*** (0.03)	0.30*** (0.04)	0.29*** (0.02)	-0.49*** (0.04)
lag.distr.fair.adv	0.65*** (0.09)	0.54*** (0.09)	0.46*** (0.09)	0.48*** (0.06)	-0.04 (0.04)
AIC	12162.64	12222.36	12374.95	24584.87	12068.03
BIC	12195.83	12255.56	12408.15	24622.23	12101.23
Log Likelihood	-6075.32	-6105.18	-6181.48	-12286.43	-6028.02
Num. obs.	1871	1872	1871	3743	1872

\*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$

## Additional inequality indexes

As stated in the main text, inequality decreases as meritocracy increases. In this section, we show that our finding is robust to the type of inequality measurement chosen. Fig. 7.9 displays the payoff inequality as measured by a number of different indexes commonly found in the literature of inequality studies Atkinson, 1970.




Figure 7.9: **Battery of indexes measuring payoff inequality over the forty rounds for perfect-, high-, low-, and no-meritocracy, respectively associated with the values of  $\sigma^2 = \{0, 3, 20, \infty\}$ . Inequality decreases with meritocracy for a large number of distinct inequality indexes. Error bars represent the 95%-confidence intervals**

## Appendix D: Implications

Our model implies that situations consistent with our model assumptions would benefit from higher degrees of meritocracy, both in terms of efficiency and in terms of equality. This positive result relies on several features of the underlying model. It is an avenue for future research to consider these generalizations. First, our model describes an *ex ante* homogeneous population. Differences in payoff are driven by differences in actions and by neutral stochastic elements alone. Heterogeneity in priority given by the matching mechanism and/or heterogeneities in the individual rates of return could influence the results. This is true for any public-goods game including the standard models with random interactions (e.g. Buckley and Croson, 2006; Fischbacher, Schudy, and Teyssier, 2014). However, it should be noted that meritocracy may actually mitigate the associated inequality problems. Second, related to heterogeneity, our model allows for no wealth creation, that is, individuals receive a new budget every period and the size of this budget is fixed and constant over time. Players cannot accumulate wealth. The role of wealth creation in public-goods games has received some attention and has been shown to lead to the emergence of different classes of contributions and income (e.g. Tamai, 2010, see also King and Rebelo, 1990; Rebelo, 1991). Under assortative matching, wealth creation can be problematic as it allows rich players to block out poor players. Third, group sizes are fixed. Alternative models have been proposed (e.g. Cinyabuguma, Page, and Putterman, 2005; Charness and Yang, 2008; Ehrhart and Keser, 1999; Ahn, Isaac, and Salmon, 2008; Coricelli, Fehr, and Fellner, 2004; Page, Putterman, and Unel, 2005; Brekke, Nyborg, and Rege,

2007; Brekke et al., 2011).

## References

- Adams, J. S. (1965). “Inequity in social exchange”. In: *L. Berkowitz (Ed.): Advances in experimental social psychology* 2, pp. 267–299.
- Ahn, T., R. M. Isaac, and T. C. Salmon (2008). “Endogenous group formation”. In: *Journal of Public Economic Theory* 10, pp. 171–194.
- Alger, I. and J.W. Weibull (2013). “Homo Moralis - Preference Evolution Under Incomplete Information and Assortative Matching”. In: *Econometrica* 81, pp. 2269–2302.
- Allchin, D. (2009). *The Evolution of Morality*. Evolution: Education and Outreach. Springer.
- Andreoni, J. (1988). “Why free ride? Strategies and learning in public goods experiments”. In: *Journal of Public Economics* 37, pp. 291–304.
- (1993). “An experimental test of the public goods crowding-out hypothesis”. In: *American Economic Review* 83, pp. 1317–1327.
- (1995). “Cooperation in public-goods experiments: kindness or confusion?”. In: *American Economic Review* 85, pp. 891–904.
- Andreoni, J. and R. Petrie (2004). “Public goods experiments without confidentiality: a glimpse into fund-raising”. In: *Journal of Public Economics* 88, pp. 1605–1623.
- Arrow, K., S. Bowles, and S. Durlauf (2000a). *Meritocracy and Economic Inequality*. Princeton University Press.
- (2000b). *Meritocracy and Economic Inequality*. Princeton University Press.

- Arrow, K. J. (1951). *Social Choice and Individual Values*. Yale, USA: Yale University Press.
- Atkinson, A. B. (1970). “On the measurement of inequality”. In: *Journal of Economic Theory* 2, pp. 244–263.
- (2012). “Public Economics after the Idea of Justice”. In: *Journal of Human Development and Capabilities* 13.4, pp. 521–536.
- Bayer, R.-C., E. Renner, and R. Sausgruber (2013). “Confusion and learning in the voluntary contributions game”. In: *Experimental Economics* 16, pp. 478–496.
- Becker, G. S. (1973). “A Theory of Marriage: Part 1”. In: *Journal of Political Economy* 81, pp. 813–846.
- Bentham, J. (1907). *An Introduction to the Principles of Morals and Legislation*. Clarendon Press.
- Binmore, K. (2005). *Natural Justice*. Oxford University Press.
- Bohm, R. and B. Rockenbach (2013). “The inter-group comparisonintra-group cooperation hypothesis”. In: *PLoS ONE* 8, p. 56152.
- Bowles, S. and H. Gintis (2011). *A cooperative specieshuman reciprocity and its evolution*. Princeton University Press.
- Brekke, K., K. Nyborg, and M. Rege (2007). “The fear of exclusion: individual effort when group formation is endogenous”. In: *Scandinavian Journal of Economics* 109, pp. 531–550.
- Brekke, K. et al. (2011). “Playing with the good guys. A public good game with endogenous group formation”. In: *Journal of Public Economics* 95, pp. 1111–1118.

- Buckley, Edward and Rachel Croson (2006). “Income and wealth heterogeneity in the voluntary provision of linear public goods”. In: *Journal of Public Economics* 90.4-5, pp. 935–955.
- Cabral, L. M. B. (1988). “ASYMMETRIC EQUILIBRIA IN SYMMETRIC GAMES WITH MANY PLAYERS”. In: *Economic Letters* 27, pp. 205–208.
- Charness, G. B. and C.-L. Yang (2008). “Endogenous Group Formation and Public Goods Provision: Exclusion, Exit, Mergers, and Redemption”. In: *University of California at Santa Barbara, Economics WP*.
- Chaudhuri, A. (2011a). “Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature”. In: *Experimental Economics* 14, pp. 47–83.
- Chaudhuri, Ananish (2011b). “Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature”. In: *Experimental Economics* 14, pp. 47–83.
- Cinyabuguma, M., T. Page, and L. Putterman (2005). “Cooperation under the threat of expulsion in a public goods experiment”. In: *Journal of Public Economics* 89, pp. 1421–1435.
- Cole, H. J., G. Mailath, and A. Postlewaite (1992). “Social Norms, Savings Behavior, and Growth”. In: *Journal of Political Economy* 100, pp. 1092–1125.
- Coricelli, G., D. Fehr, and G. Fellner (2004). “Partner Selection in Public Goods Experiments”. In: *Economics Series* 151.
- Cowell, F. (2011). *Measuring Inequality*. Oxford University Press.

- Dickinson, D. L. and R. M. Isaac (1998). “Absolute and relative rewards for individuals in team production”. In: *Managerial and Decision Economics* 19, pp. 299–310.
- Ehrhart, K. and C. Keser (1999). “Mobility and cooperation: On the run”. In: *CIRANO WP 99.s-24*.
- Erev, Ido, Eyal Ert, and Eldad Yechiam (2008). “Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions”. In: *Journal of Behavioral Decision Making* 21.5, pp. 575–597.
- Fehr, E. and C. Camerer (2007). “Social neuroeconomics: the neural circuitry of social preferences”. In: *Trends in Cognitive Sciences* 11, pp. 419–427.
- Fehr, Ernst and Simon Gächter (2000). “Cooperation and Punishment in Public Goods Experiments”. In: *Am. Econ. Rev.* 90, pp. 980–994.
- Fehr, Ernst and Klaus M. Schmidt (1999). “A Theory of Fairness, Competition, and Cooperation”. In: *Quarterly J. Econ.* 114, pp. 817–868.
- Feldman, A. (1980). *Welfare Economics and Social Choice Theory*. Boston, USA: Martinus Nijhoff Publishing.
- Ferraro, P. J. and C. A. Vossler (2010). “The source and significance of confusion in public goods experiments”. In: *The B.E. Journal in Economic Analysis and Policy* 10, p. 53.
- Fischbacher, U. and S. Gaechter (2010). “Social preferences, beliefs, and the dynamics of free riding in public good experiments”. In: *American Economic Review* 100, pp. 541–556.
- Fischbacher, Urs, Simeon Schudy, and Sabrina Teyssier (2014). “Heterogeneous reactions to heterogeneity in returns from public goods”. In: *Social Choice and Welfare* 43.1, pp. 195–217.



- Foster, D. and H. P. Young (1990). “Stochastic evolutionary game dynamics”.  
In: *Theoretical Population Biology* 38, pp. 219–232.
- Goeree, J. K., C. A. Holt, and S. K. Laury (2002). “Private costs and public benefits: Unraveling the effects of altruism and noisy behavior”. In: *Journal of Public Economics* 83, pp. 255–276.
- Greenberg, Jerald (1987). “A taxonomy of organizational justice theories”. In: *Academy of Management review* 12.1, pp. 9–22.
- Greenwood, J. et al. (2014). “MARRY YOUR LIKE: ASSORTATIVE MATING AND INCOME INEQUALITY”. In: *NBER WP* 19829.
- Grund, T., C. Waloszek, and D. Helbing (2013). “How Natural Selection Can Create Both Self- and Other-Regarding Preferences, and Networked Minds”. In: *Scientific Reports* 3, p. 1480.
- Gunthorsdottir, A. and P. Thorsteinsson (2010). “Tacit Coordination and Equilibrium Selection in a Merit-Based Grouping Mechanism: A Cross-Cultural Validation Study”. In: *Department of Economics WP* 0.
- Gunthorsdottir, A., R. Vragov, and J. Shen (2010). “TACIT COORDINATION IN CONTRIBUTION-BASED GROUPING WITH TWO ENDOWMENT LEVELS”. In: *Research in Experimental Economics* 13, pp. 13–75.
- Gunthorsdottir, A. et al. (2010). “Near-efficient equilibria in contribution-based competitive grouping”. In: *Journal of Public Economics* 94, pp. 987–994.
- Hamilton, W. D. (1964a). “The Genetical Evolution of Social Behaviour I”. In: *Journal of Theoretical Biology* 7, pp. 1–16.
- (1964b). “The Genetical Evolution of Social Behaviour II”. In: *Journal of Theoretical Biology* 7, pp. 17–52.

- Hardin, Gerrett (1968). “The Tragedy of the Commons”. In: *Science* 162, pp. 1243–1248.
- Harsanyi, J. (1953). “Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking”. In: *Journal of Political Economy* 61, pp. 434–435.
- Harsanyi, J. C. and R. Selten (1988a). *A General Theory of Equilibrium Selection in Games*. MIT Press.
- (1988b). *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.
- Hayek, F. A. von (1935). “The Nature and History of the Problem”. In: *Collectivist Economic Planning*, pp. 1–47.
- Helbing, D. (1996). “A stochastic behavioral model and a ‘microscopic’ foundation of evolutionary game theory”. In: *Theory and Decision* 40, pp. 149–179.
- Irlenbusch, B. and G. Ruchala (2008). “Relative Rewards within Team-Based Compensation”. In: *Labour Economics* 15, pp. 141–167.
- Isaac, M. and J. Walker (1988). “Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism”. In: *Quarterly Journal of Economics* 103, pp. 179–199.
- Isaac, M. R., K. F. McCue, and C. R. Plott (1985a). “Public goods provision in an experimental environment”. In: *Journal of Public Economics* 26, pp. 51–74.
- Isaac, Mark R., Kenneth F. McCue, and Charles R. Plott (1985b). “Public goods provision in an experimental environment”. In: *Journal of Public Economics* 26, pp. 51–74.

- Jones-Lee, M. W. and G. Loomes (1995). “Discounting and Safety”. English. In: *Oxford Economic Papers*. New Series 47, pp. 501–512.
- Kahneman, D. and A. Tversky (1979). “Prospect Theory: An Analysis of Decision under Risk”. In: *Econometrica* 47, pp. 263–291.
- Kandori, M., G. J. Mailath, and R. Rob (1993). “Learning, mutation, and long run equilibria in games”. In: *Econometrica* 61, pp. 29–56.
- King, Robert G and Sergio Rebelo (1990). “Public Policy and Economic Growth: Developing Neoclassical Implications”. In: *Journal of Political Economy* 98.5, S126–50.
- Lane, G. (2004). *Genghis Khan and Mongol Rule*. Greenwood.
- Ledyard, J. O. (1995). “Public Goods: A Survey of Experimental Research”. In: *in J. H. Kagel and A. E. Roth (Eds.), Handbook of experimental economics* 37, pp. 111–194.
- (1997). “Public Goods: A Survey of Experimental Research”. In: *The Handbook of Experimental Economics*. Ed. by J. H. Kagel and A. E. Roth. Princeton, NJ: Princeton University Press, pp. 111–194.
- Maynard Smith, J. and G. R. Price (1973). “The logic of animal conflict”. In: *Nature* 246, pp. 15–18.
- Mises, Ludwig von (1922). *Die Gemeinwirtschaft: Untersuchungen über den Sozialismus*. Jena, Germany: Gustav Fischer Verlag.
- Miyazaki, I. (1976). *China’s Examination Hell: The Civil Service Examinations of Imperial China*. Weatherhill.
- Nash, John (1950). “Equilibrium points in n-person games”. In: *Proc. Natl. Acad. Sci. USA* 36, pp. 48–49.
- (1951). “Non-cooperative games”. In: *Ann. Math.* 54, pp. 286–295.

- Nax, H. H., R. O. Murphy, and D. Helbing (2014). *Stability and welfare of 'merit-based' group-matching mechanisms in voluntary contribution games*.
- Nax, H. H. et al. (2013). "Learning in a Black Box". In: *Department of Economics WP, University of Oxford* 653.
- Nowak, M. A. (2006). "Five rules for the evolution of cooperation". In: *Science* 314, pp. 1560–1563.
- Ockenfels, Axel and Gary E. Bolton (2000). "ERC: A Theory of Equity, Reciprocity, and Competition". In: *American Economic Review* 90.1, pp. 166–193.
- Okun, A.M. (1975). *The Big Tradeoff*. Washington D.C.: Brookings Institution Press.
- Ones, U. and L. Putterman (2007). "The ecology of collective action: A public goods and sanctions experiment with controlled group formation". In: *Journal of Economic Behavior and Organization* 62, pp. 495–521.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge, U.K.: Cambridge University Press.
- Ostrom, Elinor (1999). "COPING WITH TRAGEDIES OF THE COMMONS". In: *Annu. Rev. Polit. Sci.* 2, pp. 493–535.
- Page, T., L. Putterman, and B. Unel (2005). "Voluntary association in public goods experiments: reciprocity, mimicry and efficiency". In: *The Economic Journal* 115, pp. 1032–1053.
- Palfrey, T. R. and J. E. Prisbrey (1996). "Altruism, reputation and noise in linear public goods experiments". In: *Journal of Public Economics* 61, pp. 409–427.

- Palfrey, T. R. and J. E. Prisbrey (1997). “Anomalous behavior in public goods experiments: how much and why?” In: *American Economic Review* 87, pp. 829–846.
- Rabanal, J. P. and O. A. Rabanal. “Efcient Investment via Assortative Matching: A laboratory experiment”. In: *mimeo*.
- Rawls, J. (1971). *A Theory of Justice*. Belknap Press.
- Rebelo, Sergio (1991). “Long-Run Policy Analysis and Long-Run Growth”. In: *Journal of Political Economy* 99.3, pp. 500–521.
- Samek, S. and R. Sheremeta (2014). “Visibility of Contributors: An Experiment on Public Goods”. In: *Experimental Economics*.
- Samuelson, P. A. (1980). *Foundations of Economic Analysis*. Cambridge, USA: Harvard University Press.
- Sen, Amartya (1970). “The Impossibility of a Paretian Liberal”. In: *Journal of Political Economy* 78.1, pp. 152–57.
- Simon, H. A. (1990). “A mechanism for social selection and successful altruism”. In: *Science* 250, pp. 1665–1668.
- Tamai, Toshiki (2010). “Public goods provision, redistributive taxation, and wealth accumulation”. In: *Journal of Public Economics* 94.11-12, pp. 1067–1072.
- Taylor, P. D. and L. Jonker (1978). “Evolutionary stable strategies and game dynamics”. In: *Mathematical Bioscience* 40, 145156.
- Tversky, A. and D. Kahneman (1991). “Loss Aversion in Riskless Choice: A Reference Dependent Model”. In: *Quarterly Journal of Economics* 106, pp. 1039–1061.
- Weibull, J. (1995). *Evolutionary Game Theory*. The MIT Press.

- Young, H. P. (1993). “The Evolution of Conventions”. In: *Econometrica* 61, pp. 57–84.
- (1998). *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press.
- Young, M. (1958a). *The Rise of the Meritocracy, 1870-2033: An Essay on Education and Equality*. Transaction Publishers.
- (1958b). *The Rise of the Meritocracy, 1870-2033: An Essay on Education and Equality*. Transaction Publishers.

# Chapter 8

# Conclusion

Maybe better names for “game theory” would be “strategy theory/strategics” or “interaction theory/interactics”. The word “game”, in everyday language, insinuates joy or playfulness, and is therefore often mis-associated with such things as the (computer) gaming industry or board games. This stands in stark contrast to the seriousness of many of the interactions that are studied using game-theoretic models such as political conflict, public goods provision, or organ transplantation markets. But the word “game” also does something useful. Namely, it captures something integral to human nature related to what has been described as *homo ludens* by Johan Huizinga (in his 1938-book), which is that humans, even in very serious situations, often behave in ways that are hard to predict because they experiment/gamble/reason in ways of strategic logic that are hard to decipher.

Indeed, standard game-theoretic solution concepts often fail to predict human behavior, overestimating either the selfishness or the strategic rationality of players. Behavioral game theory relaxes the extreme assumptions regarding selfishness and/or strategic rationality. Behavioral models of game behaviors allow humans to be driven by social preferences and norms, and/or to learn and to make mistakes as the game plays out. Outside economics, not making extreme assumptions such as infallibility would be commonsensical, and indeed laboratory and real-world behaviors provide many examples. Economics is beginning to acknowledge such behavioral components, too.

The aim of this thesis is to improve our understanding of two separate aspects fundamental to behavioral game theory. On the one hand, the thesis aims to contribute to predicting the consequences of behavioral models of game play, especially of game dynamics driven by learning. Chapters 2, 4, 5 and 7 are written to this end. On the other hand, the thesis seeks to improve our modeling foundations, that is, to know what behavioral models best describe



the deviations from standard economic predictions. Chapters 1, 3, 6 and 7 pursue this goal.

The findings of this thesis can be summarized as follows. In terms of theoretical predictions, learning dynamics are shown to approach equilibrium predictions, at least in a zonal sense, even if players have very little knowledge about the game and about other players' roles, actions and payoffs. This finding holds true in public goods environments (Chapters 2 and 7) and in market games (Chapter 4). Particularly noteworthy is the finding from Chapter 7 that humans seem responsive and capable of exploiting the meritocracy of a mechanism to coordinate to payoff-superior equilibria. Chapter 5 illustrates what kind of cooperative outcomes may be feasible in complex environments.

In terms of deducing behavioral patterns of game play from laboratory experiments and real-world data, the thesis corroborates a number of existing theories, challenges others, and, more importantly, contributes toward depicting an integrated model of context-dependent agentic heterogeneity. First, consistent heterogeneity in behaviors in the population in all experiments of the thesis is identified, not just in terms of the magnitudes of behavioral components but also in terms of the nature of their motivations. It was estimated (Chapter 1) that roughly half of the population plays or is able to learn equilibrium behavior, that roughly one third is driven by social preferences and therefore deviates from equilibrium play consistently, and that the remaining players are inexplicably inconsistent and do not learn equilibrium play. This leaves open what kind of learning behavior drives players, either predominantly or given a specific context. Indeed, this question was addressed (Chapter 3) and it was found that the kind of adjustment dynamics that describe learning behavior well depend on the precise informational context of the game, but that certain payoff-based, directional learning components tend to be robust

(precisely these components were explored theoretically in Chapters 2 and 4). Such trend-following behaviors persist with possibly grave consequences even on financial markets where extreme rationality assumptions are often made instead (Chapter 6).

Perhaps the most subtle findings of the thesis were borne out of our recent work on institutionally “meritocratic” mechanism designs (Chapter 7). Theory made the prediction that higher levels of meritocracy would increase efficiency, but at the cost of increased inequality. In reality, however, despite aggregate macro-behavior closely resembling equilibrium predictions in the higher meritocracy regimes, meritocracy increased both efficiency *and* equality. This mismatch between theory and evidence is resolved by inspection of the underlying micro-adjustments. It turns out that a supercritical number of agents cares about ‘meritocratic fairness’, and these agents adjust their own behavior in reaction to inequalities sufficiently, so that more efficient regimes are made also more equitable.

Of course, this thesis has not settled matters definitely regarding the complex subject of human interactions. Instead, the contribution of the thesis was to propose ways and methods, as part of a novel research agenda, with the aim of integrating existing theories. I intend to pursue this agenda further in the future, starting with several directions that I find most pressing and interesting. In particular, I intend to study how game structures and information contexts influence the type of reasoning humans tend to use. In parallel, I want to explore how institutions and mechanisms influence behaviors, focussing on real-world applications and laboratory experiments. Subsequently, such findings could be prove useful to help design better institutions and mechanisms.