Marco Mantovani

Essays in

Forward Looking Behavior in Strategic Interactions

UNIVERSITÀ
DEGLI STUDI
DI MILANO

ULB
UNIVERSITÉ
LIBRE
DE BRUXELLES

UNIVERSITÀ DEGLI STUDI DI MILANO

UNIVERSITÉ LIBRE DE BRUXELLES

Department of Economics and Quantitative Methods
GSSEPS - Economics XXIV cycle

Solvay Brussels School
Ph.D. program in
Economics and Management

PH.D. THESIS

Essays in Forward Looking Behavior in Strategic Interactions

Author:
Marco MANTOVANI

Supervisors:
Dr. Antonio FILIPPIN
Prof. Georg KIRCHSTEIGER

Members of the jury:
Prof. Nagore IRIBERRI
Prof. Arno RIEDL
Dr. Antonio FILIPPIN
Prof. Georg KIRCHSTEIGER

Academic Year 2011/2012

Acknowledgements

Faculty, friends, and family members have helped me to complete this dissertation. I would like to thank them for their support, despite being already aware that the customary prose of acknowledgements will not do justice to my sincere gratitude.

I would not have made it through those four years without the support of my advisors, Antonio Filippin and Georg Kirchsteiger. I could not have asked for better role models, each inspirational, supportive, and patient.

I would advise anyone willing to prove how sound his work is, to go and talk to Antonio. His frank and no-frills attitude was sometimes hard to accept, but I soon understood that was exactly the kind of training a student, and myself in particular, desperately needed. Thanks to him, I now have a pre-presentation mantra, prescribing that I should be able to explain my work to his grandmother. Despite inevitably failing in the task, I hope I will carry with me this attitude toward research that can be communicated and the most simple and intuitive way to communicate it.

Georg is one of the smartest researcher I have ever met, and of that kind that devotes all of his acumen for the relevant aspects of an argument, never indulging in showing-off over irrelevant details. Despite this, he engaged with dedication in the fine details of my writings, and they all incredibly benefited from the goad of his variable-length sequences of red question marks. In my wandering between Milan and Brussels, and more so during my hardest times, Georg's office has become a warm shelter; a place I occasionally entered anxious and lost, and regularly left reassured and enriched.

I would have troubles assessing how much I owe to Ana Mauleon and Vincent Vannetelbosch. They invited me to Core, when I was visiting the Univerité Catholique de Louvain, then welcomed me at Cerec at the Facultés Universitaires Saint-Louis. Beyond being a great source of challenging ideas, Ana and Vincent supported and sponsored me in a moltitude of occasions, proving not only excellent people to work with, but also very good friends.

Paolo Crosetto is the co-author you have been dreaming of: as sharp as easy-going, as rigorous as fun, always efficient and never apprehensive. Beyond this, I do not count the times I have turned to him as I was stuck with a code. On top of thanking him for all this, this is a promise - sounds like a threat? - to work again together in the very near future.

The faculty of Ecares warmly welcomed me in their community, providing me with a priceless opportunity to meet outstanding researchers and great people.

To the staff and students of Core and of Cerec, I am grateful for the chance to visit and be a part of your group. Thanks to those I just bumped into once at a coffee machine or in a corridor, and showed curiosity for who I was and what I was working on. I think that really makes the difference for a student (if you are wondering about it: yes, the life of a Ph.D student is that miserable!). Gilles Grandjean did much more than that, and in both places. On top of that, he let me win every single football match we played! Daria, Paolo, let your current office-mates know I am terribly jealous of them.

Those essays would not be there without the financial support from the Italian Ministry of Eucation (doctoral scholarship and PUR grant n.), the FNRS (conv n. 2.4538.07), the Max Planck Institute of Economics - Jena and the Spanish Ministry of Science and Innovation (project ECO2009-09120), which are all gratefully acknowledged.

My engagement in a duble degree left many people with lots of paperwork to do, to my exclusive advantage. It was hard and it would have been impossible without the support of the coordinator of GSSEPS doctoral school in economics, Michele Santoni, and without the administrative staff both at the Université Libre de Bruxelles and at the University of Milan.

In my commuting between Bruxelles and Milan, I became a constant and invasive presence in the houses of many (aka the-man-with-the-ugly-white-bag). Luca, Daniel, Margherita, Marco, Donatella, Renato, you have been my family when I was far away from home, and I am wordless in front of your open-handedness. I hope I will be able to do the same for you or for somebody else, sometimes (by the way, I will be back in a couple of weeks...).

I most want to thank you, Lisa. This is about enthusiasm, inspiration, love, patience, surprise. This list could be extended at will, but it would still express very little of what I feel. In other words, the n-th absent-minded look you have tolerated, was to find a suitable thanksgiving for you, and it was useless. Agata, Mattia, you have been amazing me every single day of the last four years; thank you for ruthlessly reducing my productivity, daddy's book is finished now, and I am coming home. (Well, more or less. . .)

Contents

List of Tables

List of Figures

CHAPTER 1

INTRODUCTION

Decision-making generally features a dynamic dimension, in the sense that the outcome of the decision is delayed in time and it depends on other moves that follow, may those be taken by the same individual, by others, as well as by nature. It is indeed hard to overestimate the prevalence of this type of decision settings.

Some can be represented as individual decison problems, including investment in education or the choice of a retirement plan. Others have more apparent strategic features: trading assets, entering a market, adopting a new technology, building social relations, voting.

It is crucial for both our comprehension of the aggregate outcomes, as well as to correctly predict the effects of policies, that we have good models of how the agents take decisions in those domains.

Fully rational decison making requires, at least, that the agent is conscious of the whole problem, including all available plans of action for each agent and the consequences associated to each, and chooses optimally based on his preferences over such outcomes and his beliefs about the others. Actual behavior, in particular as recorded by controlled laboratory experiments, most often fails to conform to this benchmark and bounded rationality is now widely incorporated into economic models.

In what follows, most of the efforts will be devoted to figuring out the specific challenges that decision-makers face in dynamic environments; to showing, through controlled laboratory experiments, what consequences they have on actual behavior; and to finding ways to account for those in economic models.

The main take of those exercises is that strategic thinking is bounded in a way that is specific to the dynamic dimension of the interaction. In particular, in a sequence of moves, those that are close to the current decision are the object of deep strategic consideration, whereas distant ones are barely considered, resulting in systematic deviations from both fully rational and alternative models that do not take this into account. We report data on three experiments. Despite huge differences, behavior is surprisingly consistent among them. A majority of the population seems endowed, in all of the experiments, with some form of limited look-ahead; the number of steps most subjects can look forward to, is bounded between two and four.

Most notably, we show how we can successfully account for those bounds using simple and tractable models, that are likely to produce more accurate predictions in many applications.

Foresight and farsightedness

We offer two different ways for accounting for bounded thinking in dynamic settings.

In sequential games of perfect information, we propose an out-of-equilibrium model where backward reasoning is performed only locally, and to an extent that depends on the *foresight* of the agent. An experiment on the race game shows the explicative power of the model. Using a centipede game experiment, we also test the way in which the foresight of the agents depends on the complexity of the environment. Overall, we see this model as a proper dynamic analogue of level-k models.

In complex envirnments, such as network formation, we suggest a half-way solution between the extremes of myopic and farsighted stability. Limited *farsightedness* makes both myopically and farsightedly stable outcomes more fragile, either because, in the former case, more deviations are available, or because, in the latter case, deviations are not deterred by longer inferences. This is a desirable feature, given that, in general, both approaches tend to predict too many outcomes.

The difference between foresight and farsightedness becomes all the more clear when we move to the limited versions of both. Indeed, the clarification of this distinction is one of the contribution of these essays, as the two concepts are often confused in the literature[1].

Consider a finite extensive form alternate-move game where the agent controlling each decision node can choose whether to change state or not, and to each final state corresponds an outcome for each player. Under farsightedness, the agent controlling each decision node is considered as he was comparing the final outcomes to the one corresponding to the current state. If one state is preferred to the current one, the agent looks for a path of decisions/states, where each subject choosing along the path makes the same consideration - i.e. prefers the final state to the current one. If that is the case, this path will be a farsighted improving path. Limited farsightedness would simply limit the lenght of those paths.

[1] This difference reflects the origin of the two concepts in non-cooperative and cooperative game theory.

Under (limited) foresight, the agents consider the optimal choice of the agent choosing at the end of his horizon. That is, if he prefers changing state or not, at all the possible states he may be choosing at. Taking these actions as given, he moves backwards and solves for the next-to-last stage, and so on, figuring out a chain of best responses. Clearly, such chain may not be a farsighted improving path. At the same time, a farsighted improving path may not be a chain of best responses. For example, on a chain of best replies, an agent may not prefer the final state to the current one, but realize that the former is better than the one that will prevail, if he takes another action.

The choice between the two largely depends on the context, and in particular on how fine is the information available regarding the details of the interaction. The stability approach only requires a set of states to be defined and is then independent on the game protocol. Limited backward induction is a fully strategic model and requires all the details of the extensive form of the game.

Models of strategic thinking and limited foresight

In the last twenty years, different models have addressed the failures of fully rational, game theoretic models, generally claiming more success in explaining and predicting laboratory behavior[2]. Some of them are based on the assumption that the agents can make errors (e.g. quantal response equilibrium - QRE [McKelvey and Palfrey, 1995]); others relax the equilibrium condition in that the players best reply to some (correct) aggregate statistics of the others' strategies (e.g. cursed equilibrium [Eyster and Rabin, 2005]; others relax in different ways the assumption of common knowledge of rationality (k-rationalizabilty [Bernheim, 1984], level-k models [Costa-Gomes et al., 2003; Stahl and Wilson, 1995], cognitive hierarchy models [Camerer et al., 2004]).

Though those models were meant to target normal form games, some of them proved suited to be adapted to the extensive form. For example, Ho and Su [2013] builds a dynamic level-k model, mixing elements of level-k and cognitive hierarchy models, and apply it to the centipede game; McKelvey and Palfrey [1998] propose Agent QRE as the QRE counterpart for extensive form games, by spelling out the model in terms of behavioral strategies; the analogy-based expectation equilibrium [Jehiel, 2005], is close in spirit to a cursed equilibrium, applied to multistage games of perfect information.

There has been surprisingly little effort to capture the peculiar aspects of dynamic strategic interactions, and in particular that the depth of strategic thinking can vary throughout a game tree in a way that is not possible when considering the normal form. Probably the first attempt in this direction, and indeed close to ours, is the work of Jehiel

[2]For an excellent survey, see Crawford et al. [2012]

[1995], where a notion of limited forecast equilibrium is presented for infinitely repeated alternate-move games. Every player forecasts the future actions only within a restricted horizon, and takes the action that maximizes his average payoff within it. In equilibrium, their forecasts prove correct.

Jehiel [1998b] extends the limited forecast equilibrium to repeated simultaneous move games. The utility of the players includes a term that captures the unawareness of the players about what will come beyond their horizon of forecast. The equilibrium forecasts are correct on and off the equilibrium path, in the sense they coincide with the true distribution of actions resulting form the behavioral strategies of the players. As shown by Jehiel [2001], for intermediate discount factors, full cooperation can result, in the iterated prosoner's dilemma, as a limited forecast equilibrium outcome, whereas defection cannot. A learning justification for those approaches is provided by Jehiel [1998a].

Diasakos [2008] proposes a model of limited foresight for individual decision problems. Limited foresight emerges endogenously as a solution to a two-stage optimization problem. The agent chooses his foresight, balancing between (individual) search costs, arising from the decison problem's complexity, and the benefits of deeper reasoning.

Our attempt takes advantage of different features of those models. In particular, our model of limited foresight is close in spirit to level-k models: we develop an out-of-equilibrium model, based on a hierarchy of sophistication levels, each of which chooses the action that best replies to the immediate lower level. We share with Jehiel's contributions the idea of a limited horizon whithin which the agents make their plans. The way in which we assume the agents to project, on their horizon, the consequences that are beyond their foresight, is similar to Diasakos's.

Outline of the essays

In the first essay, we present a general out-of-equilibrium framework for strategic thinking in sequential games, Limited Backward Induction (LBI). It assumes the agents to take decisions reasoning backwards on restricted game trees, according to their (limited) foresight level. We develop a simple way in which the foresight is derived as a function of the stakes of the game and its complexity, captured by an individual cost for thinking forward; we also extend the model to apply it to infinite games.

We test for LBI using a variant of the race game, where the players take turns in adding up numbers up to a final one. The player reaching this number wins a prize. This game has special features that makes it particualrly suited for identifying forward looking behavior and backward reasoning. In particular, we can separate the predictions of LBI from other models, without making any specific assumption on the preferences of the subjects and their beliefs on their opponent's strategies. In a treatment, we add a small prize, off the equilibrium path, to identify reasoning on restricted game trees.

The results provide strong support in favor of LBI, showing that most players solve for the small prize before they do for the final one. Only a small fraction of subjects play close to equilibrium. Overall the intermediate prize keeps the subjects off the equilibrium path longer than in the base game. The results cannot be rationalized using the most popular models of strategic reasoning, let alone equilibrium analysis. Remarkably, the players effort in making a decision, as recorded by the time it takes before acting, is very low until the game gets close to one of the prizes; it then peaks steeply and then decrease as it becomes clear who is going to get the prize. Most of the player are consistent with a level of foresight of two, three or four steps.

In the second essay, we test a specific implication of LBI: that the foresight of the players - i.e. the number of steps of backward induction they are able to perform - is decreasing in the complexity of the environment. We present the results of a novel experiment, using a centipede game, where we manipulate complexity by reducing the availability of information regarding the payoffs. We run three treatments featuring the same game, but where the payoffs are represented in different ways.

We show that reduced availability of information is sufficient to shift the distribution of take-nodes further from the equilibrium prediction, and similar results are obtained in a treatment where reduced availability of information is combined with an attempt to elicit preferences for reciprocity, through the presentation of the centipede as a repeated trust game.

Behavior in the centipede game has been explained either by appealing to failures of backward induction or by calling for preferences that induce equilibria consistent with observed behavior. Our results could be interpreted as cognitive limitations being more effective than preferences in determining (shifts in) behavior in our experimental centipede.

Furthermore our results are at odds with the recent ones in Cox and James [2012], suggesting caution in generalizing their results. Reducing the availability of information may hamper backward induction or induce myopic behavior, depending on the strategic environment. Most notably, both effects can be rationalized within the framework of limited backward induction.

In a nutshell, as complexity increases, the agents respond by reducing the number of steps over which they perform strategic reasoning, but still incorporate the efficiency gains that are achievable at distant nodes by projecting those payoffs on their foresight bound. Beyond a certain threshold, however, on top of being able of very limited backward reasoning, they stop considering distant payoffs, resulting in myopic behavior (“take the money and run”).

In the third essay we change completely both the environment and the framework. We investigate network formation and test for the stability notions that are there applied. Given the prevalence of interactions through social networks and the extraordinary growth of the network literature in the recent years, we see the empirical foundation and development of sensible models of network formation as mostly needed. Moreover, network formation is typically a complex environment, where bounded reasoning is most likely to bind actual behavior.

Pairwise stability [Jackson and Wolinsky, 1996] is the standard stability concept in network formation. It assumes myopic behavior of the agents in the sense that they do not forecast how others might react to their actions. Assuming that agents are perfectly farsighted, related stability concepts have been proposed. We design a simple network formation experiment to test these theories.

Our results reject both of those extreme stability notions. In particular, we show that the behavioral models thay assume are both untenable. The agents are, instead, consistent with a form of limited farsightedness. Both myopically and farsightedly stable networks are found to be fragile to farsighted deviations of short lenght (two, three steps). The selection among pairwise stable networks seems to be driven by their resilience to those deviations. Indeed, to the best of our knowledge, no other theory can account for the variance across treatments in the outcomes that we observe. Beyond this, we find support for this interpretation in the analysis of individual behavior. Low level of farsightedness appear relevant to explain the choices of our subjects.

CHAPTER 2

LIMITED BACKWARD INDUCTION

Abstract

We present a general out-of-equilibrium framework for strategic thinking in sequential games. It assumes the agents to take decisions on restricted game trees, according to their (limited) foresight level, following backward induction. Therefore we talk of limited backward induction (LBI).

We test for LBI using a variant of the race game. Our design allows to identify restricted game trees and backward reasoning, thus properly disentangling LBI behavior.

The results provide strong support in favor of LBI. Most players solve intermediate tasks - i.e. restricted games - without reasoning on the terminal histories. Only a small fraction of subjects play close to equilibrium, and (slow) convergence toward it appears, though only in the base game. An intermediate task keeps the subjects off the equilibrium path longer than in the base game. The results cannot be rationalized using the most popular models of strategic reasoning, let alone equilibrium analysis.

JEL classification: D03, C51, C72, C91

Keywords: Behavioral game theory, bounded rationality, race game, sequential games, strategic thinking, level-k.

2.1 Introduction

How do you figure out your moves in a chess game? We would say that most people think of what the other is going to do next; some think of their next move as well, and maybe of the opponent following choice. Deep consideration of further stages characterizes chess lovers and professionals[1]. This behavior is just backward induction performed on a limited number of stages. In a nutshell, that represents what we call limited backward induction (LBI). Going back to chess, it is reasonable that we will look further ahead as our king (or the opponent's) is plainly menaced. At the same time, we will able to push our reasoning deeper as the action space shrinks throughout the game, or in games where the action space is more limited, such as tic-tac-toe. This paper presents a model that catches such features of strategic reasoning in sequential games and presents a novel experiment meant to test for it.

The strategic ability of human beings, as recorded form the experimental literature, seems more limited than assumed in game theory[2]. Backward induction is no exception (e.g. Binmore et al. [2002]) and has been long criticized as well on theoretical ground[3]. In the last twenty years, different models have addressed the issue, generally claiming more success in the lab than competing fully-rational models. An excellent survey of the subject can be found in Crawford et al. [2012].

Despite those models were meant to target normal form games, some of them proved suited to be adapted to extensive form. As an example, Ho and Su [2013] builds a dynamic level-k model to be applied on the centipede game[4], while McKelvey and Palfrey [1998] proposed the agent quantal response equilibrium (AQRE) as the QRE counterpart for sequential games. An attempt closer to ours is that of Jehiel [1995]; his limited forecast equilibrium is close in spirit to our approach, although it sticks to equilibrium analysis. An independent attemp, similar to ours, is being carried out at the moment by Roomets [2010]. Despite many similarities, in his paper, the level of foresight is exogenously given and not endogenous to the game, as in ours; the way in which intermediate payoffs are derived is largely unspecified; most notably, in Roomets [2010] there is no experimental test of the model.

Beyond this, no model of strategic thinking attempted to address the specific aspects of dynamic strategic environments, which is the goal of the present paper. In carrying out the task, we retain the intuition underlying backward induction, but we limit the number

[1] The literature on chess heuristics is vast and spans from artificial intelligence to psychology. See for example Reynolds [1982] and De Groot et al. [1996]

[2] Excellent surveys can be found in Kagel and Roth [1995], Smith [1994] and Selten [1998]

[3] Aumann and Binmore gave life to a famous crosstalk on the subject. See Aumann [1996] and Binmore [1996]

[4] Kawagoe and Takizawa [2012] study classic results on the centipede game with a similar model.

of stages on which it is performed. As such, LBI is best suited for studying multistage games with observable actions, characterized by perfect information.

Under LBI, a player faces a reduced game tree, called the limited-foresight game (LF-game), that encompasses only the stages of the game that are closer to the current decision node. For each (pseudo-)terminal history of the LF-game, intermediate values are determined, based on the final payoffs that are consistent with it; finally, actions consistent with a subgame perfect equilibrium (SPE) in the LF-game are taken. The higher the level of foresight, the more stages are included in the LF-game so that, in the limit, this coincides with the whole game and the actions taken are consistent with subgame perfection.

We use a variant of the Race game (also known as the Game of 21) to identify LBI. In this simple game players alternate to choose numbers within a range; those are summed up, until a certain target number is reached. The player who reaches it wins a prize, the other loses. One of the players has an advantage at the beginning of the race, and can secure the victory of the prize. This possibility transfers to the other player in case of error. Each player has a family of dominant strategies, which can be identified by backward induction and should be played in any SPE. Level-*k* players [Costa-Gomes et al., 2003] should play consistent with equilibrium.

Previous results [Dufwenberg et al., 2010; Gneezy et al., 2010; Levitt et al., 2011] show little compliance with equilibrium predictions. The subjects seem to discover the solution as they play and find it hard to substitute a subgame with its outcome, even after gaining experience. We replicate those results in a base treatment and show how they are due to a LBI type of reasoning, using a second treatment. The treatment variable is the presence of a small prize on the path to the final target. This manipulation is suited to identify reasoning on a restricted game tree (the LF-game). We also introduce, in both treatments, the (incentivized) possibility to claim victory of any prize at any time in the game; this helps us tracking which prize the subjects are targeting beyond their observed actions.

Our results are stark. Most subjects solve for the trap prize before they do for the final one. This is consistently witnessed by both their actions and their claims, as their claims of the trap prize anticipate those for the final one and subjects reach the path to the former earlier than that to the latter. The timing of their decisions shows how reasoning efforts emerge only when the game approaches one of the prizes - i.e. when this enters the LF-game. In the presence of the small prize, the players stay off the equilibrium path longer that when the small prize is absent. On aggregate, we provide strong support in favor of a LBI type of reasoning. The majority of the population proves able to run no more than two or three stages of backward induction in our race game, whilst only a small fraction of the subjects play consistently with equilibrium.

We propose a formal model of LBI. The foresight of an agent is determined after considering how relevant the stakes are - i.e. how beneficial strategic reasoning can be - and how complex the game is - i.e. how costly strategic reasoning can be. To build the LF-game, he then needs to assign payoffs to intermediate nodes. There are classes of games where this exercise has an intuitive solution, and most notably, in games where stage-payoffs are gained. Under this condition, LBI can tackle infinite games and encompasses perfect myopia as a special case.

More generally, we specify a class of procedures that extract the intermediate values for a pseudo-terminal history of the LF-game, by projecting the payoffs that arise beyond the foresight of the agent, and are consistent with that pseudo-terminal history being reached. In the chess example, a valuation of an intermediate node would depend mostly on the value assigned to the single pieces, this being associated to the likelihood of winning after losing that piece[5].

We find our attempt to have the same flavor of level-*k* models. Indeed, LBI types depend on the depth of their strategic considerations and each type best responds to the type which is one step lower in the hierarchy. The chain of best replies is anchored to the behavior of a non-strategic type, which, in our case, is represented by the agent choosing before the foresight bound. Despite those similarities, LBI predictions will generally diverge from the level-*k* ones in many contexts, as made clear by our design. Both level-*k* and LBI constitute out-of-equilibrium models of behavior and should be understood to capture initial responses to a game.

The paper proceeds as follows: section 2.2 provides an informal description of LBI and draw some didactic examples ; section 2.3 offers a formal model of LBIsection; 2.4 introduces the race game and the relative experimental literature; section 2.5 presents the experimental design and procedures; section 2.6 shows results and section 2.7 concludes.

2.2 A sketch of Limited Backward Induction

Consider[6] the four-stage game in Figure 2.1, where each outcome $\mathbf{a}, ..., \mathbf{p}$ is a vector in $\mathbb{R}^2$, identifying von Neumann-Morgenstern utilities for each player. According to backward induction, player 1 knows what player 2 is choosing after every history in stage four. He can substitute the preferred outcomes to the decision nodes of stage four and roll over this reasoning to determine what actions are chosen by himself in stage three and by

[5] An appropriate evaluation will depend on the position of the pieces as well, however most chess manuals report standardized pieces' values [Capablanca, 2006, e.g.] and those are regularly used by computerized chess players [Levy and Newborn, 1991].

[6] See section 2.3 for a formal tractation.

player 2 in stage two. Finally, he picks his best reply to those profiles in stage one, which results in a SPE of the game.

Suppose player 1 is not able to run backward induction from the terminal histories of the game, because of limited foresight. The dashed line in Figure 2.1 represents his foresight bound, implying he best replies to what he believes the next player is choosing, without conditioning on the following moves. In the terminology of LBI, his level of foresight is two: the LF-game he can solve includes two stages. We label this type F_2.

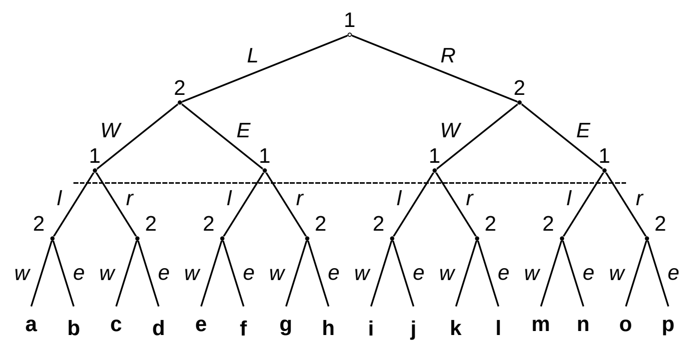

Figure 2.1: A four-stages sequential game with a foresight bound

In order to figure out what his opponent is going to do in stage 2, he needs some intermediate valuations for the pseudo-terminal histories of the LF-game. Provided that he knows the final outcomes of the game, he uses this information to retrieve the intermediate valuations. Each payoff of the LF-game will be derived as a projection of the payoffs of the complete game that are consistent with each pseudo-terminal history. As an example, the payoffs considered in the LF-game at the pseudo-terminal history (L, W) will be a function of $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}$, those after (L, E) of $\mathbf{e}, \mathbf{f}, \mathbf{g}, \mathbf{h}$, and so on. If player 1 was of type F_3 - i.e. he had a foresight of three stages - the dashed line in figure 2.1 would move one stage downwards. The LF-game would be larger and its payoffs would be derived from smaller sets of final outcomes.

We will not propose a one-fits-all solution to the problem of how to project terminal payoffs on the terminal histories of the game. In most applications, the obvious choices will be simple functions like the average, or the median point in the range of available payoffs, which can be considered as baseline hypothesis. The degrees of freedom that are left to the model, are similar to those that arise in the definition of the L_0 player in level-k models, or with the choice of the noise element in QRE. Contrary to those, however, in LBI the context will generally provide an intuition or a theoretical guide to which function to choose, as will be clear from the examples of section 2.2.1. Most importantly, the proper definition of the projection function will be completely irrelevant in our experimental test. We will discuss further the issue of intermediate payoffs in the formal set up (see section 2.3.3).

To fix ideas, assume player has a foresight of k stages; call the decision rule of this player F_k. By construction, the LF-game he is facing in the first stage includes k stages. Reasoning backward, his implicit belief is that the player choosing at the last stage of the LF-game is of level F_1. This player, who serves as anchor of the LBI reasoning, is non-strategic, in the sense that he does not perform any strategic reasoning. The player controlling the next-to-last decision nodes is believed to act as level F_2, and so on. Thus, player one will best responds to the following player, assumed to choose as F_{k-1}, who best responds to $F_{k-2}, \ldots,$ who best responds to F_1.

Provided that the foresight bounds of players choosing sequentially do not coincide, beliefs about the next players' chosen actions will generally prove incorrect. The actions chosen by one single player need not be consistent one with the other. Those observations clarify that LBI is an out-of-equilibrium model of the initial responses to a game, by untrained subjects.

The sketch of LBI we have given is perfectly sufficient to understand the experimental part of the paper. A reader not interested in the examples and in the formal set up may then want to jump to section 2.4, which is self-contained.

2.2.1 Examples

We here sketch how LBI can be applied to a couple of classical examples. We briefly draw predicted behavior and compare it to the experimental evidence.

Centipede game The centipede game (see Figure 2.2) has long been a major workhorse for investigations of backward induction, as for its simple sequential structure[7]. It is a two-player, finite sequential game in which the subjects alternate choosing whether to end the game ("take") or to pass to the other player ("pass"). The payoff from taking in the current decision node is greater than that received in case the other player takes in the next one, but less than the payoff earned if the other player were to pass as well. The player making the final choice is paid more from taking than from passing, and

[7]With respect to the motivations underlying the players' choices, the centipede game proves much less simpler than it may seem at a first look. We do not think it is a good test of backward induction, and for this reason we chose a different game for our experiment.

would therefore be expected to take. Iterating this argument, backward induction leads to the unique subgame perfect equilibrium: the game is stopped at the first decision node, implying a huge efficiency loss.

Experimental evidence has shown little compliance with SPE in the laboratory[8]. The typical results feature a bell-shaped distribution of endnodes. Figure 2.3 depicts the results from the six-leg centipede in the seminal paper of McKelvey and Palfrey [1992]. Beyond the failure of SPE, robust findings show that longer games result, ceteris paribus, in higher endnodes.

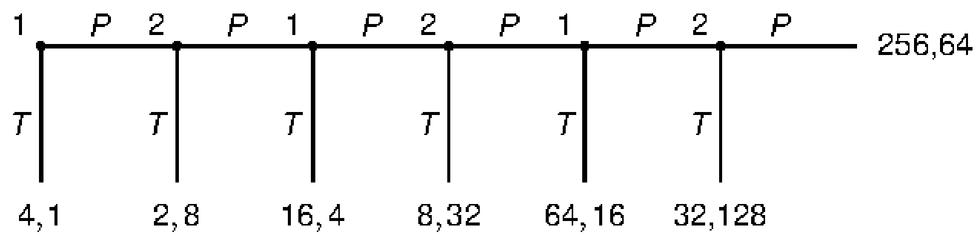

Figure 2.2: The six-legs centipede game in McKelvey and Palfrey [1992]

Under LBI, a subject playing the centipede in Figure 2.2 does not take until the terminal histories are included in his LF-game, as he expects the opponent (and himself) to pass (P) in the following decision nodes. Any function of the payoffs that follow some decision node, satisfying some basic axioms[9], will give a higher value to P than to T at the foresight bound[10]. As soon as terminal histories are included in the LF-game of a player, he takes. This happens later in the game, ceteris paribus, if the centipede features more decision nodes.

Note that, for an agent to choose T, it is not necessary, in general, that the terminal histories of the centipede game are included in the LF-game. The point, at which a player with a certain foresight takes, crucially depends on the progression of the payoffs[11]. A population of F_2 and F_3 (and, possibly F_4) describes the initial behavior of the majority of the experimental subjects in most standard experimental centipede games we are aware of [12].

Sequential bargaining In sequential bargaining (see Figure 2.4), two players must agree on the division of a cake that shrinks every time they do not find an agreement (generally

[8]See, for example: Levitt et al. [2011]; McKelvey and Palfrey [1992]; Palacios-Huerta and Volij [2009]; Zauner [1999]

[9]See section 2.3.3 for a discussion of those.

[10]For example, the average of the n payoffs that come after a payoff of x is in the form $(1/2 + 4 + 2 + 16...)x/n > x$.

[11]See Crosetto and Mantovani [2012] for a discussion of the issue.

[12]See Kawagoe and Takizawa [2012] for a presentation of initial response results in the centipede game.

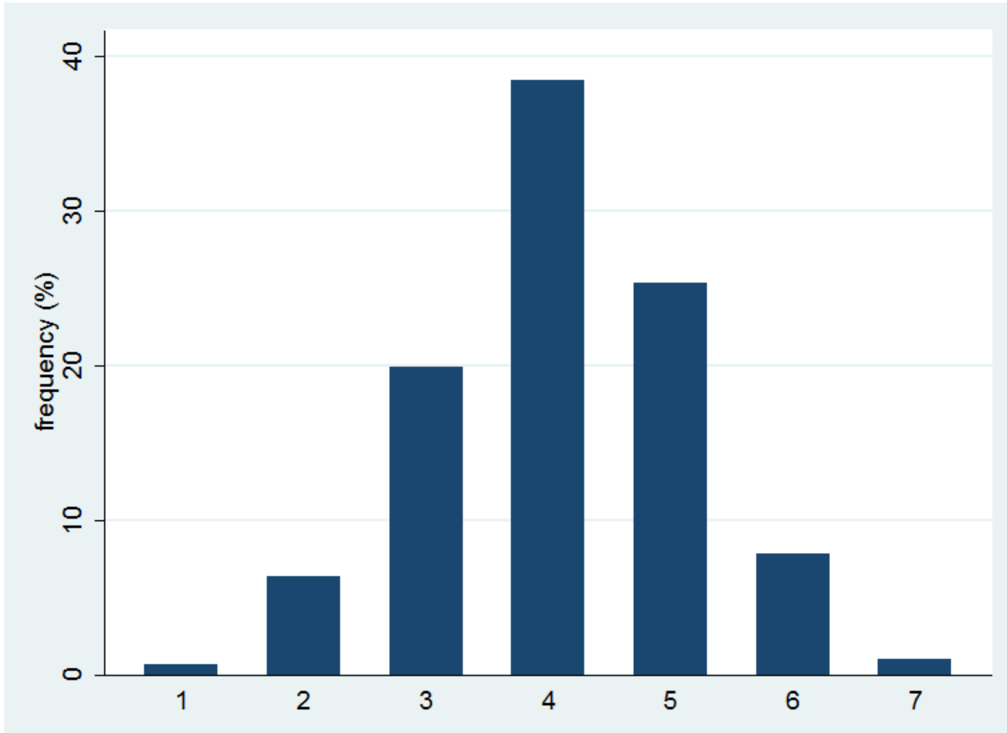

Figure 2.3: Distribution of endnodes in McKelvey and Palfrey [1992]

intended as the effect of a discount factor). The players alternate in making an offer, with the other player either accepting, in which case the game ends and the proposed split is implemented, or rejecting, in which case the cake is reduced and they move to a new stage, switching roles. In finite bargaining, the last round is an ultimatum game. Its solution provides the minimal accepted offer in the previous stage. This reasoning can be iterated backwards up to the first stage. In the unique SPE, in every stage, the proposer submits the minimal accepted offer, the responder accepts any offer weakly higher than the minimal acceptable one, so that the first offer is accepted and the game stops.

Broadly speaking, the existing experimental evidence[13] shows that offers are, on average, more generous than in the SPE, and those offers that are close to equilibrium are often rejected. In general, and contrary to the theory, the first offers are relatively stable, independently of the number of bargaining stages[14].

Consider the game in Johnson et al. [2002], in Figure 2.4; the cake is initially worth \$5 and it is halved at every new round, up to the third. They find a first round average offer

[13] See, for example, Binmore et al. [2002]; Bolton [1991]; Harrison and McCabe [1996]; Johnson et al. [2002]; Ochs and Roth [1989]

[14] The number of stages is the major determinant of who retains more bargaining power, by determining the roles in the last round, which is an ultimatum game.

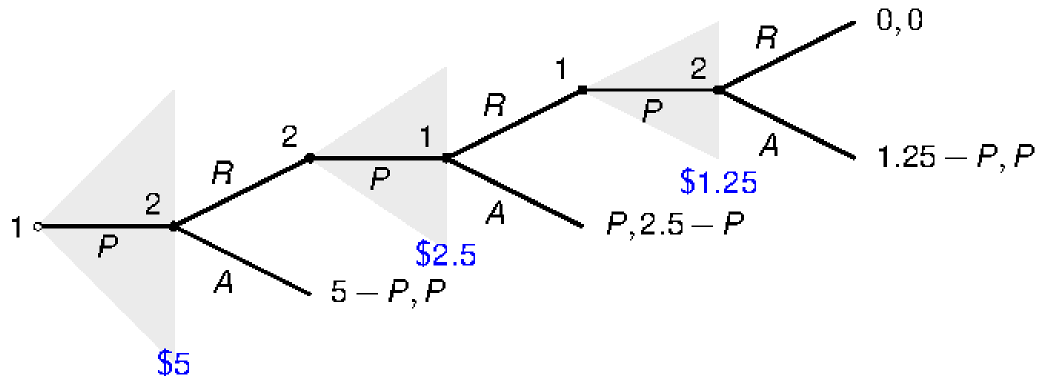

Figure 2.4: The sequential bargaining game in Johnson et al. [2002]

of around \$2.10, against a SPE of \$1.26; half of the offers below \$2 were rejected. They also report on a treatment where other-regarding preferences are switched off, using robots as opponents: the average first-round offer declines to \$1.84.

Under LBI, a first mover that does not consider the last round of bargaining within his LF-game understates the minimal accepted offer in the second round of bargaining, which in turns implies he make a better offer than in the SPE. Given that offers are a compact set, we assume the players to use the average value of the set, to compute the payoffs of the LF-gameThis would be reasonable also in the case of a finite set of possible offers: the simple average of the terminal payoffs implies an over-representations of all the zeros that follow a rejection. See section 2.3 for a discussion.. A first mover whose LF-game includes only the next round of bargaining, assumes intermediate values following a rejection in round two of \$0.625 each[15] Reasoning backwards, he expects, in the second round, player 2 to offer $P \cong \$0.63$, and keep \$1.83 for himself. The offer that best replies to this belief is $P \cong \$1.87$ in the first round. This prediction perfectly matches the average offer in the robot treatment of Johnson et al. [2002][16].

Moreover, given a distribution of LBI types, the length of the game does not affect the behavior of all the first movers whose foresight does not reach the terminal histories.

[15].

[16]The fact that in human vs. human treatments offers are even higher, suggests that other-regarding preference play a role beyond cognitive limitations in this game.

2.3 A formalization of LBI

2.3.1 General notation

Take a set $\mathcal{I}$ of players, $\#\mathcal{I} = I$ finite, playing a multistage game of perfect information with $T+1$ stages ($t = 0,1,...,T$), with T finite[17]. As usual, a history at the beginning of stage t is a collection of actions in the form $h^t = (a^0, a^1, ..., a^{t-1})$. Let $A_i(h^t)$ be the finite set of feasible actions for player i when history is h^t and H^t the set of all histories at stage t, with H^{T+1} the set of terminal histories. Then: $A_i(H^t) = \bigcup_{h^t \in H^t} A_i(h^t)$. Recall that perfect information implies that at every stage t and history h^t for exactly one player it holds that $\#A_i(h^t) > 1$. Function $l : H \setminus H^{T+1} \to \mathcal{I}$ is the mapping of who moves at each non-terminal history. Utilities are in the form $u_i : H^{T+1} \to \mathbb{R}$.

This defines a game $G = \left(\mathcal{I}, H, l, (u_i)_{i \in \mathcal{I}}\right)$. We introduce a property that will be useful in what follows.

Definition 1. *Payoffs are said to satisfy Additive Separability (AS) if:*

$$u_i = \sum_{t=0}^{T} \pi_i(a^t),$$

where $\pi_i(a^t)$ is the single stage payoff resulting from actions consistent with h^{T+1}.

In other words, a game displays additive separability of payoffs if utilities can be represented as the sum of the payoffs gained along the game[18].

For any history h^t, let G_{h^t} be the game that starts at h^t; $H^{T+1}|_{h^t} : \{h^{T+1} = (h^t, a^{t+1}, ..., a^T)\}$ will be the corresponding set of terminal histories. With a slight abuse of notation let $\pi_i(h^{\bar{t}}) = \sum_{t=0}^{\bar{t}} \pi_i(a^t)$, and,

$$u_i^{h^t} = \{u_i(h^{T+1}) \text{ s.t. } h^{T+1} \in H^{T+1}|_{h^t}\}$$

and, for $t^2 \geq t^1$,

$$\pi_i^{H^{t^2}|_{h^{t^1}}} = \{\pi_i(h^{t^2}) \text{ s.t. } h^{t^2} \in H^{t^2}|_{h^{t^1}}\}$$

That is, respectively, the set of terminal utilities that are viable after h^t and the set of cumulative payoffs at stage t^2, that are viable after h^{t^1}.

[17] We present the finite case to keep things simple. Our approach easily fits an infinite number of stages, provided that the players are meant to represent it as a finite game.

[18] For simplicity, we avoid talking about discounting. Nothing would change in the model if we included it.

2.3.2 Sight and foresight

We turn to the players' *sight* and *foresight*. Given any $L \in \mathbb{N}^+$, an agent with *sight* S_L at history h^t sees all histories in $\{H^l|_{h^t}\}_{l=t+1}^{t+L}$ and the corresponding payoffs. His *sight* captures his understanding of the game, with no relation to strategic considerations. Note that his *sight* reaches the terminal histories only if $t+L \geq T+1$; when this is not the case, we will talk about cognitive-terminal histories refering to the set $H^{t+L}|_{h^t}$; if $L \geq T+1$ the player sees all relevant information about the game from the very first stage. With a slight abuse of notation we will interpret $t+L$ as $\min\{t+L, T+1\}$, to avoid specifying such minimum every time.

If the payoffs satisfy additive separability, the *sight* of a player always provides information about utilities to players, the same is not true if this property is not satisfied. More precisely, given AS, a *sight* S_L identifies a game, and allows the player to build the LF-game, according to the description that follows. Absent AS, only a game form is specified, with no possibility to build the LF-game, except for the case when $t+L \geq T+1$.

To clarify the point, imagine you are playing chess knowing all the rules, except that the first who checkmates wins the match. There is no way of playing meaningfully without information about the payoffs. Suppose now you are playing tennis, knowing all the rules except for the match-winning rule[19]. You can still play meaningfully and have fun since you know how to score and to win games.

The usual way in which a game is presented, absent AS, implies, *ipso facto*, that $L \geq T$ and thus the sight of players reaches the terminal histories [20]. In particular, in economic experiments, the subjects are generally informed of the final payoffs, and the experimenter makes sure that they understood, at least, that point. This is not granted in real life games, for the cases where the payoffs arise stage after stage. In what follows we assume that the *sight* of the players always reaches the terminal histories when AS is not satisfied.

The *sight* of a player is exogenously given. His *foresight* represents the depth of his strategic thinking. It is derived as a function of the stakes of the game and its complexity using the information provided by his *sight*. That is both the stakes $D_{iL}^{h^t}$, and the complexity, $C_{iL}^{h^t}$, are assessed according to S_L. We denote the level of foresight of a player with:

$$K = \left\lfloor f(D_{iL}^{h^t}, C_{iL}^{h^t}) \right\rfloor$$

and we denote with F_K a player with foresight of K steps that uses LBI. We will as-

[19] For example, you do not know how many games are there in a set, and how many sets are needed to win the match

[20] See Crosetto and Mantovani [2012] for a discussion and an experiment on representation effects in the centipede game

sume $f_{C^{h^t}_{iL}} < 0$, $f_{D^{h^t}_{iL}} > 0$, $f_{D^{h^t}_{iL}D^{h^t}_{iL}} < 0$, $f_{C^{h^t}_{iL}C^{h^t}_{iL}} > 0$; in other words, the level of foresight is increasing concave in the stakes and decreasing convex in complexity. As $K \leq L$, the level of foresight is always finite.

A player with *sight* S_L engages in more steps of strategic reasoning the higher the stakes of the game, and the lower its complexity. Though intuitively appealing, it is indeed challenging to formally define stakes and complexity, not to mention any specific functional form. Consistent with the purpose of proposing a general framework, we define stakes as a function of the payoffs within the sight of the agent, assuming weak monotonicity in the payoffs and in their variance, and complexity as function of the number of cognitive-terminal histories in his sight, plus an individual cost parameter, increasing in both terms. That is

$$D^{h^t}_{iL} = d(\pi_i^{H^{t+L}|_{h^t}})$$

and

$$C^{h^t}_{iL} = e(c_i, \#H^{t+L}|_{h^t})$$

Absent AS, the above definitions reduce to $D^{h^t}_i = d(u_{h^t})$ and $C^{h^t}_i = e(c_i, \#H^{T+1}|_{h^t})$. F_K determines the set of the histories of the LF-game, defined as

$$H_{h^t,K} = \{H^{t+k}|_{h^t}\}_{k=0}^{K}$$

The pseudo-terminal histories of the LF game are then $H^{t+K}|_{h^t}$.

A natural extension of this definition of *foresight* is to take into account how stakes and complexity vary over the action space of the subject, letting F_K vary over the set $H^{t+1}|_{h^t}$. In general, simplicity reasons suggest not to consider such a case, but one should keep in mind this is possible, for the cases where those differences are likely to be relevant.

2.3.3 Intermediate payoffs and the LF-game

Let H^t_i be the set of histories controlled by player i, $H^t_i = \{h^t \text{ s.t. } \#A_i(h^t) > 0\}$. At each node in this set, player i, with sight S_L, is characterized by a level of foresight, F_K. This level specifies, at every history, a restricted game form $(\mathcal{I}, H_{h^t,K}, l)$. To complete the LF-game, utilities must be defined over the pseudo-terminal histories $H^{t+K}|_{h^t}$. Let those be a function in the form:

$$v_i : \begin{cases} \{\pi_i^{H^{t+L}|_{h^t}}\}_{h^t \in H \setminus H^{T+1}} \to \mathbb{R} & \text{if AS satisfied} \\ \{u_i^{h^t}\}_{h^t \in H \setminus H^{T+1}} \to \mathbb{R} & \text{otherwise} \end{cases} \tag{2.1}$$

In other words, the intermediate utilities are functions of the set of consequences that are consistent with the pseudo-terminal history they refer to. Those consequences are the terminal utilities, if AS is not satisfied, but can be cognitive-terminal utilities if AS is satisfied. In the latter case, the players base their evaluation of pseudo-terminal histories using the payoffs that are viable from those histories on, and that fall within their sight S_L. In principle, different subjects may have different mappings and the same subject may use different mappings for different pseudo-terminal histories. Despite this, we will not consider these possibilities in the general case.

Let $v \in \mathbb{R}^n$ be a vector of the intermediate utility for each player; the LF-game at history h^t, controlled by an agent with foresight F_K, is:

$$G_{h^t,K} = (\mathcal{I}, H_{h^t,K}, l, v)$$

In applications one wants to make intermediate utilities operational. Up to now, we only constrained them to depend only on payoffs consistent with the pseudo-terminal history they refer to. We put forward a set of properties, indeed quite standard, that must be satisfied by any specification of the intermediate utilities[21]. We then show some examples, having intuitive applications, that satisfy those properties.

Definition 2. *Given any two histories $h, h' \in H^t$ such that, $min\{u_i^h\} \geq max\{u_i^{h'}\}$, intermediate utilities satisfy **dominance** iff $v_i(u(h)) \geq v_i(u(h'))$.*

Dominance states that if all the utilities viable after a certain history are higher than those after another history, then the former should be preferred to the latter. A stronger version is the following, postulating that a history is preferred over another if each utility, that is viable after the former, beats the corresponding one, that is viable after the latter, after ordering both sets in the same way. Given u_i^h and $u_i^{h'}$ with the same cardinality, let $\Delta u_i^{h,h'}$ be the set containing the pairwise differences between the ordered elements of the two sets.

Definition 3. *Given any two histories $h, h' \in H^t$ such that $\#H^{T+1}|_h = \#H^{T+1}|_{h'}$ and $\Delta u_i^{h,h'} > 0$, intermediate payoffs satisfy **stage monotonicity** iff $v_i(u(h)) \geq v_i(u(h'))$.*

Extension monotonicity states that, if the set of utilities viable after some history, is the same as those after another history, plus some utilities that are less valued (more valued) than the previous ones, than the latter (former) history is preferred to the former (latter).

[21] The notation is suited for cases where the sight of the agents encopasses the terminal histories. The reader can easily check that it transfers to the other cases, though at the cost of becoming more cumbersome

Definition 4. *Given any two histories $h, h' \in H^t$ such that $u_i^h \subseteq u_i^{h'}$, intermediate payoffs satisfy **extension monotonicity** iff:*

$$i.\ max(u_i^{h'} \setminus u_i^h) \leq min(u_i^h) \Rightarrow v_i(u(h)) \geq v_i(u(h'))$$

$$ii.\ min(u_i^{h'} \setminus u_i^h) \geq max(u_i^h) \Rightarrow v_i(u(h')) \geq v_i(u(h))$$

Examples. *The following intermediate payoffs satisfy definitions 2-4:*
Simple average: $v_i(u_i^h) = \frac{\sum_{u_i \in u_i^h} u_i}{\# u_i^h}$
Range average: $v_i(u_i^h) = \frac{max(u_i^h) - min(u_i^h)}{2}$
Random choices: $v_i(u_i^h) = E_r(u_i^h)$*, where* E_r *assigns equal probability to each action in the action set at every node before the pseudo-terminal history.*
Maximum(Minimum): $v_i(u_i^h) = max(u_i^h) \quad (min(u_i^h))$

As we deny proposing a one-fits-all solution, each operational solution needs to be justified within a context. In general, we favor the use of the most intuitive solutions, such as the simple average or the range average. For example, in sequential bargaining, if the set of possible offers if finite, the simple average counts one zero for each possible offer (in case the offer is rejected). This over-representation of zeros may discourage the use of the simple average.

On the LF-game the agents take decisions following backward induction. The agent choosing at h^t starts by finding the optimal actions of the subjects that controls the histories $H^{t+K-1}|_{h^t}$. He then moves to $H^{t+K-2}|_{h^t}$, taking the optimal actions in the following stage as given. And so on, until he reaches h^t. is actions are then consistent with a subgame perfect equilibrium in the LF-game. The following statements are true.

Proposition 1. i. *For any F_K, LBI always prescribes at least one action.*

ii. *If $t + K \geq T + 1$ the actions prescribed by LBI are all and only those that are part of a subgame perfect equilibrium strategy of the game.*

iii. *For $K \to \infty$, LBI prescribes all and only the actions that are consistent with a subgame perfect equilibrium of the game.*

The proofs are self-evident and are omitted. Given a population featuring a certain distribution of levels of foresight, the actions consistent with LBI become closer to equilibrium as we approach the end of the game, matching a common experimental finding, perfectly depicted by our experiment. As the agents gain experience, their level of foresight can increase and their actions will converge toward SPE[22].

[22] This does not imply that the agents will improve their level of foresight and converge to SPE in all situations.

2.3.4 Strategic reasoning in LBI and in level-k models

There are many similarities between LBI and level-*k* models. Both are first-response models. Their scope is to explain the out-of-equilibrium behavior of untrained subjects and to predict to which equilibrium the subjects are likely to converge, in case they do.

LBI and level-*k* models are based on a hierarchy of types and the decision rule is such that each type best replies to the next lower one. Finally, this chain of best replies is anchored to a non-strategic type.

However, LBI distinguishes from level-*k* in many aspects. A level-*k* agent knows the actions chosen by the agents with level between 0 and $k-1$, and best replies to the $k-1$ action. When applied to dynamic contexts [Ho and Su, 2013; Kawagoe and Takizawa, 2012] the players specify an action plan for the whole game, consistently with their belief about the other's type.

Under LBI, the agent of type F_K applies backward induction on a restricted game tree. In doing so, he acts as if he imposed decreasing levels of foresight on the agents that control the following nodes: the next player to be active is assumed to act as a F_{K-1} type, the following as a F_{K-2} type. The agent controlling a node at $K-1$ stages of distance from the current one perform no strategic reasoning, in the sense that he does not consider strategically the other choices.

On the one hand, this implies that, within the LF-game, the player assumes the others' beliefs to be consistent with his owns. A player of level-*k* assumes the opponents are of level $k-1$, and they believe the others are of level $k-2$. On the other hand, the agents are not aware of their own type under LBI, in the sense that they impose a lower level of foresight on themselves, when considering their choices at future nodes. This entails that the actions of the same player need not be consistent throughout the game. New information is taken into account as the player explores new portions of the game tree, and the action plan changes accordingly.

The different features of LBI with respect to level-*k* models make the two models best suited for different situations. Specifically, LBI is meant to address those games where the dynamic aspects are salient and the foresight of the players is likely to bind their strategic reasoning. Level-*k* is not suited for those settings, as already clear from our experiment.

2.4 The Race Game experiment

In all of the previous examples, it is hard to disentangle the impact of limited cognition from other aspects of decision making, such as other-regarding preferences, reciprocity or efficiency considerations. To test for LBI, we design a novel experiment featuring a race game, a sequential, perfect-information game. As it will be clear, this game has nice features for our scope.

First, it is a zero-sum game, with only two possible outcomes; this implies that we can overlook all preference-related aspects of the decision problem. Second, each player has a set of weakly dominant strategies, which mitigates the problem of not observing the players' beliefs about others' strategies. Third, the set of outcomes that is consistent with some history is identical across all histories within the same stage; this means that the function that projects the terminal payoffs on the pseudo-terminal histories is completely irrelevant. Those features make it possible to observe the pure effect of limited strategic ability, and, in particular, of those limits that are specific to dynamic strategic settings.

2.4.1 Base Game

In the race game, two players take turns choosing a number of steps, an integer within a range $1,...,k$. The steps chosen are summed up: assuming the players start at position one, the position at some stage s of the game is given by total number of steps taken in stages $1,...,s$ plus one. When a player reaches the target number M, he wins a prize, the other loses (and gets nothing).

Any race game can be solved backwards: a player easily wins from positions $M - k,...,M-1$; thus a player choosing at $M-(k+1)$ is meant to lose. This position can be reached from $M-(2k+1),...,M-(k+2)$, meaning that a player choosing at $M - 2(k+1)$ is meant to lose. This reasoning can be iterated back to position one, unveiling a sequence of losing positions[23]. An agent that is able to reach with his choice any of these positions, is able to secure the victory of the game, by reaching the subsequent losing positions in his following decision nodes.

Formally, the set of losing position is $\mathcal{L} = \{t \in \mathcal{T} : t = M - (ik+i)$, for some $i = 1,2,...\}$, where $\mathcal{T}$ is the ordered set of all positions. The set of winning position is then: $\mathcal{W} = \mathcal{T} \setminus \mathcal{L}$. The game displays a set of (weakly) dominant strategies for both players, prescribing to reach the closest losing position whenever possible (and choose whatever number at losing positions). If $1 \in \mathcal{L}$, player 2 has an advantage in the sense that he wins the game according to any of his dominant strategy. If $1 \notin \mathcal{L}$, player 1 has the advantage. Whenever a player plays an action in a dominated strategy, the advantage transfers to the other player.

We refer with $G(k, M)$ to the race game with M position and a choice set $\{1,..,k\}$

[23]We choose the term losing position, following Gneezy et al. [2010]. It comes form the fact that an agent that chooses there is meant to lose the game, so that an agent reaching a losing position is, actually, winning.

2.4.2 Previous experiments

The race game has been the object of a series of recent investigations. Dufwenberg et al. [2010] study the games $G(3{,}21)$ and $G(2{,}6)$. The subjects play both games, but the order in which those games are played is varied. They ask whether solving a simpler game helps in tackling a similar but more complex one, finding support for this hypothesis.

Gneezy et al. [2010] analyze $G(3{,}15)$ and $G(4{,}17)$. They find that subjects switch to backward analysis as they gain experience. Their understanding of the game proceeds from the final positions: the losing positions are discovered sequentially, starting from the one that is closer to the final position.

Levitt et al. [2011] have $G(9{,}100)$, $G(10{,}100)$ and a centipede game played by chess players. Consistently with previous studies, they find the players' actions to be, on average, closer to the dominant strategy as the game proceeds. Also, a relatively minor change in the game - changing k from 9 to 10 - has a major impact on the performance of the players[24]. Interestingly, the performance in the race game is strongly correlated to the players' ranking as chess players, which is not true with respect to the centipede game[25]. This suggests that the race game is able to capture the pure ability to backward induct, which makes it an ideal set up to test strategic reasoning in sequential games.

2.5 Design

On aggregate, previous results indicate that individuals are unable to figure out their dominant strategy from the beginning; rather, they discover it as they gain experience, starting from the actions closer to the end. This observation, despite being consistent with the LBI hypothesis, leaves open the question of whether the subjects actually reason only on a limited number of steps ahead of the current decision node and if they do so consistently with backward induction.

Under limited backward induction, the players solve only a reduced game that includes the stages that are closer to the current decision node. To identify this behavior we need to show that the agents (*i*) reason backward, and (*ii*) do so on a reduced game tree. In particular, there are two alternative explanations that we must disentangle from LBI: under the first, the players perform backward induction from the terminal histories, but stop the iterative process after some steps[26]; under the second, behavior is fully driven by beliefs about the others playing a dominated strategy longer than they do.

[24]In $G(9{,}100)$ the set of losing positions is $\mathcal{L} = \{10,20,30,40,50,60,70,80,90\}$. It is a focal sequence, and is easier to see, with respect to that in $G(10{,}100)$, which is $\mathcal{L} = \{1,12,23,34,45,56,67,78,89\}$.

[25]Palacios-Huerta and Volij [2009] report partially different findings with respect to the centipede game.

[26]Gneezy et al. [2010] seem to favor this interpretation.

We introduce two modifications to the base game to disentangle LBI behavior.

The trap prize: we add a small prize p at an intermediate position $m \notin \mathcal{L}$. Winning p gives your opponent the chance of winning P. Intuitively, p allows indentifying the reduced game trees that include m as pseudo-terminal histories. Note that p implies there is no longer a dominant strategy: an agent could try to win both prizes, or only one, depending on his beliefs about the strategic ability of his opponent. As a consequence, when we observe an agent playing consistently with SPE on the reduced game tree, we could not distinguish between LBI and "confident" behavior. We achieve this distinction with the following device.

Claims: we allow players to claim they are going to win p and P in any position, and independently of who is moving. One can claim both prizes at the same time; a claim cannot be withdrawn. The players get no feedback on the claims of their opponent and those are not affecting their payoff; so the claims are non-strategic. However, one's own claims are payoff-relevant: by claiming P (p) at position t an agent gets $M - t$ ($m - t$), on top of the prize, in case his claim is realized; otherwise, he gets a fine, F . We use claims to track what the players are targeting along the game.

We call $\mathcal{L}_p$ and $\mathcal{W}_p$ the set of losing and winning positions on the path to p; $\mathcal{L}|_m \subseteq \mathcal{L}$ and $\mathcal{W}|_m \subseteq \mathcal{W}$ indicate the winning and losing positions toward P, restricted to positions higher than m. Obviously, one of two players will have an initial advantage to get p. A generic modified race game will be identified as $G(k, M, m, P, p, F)$.

2.5.1 Parameters and treatments

In accordance to the latter convention, we will denote a game with no trap prize as featuring $m = 0$ and $p = 0$. We investigate the games $G_0(6, 66, 0, 100, 0, -15)$ and $G_1(6, 66, 40, 100, 30, -15)$.

G_0 is just a base game with payoff-relevant claims. The set of losing positions is

$$\mathcal{L} = \{3, 10, 17, 24, 31, 38, 45, 52, 59, 66\}.$$

The game displays first mover advantage: in every SPE, player 1 wins P and claims it as soon as the game starts ($t = 1$).

The same is true for G_1. The set of losing positions on the path to p is

$$\mathcal{L}_p = \{5, 12, 19, 26, 33, 40\}.$$

Player 1 has an advantage to win p, as well. However, in any SPE player 1 wins P and player 2 wins p. Each player claims victory of his respective prize in the first position ($t = 1$).

We run two treatments, $T0$ and $T1$, featuring G_0 and G_1, respectively. Each subject participated only in one of those treatments (*between* protocol). Subjects played 8 race games of the relevant type changing partner and role (Player 1 and Player 2) in every repetition (perfect strangers matching). $T0$ is used as a benchmark for actions and claims related to the large prize P. We let the player take all the time they need to take each single action, and we do not impose any time constraint on them, so that their strategic ability is not biased by time pressure.

2.5.2 Equilibrium and out-of-equilibrium behavior

In any SPE equilibrium of the race game, the player who has the initial advantage wins P; this holds true both in the base and in the modified game. When a small prize is available, the player at disadvantage with respect to P, wins p, no matter if he has an initial advantage toward it or not. Each player claims the prize he is going to win in equilibrium at the initial position.

As it looks like a natural alternative to LBI, we now consider a dynamic level-k model for the race game, as in Kawagoe and Takizawa [2012]. Level zero identifies a random player, normally assumed to be fictitious - i.e. it exists only as a belief in the mind of the higher-level players. Each level, l, believes the others to be of level $l-1$, and best replies to their actions. Following Ellingsen and Ostling [2010], we assume that a player observing an action that is inconsistent with his beliefs, revises them assuming the opponent to be of the highest, among the levels lower than his own, that makes his inference consistent. Since any action profile is played with positive probability by L_0, a player can always hold a belief that is consistent with the current history. Noting that L_0 plays with positive probability one of his dominant strategies, for $k \geq 1$, L_k never plays a dominated strategy. Thus, in the standard race ($p = 0$) every level should mimic SPE.

When we add p, L_1 may try to win both prizes. Intuitively, this happens if the probability that a random player will pick by chance all $t \in \mathcal{L}|_m$ is sufficiently low[27]. Thus, if the action space, k, the distance between the prizes, $M - m$, and p are relatively large, an L_1 player targets the positions in $\mathcal{L}_p$, switching to those in $\mathcal{L}|_m$ for $t > m$. In the other cases, his play mimics SPE. The higher levels play as in SPE, regardless of the parameters, unless they end up believing their opponent is L_0, in which case they mimic L_1. The only type that tries to win both prizes is the less sophisticated one, L_1, and he should fail achieving them (as long as L_0 players do not exist). Most notably, the players should pass through all losing positions, $\mathcal{L}_p$ and $\mathcal{L}$.

LBI provides a different perspective. As the players move to higher positions, they

[27] How to compute this probability depends on the type of randomness assumed for L_0. See Kawagoe and Takizawa [2012] for a discussion of the issue.

discover new portions of the game tree. If no history corresponding to winning a prize is included in the LF-game, a player cannot distinguish between any two positions in terms of chances of winning. As a consequence, the same intermediate payoff is associated to all pseudo-terminal histories. In the base game, as far as a node corresponding to position M enter his foresight, he will switch to actions consistent with his dominant strategy - i.e. try to reach the positions in $\mathcal{L}$. The higher his foresight, the sooner this will happen, the more chances of winning the player will have. For example, F_1 realizes how to win only when the distance to the prize is lower than k, F_2 when it is lower than $2k$, and so on. A player is expected to claim the prize as soon as the positions associated to it fall within his foresight, and conditional on being at a winning position.

A similar reasoning applies to the modified game. What changes is that position m falls in the LF-game before M. Thus, the players start targeting $\mathcal{L}_p$ before $\mathcal{L}$, and claiming p before P. This latter feature distinguish behavior consistent with LBI from that of players who can reason backwards for a limited number of steps, but do so considering the whole game. In both games, the players spend most of the time off the equilibrium path and converge to losing positions only when a prize is approached.

We next show that we can properly identify LBI behavior using actions and claims. Consider a player that was not able to solve the whole game from the start. At some position $t \leq m$, he discovers the solution, reasoning backwards from the terminal histories. Now he knows how to get p and P. Depending on his belief about the strategic ability of his opponent, he will either target P or *p and then P*. In the latter case, we may misinterpret as LBI a behavior which is not. It is driven by the player believing the probability that his opponent solves the game while $t \leq m$ to be sufficiently low. For obvious reasons, we refer to this type of behavior as "*Confident*".

Claims turn out useful here, as we cannot disentangle *Confident* behavior and *LBI* from the actions. We can show that, given our parameters, a *Confident* player should claim P as soon as he learns the solution.

Assume that at some $\bar{t}$, $\bar{t} < m$ and $\bar{t} \in \mathcal{W} \cap \mathcal{W}_p$ a player, j, discovers the solution. He has, basically, three options: target and claim P (S_1); target p and P, claim p now and P only when sure of getting it (S_2); target p and P, claim p and P (S_3). The payoff from S_1 is

$$\pi_1 = P + (M - \bar{t})$$

The payoff from S_2 is

$$\pi_2 = p + (m - \bar{t}) + q(P + (M - m - k_{opp,m}))$$

where q is the probability that the opponent does not solve the game within m and $k_{opp,m}$

is the action of the opponent at m[28]. The payoff from S_3 is

$$\pi_3 = p + (m - \bar{t}) + q(P + (M - \bar{t})) + (1 - q)F$$

We show that there is no belief that sustains S_2 with our parameters, so that can S_2-consistent behavior (acting towards p and claiming only p) can be attributed only to LBI.

Proposition 2. *For a risk neutral agent, $S_2 \succeq S_3$ and $S_2 \succeq S_1$ if and only if:*

$$\frac{\Delta_P + \Delta_M}{P + \Delta_M - k_{opp,m}} \leq q \leq \frac{-F}{m + k_{opp,m} - \bar{t} - F} \tag{2.2}$$

where $\Delta_P = P - p$ and $\Delta_M = M - m$.

Corollary 1. *For $\bar{t} < 37$ there exists no q such that $S_2 \succeq S_3$ and $S_2 \succeq S_1$.*

This and the following proofs are immediate and are thus omitted. Intuitively, to choose S_2, a player needs to believe he has enough chances of winning P after winning p, but not that many so as to induce him to claim P immediately. With our parameters, the above interval for q does not virtually exist: even taking the minimum possible $k_{opp,m}$, 1, which corresponds to the largest possible interval, such a probability exists only for $\bar{t} \geq 37$; in this case the interval of beliefs that sustain S_2 is $q \in (0.77, 0.79)$.

The following two propositions regard players displaying constant absolute risk aversion (CARA), and constant relative risk aversion (CRRA), respectively. They show that risk aversion is not a major concern in this context. The reason for this is that S_3 is more risky than S_2, but S_2 is more risky than S_1: as a consequence, for risk averse players, both the upper and the lower bound of the above interval move in the same direction.

Proposition 3. *Consider an agent, whose utility function is $U(x) = -e^{-\alpha x}$, $\alpha > 0$, featuring CARA. Then $S_2 \succeq S_3$ and $S_2 \succeq S_1$ if and only if:*

$$\frac{1 - e^{-\alpha(\Delta_P + \Delta_M)}}{1 - e^{-\alpha(P + \Delta_M - k_{opp,m})}} \quad \leq \quad q \quad \leq \quad \frac{1 - e^{-\alpha F}}{1 + e^{-\alpha(P + M - \bar{t})} - e^{-\alpha(P + \Delta_M - k_{opp,m})} - e^{-\alpha F}} \tag{2.3}$$

Now let u_s^i, $i \in \{1, 2, 3\}$, $s \in \{g, b\}$, be the (rescaled) utility of a CRRA agent, when his strategy is S_i, conditional on state s. In case the opponent does not solve the game within

[28] More precisely, q is the probability that the opponent, choosing at m, will not reach a position in $\mathcal{L}$ (for whatever reason).

m, $s = g$; otherwise, $s = b$. Let ρ be the coefficient of CRRA. Then:

$$\begin{aligned}
u^1_{g,b} &= (p + m + \Delta_P + \Delta_M - \bar{t})^{1-\rho} \\
u^2_g &= (2p + m + \Delta_P + \Delta_M - \bar{t} - k_{opp,m})^{1-\rho} \\
u^2_b &= (p + m + -t)^{1-\rho} \\
u^3_g &= (2p + 2m + \Delta_P + \Delta_M - 2\bar{t})^{1-\rho} \\
u^3_b &= (p + m - \bar{t} + F)^{1-\rho}
\end{aligned}$$

Proposition 4. *Consider an agent, whose utility function is* $U(x) = \frac{x^{1-\rho}}{1-\rho}$, $\rho > 0$ *featuring CRRA, then* $S_2 \succeq S_3$ *and* $S_2 \succeq S_1$ *if and only if:*

$$\frac{u^1_{g,b} - u^2_b}{u^2_g - u^2_b} \leq q \leq \frac{u^3_b - u^2_b}{(u^2_g - u^2_b) - (u^3_g - u^3_b)} \tag{2.4}$$

Corollary 2. *If an agent is not able to reach m in a single move:*

i *Under CARA, for* $\alpha < 0.8$ *there exists no q such that* $S_2 \succeq S_3$ *and* $S_2 \succeq S_1$*;*

ii *Under CRRA, for* $\rho < 0.5$ *there exists no q such that* $S_2 \succeq S_3$ *and* $S_2 \succeq S_1$*.*

This implies that, until very close to m, under no reasonable parameter of risk aversion there exists a belief sustaining S_2[29]. Moreover, even moving closer to m and admitting higher risk aversion levels, the interval of beliefs that sustain S_2 remains virtually irrelevant.

In Figure 2.5 the bounds of the interval for q, in the case of CARA (circles) and CRRA (stars), are plotted against the coefficient of (constant or relative) risk aversion. To sustain S_2, q must be higher than the solid line, and lower than the dashed line. It is assumed that $\bar{t} = 36$ and $k_{opp,m} = 1$, giving the interval the highest chances to exist and the largest magnitude. As shown, for low levels of risk aversion, there exists no belief supporting S_2. For intermediate values, a tiny interval, smaller than 0.06, appears[30].

Overall we might confuse LBI and confident behavior for claims only under the following conditions: the claim happens within a limited number of positions, in the neighborhood of m (*a*); under extremely restrictive beliefs about the strategic ability of the opponent, corresponding to small range (≤ 0.06) of probabilities for the opponent failing to solve the game at m around $0.8 - 0.9$ or above (*b*); under restrictive beliefs about the choice of the opponent at m ($k_{opp,m} = 1$) (*c*); for very high values of risk aversion (*d*).

[29] There is no consensus on the estimation of the coefficients of relative and absolute risk aversion. However most experimental and field studies agrees on average coefficients that are below our thresholds [see Harrison and Rutström, 2008]

[30] If we let risk aversion grow beyond the bounds in the figure, this interval would eventually shrink.

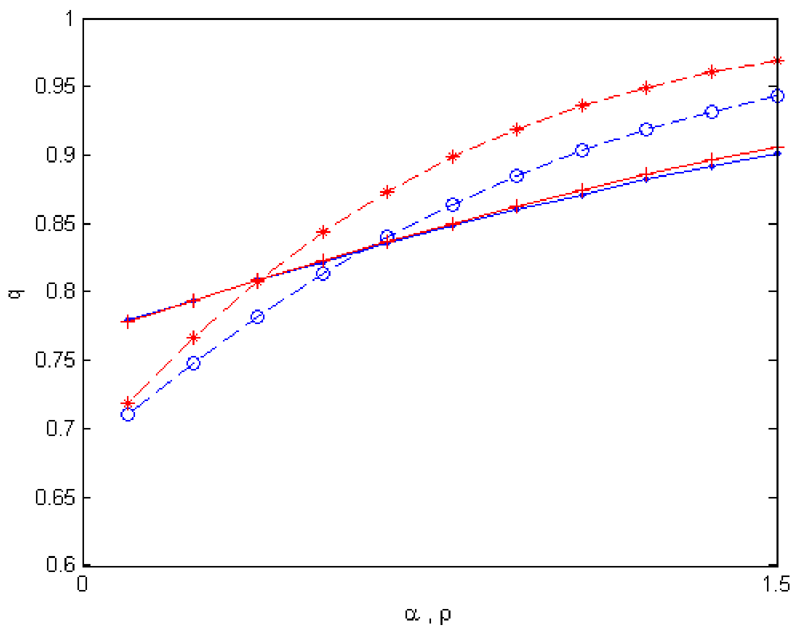

Figure 2.5: Risk aversion and beliefs sustaining S_2: CARA (red), CRRA (blue)

Given the low probability of (*a*),(*b*),(*c*) and (*d*) occurring at the same time[31], we claim that our interpretation of an S_2-consistent strategy as due to LBI, to be sufficiently sound.

Summing up, in equilibrium, we should observe only positions in $\mathcal{L}$ and $\mathcal{L}_p$ in G_1 and no error in either treatment; Level-k agents should mimic this behavior, although we may observe some players (L_1) claiming both prizes. A more relaxed version would see the subjects learning a strategy for the *whole* game as they play it. In other words we should see portions of an equilibrium strategy being played as the game proceeds. Even here, we would expect errors toward P to fade away *before* errors toward p, and P being claimed before p. Let us call this *equilibrium-like* behavior. *Confident* players - actually only those whose beliefs prove correct - might invert this ordering with respect to actions, but not with respect to claims: we should see errors toward P decline later than errors toward p, but prizes being claimed at the same time.

Under LBI, both claims and actions should display a reverse timing with respect to equilibrium and equilibrium-like behavior. We say that a player makes a *p-error* (*P-error*), if, choosing at a position in $\mathcal{W}_p$ ($\mathcal{W}$), he does not reach a position in $\mathcal{L}_p$ ($\mathcal{L}$); that is, if he does not exploit his advantage toward winning a prize. We state the following main hypothesis, within $T1$:

[31]There also seems to be a contrast in the required pair of high confidence (*b*, *c*) and high risk aversion (*d*).

Hypothesis 1. *The rate of p-errors decreases earlier than that of P-errors.*

Hypothesis 2. *Prize p is claimed before prize P.*

Through repetitions, the subjects are expected to learn the strategic features of the game (and, eventually, fully solve it). Error rates should then decrease and claims become more precise.

Comparing $T0$ and $T1$, we expect the presence of the trap prize to induce higher P-error rates in all positions $t < m$. This conclusion stems from two considerations. First, a successful targeting of p induces more P-errors, even when compared to random behavior. Second, the presence of p makes the game more complex, resulting in a lower strategic performance of the subjects[32].

Hypothesis 3. *The rate of P-errors is higher in $T1$ than in $T0$, due to the trap effect of p.*

2.5.3 Procedures

The experiment took place at the EELAB of the University of Milan-Bicocca on June 15th, 2012. The computerized program was developed using *Z-tree* [Fischbacher, 2007]. The subject display was as similar as possible to the one used by Gneezy et al. (2010). We run 4 sessions with 24 subjects per session, for a total of 96 participants, equally split across tratments. Participants were undergraduate students from various disciplines,[33] recruited through an announcement on the EELAB website.

Instructions were read aloud (see Appendix A for an English translation of the instructions). Participants filled in a control questionnaire to ensure everybody understood the instructions before starting the experiment.

Sessions took on average 70 minutes, including instructions, control and final questionnaire phases.

During the experiment subjects earned Experimental Currency Units. At the end one game was selected at random for each couple and subjects were paid the points they earned in that game only, according to an exchange rate of $1€ = 10ECU$. Average payment was 11.10€ with a minimum of 2.50€ and a maximum of 25.40€. Subjects received an initial endowment of 4€ that could be partially spent to pay fines in case of bankruptcy during the experiment.

[32]To the extreme, a subject may adopt a "one-problem-at-a-time" approach, starting to think about P at m. Quite remarkably, the working paper version of Gneezy et al. [2010] was titled "I will cross that bridge when I come to it".

[33]Sociology, economics, business, psychology, statistics, computer science, law, biology, medicine, mathematics, pedagogy and engineering.

2.6 Results

To facilitate the presentation of the results, we partition the set of positions into intervals. Each interval is formed by all the positions within two losing positions, excluding the lower and including the upper bound. In $T1$, we have a similar partition for the small prize, p, over positions $1,...,40$. Given that the sets $\mathcal{L}$ and $\mathcal{L}_p$ are disjoint, the intervals for P and p do not perfectly overlap[34] and should be kept distinguished. If not differently specified, intervals should be understood with reference to the prize they are specific for.

2.6.1 Errors

We start by tracking P- and p-errors over intervals. Figure 2.6 reports the rate of errors in the two treatments and for both prizes, distinguishing between repetitions 1-4 (a) and 5-8 (b). It represents the fraction of the subjects that did not reach the upper bound of the interval. Choices taken at $\mathcal{L}$ and $\mathcal{L}_p$ are excluded from the computation of P- and p-errors, respectively, since, by definition, there is no correct action available at those positions. Note that the simple possibility to make two different kinds of error, in $T1$, is not sufficient to inflate the rate of errors, since an action that reaches $\mathcal{L}$ ($\mathcal{L}_p$) counts as a p-error (P-error), *only if* it is not taken at a position in $\mathcal{L}_p$ ($\mathcal{L}$).

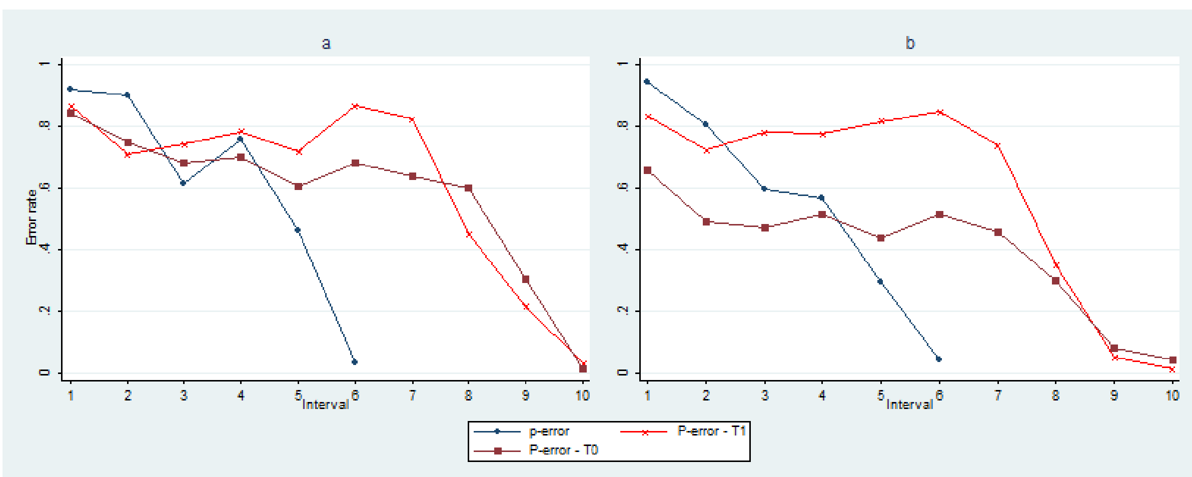

Figure 2.6: Error rates over intervals: first (a) and last (b) four repetitions

In the first repetitions, the rate of P-errors follows a similar pace in $T0$ and $T1$, though it is slightly higher in the latter between interval 3 and 7. Around eighty percent[35] of the

[34]In particular, the upper bound of an interval that refers to P is two positions lower that the corresponding interval that refers to p.

[35]An agent choosing at random makes an error 83 percent of the time.

subjects choosing at winning positions makes an error in the first interval. This percentage decreases slightly in $T0$, reaching 60 percent at interval 8. It remains more or less stable in $T1$ until it drops sharply after interval 7. The rate of p-errors starts above 90 percent, it drops then significantly, reaching 45 percent in interval 5. In the last four repetitions, the performance of the subjects improves significantly in $T0$, though it remains above 40 percent until interval 7. In $T1$ the rate of P-errors is unchanged, with respect to the first repetitions, until interval 7, after which is substantially lower. The performance with respect to prize p shows a relevant improvement.

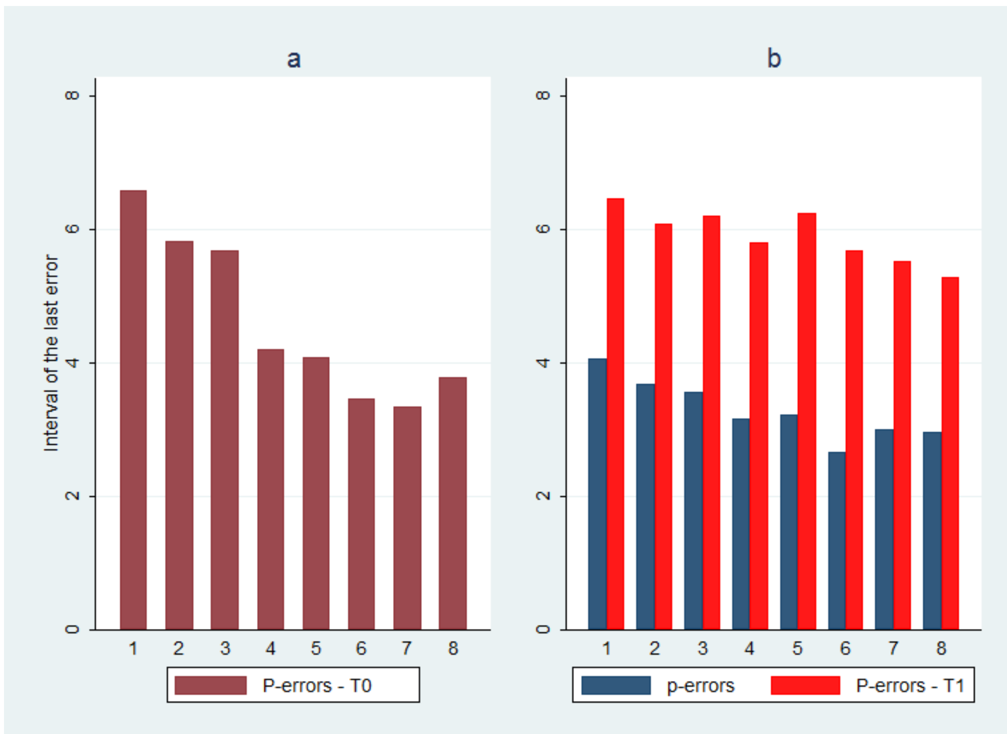

Figure 2.7: Average last error over reetitions: T0 (a) and T1 (b)

A more precise measure of the moment where a subject understands the solution of the game is to identify the last interval where he makes an error. It should be noted that when an agent stops making errors, we cannot register any error on the part of his opponent. As a consequence the last errors we record and their distribution should be understood as a lower bound for the real ones.

Figure 2.7 shows the average last error for each repetition. As before, the performance improves sharply in $T0$, passing from above the sixth to below the fourth interval. Starting from a similar level in the first repetition, the improvement is smaller in $T1$. The average last p-error is consistently below the average last P-error in $T1$, with a difference of up to three intervals.

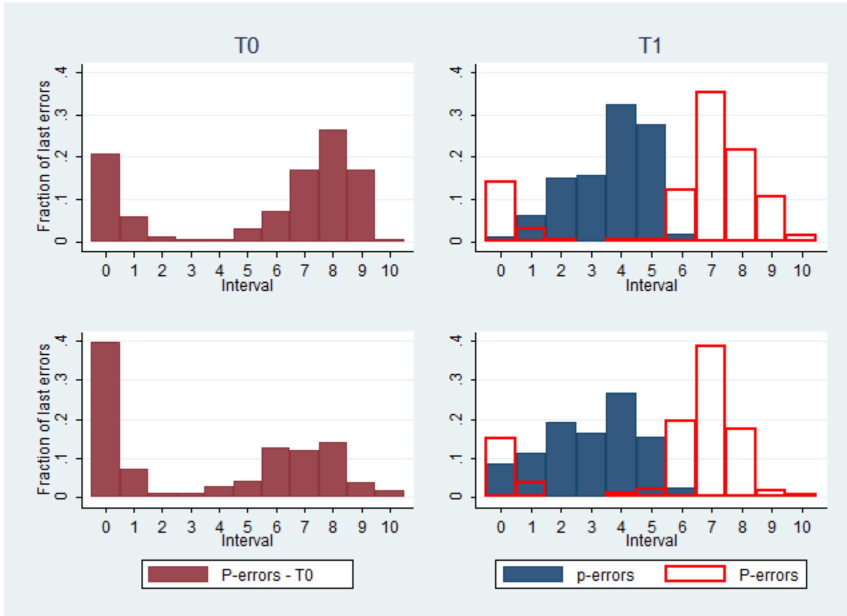

Figure 2.8: Distribution of last errors: first (top) and last (bottom) four repetitions

Figure 2.8 reports the distributions of the last errors, by treatment and for the first and last repetitions. The same distributions are depicted in Figure 2.9 only for those who claimed the prize for which the last error is recorded. In the latter case the distribution shifts slightly toward lower intervals, without affecting the general picture. The distributions for prize P are bimodal. A fraction of the subjects does not make any error, or do so only in the first interval. The rest of the population stops making errors only close to the final position. Those, who constitute a majority of the population, display a bell shaped distribution, peaking around intervals six, seven and eight.

We reject the null of an equal distribution between $T0$ and $T1$ using a Kolmogorov-Smirnov two-sample test (KST), both for the raw distributions ($D = 0.2188$, P-val $= 0.00$) and for those restricted to the claiming subjects ($D = 0.2114$, P-val $= 0.00$). A Mann-Whitney rank sum test (MWRST) confirms that the average last P-error is different in the two treatments ($z = -2.766$, P-val $= 0.00$). Indeed the fraction of subjects not making errors is apparently higher in $T0$. Moreover, the learning effect is much higher: around 40 percent of the subjects show a perfect play in the last repetitions of $T0$[36]. We observe smaller differences across repetitions in $T1$, concentrated in the right-hand side of the distributions.

[36]Recall that when a player does not make any error, we cannot register any error for his opponent as well.

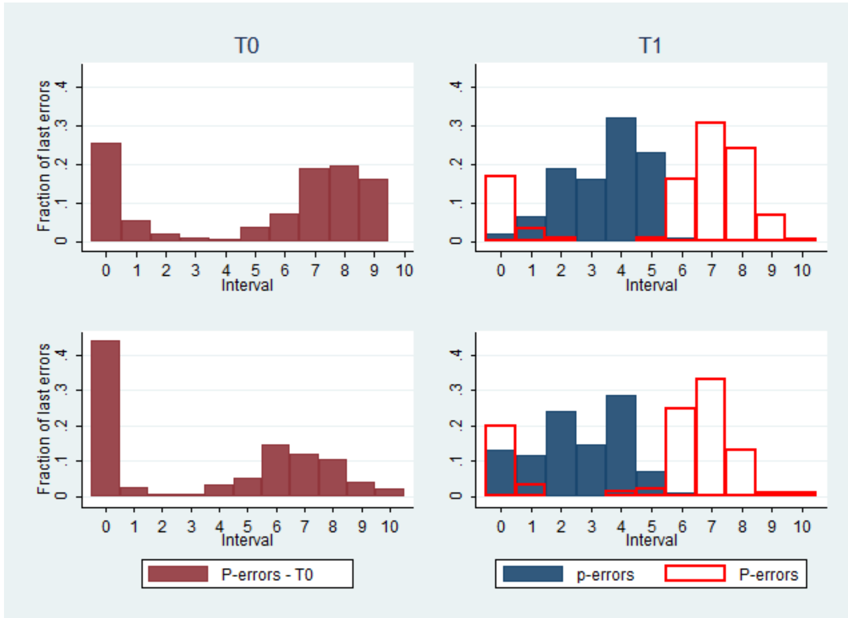

Figure 2.9: Distribution of last errors, claimed prizes only: first (top) and last (bottom) four repetititions

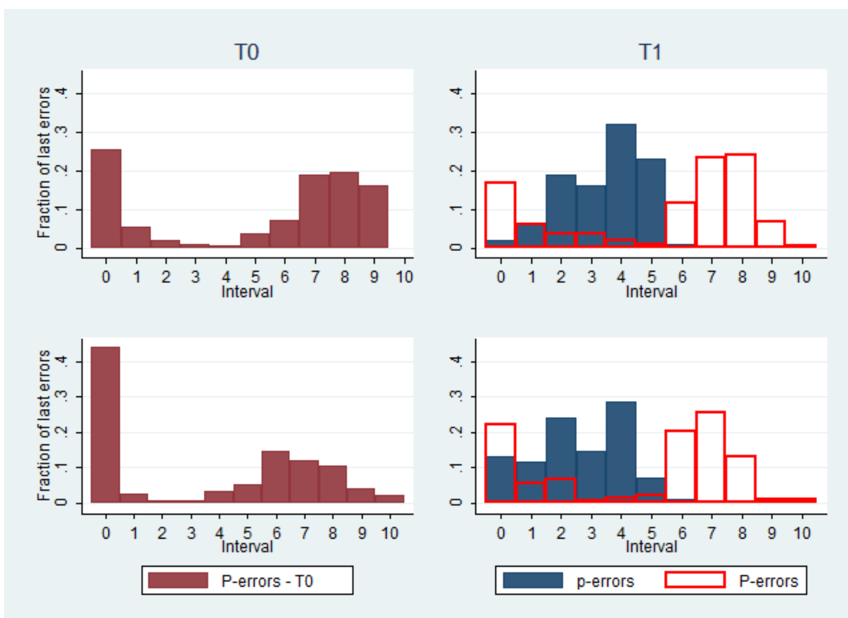

Figure 2.10: Distribution of adjusted last errors, claimed prizes only: first (top) and last (bottom) four repetititions

The distribution of the last p-errors shows again a majority of the subjects making

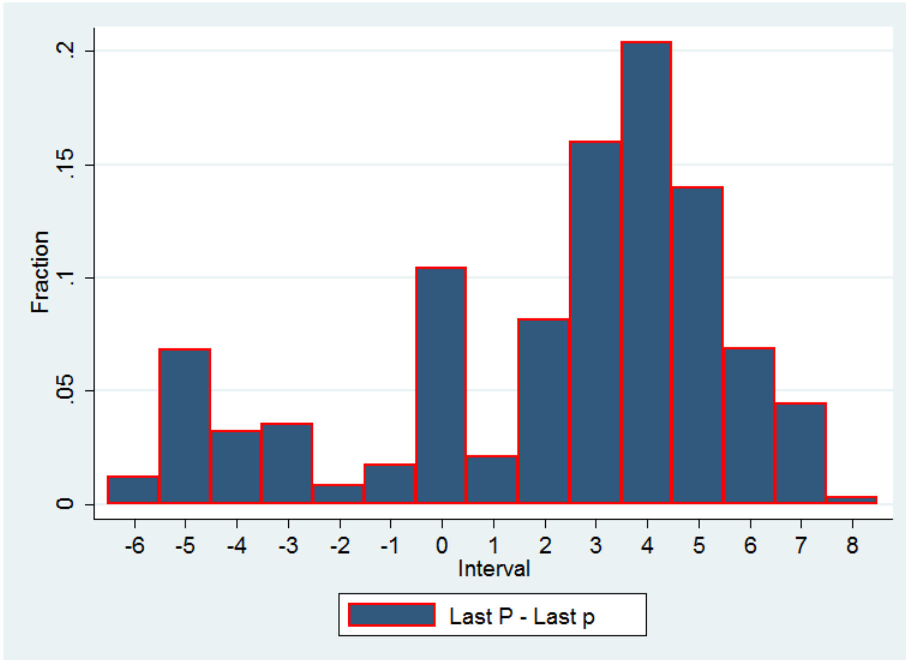

Figure 2.11: Distribution of differences between the last P-error and last p-error

the last error less than three intervals before the prize. We observe only a tiny fraction of players not making any p-error. Within $T1$, improvements in understanding how to win p capture most of the learning effect.

That most subjects make the last P-error after they make their last p-error is apparent in Figure 2.11, where the distribution of the difference for each individual is plotted.

Given that, in $T1$, a player can make a P-error even if he understood how to win P, in case he was aiming at winning both prizes, we check if the differences across treatments are not due to this effect. In case a subject claimed both prizes[37] and won prize p, without making any P-error after that, we attribute to him a last P-error equal to his last p-error. The new distributions are shown in Figure 2.10. Despite we do observe a better performance with respect to the base distributions, the effect is minor and does not affect the difference with respect to $T0$ (KST: $D = 0.1458$, P-val < 0.01; MWRST: $z = -1.873$, P-val $= 0.06$).

Our results are highly consistent with previous investigations. Despite a number of subjects play (or learn to play) close to equilibrium, most of them find it difficult to solve the game backwards, even after gaining significant experience. In the early intervals, their actions cannot be distinguished from random play. A majority of the subjects discover how to win a prize only when they are at a distance of at most three intervals from

[37]Recall the discussion in section 2.5.2

it.

We find strong evidence supporting hypothesis 3 and 1. The subjects manage to reach the positions in $\mathcal{L}_p$ before those in $\mathcal{L}$. To check whether this behavior is due to LBI reasoning we analyze claiming behavior.

2.6.2 Claims

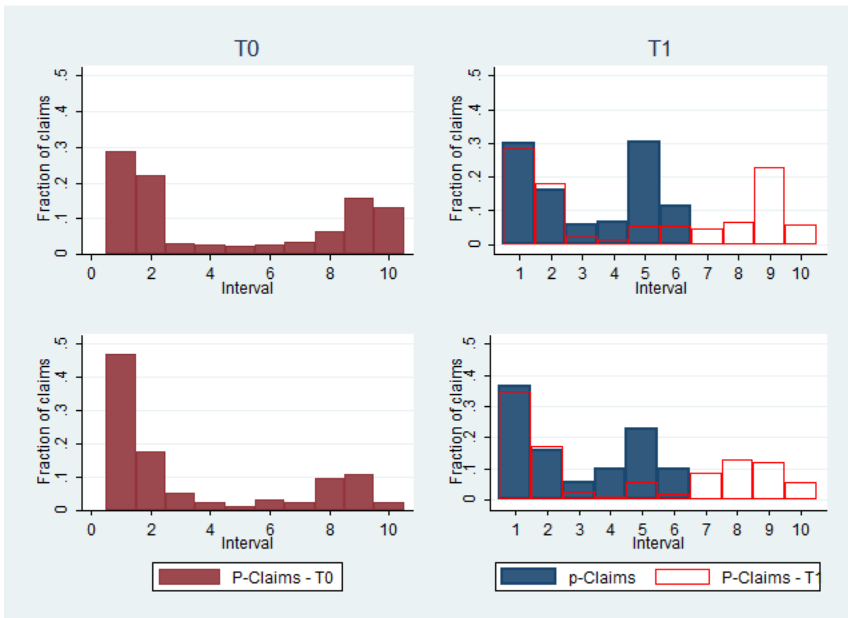

Figure 2.12: Distribution of claims: first (top) and last (bottom) four repetitions

The subjects made a wide use of the claiming device. Around 65 and 67 percent of the subjects claimed they would have won prize P in $T0$ and $T1$, respectively; 60 percent claimed prize p in $T1$. Some claims were unwarranted, and, indeed, a fine was imposed on 29 percent of both the p-claims and the P-claims; thus, most claims ended with the claiming player winning the prize.

Figure 2.12 displays the distribution of claims over intervals; the upper and the lower panels display results for repetitions 1-4 and 5-8, respectively, for $T0$ (left) and $T1$ (right). The distributions of P-claims tracks the pace of that of last errors, separated between early (interval 1-2) and late claimers (intervals 7-9) . The comparisons between treatments and prizes that we made for errors, holds basically unchanged for the claims. In particular, the distribution of p-claims is first order stochastically dominant with respect to that of P-claims.

With respect to last errors, the proportion between the early and late claimers shifts

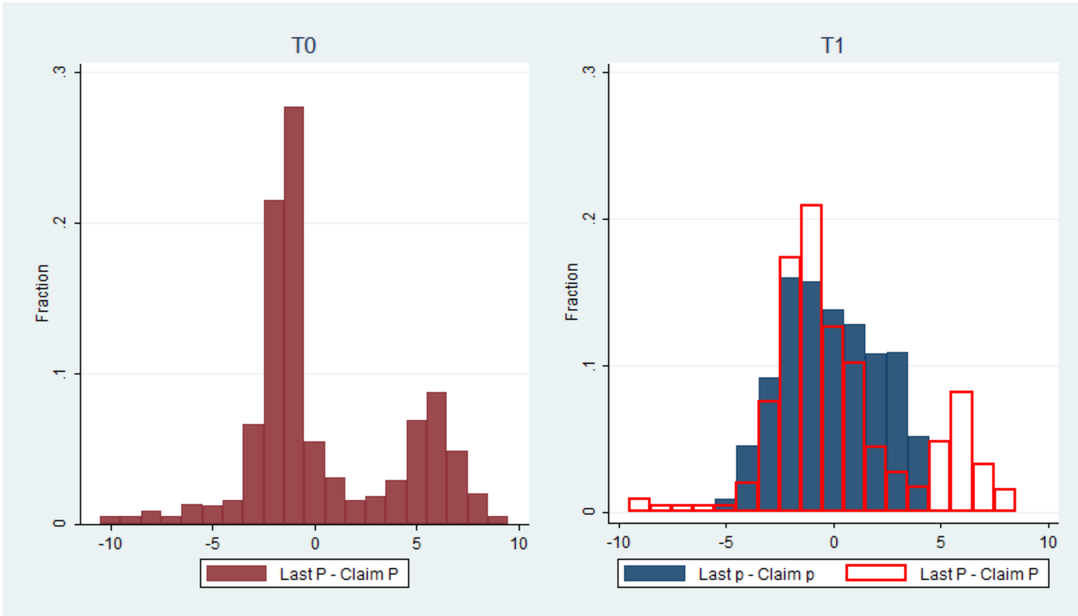

Figure 2.13: Distribution of differences between the timing of the last error and that of the claim

in favor of the former. This tendency to bet on one's ability to solve the game before the opponent is documented in Figure 2.13: many players claim a prize before they make their last error, in both treatments. However, a majority of the claims happen at a distance of less than one interval from the last error.

Given the previous remarks, we look at claims that are *perfect*, in the sense that the player does not make any error after he claims. As with the last errors, we correct for the possibility that a player makes an error consciously, as he thinks he can win both prizes: in $T1$, we take a P-claim to be perfect if an agent claim both prizes, does not make any p-error after claiming P and does not make any P-error after winning p. Figure 2.14 shows the fraction of those claims on the total for each interval; the two upper panels regard P-claims in $T0$ (left) and $T1$ (right), the lower panel regards p-claims. Late claims are more likely to be perfect: in the three intervals prior to a prize, virtually all claims are perfect. This fraction drops significantly if we move further form the prize, and particularly so for p-claims. The distribution of perfect claims (Figure 2.15) does not present any surprise.

Overall, the evidence is consistent with hypothesis 2: prize p is claimed before prize P. Combined with the results about errors, this provides strong support for the LBI hypothesis. It should be noted that, as with errors, repetitions drive behavior apart in the two treatments.

We close this section with some insights from the timing of the decisions of the sub-

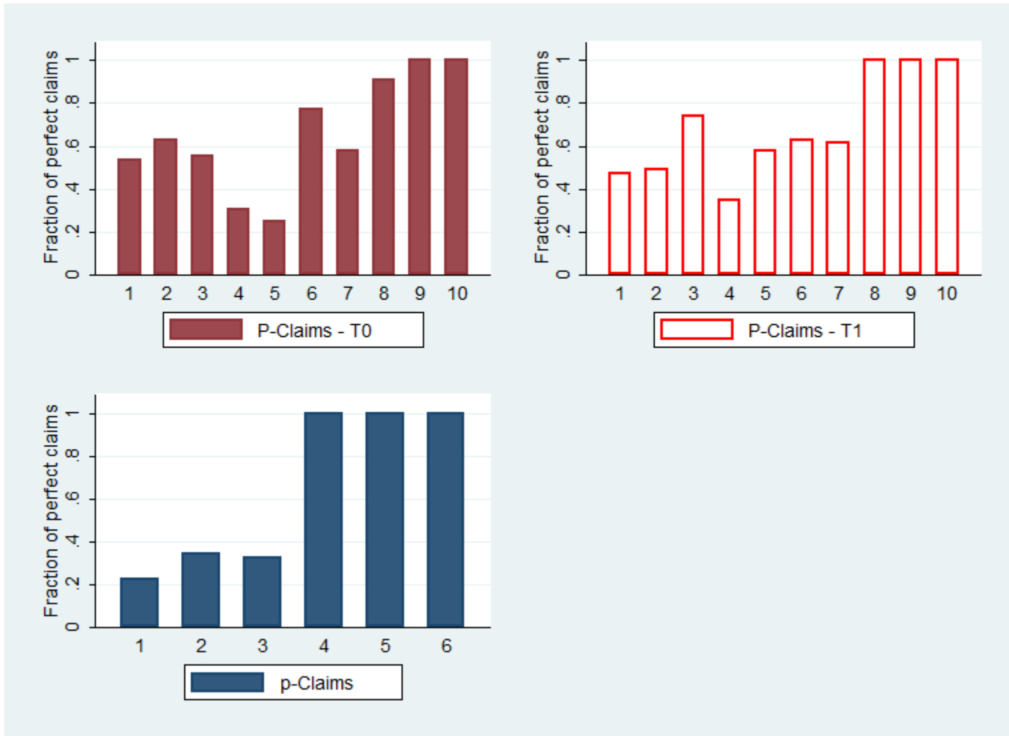

Figure 2.14: Fraction of claims that are perfect, over repetititions

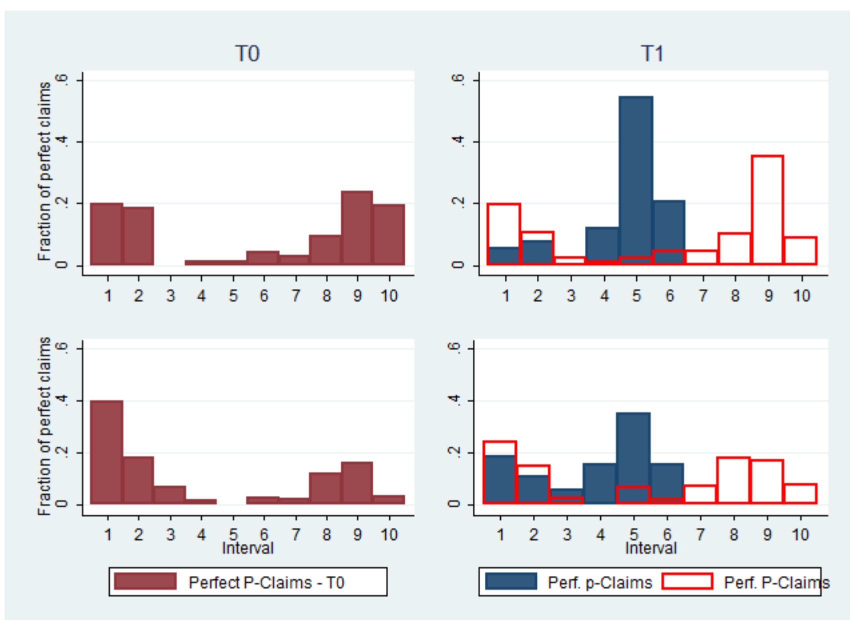

Figure 2.15: Distribution of perfect claims: first (top) and last (bottom) four repetititions

jects.

2.6.3 Timing

We wanted the players to take the time they needed to reason and reach a decision without the pressure of time constraints. We then analyze how the subjects used their time.

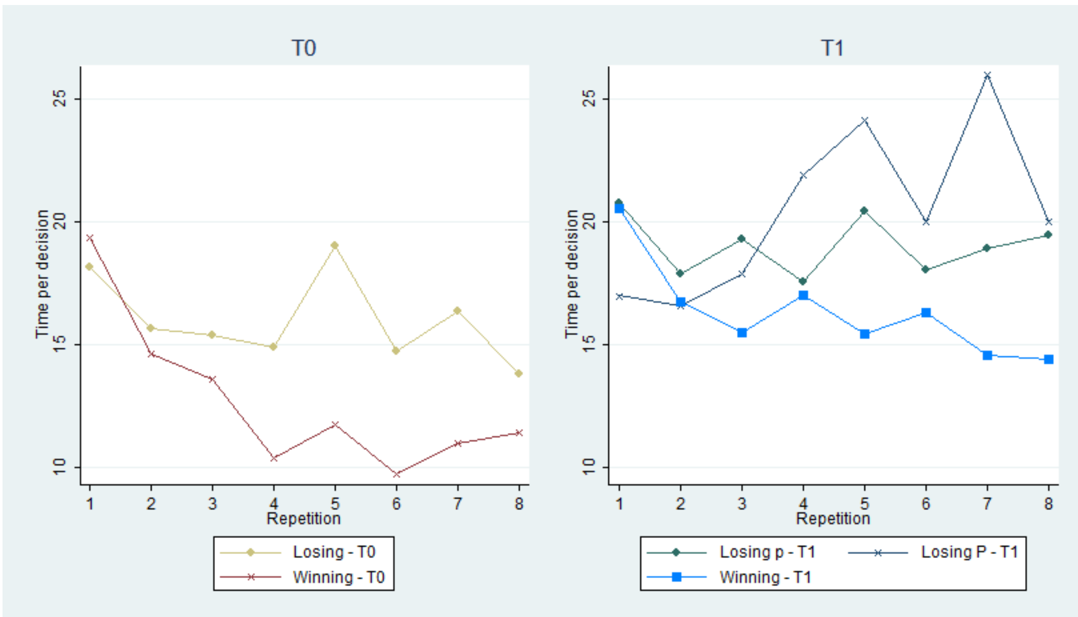

Figure 2.16: Average time for each decision over repetitions

Figure 2.16 displays the average number of seconds the agents took to choose how many steps to take. The two panels refers to $T0$ (left) and $T1$ (right). Different averages are computed for winning and losing positions, so that we observe, for $T0$, the mean for $\mathcal{L}$ and $\mathcal{W}$, and for $T1$, the mean for $\mathcal{L}$, $\mathcal{L}_p$ and $\mathcal{W} \cap \mathcal{W}_p$.

In the first repetition the averages are indistinguishable and around twenty seconds for each decision. The time for each choice drops progressively through repetitions only in the case of winning positions, so that deciding at losing positions takes relatively more time with respect to winning positions as the game is repeated; this result is consistent with the evidence in Gneezy et al. [2010].

We have a more precise picture of where it takes longer to decide by tracking the average time per decision for the losing positions that are closer to the prizes. The corresponding figures are shown in figure 2.17, with the upper panels referring to prize P in $T0$ (left) and $T1$ (right) and the lower panel referring to prize p.

The losing position that is closest to a prize trigger a decision in about the same time

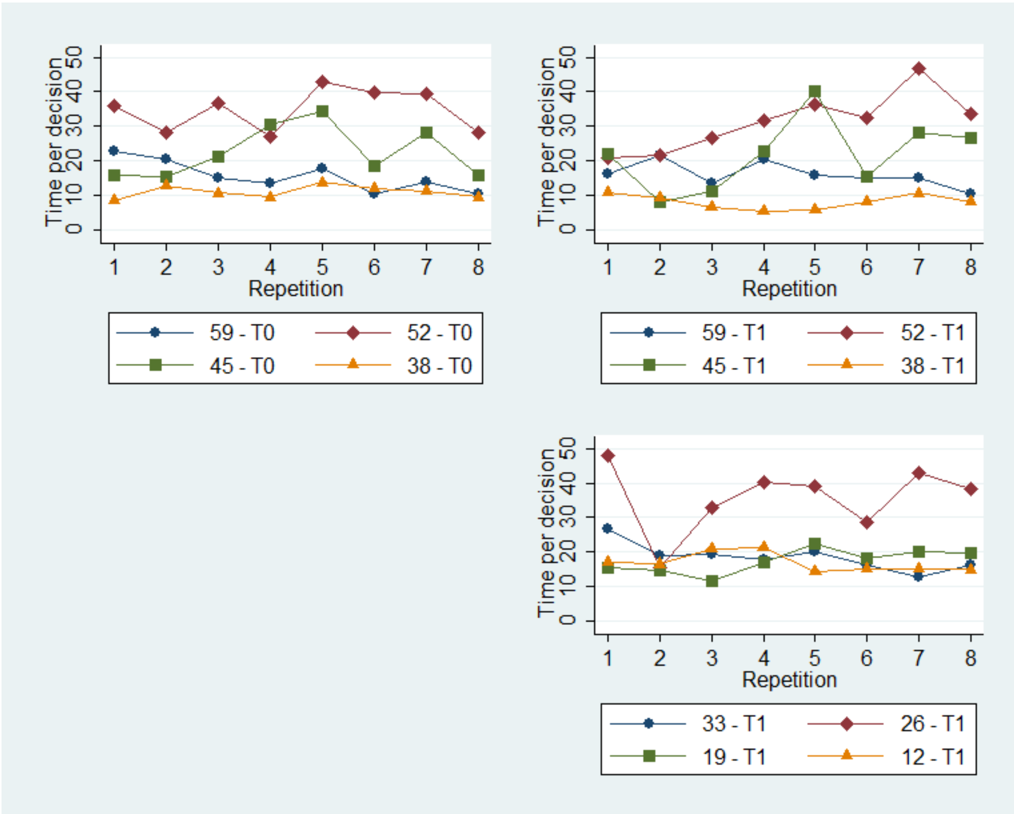

Figure 2.17: Average time for each decision in losing positions

as in winning positions. The same can be said of the fourth less distant losing position[38]. Most of the difference, between losing and winning positions, comes from the choices at the two losing positions in between. The time needed to reach a decision in the second closest losing position to the prize increases in the first repetitions reaching averages of more than forty seconds, and giving slight signs of decrease in the last repetition. For prize P, the time per decision at the third closest losing position to the prize traces the previous, with a couple of repetitions of delay.

An immediate interpretation for those results is that the losing subjects realize they are doomed to lose and check carefully if that is indeed the case. There is no sign of this behavior further than three steps from a prize.

We can have a broader view of how the timing of decisions evolves throughout the game by looking at Figure 2.18, where the average number of seconds per decision is plotted against the set of positions, for both treatments. In $T0$, the graph remains flat, at around ten seconds, until position 40; it then shows a steep increase, reaching thirty

[38]And for all farther positions (not shown in figure 2.17).

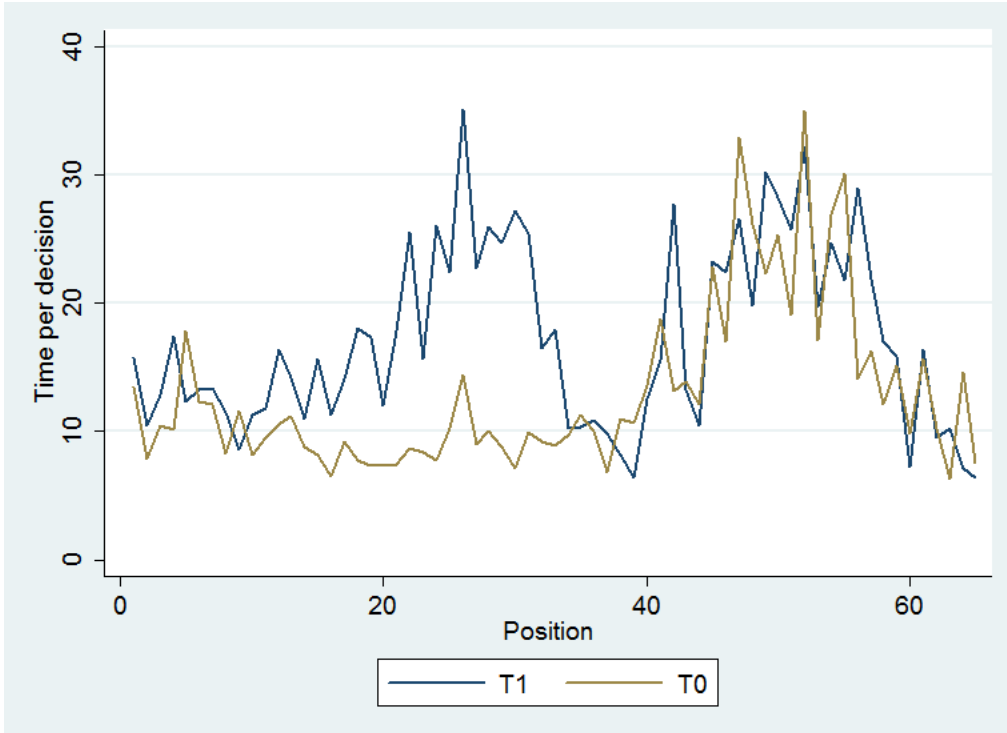

Figure 2.18: Average time for each decision over positions

seconds around 50, followed by a symmetric drop, and is back to around ten seconds by 60. $T1$ displays an identical pace after 40. However it shows another steep increase, followed by a symmetric drop, between positions 20 and 40. It is clear from the figure that both the magnitude and the length of both the early and the late raise, in the effort made to take a decision, are similar, if not identical.

It should be noted that equilibrium reasoning seems to imply that all the strategic effort is made at the beginning of the game, after which players just follow the planned strategy. Indeed, no theory that does not target specifically the dynamic structure of the game can predict such a pace for the timing of decisions. On the other hand LBI predicts exactly the pace we observe.

At the beginning of the game, the LF-game provides no information on how to choose a meaningful action, as the probability of winning seems unaffected by the current choice (without conditioning on future choices). As a consequence, there is no scope for in-depth consideration of different alternatives, and most players choose basically at random[39]. When, as the game proceeds, a prize is included in the LF-game, it becomes

[39]Recall that, indeed, the error rate in the first intervals are very close to those obtained by fully random players.

worth reasoning on which action to pick. This seems to happen on average between three and two intervals of distance from each prize, which is perfectly consistent with the figures shown for errors and claims, and matches the observation of one versus two peaks in $T0$ and $T1$, respectively.

Overall we find evidence in favor of every single hypothesis derived from LBI. Most players put effort in deciding only when a prize approaches. As a consequence the error rate for the small prize declines before that of the large prize and the former is claimed before the latter, on average. We also find evidence of some equilibrium players and of learning towards equilibrium play (in $T0$). There seems to exist a huge divide between those latter players and the rest. This is proven also by the answers given to a questionnaire we provided at the end of the experiment, including some questions about the chosen strategy. Around fifteen percent of the subjects identified the set of losing positions as the guide for their strategy. Half of them stated as "strategic" only a subset of the last three losing positions and around one third clearly stated they were trying to move as quickly as possible to the "hot-spots", close to the prizes.

Our results suggest there is a huge scope for further developing LBI, which we started doing in section 2.3.

2.7 Conclusion

The paper presents a general framework of out-of-equilibrium behavior in sequential games, limited backward induction, including a novel experiment aimed at testing its predictions.

Under LBI, the agents take decisions according to backward induction over restricted decision trees, the LF-game, the dimension of which is determined by the level of foresight of the players. The framework is flexible and applies to all sequential games with perfect information, including infinite games. As we let the level of foresight grow, in the limit, LBI mimic subgame perfection. On the other side, it encompasses perfect myopia as a special case.

The experiment is based on the race game, which, in our view, is the best setting to study strategic thinking in sequential games. A small trap prize, off the equilibrium path, and the possibility to claim prizes, allow us to gather a number of new insights. As already noted by Gneezy et al. [2010], we find that backward reasoning is the main cognitive procedure used by the subjects. However this procedure does not proceed backwards from the terminal histories, but is rather routinely performed on the few stages that are closer to the current decision node, as predicted by LBI.

This conclusion is supported by the analysis of errors, claims and timing of decisions.

Effort in decision making increases only when a prize is approached. The path to the trap prize is discovered and pursued before that to the final prize. The trap prize keeps the subjects off the equilibrium prize longer than in the base game.

Beyond its clear success in this experiment, LBI can explain aggregate behavior in other contexts, such as, for example, the centipede game, and sequential bargaining.

LBI type of reasoning bares important messages for real life decisions that include planning or anticipating the choices of other people in the future. First, the longer the chain of backward reasoning needed to reach the optimal solution, the more likely errors are and the more experience is needed to reach it. Second, a myopic bias may emerge in situations where individual plans involve choices that have different consequences at different points in time. This may be misunderstood as stemming from low discount factors, despite being due only to limited strategic thinking. These remarks have relevant consequences in many applications, such as, for example, retirement decisions and asset bubbles.

APPENDIX A

EXPERIMENTAL INSTRUCTIONS

Welcome to this experiment in decision-making. You will receive 3 Euros as a show-up fee. Please, read carefully thesse instructions. The amount of money you earn depends on the decisions you and other participants make. In the experiment you will earn ECU (Experimental Currency Units). At the end of the experiment we will convert the ECU you have earned into euros according to the rate: **1 Euro = 10 ECU**. You will be paid your earnings privately and confidentially after the experiment. **Throughout the experiment you are not allowed to communicate with other participants in any way. If you have a question please raise your hand. One of us will come to your desk to answer it.** [Between square brackets, we report the instructions specific to T1]

The game

- You will play a game with two players, *P* and *Q*.
- The players decide **sequentially**: they will take turns, one after the other. Each decision consists in a number of steps, between 1 and 6 (included).
- You will start at **position 1**. *P* is the first to decide.
- At the beginning, *P* chooses a number of steps between 1 and 6. Summed to the initial position, those steps determine a new position (example: *P* chooses 3; new position = 1+3 = 4).
- Then *Q* chooses a number between 1 and 6. those are summed to the position reached by *P* (example, follows: *B* chooses 5; new position = 4 + 5 = 9). And so on.
- The game ends when one of the players reaches position 66 with his decision.
- The players are always informed of the current position.

Prizes

- [When a player reaches **position 40** with his choice, he obtains the **prize A,** valued 30 ECU].
- When a player reaches **position 66** with his choice, he obtains the **prize [B], valued 100 ECU**.
- At any time you can claim you are going to win the prize [A or the prize B; you are allowed to claim both prizes].
- If a player obtains the prize he has claimed, he earns, **on top of the prize**, a number of ECU equal to the **difference between 66 [the position of the prize] and the position where he has declared to win it** (example: P declares at position 60 he is going to win the prize, and then wins; he receives 6 ECU on top of the prize [P declares at position 35 he is going to win prize A, and then wins; he receives 5 ECU on top of the prize]).
- If a player does not win a prize he has claimed, he gets a fine worth 20 ECU.
- When a player declares he is going to win [a prize], his opponent is not informes and can himself declare he is winning [the same prize].
- The number of ECU earned are the sum of the prize[s] and the adjunctive ECU obtained, minus the fine[s].

Structure of the experiment

- You will play 8 rounds of this game.
- You will play alternatively as player P and Q; this means you will choose alternatively as first and second.
- In every new round you will play agianst a **new** partner, chosen at random between the other participants.
- You will never play twice with the same partner.
- Two of your opponents will never play one against the other.

Earnings

- **Only one out of the eight rounds will be paid to you**.
- At the end of the experiment, one number between 1 and 8 will be selected at random by the computer, and the corresponding game will be paid.
- You will be informed of the chosen game, of your final payoff in ECU and of the corrosponding Euro.

Concluding remarks

You have reached the end of the instructions. It is important that you understand them. If anything is unclear to you or if you have questions, please raise your hand. To ensure that you understood the instructions we ask you to answer a few control questions. After everyone has answered these control questions correctly the experiment will start.

CHAPTER 3

AVAILABILITY OF INFORMATION AND REPRESENTATION EFFECTS IN THE CENTIPEDE GAME[1]

Abstract

The paper presents the results of a novel experiment testing the effects of environment complexity on strategic behavior, using a centipede game.

Behavior in the centipede game has been explained either by appealing to failures of backward induction or by calling for preferences that induce equilibria consistent with observed behavior. By manipulating the way in which information is provided to subjects we show that reduced availability of information is sufficient to shift the distribution of take nodes further from the equilibrium prediction. Similar results are obtained in a treatment where reduced availability of information is combined with an attempt to elicit preferences for reciprocity, through the presentation of the centipede as a repeated trust game.

Our results can be interpreted as cognitive limitations being more effective than preferences in determining (shifts in) behavior in our experimental centipede game. Furthermore our results are at odds with the recent findings in Cox and James [2012], suggesting caution in generalizing their results. Reducing the availability of information may hamper backward induction or induce myopic behavior, depending on the strategic environment.

JEL classification: C72, C73, C91

Keywords: Centipede; Backward Induction; Representation Effects.

[1] Joint work with Paolo Crosetto (Max Planck Institute of Economics, Strategic Interaction Group, Jena)

3.1 Introduction

The effects on observed behavior of apparently superficial changes in presentation are generally referred to as framing effects. Their existence[2] suggests that the game agents play is hardly ever identical to the canonical representation assumed by the experimenter. There are two layers of the subject's representation that can be affected by those changes: in some cases, an institutional format may elicit preferences that another does not; in others, the institutional format affects the players' understanding of the structure of the game. In terms of extensive form games, utilities only are affected in the former case, the game form and, as a consequence, utilities in the latter. Obviously, both mechanisms may be at work at the same time.

We perform two institutional manipulations on the centipede game to gather insights on the commonly observed patterns of behavior in this game. In particular, by manipulating the presentation of information about payoffs, we achieve a preference-neutral and a preference-non-neutral variation on the standard game, which we use to identify what is effective in shifting aggregate behavior in the game, distinguishing between preference-related and cognitive factors. As our manipulated institutional formats are more complex than the standard format, we can isolate the effects on behavior of (marginal) increases in complexity in a simple sequential game.

The centipede game [Rosenthal, 1981] has attracted experimental investigation mainly due to its counterintuitive theoretical prediction. The original centipede game is a two-player, finite sequential game in which the subjects alternate choosing whether to end the game ("take") or to pass to the other player ("pass"). The payoff from taking in the current decision node is greater than that received in case the other player takes in the next one, but less than the payoff earned if the other player were to pass as well. The player making the final choice is paid more from taking than from passing, and would therefore be expected to take. Iterating this argument, backward induction leads to the unique subgame perfect equilibrium: the game is stopped at the first decision node.

Starting from the first experimental evidence [Fey et al., 1996; McKelvey and Palfrey, 1992], studies have found that players fail to comply with this extreme unraveling prediction, even after a number of repetitions.

Probably due to the combination of the simplest possible sequential structure, a clear-cut equilibrium prediction, and a still rich and subtle strategic environment, the centipede has become a workhorse for theory testing. As simple as it may seem, the identification of the motivations underlying behavior in the centipede turns out to be a chal-

[2]See, in general, Tversky and Kahneman [1981]; for an application to games, see Devetag and Warglien [2003]; Kreps [1990].

lenging task. The list of possible reasons why players may take actions that diverge from subgame perfect equilibrium turns out to be long and often twisted.[3] Broadly speaking, we can identify different families of explanations regarding the roots of deviations from equilibrium, depending on whether they rely on preferences [e.g., Dufwenberg and Kirchsteiger, 2004], on bounded strategic thinking [e.g., Kawagoe and Takizawa, 2012; Palacios-Huerta and Volij, 2009][4] or on a combination of the two [e.g., Maniadis, 2011; McKelvey and Palfrey, 1992; Zauner, 1999].[5]

In a recent paper, Cox and James [2012] found that a strategically irrelevant manipulation of the institutional format by which two, otherwise identical, centipede games are represented can have a significant impact on behavior. In particular, they found that framing the game as a sequential auction, where the players are informed about the payoffs if buying in the current node but have to compute the payoffs for future stages by themselves, triggers an unprecedented proportion of behavior observationally equivalent to subgame perfect strategies. They interpret this finding as an instance of myopia arising from making information about the game less available.

Others, preference-non-neutral manipulations that were used on different games can also be applied to the centipede game. In an early example of preference-eliciting institutional manipulation, Evans [1966] and Pruitt [1967] presented results on the decomposed Prisoner Dilemma. In their experiments, a standard PD is compared to a decomposed version where one player choosing a strategy directly determines an allocation to both players, which is then summed to the allocation chosen by the other player. In general, the latter presentation achieves significantly higher cooperation: the presentation of the PD as a sort of simultaneous trust game, which makes the give-and-take nature of the game salient, elicits preferences, most likely related to reciprocity, that the traditional version does not.

We exploit the two abovementioned institutional formats to investigate the role of preferences and cognitive limitations in shaping taking behavior in the centipede game. In our baseline standard treatment (*Tree*), the players are shown the standard game tree displaying the final payoffs at every terminal history. The first manipulation (*Formula*) is preference neutral and traces the *Clock* treatment in Cox and James [2012]: the players are informed only about the progression of the payoffs throughout the game; as they proceed they are told the final payoffs were the game to end at that node, but have to compute the

[3] Levitt et al. [2011] provide a nice example of such a list. A partial attempt to disentangle those reasons can be found in Atiker et al. [2011].

[4] This category actually includes departures from common knowledge of rationality (or incorrect beliefs) and correct beliefs but an imperfect best reply.

[5] Other relevant papers featuring theoretical and experimental analyses on the centipede are Nagel and Tang [1998]; Ponti [1996]; Rapoport et al. [2003].

final payoffs for future decision nodes (if they so wish). The second manipulation (*Decomposed*) is identical to that of Pruitt [1967]: the payoffs are decomposed in stage payoffs, so that every pass entails some losses for the passing player and some gains for the other player. To compute the final payoffs, the players need to sum up the stage payoffs to each terminal history. As before, they are informed about the final payoffs, were the game to end at the current node. The final payoffs, the rules of the game and their description, together with all other details of the design are identical across treatments, and exactly the same amount of information is available to the players, although presented in a different way.

Following Cox and James [2012], *Formula* could elicit myopia due to information being less available: facing some higher complexity, the players would focus only on the closest decision nodes and not consider the possible gains from passing. Note however that, in principle, the more complex environment could trigger an opposite effect: a player could find it harder (or more costly) to perform backward induction and pass through more nodes as she fails longer to recognize the strategic structure of the game.

Those considerations apply as well to treatment *Decomposed*. On top of that, *Decomposed* could elicit preferences for reciprocity as the game is represented as a repeated trust game. If this is the case, and assuming additivity for the preference and the cognitive effect, then players should take later in *Decomposed* with respect to *Formula*.

We find two main novel results. With respect to the base treatment, both institutional transformations achieve later take nodes which are furher away from the theoretical prediction: apparently, making information less available makes it more difficult for subjects to understand the strategic structure of the game,[6] with no evidence of myopia. We observed no difference between the preference-neutral *Formula* and the preference-non-neutral *Decomposed*: though we cannot properly separate cognition- and preference-based effects in *Decomposed*, it looks as if preference elicitation is ineffective in pushing the take nodes further away.

More notably, the first result is sharply at odds with results in Cox and James [2012]: although we perform the same manipulation, our subjects take later where theirs take earlier. We interpret this gap as stemming from relevant differences in the base game: their centipede is extremely competitive and already complex in the tree format, whereas ours is a more standard, less competitive and simple game. Thus it looks reasonable that a reduction in the availability of information induces, in the former environment, no use of the information about distant nodes, resulting in myopic early takes, and only hampers backward induction (or reduced use of the information about distant nodes), in turn resulting in late takes in the latter. This apparent conflict suggests cautiousness

[6]Or, change the beliefs obout the others' ability to understand the strategic structure of the game.

about generalizations of the effects of complexity on strategic behavior and elicits new fascinating research questions on the topic.

The paper is organized as follows. In the next section, we describe the experimental design and present our main hypotheses. The actual implementation of the design in the lab is detailed in Section 3.3. Section 3.4 describes the results, and Section 3.5 concludes.

3.2 Experimental Design

We implement a twelve-legged centipede, with actions labeled "Stop" and "Continue". Terminal histories are ordered and assigned a number between 1 and 13 (Stop at first node: 1; ... ; Always continue: 13). The aggregate payoff at each terminal history is worth 5 times the corresponding number; the player choosing "Stop" gathers $\frac{4}{5}$ of the total value, while the opponent gathers the remaining $\frac{1}{5}$ (see Fig. 3.1).

The length and the linear increase in the joint payoffs distinguish our game from the most exploited experimental centipedes.

The length of the game is meant to allow more room for responses to relatively minor treatment variations to emerge and to enhance the relevance of sequential reasoning.

We chose an arithmetic progression with respect to the more common, geometric one [as in McKelvey and Palfrey, 1992] for two main reasons. The first, specific to our design, is that a linear increase (as a function of the decision node) makes the underlying formula easy to convey also to subjects with potentially low numeracy skills.[7] The second, more general, is to avoid the unpleasant choice the experimenter faces with geometric centipedes between a very short game, an exchange rate that makes initial payoffs economically irrelevant, or a geometric factor that makes the progression at first nodes virtually flat.[8] In our setting, it is possible to keep the range of payoffs in line with the literature while providing economic relevance to choices at all decision nodes, including the first ones. Our choice allows us to show payoffs directly in euro, with the first decision node entailing a payoff of $(4, 1)$ euro for the player controlling the node and the opponent, respectively, and a payoff of $(52, 13)$ euro if both players choose "Continue" at all decision nodes.

In this general framework, we implement three different ways of conveying the payoff information:

---

[7] A pilot featuring a geometric progression was run, but proper understanding of the treatment "Formula" proved difficult, undermining the comparability of the results. All data and materials are available upon request.

[8] Rapoport et al. [2003] avoid the problem for a limited number of subjects in their "high stakes" treatment, bearing the risk of a potentially explosive budget. More commonly, the increase in payoffs at the first decision nodes is in terms of cents.

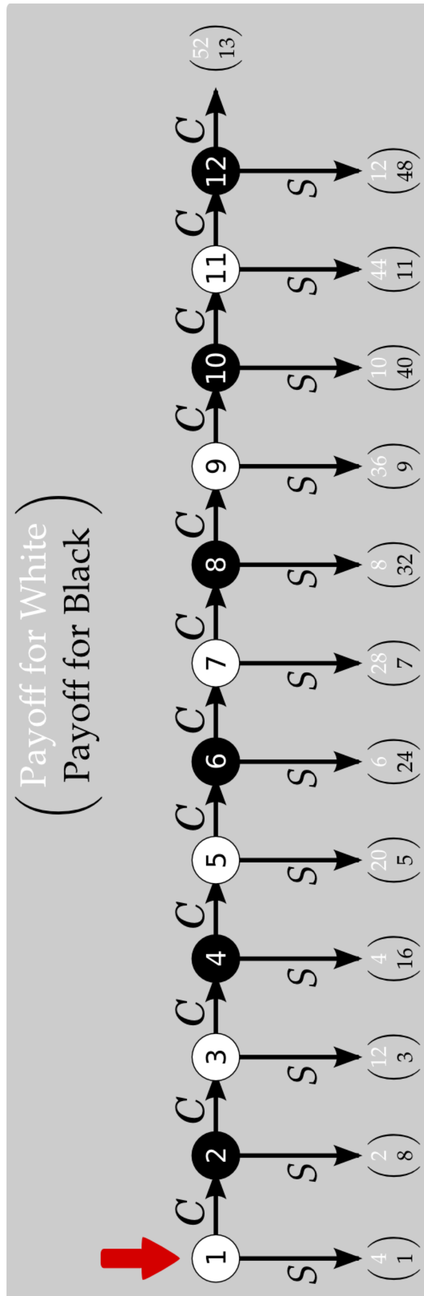

Figure 3.1: The game representation in the *Tree* condition, payoffs in euro

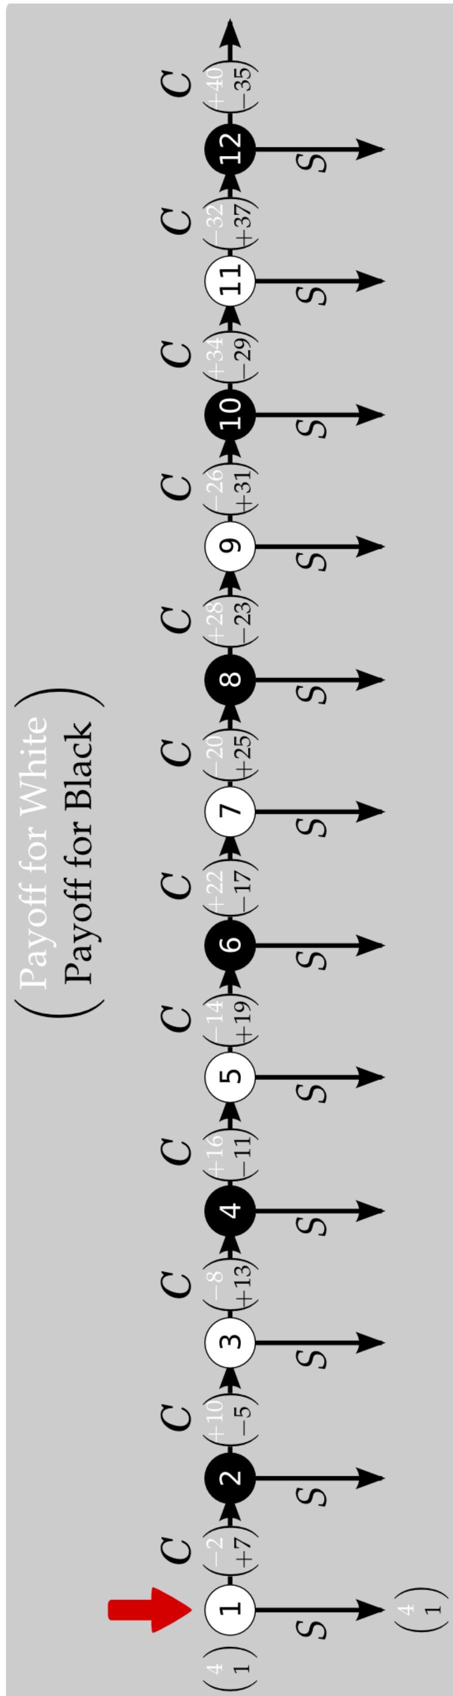

Figure 3.2: The game representation in the *Decomposed* condition, payoffs in euro

Tree: as is standard in the literature, subjects are shown the game tree that reports at each terminal history the final payoffs accruing to both players. The tree, as shown to subjects, can be seen in Figure 3.1. This condition replicates the standard way[9] to convey the centipede game in experiments.

Formula: the subjects are not shown the tree but only the formula to compute the payoffs. In particular, subjects are told that, when one player chooses "Stop", she earns four times the number of the current decision node, while the other earns an amount equal to the number of the decision node.

Decomposed: the subjects are shown the game tree, but, instead of final payoffs, the stage-payoffs, i.e., the variations with respect to the currently earned payoff, are shown for each decision node. The tree, as shown to the subjects of the *Decomposed* condition, can be seen in Figure 3.2.

Thus the *Tree* and the *Decomposed* conditions conveyed infomation using a comprehensive visual representation of the game.[10] In the *Formula* condition, all information was conveyed by means of words.[11]

It should be noted that the players were given exactly the same amount of information under all treatment conditions, the only difference being its availability: in the *Formula* and *Decomposed* conditions, players have to compute endgame payoffs for future stages on their own. Given our payoffs, this step is, however, minimally demanding: it requires the computation of the four-times table or of simple integer sums, respectively.

Beyond being less available, the *Decomposed* structure presents the payoffs with a *give-and-take* frame, underlying the intrinsic nature of repeated trust game of the centipede and possibly eliciting reciprocal behavior.

The proposed game is the same in all treatments and the presentation variations are minimal. Considering these features, combined with the well-known learning dynamics in the centipede game, we opted for a pure between subjects design.[12]

[9] In particular, the figure is identical to that in Palacios-Huerta and Volij [2009] and Levitt et al. [2011].

[10] The subjects, identified by color, were shown the full length of the tree and (final or stage) payoffs at each node. Moreover, every decision node was numbered and intuitively assigned to a player/color. The images in Figures 3.1 and 3.2 were both given to the subjects in a printed version as part of their instructions and presented on screen at every decision node; in the screen version, the red arrow would move to indicate the current decision node; moreover, all past decision nodes would gray out on screen. Both active and inactive players were shown the same set of pictures, the difference being that the inactive player faced no choice but was reminded of the choice that the matched player was considering at that moment.

[11] The part of the screen regarding the current decision node was identical to the *Tree* condition; with respect to the latter, a description of the rules of the game (including the formulas to compute the payoffs) took the place of the visual representation.

[12] The between subjects is a robust choice if the samples for the two treatments do not differ in underlying characteristics. This can be guaranteed either by a high number of subjects, or, alternatively, relying on subject's randomization. We chose to enlist a mid-sized sample but introduced several controls that allowed us to check whether a set of relevant subject charachteristics (age, gender, risk and trust attitudes) showed

Within each experiment, subjects repeated the game 12 times in a perfect stranger matching, implemented by using the turnpike protocol. This matching allows us to assure subjects that they will never play the same partner twice, and that their partners will never play one another, thus ensuring absence of contagion effects. Repetitions were meant to allow for learning, though still focusing on first response behavior. We also chose to keep the roles fixed across repetitions to restrict the confounding effects of identification.

In the following we formuilate our hypothesis. A first set of them regards the effect of the availability of information, thus comparing the *Formula* and *Decomposed* conditions to the *Tree* condition. As mentioned in the Introduction, one possibility is that the reduced availability of the consequences of passing in the *Formula* and *Decomposed* treatments may trigger myopic behavior (or beliefs of myopic behavior): subjects do not use information about efficiency gains and focus on immediate decision nodes, taking as early as possible. On the other hand, information being less available may not induce subjects to disregard it but only hamper their ability to reason backwards.[13] If that was the case, we would observe later "Stop" decisions.

Hypothesis 1.1. In conditions *Formula* and *Decomposed*, the subjects choose "Stop" earlier than in the *Tree* condition, due to myopia.

Hypothesis 1.2. In conditions *Formula* and *Decomposed*, the subjects choose "Stop" later than in the *Tree* condition, due to hampered backward induction.

Besides these cognitive effects, the *Decomposed* treatment should elicit more reciprocal behavior, resulting in the subjects passing longer in the game with respect to the *Formula* condition.

Hypothesis 2. In condition *Decomposed*, the subjects choose "Stop" later than in the *Formula* condition, due to enhanced reciprocity.

3.3 Experimental Procedure

The computerized experiment was run in Jena in June 2012, involving 210 subjects distributed over 8 experimental sessions. Seventy-two subjects took part in the baseline *Tree* sessions; a further 74 subjects participated in the *Formula* and 64 in the *Decomposed* conditions. The experiment lasted about 1 hour, and average payoff across all sessions and

any particular bias across treatments.

[13] Our manipulations are close to cognitive load experiments [Cappelletti et al., 2011; Swann, 1990; Shiv and Fedorikhin, 1999, e.g.] in that we manipulate the level of cognition imposing or not imposing (computational) burdens on otherwise identical tasks. The hypothesis that reducing the availability of information may reduce subjects' strategic ability to reason backwards is consistent with the results in this literature, as reported by Devetag and Warglien [2003] and Duffy and Smith [2012].

conditions amounted to 11.8 euro, including a 2.5 euro show-up fee.

All sessions followed an identical procedure. After subjects were allowed into the lab, instructions were read aloud and extra time was given to the subjects to go through them on their own. Then all subjects had to correctly answer a set of control questions before being allowed to proceed. The number of mistakes recorded in the questions, and the time needed to clear the control questions screen, were both recorded and used as an objective measure of the complexity of the treatment. During this phase, subjects could – and many did – ask help from the experimenters with going through the control questions.

After all subjects had cleared their control questions, the experiment started. Subjects were randomly assigned to their roles ("White" or "Black"), randomly matched, and proceeded to play the game. The same game was repeated 12 times, in a perfect stranger matching design. The pairs were allowed to proceed each at their own pace within the 12 decision nodes of the game but had to wait for all the other pairs between repetitions.

After completing the 12^{th} repetition, subjects were paid according to the results of a randomly drawn repetition, and were asked to fill in a questionnaire. We gathered qualitative information about the expectations from the game and the opponent, the strategy followed, and the belief on the opponent's behavior. Moreover, we elicited self-reported quantitative measures of trust and risk aversion [using the SOEP German Panel trust and risk questions. For the risk question, see Dohmen et al., 2011] and of the perceived complexity of the task.

The experiment was conducted in German. The English version of the experimental instructions is available in Appendix B.[14]

3.4 Results

3.4.1 Aggregate behavior

Consistently with the bulk of the literature on the centipede game, the players did not adhere to the Subgame Perfect Nash Equilibrium but played on into the game. Moreover, there was some unraveling of the game: in all conditions, the average endnode became significantly lower with the repetitions (WRST, repetition 12 *vs.* repetition 1: *Tree*, *Formula* and *Decomposed*, all p-values < 0.001). This trend is monotone and qualitatively similar in all conditions,[15] with the partial exception of *Decomposed* where unraveling is reversed in the last two repetitions, in which the average endnode slightly (though not significantly)

[14] The original German instructions, along with the experimental software [developed using zTree, Fischbacher, 2007] and the raw data from the experiment, are available upon request.

[15] Average reduction by repetition: 0.21, 0.24 and 0.22 in *Tree*, *Formula* and *Decomposed*, respectively

increased.[16] The average endnodes by repetition and treatment are summarized in Figure 3.4; the distribution of endnodes in the first and second 6 repetitions for all conditions is instead represented in Figure 3.3.

Ã—	Repetition											
Ã—	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
Tree	4.54	4.43	4.24	4.16	3.94	3.67	3.43	3.13	3.03	2.76	2.62	2.27
Formula	5.36	5.27	4.94	4.67	4.39	4.22	3.89	3.67	3.5	3.28	3.02	2.69
Decomposed	5.43	5.15	5.06	4.65	4.62	4.41	3.84	3.62	3.46	3.18	3.22	3.31

Table 3.1: Mean endnode by treatment and repetition

Result 1: *In all conditions, the players do not adhere to the SPNE, reaching, on average, slightly more than a third of the game in the first stages. We observe slow but constant unraveling of the game toward the SPNE as repetitions are played.*

It should be noted that, with respect to the bulk of existing literature, the distance from equilibrium in our experiment is, on average, relatively low. Although it is hard to perform a direct comparison, this is consistent with Rapoport et al. [2003], in which imposing relatively high stakes from the first decision nodes resulted in closer-to-equilibrium play.

3.4.2 Treatment effects and test of hypotheses

In the following, we analyze treatment effects by making use of the hypotheses laid out in Section 3.2.

First both *Formula* and *Decomposed* result in *later* take nodes with respect to the baseline *Tree*. We hence find support for a lower incidence of backward induction (Hypothesis 1.2) and have to reject instead that choices are driven by myopia (Hypotheses 1.1).

When comparing the *Tree* and *Formula* conditions, we find a significant and strong treatment effect. In the *Formula* condition subjects stop the game about $\frac{2}{3}$ of an endnode later than in the *Tree* condition. This is true both when computing the overall mean across all repetitions (4.08 vs. 3.52, WRST p-value < 0.001) and when considering each single repetition: the average endnode of *Formula* is stably more than half a stage above *Tree* in each period, though not always significantly different (WRST, p-value < 0.05 in all but repetitions 1,11 and 12).

[16]More tests: WRST, repetition 6 *vs.* repetition 1: *Tree* p-value $= 0.024$, *Formula* p-value $= 0.014$, *Decomposed* p-value $= 0.198$; repetition 12 *vs.* repetition 7: *Tree* and *Formula*, p-values < 0.001, *Decomposed* p-value $= 0.06$. As discussed below, *Decomposed* generally shows a higher variance in behavior, which explains why significance is harder to achieve there.

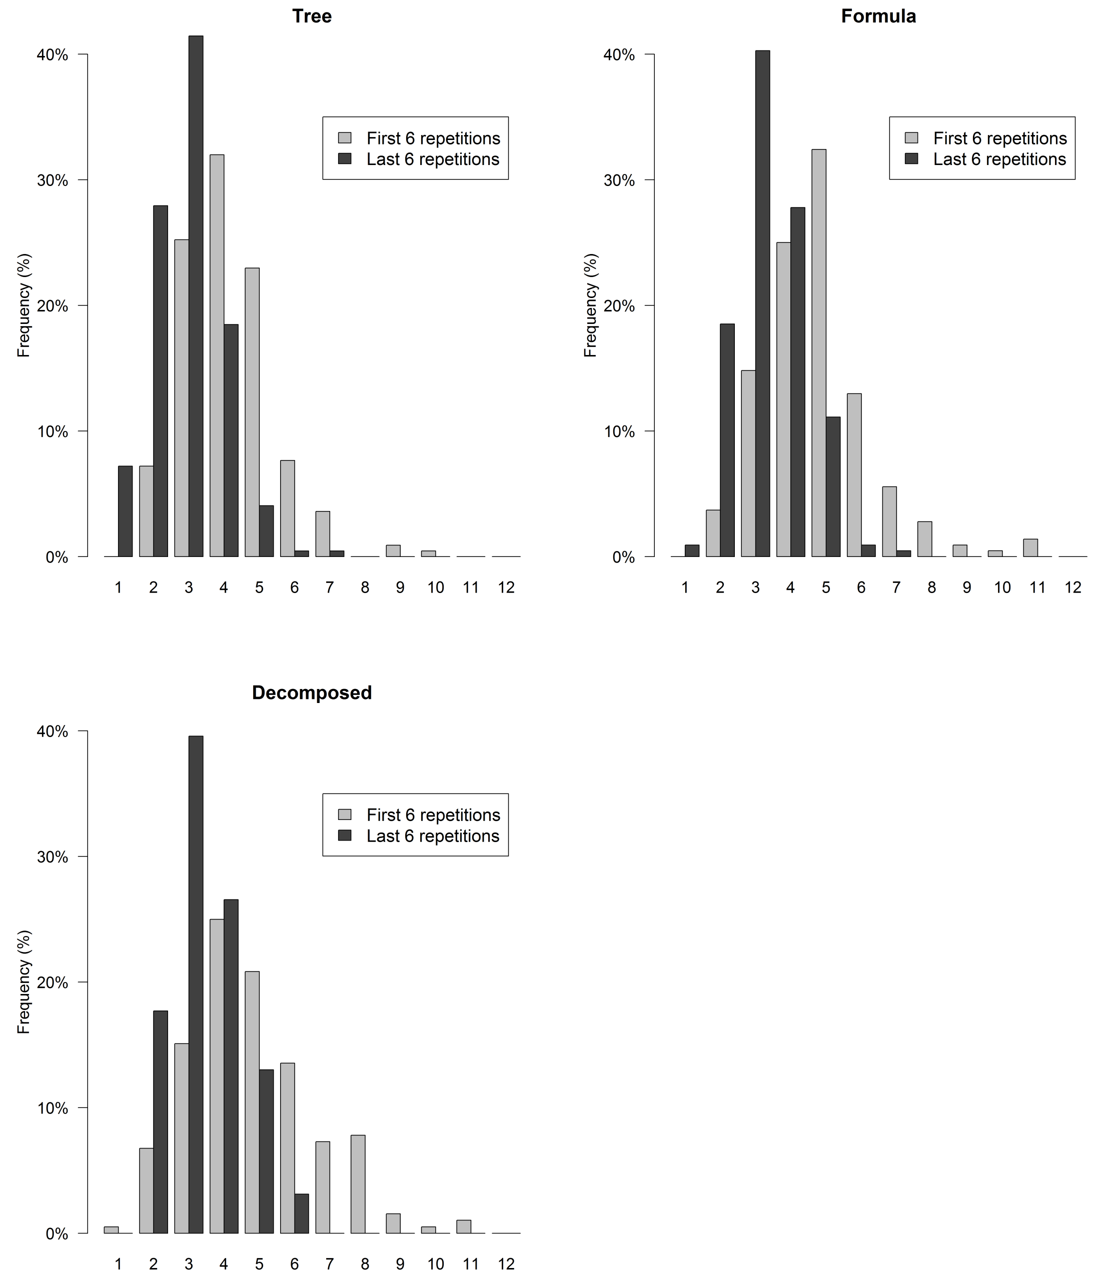

Figure 3.3: Endnode in the first and second half, by treatment

Moreover, a paired histogram of the distribution of endnodes in both conditions (Fig. 3.5) readily shows that the distribution for the *Formula* condition is shifted to the

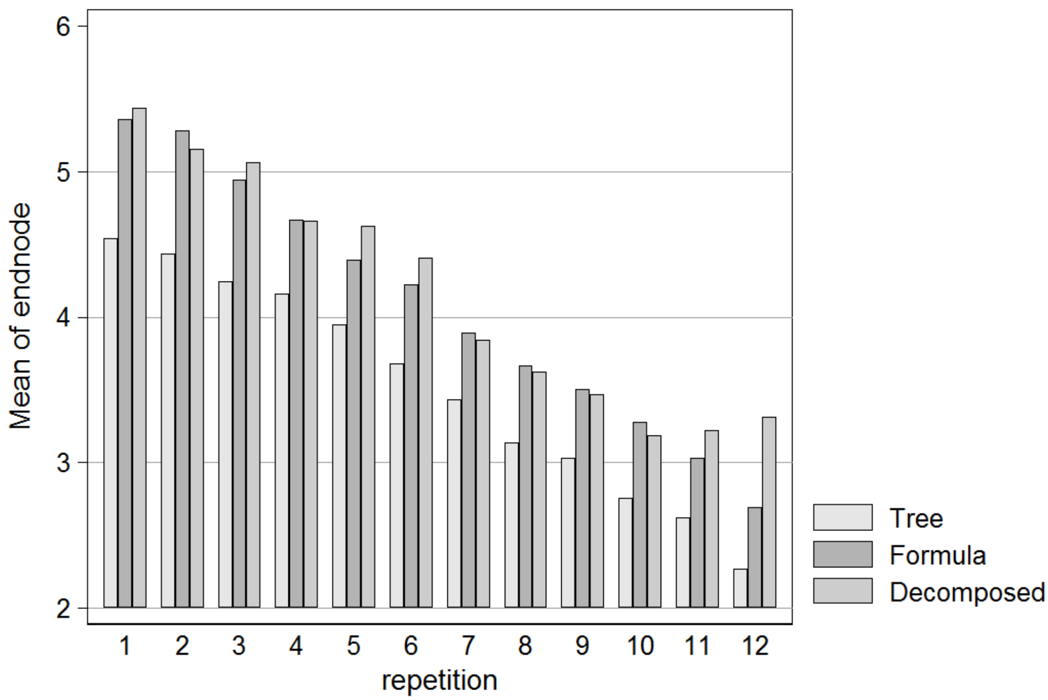

Figure 3.4: Mean endnode by treatment and repetition

right with respect to the *Tree* distribution; testing equality in distribution (Kolmogorov-Smirnov 2-sample test, p-value < 0.001) confirms the significance of the difference.

Result 2: *In the Formula condition, subjects exit significantly later than in the Tree condition.*

The comparison between the *Tree* and *Decomposed* conditions reveals a similar pattern to the one between *Tree* and *Formula* but with slightly less statistical significance. This is due to the fact that the variance of behavior is much higher in the *Decomposed* condition, especially in the first repetitions (see Fig. 3.3), possibly reflecting the higher self-reported and objective difficulty encountered by subjects in understanding the game (see below).

In the *Decomposed* condition, the average endnode is about $\frac{2}{3}$ of an endnode higher with respect to the *Tree* sessions, considering the overall average (4.17 vs. 3.52, WRST, p-value < 0.001;), but it is statistically significantly higher only in repetitions 3,5,6,11, and 12.

A paired histogram of the distribution of endnodes (Fig. 3.5) readily shows that the distribution in *Decomposed* stochastically dominates that in *Tree*; this is confirmed by a KS test (p-value < 0.001).

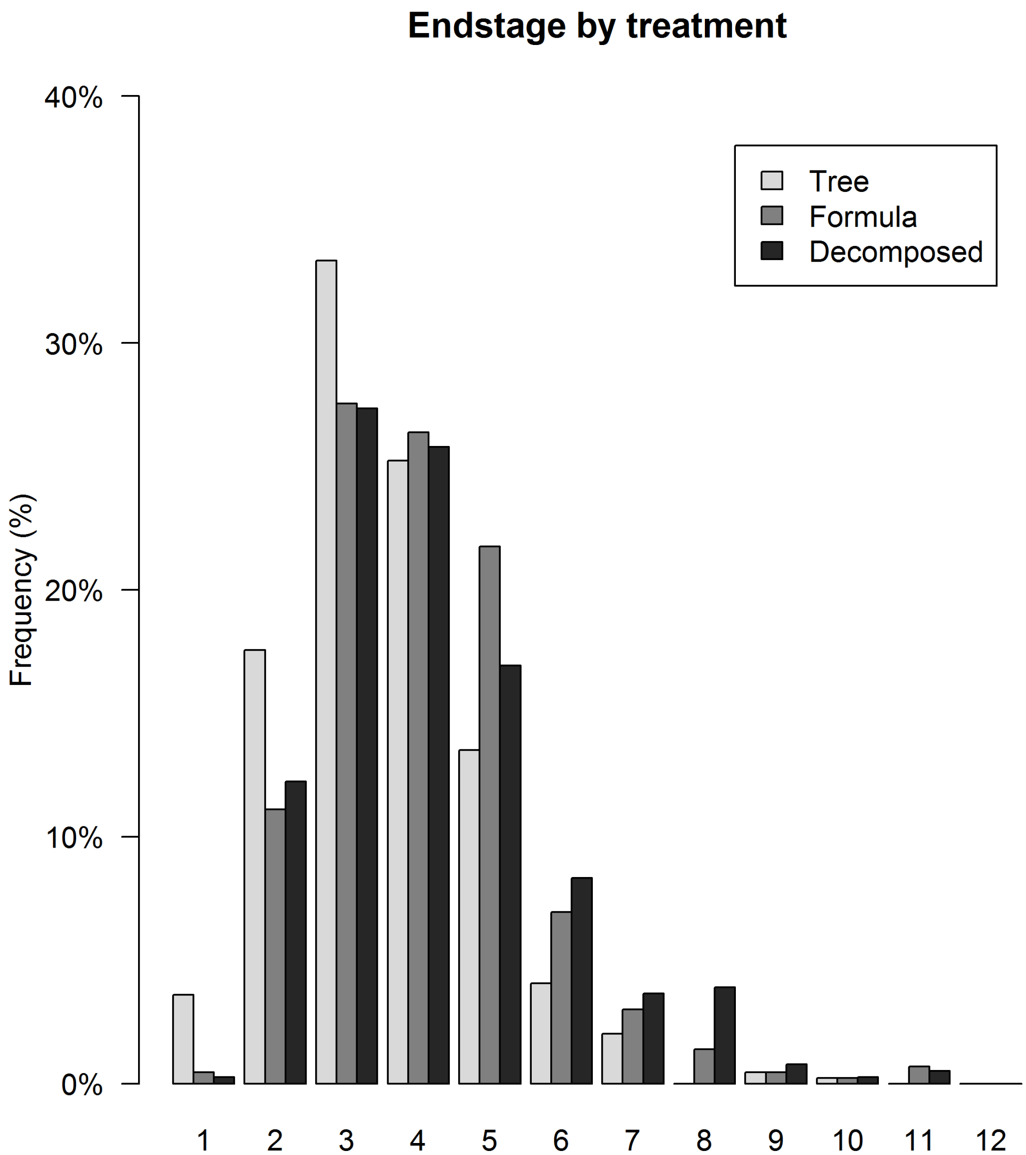

Figure 3.5: *Tree vs. Formula*

Result 3: *In the Decomposed condition, subjects exit significantly later than in the Tree condition.*

We find no support for Hypothesis 2: there is no statistical difference between the

Formula and *Decomposed* conditions, once we discount the higher initial variance of the *Decomposed* condition. The mean endnode across all repetitions is not statistically different (WRST, p-value $= 0.796$). Moreover, the endnode is not statistically different in any of the single repetitions (WRST, all p-values > 0.356) apart from the last, in which unraveling stops in *Decomposed* but continues in *Formula* (WRST, p-value $= 0.006$).

In distribution, the two conditions are not statistically different (KS, p-value $= 0.728$).

Result 4: *The Formula and Decomposed conditions do not differ statistically.*

In Figures 3.3, 3.4 and 3.5 we observe an impressive similarity between *Formula* and *Decomposed*. Nevertheless, in the latter we observe a higher variance, concentrated especially in the first repetitions. This is likely related to the higher level of perceived complexity, as documented next.

3.4.3 Controls

The above results could be due to systematic differences in the composition of subjects taking part in the between treatments. Moreover, the questionnaire answers and the statistics gathered on the control questions allow us to see if and to what extent the treatment differences can be ascribed to comprehension problems. This section addresses these issues.

First, treatments did not differ for all the characteristics that we controlled for (age, gender, attitudes toward risk and trust). Treatments did not differ in terms of trust (WRST, all p-values > 0.12) and risk attitudes (WRST, all p-values > 0.08) of the subjects involved. The composition of the treatment also did not differ statistically by gender (WRST, all p-values > 0.64) and age (WRST, all p-values > 0.38). Hence, the treatment effects cannot be said to depend on heterogeneity in the observed subjects' characteristics.

Result 5: *Participants in the different treatments do not differ, on average, by age, gender, attitudes to risk, and indicators of trust in others and the society at large.*

In order to evaluate the complexity of each treatment, we both directly asked subjects to rate the perceived complexity and measured the number of errors in the answers to the control questions and the time spent completing the control questions screen. The game was significantly more difficult to understand for subjects in the *Decomposed* condition (Table 3.2), while there was no significant difference between *Formula* and *Tree* in both self-reported complexity (Wilcoxon Rank Sum Test, p-value $= 0.444$) and the number of errors (WRST, p-value $= 0.253$); in the *Tree* condition, though, subjects answered the control questions significantly faster than in the *Formula* condition (WRST, p-value $= 0.007$). On the other hand, *Decomposed* proved significantly more complex in all indicators with respect to both *Tree* (WRST p-values: complexity = 0.034, errors 0.068, time 0.000) and *Formula* (WRST p-values: complexity = 0.003, errors 0.005, time 0.000).

Ã—	N	Complexity (0-10)	Errors (num)	Time (sec)
Tree	74	2.32	0.51	104
Formula	72	2.44	0.55	148^{*}
Decomposed	64	$2.89^{*,**}$	$0.95^{*,**}$	$257^{*,**}$

significant with respect to: * row above; ** two rows above

Table 3.2: Self-reported and objective measures of complexity

Result 6: *The Decomposed condition is more difficult to understand than both the Tree and Formula conditions, taking into account both self-reported and objective measures of complexity.*

3.4.4 Discussion

Our results are small in magnitude but significant and robust, especially when compared to our minor treatment variations: our subjects are all playing exactly the same game, having the chance of experiencing it 12 times, but despite this, differences persist consistently across repetitions.

Result 2 shows that a simple reduction in the availability of information can shift take nodes further away from the equilibrium with no sign of convergence through repetitions. Cox and James [2012] found exactly the opposite, performing the same manipulation: their centipede game is presented either in tree format or as a sequential Dutch auction, where subjects know the current price and are informed about future price decrements. Their result is interpreted as an instance of myopia, i.e., not using information about future nodes, while we interpret our result as evidence of more limited backward induction, i.e., reduced use of information about future nodes.

The apparent conflict can be defused by considering differences in the base game. Cox and James [2012] use an incomplete information game which is strategically identical to a centipede game under any belief about the opponent's payoffs. Moreover, the player who does not take always earns a payoff of zero, while the increase in the payoff, for the player who takes is relatively low. Those elements build up a setting that is both extremely competitive (strict efficiency gains are not possible) and complex even in the standard tree format. Facing a further increase in complexity due to the reduced availability of information, subjects stop exploring the strategy space deep into the game and just "take the money and run."[17] The same effect is not granted under games that are cognitively

[17]The manipulation in Cox and James [2012] also includes a language shift: in the auction, in order to

less demanding and exert less competitive pressure, as it is the case in our centipede game. Subjects are still affected by reduced availability of information, as they find it harder to reason backward and reduce the depth of their strategic thinking. However, this results in later take nodes.

Both effects - i.e. myopic behavior and hampered backward induction - can be rationalized following LBI (see Section 2.3). Starting from a reasonably simple situation, an increase in complexity reduces the foresight of the agent, inducing him to pass longer through the game. As complexity increases, at some point, it may become costly even to retrieve information about the payoffs at distant node. In this case, it is the *sight* of the player that shrinks, which, in the centipede game, implies that the agent does not consider (some of) the efficiency gains that are possible. This, in turns, may induce him to act myopically.

If this interpretation holds, it suggests cautiousness in generalizing the effects of institutional format manipulations on strategic reasoning: behavior may react in different ways, depending on the underlying strategic environment. In particular, consistently with the results in Devetag and Warglien [2003], the observer should consider whether the game is complex enough for a marginal increase in the cognitive load to be able to trigger a shift to a simple heuristic (e.g., myopia) or just throw sand in the gearbox of strategic thinking.

Since in *Decomposed* we may be eliciting preferences for reciprocity, while reducing the availability of information, an immediate interpretation, combining results 3 and 4, is that cognitive limitations are more effective than preferences in shifting behavior in the centipede game. Indirectly, this would question interpretations of the results in the standard centipede game as driven by preferences, given that we know the same manipulation to shift behavior in games where preferences for reciprocity are relevant. However, we should be cautious with respect to this interpretation as it relies on a series of reasonable but additional hypotheses; namely that the effects of preference elicitation and reduced availability of information are additive and that preferences are not endogenously affected by marginal (pure) increases in complexity.

3.5 Conclusion

The failure of subgame perfect equilibrium in the centipede game has attracted a number of scholars, their explanations focusing either on cognitive limitations that hamper backward induction or on preferences that mandate different equilibrium strategies.

take, the player must “Acquire” a good at a certain “Price”, with the payoff being the difference between his private value for the good and the realized price.

In this paper, we made small institutional changes to a centipede game that vary the way in which information is provided to the subjects, performing a preference-neutral and a preference-non-neutral manipulation. We show that making information about future payoffs less available is sufficient, on average, to significantly delay the decision to take. As this can be attributed to a more limited ability to backward induct, this result supports the potential of cognitive limitations to determine behavior in the centipede game.

On the other hand, highlighting the repeated-trust-game nature of the centipede game, i.e., by presenting the payoffs in a *give-and-take* frame, has apparently no further effect.

Our results are starkly at odds with those in Cox and James [2012], where performing a manipulation similar to our preference-neutral one significantly anticipated the decison to take. Given that our baseline game widely differs from theirs – woth our game presenting a much simpler strategic environment – this conflict suggests that reducing the availability of information can hamper backward induction (i.e., cause reduced use of the information about some future nodes) or induce myopic behavior (i.e., cause nonuse of the information about some future nodes), depending on the circumstances. Exploring which factors lead to which of the two outcomes is an exciting research question to be explored by future work.

Appendix B

Experimental Instructions

In the following, the English instructions for condition “Tree” are reported. In brackets are detailed the changes made to adapt the instructions to condition “Formula” (F) and “Decomposed” (D). The original German instructions are available upon request.

Introduction: common to all conditions

Welcome and thanks for your participation to this experiment. Please remain silent and switch off your mobile phone. Please do not talk and raise your hand if there are any specific questions during the experiment: an experimenter will come to your place and answer your concerns individually. If you violate these rules, we will have to exclude you from the experiment and all payments.
You receive a 2.5 euro show-up fee for taking part in the experiment. Please read the following instructions carefully. Prior to the experiment, you will have to answer a few questions testing your comprehension of these instructions. Please note that, for convenience, the instructions are written in male gender, but refer to both genders equally.
During the experiment you are going to use ECU (Experimental Currency Units). At the end of the experiment, earned ECU will be converted into euros at an exchange rate of

1 euro = 1 ECUs.

You will take part in a game played by two persons, white and black. You will be randomly assigned the role of white or black, which you will keep for the whole experiment. The game consists of 12 ordered decision rounds (first round: round=1, ..., last round: round=12). The players play sequentially. When it is his turn to play, each player can choose between STOP and CONTINUE.

If a player chooses STOP, the game ends.

If a player chooses CONTINUE, the game continues, and the other player faces a choice between STOP and CONTINUE.

White plays first; if he chooses STOP, the game ends, but if he chooses CONTINUE, black is called to play and decide whether to STOP or CONTINUE, and so on. Thus each player has at most six choices, with white choosing at round 1, 3, 5, 7, 9, and 11 and black choosing at round 2, 4, 6, 8, 10, and 12. The sequence of choices continues until one player chooses STOP. If both players choose CONTINUE in every decision round the game ends at round = 13.

Payoff information: different across conditions

Tree

Below you can see a representation of the game. The game starts from the utmost left. The color of the circles identifies which player has to decide; the numbers in the circle represent the decision round; the numbers in the brackets represent the final payment, in ECU, obtained by each action. In white you see the payoff of white, in black the payoff of black.

[The image shown to the subjects is reproduced above in Figure 3.1]

Formula

When a player chooses STOP at round = r, the value for him is 4 times the current round, that is:

$$V_{STOP} = 4 \cdot r$$

The value for the other player is 1 times the current round, that is

$$V_{OTHER} = 1 \cdot r$$

Decomposed

Below you can see a representation of the game. The game starts from the utmost left. The color of the circles identifies which player has to decide; the numbers in the circle represent the decision round; the numbers in the brackets represent the change in payments, in ECU, on top of what you have already earned, resulting from each action. The amount you have earned so far will always be visible on your screen. In white you see the payoff of white, in black the payoff of black.

[The image shown to the subjects is reproduced above in Figure 3.2]

Actual play of the game and payment (differences in brackets)

When it is your turn to play, you will see a screen that:

1. reminds you of the current round of the game,
2. shows you the amount you and your partner earn if you choose STOP and
3. asks you to choose between STOP and CONTINUE.

You have 30 seconds to reach a decision. You can revise your choice at any time within the 30 seconds. The choice is final when you press OK.
When it is not your turn to play, you will see a screen that:

1. reminds you of the current round of the game and
2. shows you the amount you and your partner earn if your partner chooses STOP

Your partner has 30 seconds to make a decision as well. The game continues until one player chooses STOP or if the last decision round {*Tree, Decomposed*: on the right of the above representation} is reached.

{*Tree*: When the game finishes, payoffs are assigned according to the values in the picture above. You will be paid according to the values that appear at the point in which the game stops.}

{*Formula*: When the game finishes, payoffs are assigned according to the formula detailed above. You will be paid according to the decision round in which the game stops.}

{*Decomposed*: You start with a payoff of 4 if you are white, 1 if you are black. After each decision, your earnings will be updated according to the values that appear in the picture above. You will be paid what you have earned up to the point at which the game stops.}

You will play the game 12 times. Each time, you will form a couple with a new player chosen at random from the other participants in this room. You will never play the same partner twice. Your partners will never play one another.
Only one game of the 12 you play will be paid. At the end of the experiment, one number between 1 and 12 will be selected at random by the computer, and the corresponding game will be paid.
For the chosen game the result of you and your partner's action will be shown on the screen, and your final payoff will be computed.
Should you have any questions, please raise your hand now. An experimenter will come to your place and answer your questions in private.

CHAPTER 4

LIMITED FARSIGHTEDNESS IN NETWORK FORMATION [1]

Abstract

Pairwise stability Jackson and Wolinsky [1996] is the standard stability concept in network formation. It assumes myopic behavior of the agents in the sense that they do not forecast how others might react to their actions. Assuming that agents are perfectly farsighted, related stability concepts have been proposed.

We design a simple network formation experiment to test these extreme theories, but find evidence against both of them: the subjects are consistent with an intermediate rule of behavior, which we interpret as a form of limited farsightedness. On aggregate, the selection among multiple pairwise stable networks (and the performance of farsighted stability) crucially depends on the level of farsightedness needed to sustain them, and not on efficiency or cooperative considerations. Individual behavior analysis corroborates this interpretation, and suggests, in general, a low level of farsightedness (around two steps) on the part of the agents.

JEL classification: D85, C91, C92

Keywords: Network formation, experiment, myopic and farsighted stability.

[1] Joint work with Georg Kirchsteiger (ECARES, Université Libre de Bruxelles, CEPR), Ana Mauleon and Vincent Vannetelbosch (CORE, Université catholique de Louvain, CEREC, Facultés Universitaires Saint-Louis)

4.1 Introduction

The network structure of social interactions influences a variety of behaviors and economic outcomes, including the formation of opinions, decisions on which products to buy, investment in education, access to jobs, and informal borrowing and lending. A simple way to analyze the networks that one might expect to emerge in the long run is to examine the requirement that individuals do not benefit from altering the structure of the network. Any such requirement must answer the question of how individuals assess those benefits.

An extreme answer to this problem is to assume perfect myopia on the part of the agents, as in the pairwise stability notion, defined by Jackson and Wolinsky [1996]. A network is pairwise stable if no individual benefits from severing one of her links and no two individuals benefit from adding a link between them, with one benefiting strictly and the other at least weakly. Individuals are myopic, and not farsighted, in the sense that they do not forecast how others might react to their actions. Indeed, the adding or severing of one link might lead to subsequent addition or severing of another link, and so on. For instance, individuals might not add a link that appears valuable to them given the current network, as that might induce the formation of other links, ultimately leading to lower payoffs for the original individuals.

The von Neumann - Morgenstern pairwise farsightedly stable set (VNMFS) of networks predicts which networks one might expect to emerge in the long run when individuals are farsighted. As the other approaches to farsighted stability[2], it incorporates the assumption that agents are perfectly farsighted, meaning they can consider sequences of reactions to their moves of any lenght. As this constitutes the exact opposite of perfect myopia, there appears to be an unbridged gap between those extreme theories.

A notable exception is the work of Dutta et al. [2005], which allows for different degrees of farsightedness. In their equilibrium concept, for a dynamic Markovian process of network formation[3], farsightedness is captured by a discount factor, that applies to the stream of future payoffs. As such, it entangles patience and farsightedness. Moreover, their dynamic equilibrium model is hardly comparable to the static stability notions[4], in particular for intermediate values of farsightedness[5].

[2]See the work of Chwe [1994], Xue [1998], Herings et al. [2004, 2009], Mauleon and Vannetelbosch [2004], Page et al. [2005], and Page and Wooders [2009].

[3]See Konishi and Ray [2003] for a similar approach to the formation of coalitions.

[4]There are some random dynamic models of network formation that are based on incentives to form links such as Watts [2002], Jackson and Watts [2002], and Tercieux and Vannetelbosch [2006]. These models aim to use the random process to select from the set of pairwise stable networks.

[5]A discount factor of zero, properly corresponds to myopia. At the same time, we argue that a discount factor of one leads the process close to one in which people only care about the end state, as in the notions of farsighted stability. For intermediate values, the stream of future payoffs matters in a way that cannot be

In our paper we test the myopic and the (possibly limited) farsighted types of behaviors in the context of network formation and compare the stability notions that are based on them. Network formation is hard to study in the field, as many potentially conflicting factors are at work. Consequently, we run laboratory experiments. To the best of our knowledge, this constitutes the first experimental test of farsightedness versus myopia in network formation.

In the experiment, groups of four subjects had to form a network. More specifically, they were allowed sequentially to add or sever one link at a time: a link was chosen at random and the agents involved in the link had to decide if they wanted to form it (if it had not been formed yet) or to sever it (if it had been already formed). The process was repeated until all group members declared they did not want to modify the existing network. In all of the three treatments, the payoffs were designed such that a group consisting of myopic agents would never form any link. The treatments are characterized by slight manipulations of the payoffs, resulting in networks in VNMFS sets featuring different properties.

In treatment 1, a group composed of farsighted agents would form the complete network. This network provides the players with equal payoffs, is strongly stable, in the sense that no coalition can improve upon it, and features no farsighted deviations. Thus, beyond being VNMFS, the complete network can be seen as attractive in many ways[6]. In the other two treatments we vary those features to ascertain their contribution to the stability of an outcome.

A group composed of farsighted agents would form a triangle "club" network[7] or a line network among all the players, in treatments 2 and 3, respectively. In both, the payoffs are unequal, with the disadvantaged players earning around half the payoffs of the others. We remove strong stability in treatment 2, as a coalition of three players can improve upon the networks in the VNMFS. In treatment 3 the networks in the VNMFS are strongly stable, but feature a farsighted deviation in two steps [8]. We derive across-treatment hypothesis based on those properties.

In all the treatments farsighted stability refines the set of pairwise stable networks (PWS) by selecting the (unique) Pareto dominant network within the set of PWS. Note, however, that the underlying behavioral assumptions of both notions - myopia versus farsightedness - are at odds with each other, providing us with general within-treatment hypothesis.

---

captured by static stability notions.

[6]The complete network can be a focal point in itself - only for being the complete network.

[7]A network formed by a single clique (complete sub-network) of three players.

[8]Also the VNMFS set in T2 features farsighted deviations, but these are both "weak" and longer. See the discussion in Section 4.3.2

On aggregate, 75 percent of the network finally reached are pairwise stable. In treatments 1 and 2 most of the groups (up to 70 percent of the overall population) reach a VNMFS set, supporting farsighted network formation. In treatment 3, only one out of five groups reach a VNMFS set, with half of the groups ending the game in the empty network. In this treatment, VNMFS sets are accessed almost as often as in the other treatments, but, after some time, most groups leave them.

Given the properties of the VNMFS sets, this asymmetric result is inconsistent with strong stability - present in treatment 1 and 3, absent in treatment 2 - and can not be attributed to the inequality in the payoffs - equal in treatment 1, unequal in treatment 2 and 3. Nor it can be explained by other refinements of pairwise stabilility, such as Nash stability, or Pareto dominance - both present in all treatments. It is, however, perfectly consistent with the hypothesis derived from limited farsightedness.

We then show that individual behavior supports the interpretation of the aggregate results as an instance of limited farsightedness. Subjects respond to myopic incentives as well as to farsighted improving paths of short length. As a consequence if a farsightedly stable outcome features a farsighted deviation of limited length, the subjects are likely to follow it: they do not recognize the full chain of reactions that would prevent a fully farsighted agent to deviate.

Consequently, neither perfect myopia nor perfect farsightedness seem to be good models of actual behavior. A model of limited farsightedness would be a valuable development in network formation.

The number of experiments addressing networks and network formation is rapidly increasing.[9] Relatively few of them, however, deal with pure network formation, intended as a setting where no strategic interactions take place on the network once it has been formed. Among the notable exceptions stand the experiments of Goeree et al. [2009] and Falk and Kosfeld [2012]. They investigate the predictive power of a strict Nash network in the framework of Bala and Goyal [2003]. They find low support for this concept when the Nash network is asymmetric and the agents homogeneous. The main difference with our design is that they consider a model with unilateral link formation and apply non-cooperative solution concepts, while in our context of bilateral link formation those concepts provide implausible predictions [see Bloch and Jackson, 2006].

Closer to our approach is the work of Ziegelmeyer and Pantz [2005], where R&D networks in a Cournot oligopoly are investigated. Their results generally support pairwise stability. In their design pairwise stable networks are also farsightedly stable and thus there is no tension between myopia and farsightedness.[10]

[9] See Kosfeld [2004] for a partial survey.

[10] They observe huge differences between the case in which the Cournot profits are considered as exoge-

Finally, Berninghaus et al. [2011] address limited forward-looking behavior with an experiment on network formation. Relevant features distinguish our work from their model: (*i*) they assume unilateral link formation; (*ii*) players play a coordination game on the endogenously formed network and thus the assumption on the beliefs about this latter game affects the predictions; (*iii*) the forward-looking notion they consider relates specifically to the interaction between the linking strategies and the strategies in the coordination game. So their experiment combines a test of network formation and strategic behavior in the coordination game, while our paper is the first to directly investigate farsightedness and myopia in a network formation context unaffected by any other considerations.

The paper is organized as follows. In Section 4.2 we introduce the necessary notation and definitions. Section 4.3 presents the experimental design and procedures. Section 4.4 reports the experimental results. Section 4.5 concludes.

4.2 Networks: notation and definitions

Let $N = \{1, \dots, n\}$ be the finite set of players who are connected in some network relationship. The network relationships are reciprocal and the network is thus modeled as a non-directed graph. Individuals are the nodes in the graph and links indicate bilateral relationships between individuals. Thus, a network g is simply a list of which pairs of individuals are linked to each other. We write $ij \in g$ to indicate that i and j are linked under the network g. Let g^N be the collection of all subsets of N with cardinality 2, so g^N is the complete network. The set of all possible networks or graphs on N is denoted by $\mathbb{G}$ and consists of all subsets of g^N. The network obtained by adding link ij to an existing network g is denoted $g + ij$ and the network that results from deleting link ij from an existing network g is denoted $g - ij$. We say that g' is adjacent to g if $g' = g + ij$ or $g' = g - ij$ for some ij. Let us denote with A_g the networks that are adjacent to g so that $A_g = \{g' \mid g' = g + ij \vee g' = g - ij$, for some $ij\}$, and let $\bar{A}_g$ be its complement.

The material payoffs associated to a network are represented by a function $x : \mathbb{G} \to \mathbb{R}^n$ where $x_i(g)$ represents the material payoff that player i obtains in network g. The overall benefit net of costs that a player enjoys from a network g is modeled by means of a utility function $u_i(g) : \mathbb{R}^n \to \mathbb{R}$ that associates a value to the vector of material payoffs associated to network g . This might include all sorts of costs, benefits, and externalities.

Let $N_i(g) = \{j \mid ij \in g\}$ be the set of nodes that i is linked to in network g. The *degree* of

nously given and identified with the payoffs of the players in the network, and the case in which players play the production stage after forming the network. This supports pure network formation as the cleanest setting to study network formation.

a node is the number of links that involve that node. Thus node i's degree in a network g, denoted $d_i(g)$, is $d_i(g) = \#N_i(g)$. Let $S_k(g)$ be the subset of nodes that have degree k in network g: $S_k(g) = \{i \in N \mid d_i(g) = k\}$ with $k \in \{0, 1, ..., n-1\}$. The *degree distribution* of a network g is a description of the relative frequencies of nodes that have different degrees. That is, $P(k)$ is the fraction of nodes that have degree k under a degree distribution P, i.e., $P(k) = (\#S_k(g))/n$. Given a degree distribution, $\overline{P}$, we define a *class of networks* as $C_{\overline{P}} = \{g \in \mathbb{G} \mid P(k) = \overline{P}(k), \forall k\}$. A class of networks is the subset of $\mathbb{G}$ with the same degree distribution.

Consider a network formation process under which mutual consent is needed to form a link and link deletion is unilateral. A network is pairwise stable if no player benefits from severing one of their links and no other two players benefit from adding a link between them, with one benefiting strictly and the other at least weakly. Formally, a network g is pairwise stable if

(i) for all $ij \in g$, $u_i(g) \geq u_i(g - ij)$ and $u_j(g) \geq u_j(g - ij)$, and

(ii) for all $ij \notin g$, if $u_i(g) < u_i(g + ij)$ then $u_j(g) > u_j(g + ij)$.

A network g' defeats g if either $g' = g - ij$ and $u_i(g') > u_i(g)$ or $u_j(g') > u_j(g)$, or if $g' = g + ij$ with $u_i(g') \geq u_i(g)$ and $u_j(g') \geq u_j(g)$ with at least one inequality holding strictly. Pairwise stability is equivalent to the statement of not being defeated by an adjacent network. Agents are assumed to consider only their own incentives when making their linking choices and not that of the others. In particular, agents do not take into account the likely chain of reactions that follow an action, but only its immediate profitability. Thus, PWS implicitly assumes myopic behavior on the part of the agents.

Farsightedness captures the idea that agents will consider the chain of reactions that could follow when deviating from the current network, and evaluate the profitability of such deviation with reference to the final network of the chain of reactions. As a consequence, a farsighted agent will eventually choose against her immediate interest if she believes that the sequence of reactions that will follow her action could make her better off.

A *farsighted improving path* is a sequence of networks that can emerge when players form or sever links based on the improvement the end network offers relative to the current network. Each network in the sequence differs by one link from the previous one. If a link is added, then the two players involved must both prefer the end network to the current network, with at least one of the two strictly preferring the end network. If a link is deleted, then it must be that at least one of the two players involved in the link strictly prefers the end network. We now introduce the formal definition of a farsighted improving path.

Definition 1. *A farsighted improving path from a network g to a network $g' \neq g$ is a finite sequence of graphs $g_1, \ldots, g_K$ with $g_1 = g$ and $g_K = g'$ such that for any $k \in \{1, \ldots, K-1\}$ either:*

(i) $g_{k+1} = g_k - ij$ *for some ij such that $u_i(g_K) > u_i(g_k)$ or $u_j(g_K) > u_j(g_k)$ or*

(ii) $g_{k+1} = g_k + ij$ *for some ij such that $u_i(g_K) > u_i(g_k)$ and $u_j(g_K) \geq u_j(g_k)$.*

If there exists a farsighted improving path from g to g', then we write $g \to g'$. For a given network g, let $F(g) = \{g' \in \mathbb{G} \mid g \to g'\}$. This is the set of networks that can be reached by a farsighted improving path from g. The von Neumann-Morgenstern pairwise farsightedly stable set is obtained by introducing the notion of farsighted improving path into the standard definition of a von Neumann-Morgenstern stable set. In other words, we define a set of networks G to be von Neumann-Morgenstern pairwise farsightedly stable (VNMFS) if there is no farsighted improving path connecting any two networks in G and if there exists a farsighted improving path from any network outside G leading to some network in G. Formally,

Definition 2. *The set of networks G is a von Neumann-Morgenstern pairwise farsightedly stable set if*

(i) $\forall g \in G, F(g') \cap G = \varnothing$ *(internal stability) and*

(ii) $\forall g' \in \mathbb{G} \backslash G, F(g') \cap G \neq \varnothing$ *(external stability).*

Although the existence of a VNMFS set is not guaranteed in general, when a VNMFS set exists it provides narrower predictions than other definitions of farsighted stability, a feature that is particularly welcome in experimental testing. For instance, a VNMFS set is always included within the pairwise farsightedly stable sets, as defined by Herings et al. [2009].[11]

We now turn to individual behavior. We provide a comprehensive evaluation of the players' actions by assessing their consistency with progressive levels of farsightedness. The definition states that an action prescribing to form (break) a link that is not formed (has been formed) is consistent with farsightedness of level k, if building (breaking) the link lies on a farsighted improving path of lenght smaller or equal than k. An action prescribing not to form (keep) a link that is not formed (has been formed) is consistent with farsightedness of level k if forming (breaking) the link does not lie on a farsighted

[11] A set of networks $G \subseteq \mathbb{G}$ is *pairwise farsightedly stable* if (i) all possible pairwise deviations from any network $g \in G$ to a network outside G are deterred by a credible threat of ending worse off or equally well off, (ii) there exists a farsighted improving path from any network outside the set leading to some network in the set, and (iii) there is no proper subset of G satisfying Conditions (i) and (ii).

improving path of lenght smaller or equal than k. Let the length of a path be the number of steps in the sequence. Call $\mathcal{P}_g^k$ a generic farsighted improving path of length k, starting from network g, and $\{\mathcal{P}_g^k\}$ be the set containing all such paths.[12] At time t the link ij is selected, the action of agent i is $a_{it} \in \{0,1\}$, where 0 means not to form (to break) the selected link ij, and 1 means to form (to keep) the link ij.

Definition 3. *An action a_{it} is consistent with farsightedness of level k if either*

$$\text{(i) } ij \notin g_t \text{ and } ((\exists l \leq k \text{ and a } \mathcal{P}_{g_t}^l \in \{\mathcal{P}_{g_t}^l\} \text{ s.t } g_t + ij \in \mathcal{P}_{g_t}^l) \text{ and } a_{it} = 1) \vee$$
$$((\nexists l \leq k \text{ and a } \mathcal{P}_{g_t}^l \in \{\mathcal{P}_{g_t}^l\} \text{ s.t } g_t + ij \in \mathcal{P}_{g_t}^l) \text{ and } a_{it} = 0)$$

or

$$\text{(ii) } ij \in g_t \text{ and } ((\exists l \leq k \text{ and a } \mathcal{P}_{g_t}^l \in \{\mathcal{P}_{g_t}^l\} \text{ s.t } g_t - ij \in \mathcal{P}_{g_t}^l) \text{ and } a_{it} = 0) \vee$$
$$((\nexists l \leq k \text{ and a } \mathcal{P}_{g_t}^l \in \{\mathcal{P}_{g_t}^l\} \text{ s.t } g_t - ij \in \mathcal{P}_{g_t}^l) \text{ and } a_{it} = 1)$$

As they are equivalent, we call myopic an action that is consistent with farsightedness of level one - i.e. one that looks at the profitability of adjacent networks. Two aspects in this definition should be noted. First, an action that aims at changing the current network and is consistent with some level of farsightedness, including myopia, is also consistent with higher levels. Second, for an action that does not change the current network, we implicitly impose a strong assumption on farsighted behavior: that a farsighted agent should always take a profitable deviation, if available.

Indeed, given that the building blocks of farsightedness are sequences of networks, farsighted behavior is unambiguously defined only if a choice aims at changing the current network. When it does not, we are forced either to draw some further assumptions or give up categorizing those choices. In the statistical analysis of individual behavior we pursue both of the alternatives.

4.3 Experimental design and procedures

4.3.1 The game

We consider a simple dynamic link formation game, almost identical to that proposed by Watts [2001]. Time is a countable infinite set: $T = 0, 1, ..., t, ...$; g_t denotes the network that exists at the end of period t. The process starts at $t = 0$ with $n = 4$ unconnected players

[12]Note that a path of length k will have a sequence of $k+1$ networks.

(g_0 coincides with the empty network, $g^{\emptyset}$). The players meet over time and have the opportunity to form links with each other.

At every stage $t > 0$, a link ij_t is randomly identified to be updated. At $t = 1$ each link from the set g^N is selected with uniform probability. At every $t > 1$, a link ij from the set $g^N \setminus ij_{t-1}$ is selected with uniform probability. Thus, a link cannot be selected twice in two consecutive stages. If the link $ij \in g_{t-1}$, then both i and j can decide unilaterally to sever the link; if the link $ij \notin g_{t-1}$, then i and j can form the link ij if they both agree. Once the individuals involved in the link have taken their decisions, g_{t-1} is updated accordingly and we move to g_t. All group members are informed about both the decisions taken by the players involved in the selected link and the consequences on that link. They are informed through a graphical representation of the current network g_t and the associated payoffs. After every stage all group members are asked whether they want to modify the current network or not. If they unanimously declare they do not want to, the game ends; otherwise, they move to the next stage.[13] To ensure that an end is reached, a random stopping rule is added after stage 25: at every $t \geq 26$ the game ends anyway with probability 0.2.

The game is repeated three times to allow for learning: groups are kept the same throughout the experiment. Group members are identified through a capital letter (A, B, C or D). These identity letters are reassigned at every new repetition.

A vector of payoffs is associated to every network: it allocates a number of *points* to each player in the network. The subjects receive points depending only on the final network of each repetition. Thus, their total points are given by the sum of the points achieved in the final networks of the three repetitions. At the end of the experiment the points are converted into Euro at the exchange rate of 1 Euro = 6 *points*.

The subjects are informed about the payoffs associated to every possible network and know the whole structure of the game from the beginning. Before starting the first repetition the participants have the opportunity of practicing the relation between networks and payoffs and the functioning of the stages through a training stage and three trial stages.

4.3.2 Treatments and hypothesis

Since $n = 4$, it follows that $\#g^N = 6$ and $\#\mathbb{G} = 64$.

We run three treatments (T1, T2, T3) where we manipulate the payoffs in some networks to obtain VNMFS set(s) with different properties. Figures 4.1, 4.2 and 4.3 display the payoffs that were used in the three treatments for each class of networks, $C_{\bar{p}}$. The

[13] Subjects are informed about the outcome of the *satisfaction* choices - i.e. end of the repetition or not - but not about individual choices.

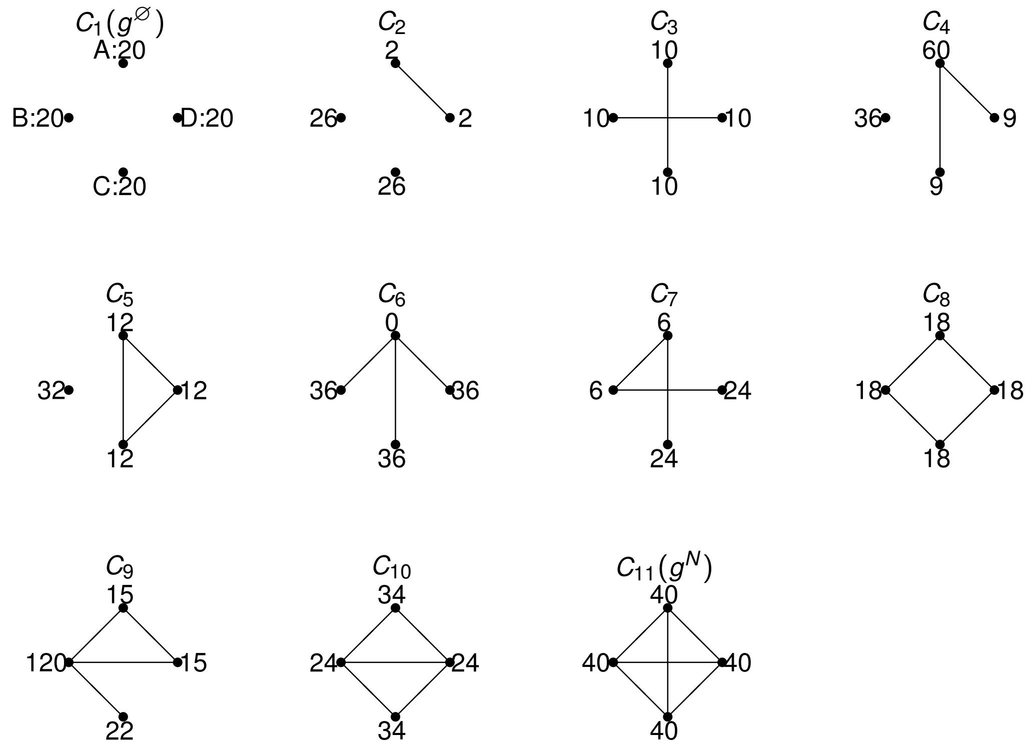

Figure 4.1: Payoffs for T1

function of material payoffs satisfies anonymity[14] and then, this representation is sufficient to assign a payoff to each player in each possible network configuration. These numbers were chosen in order to provide the resulting predictions with a set of nice properties for each treatment that are described below and are summarized in Table 4.1.[15]

The empty network, $g^{\varnothing}$, and the the four networks in class C_5 are PWS in all treatments. These are the only PWS networks in T2, whereas g^N is also PWS under T1, and the networks in C_7 in T3. Furthermore, in T1 and T3, in every network in C_5, the connected agents can improve their situation by cutting both of their links. These networks (contrary to the others PWS) are not Nash stable in the terminology of Bloch and Jackson [2006].[16]

[14]Anonimity holds if payoffs in a network are assigned to each player independently of his or his partners' identity.

[15]In general, the following considerations are valid for self-regarding agents. In some cases they hold for other-regarding preferences (for an overview, see Sobel [2005]). Most notably, in T1, assuming the inequity model of Fehr and Schmidt [1999] does not affect our predictions.

[16]Pairwise Nash stability is a refinement of both pairwise stability and Nash stability, where one requires

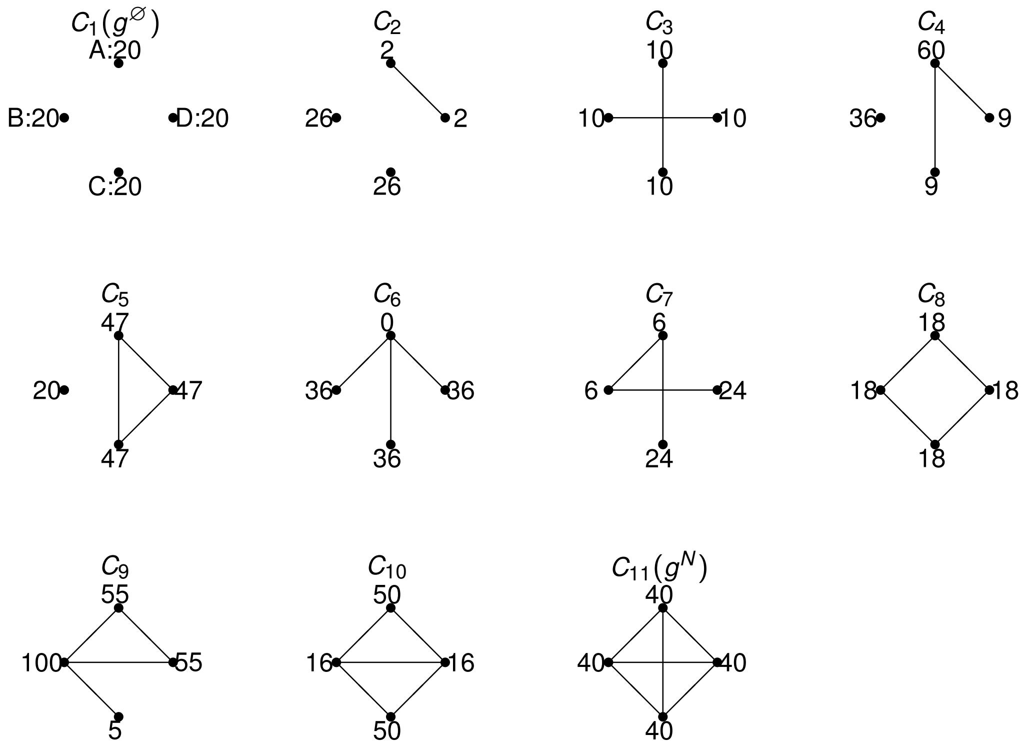

Figure 4.2: Payoffs for T2

In all treatments, all groups start at $g^{\varnothing}$, and then groups composed of myopic players are expected not to move from $g^{\varnothing}$. This prediction is robust to errors. A sequence of at least three (T1) or two (T2, T3) links added consecutively by error is needed in order to change the prediction for myopic agents. In both cases, these sequences of events are highly unlikely, and our prediction for a myopic group of players is to end up in $g^{\varnothing}$.

To identify the VNMFS sets, we need to compute $F(g)$ for every g. We can prove the following results.

Proposition 1. *Consider a set of four self-regarding agents ($u_i(g) = x_i(g)$). Then,*

i *in T1 the set $G = \{g^N\}$ is the unique VNMFS set.*

ii *in T2 the set $G = \{g|g \in C_5\}$ is the unique VNMFS set.*

---

that a network be immune to the formation of a new link by any two agents, and the deletion of any number of links by any individual agent.

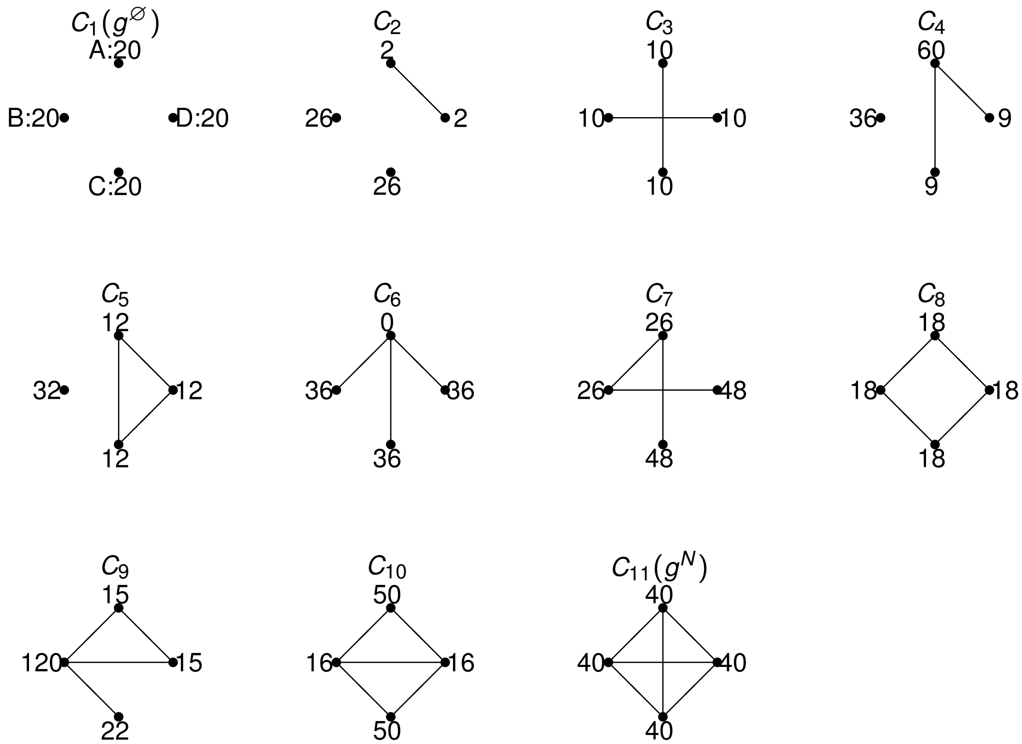

Figure 4.3: Payoffs for T3

iii *in T3 a set G is a VNMFS set if and only if* $G = \{g \mid g \in C_7$ *and* $, d_i(g) = d_i(g'), \forall i \in N, g' \in G\}$.

The proof of this proposition can be found in Appendix C.

In T1 and T2 there is a unique VNMFS set: the complete network (i) and the set composed of the four networks in C_5 (ii), respectively. We will refer to the latter as club networks. In T3 there are six VNMFS sets. Their union is C_7, i.e. it encompasses all line networks. Each set consists of a pair of line networks with identical degree distribution (iii).[17]

We expect a group composed by farsighted agents to end up in a network included in some VNMFS set. This prediction is robust to errors in the sense that the farsighted

[17] The pair of line networks in a VNMFS, are equal up to a single permutation of players with the same degree. For example, there are two networks in C_7 where A and B have 2 links each, call them g and g'. A and B are linked to one another in both networks, but A will be linked to C, and B to D, in g; vice versa in g'. The set $\{g, g'\}$ is a VNMFS.

prediction does not depend on the starting point: from any other network, there is a farsighted improving path leading to a network in G.

Let $frac_{MYO}(T_i)$ and $frac_{FAR}(T_i)$ be the fraction of groups ending in the myopic and in the farsighted prediction, respectively, in treatment i. We state the following, mutually exclusive, hypothesis, regarding perfect myopia and farsightedness.

	PWS	VNMFS	Myopic Prediction	Farsighted Prediction	Unequal Payoffs	Strongly Stable	Farsighted Deviations
T1	$g^{\varnothing}, C_5^*, g^N$	$\{g^N\}$	$g^{\varnothing}$	g^N	No	g^N	–
T2	$g^{\varnothing}, C_5$	$\{g \mid g \in C_5\}$	$g^{\varnothing}$	C_5	Yes	–	Three steps**
T3	$g^{\varnothing}, C_5^*, C_7$	$\{g, g' \mid g, g' \in C_7$ and $d_i(g) = d_i(g'), \forall i \in N\}$	$g^{\varnothing}$	C_7	Yes	$g \in C_7$	Two steps

* not Nash stable.

** Weak deviation, based on indifference breaking rule.

Table 4.1: Summary of treatment properties and predictions

Hypothesis 1. *(Myopia) In all treatments, a relative majority of the groups end the game in $g^{\varnothing}$, and, in particular, for $i = 1, 2, 3$:*

$$frac_{MYO}(T_i) > frac_{FAR}(T_i).$$

Hypothesis 2. *(Farsightedness) In all treatments, a relative majority of the groups end the game in a VNMFS, and , in particular, for $i = 1, 2, 3$:*

$$frac_{MYO}(T_i) < frac_{FAR}(T_i).$$

In our experiment, if a network is in a VNMFS set, it is also PWS. Even myopic agents will stay at the farsighted stable network once it is reached. Therefore, one cannot find direct experimental evidence against PWS as opposed to farsighted stability. But our experiment discriminates between the different behavioral models that lie behind both stability concepts. In this way our experiment can provide evidence in favor or against the farsighted models of network formation in cases where they refine PWS.

The payoffs guarantee that the predicted networks are essentially unique, in the sense that the networks included in a VNMFS set are isomorphic. Moreover, the predicted networks are not strongly efficient in the sense of Jackson and Wolinsky [1996][18] nor Pareto dominant. Previous experimental studies have shown that efficiency considerations can drive individual's behavior (see Engelmann and Strobel [2004]. But generic efficiency arguments could not explain if a network in some VNMFS set or $g^{\varnothing}$ were observed in the

[18] A network $g \in \mathbb{G}$ is *strongly efficient* if $\sum_{i \in N} x_i(g) \geq \sum_{i \in N} x_i(g')$ for all $g' \in \mathbb{G}$.

experiment. The networks included in VNMFS sets are (weakly) Pareto dominant within the set of pariwise Nash stable networks[19].

On top of these general hypothesis, the different VNMFS sets differ on three important properties, providing some testable across-treatment hypothesis (see Table 4.1).

First, the payoffs are equal in the VNMFS set in T1 (g^N) and unequal in T2 (C_5) and T3 (C_7). In the latter, the players gaining more obtain around twice as much as the least well off. Under both conditions, the disadvantaged players can lead the group to leave the VNMFS set, if they so wish, by severing a link in T3, by adding a link in T2[20]. If other-regarding preferences are sufficiently strong, the VNMFS sets could be less stable in T2 and in T3, with respect to T1.

Hypothesis 3. *The fraction of groups ending the game in a VNMFS set is higher if the networks that are there included feature equal payoffs for the players. Thus:*

i $frac_{FAR}(T_1) > frac_{FAR}(T_2)$, *and*

ii $frac_{FAR}(T_1) > frac_{FAR}(T_3)$.

Second, we also consider stability against changes in links by any coalition of individuals - i.e. look for strongly stable networks (immune to coalitional deviations). In T1 and T3 the networks included in VNMFS sets are also strongly stable. This is not true in T2, where strongly stable networks fail to exist.[21] In this view the VNMFS set seems more robust in T1 and in T3 than in T2.

Hypothesis 4. *The fraction of groups ending the game in a VNMFS set is higher if the networks that are there included are strongly stable. Thus:*

i $frac_{FAR}(T_1) > frac_{FAR}(T_2)$, *and*

ii $frac_{FAR}(T_3) > frac_{FAR}(T_2)$.

Finally, the networks belonging to the VNMFS sets differ with respect to the presence and length of farsighted deviations leaving the set. Table 4.2 provides an overview and an example for each treatment. In T1, there are no farsighted improving paths leaving the complete network ($F(g^N) = \varnothing$). In T2, $F(g \in C_5) = \{g' \mid g' \in C_9 \wedge g' \notin A_g\}$. This means that there are farsighted improving paths leaving the VNMFS set and leading to

[19] Recently, Carrillo and Gaduh [2012] suggested that the players are able to select the PWS networks that are Pareto dominant. Our results show that Pareto dominance is not a sufficient criterion to select among PWS networks.

[20] Despite needing the agreement of his partner to add a link, adding a link in C_5 is highly beneficial to the already connected agents, so that they are likely to agree on that.

[21] As shown by Jackson and Van Den Nouweland [2005] this is equivalent to an empty core in the derived cooperative game.

networks in C_9 that are not adjacent to the initial network g. The path is built as follows: from C_5 players move to C_9, then to C_{10} and finally back to another network in C_9. This path relies on the indifference-breaking convention stating that, in C_9, a player with two links is willing to build another one in order to (move to C_{10} and then) be exactly in the same situation in another network in C_9. As such, this is a "weak" deviation, that is unlikely to succeed [22]. Finally, $F(g \in C_9)$ includes, beyond the neighboring network in C_5 and the other networks in C_9, only the networks in C_4, reached with a four-step farsighted improving path, implying that even groups that leave the VNMFS set for C_9 are somewhat stuck there.

In T3 there are two-steps farsighted improving paths from any network in a VNMFS set to any network in another VNMFS set[23]. Namely, one of the players with two links cut any of his existing links ($C_7 \to C_3$ or $C_7 \to C_4$). From there, another link will be added leading back to C_7, but in a network where the initial deviator is better off (because he has only one link). After the first move away from C_7 is made, other (short) deviations are feasible, driving the group away from the VNMFS (and, most notably, toward $g^{\varnothing}$). Those differences bare little meaning in the context of perfect farsightedness. However, to the extent that the agents may be bounded in their ability to pursue farsighted deviations, the VNMFS set seems more robust in T1 and in T2 than in T3.

Hypothesis 5. *The fraction of groups ending the game in a VNMFS set is higher if the networks that are there included are robust to short farsighted deviations. Thus:*

i $frac_{FAR}(T_1) > frac_{FAR}(T_3)$, *and*

ii $frac_{FAR}(T_2) > frac_{FAR}(T_3)$.

4.3.3 Experimental procedures

The experiment took place at the EELAB of the University of Milan-Bicocca in June 2010 (T1) and April/May 2012 (T2,T3). The computerized program was developed using *Z*-tree [Fischbacher, 2007]. We run 16 sessions for a total of 288 participants and 72 groups. Those corresponds to 36 independent observations for T1, and 18 independent observations for T2 and T3. Table 4.3 summarizes sessions' details. Participants were undergraduate students from various disciplines,[24] recruited through an announcement on the EELAB website. No subject participated in more than one session.

[22] One may question how reasonable it is to keep the same indifference-breaking conventions in the case of farsighted moves as in the myopic case.

[23] There are other farsighted deviations, longer than four steps.

[24] Sociology, economics, business, psychology, statistics, computer science, law, biology, medicine, mathematics, pedagogy and engineering.

	Origin	Destination	Lenght	Example
T1	g^N	–	–	–
T2	C_5	C_9	3	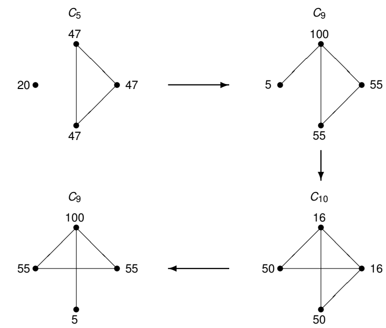
T3	C_7	C_7	2	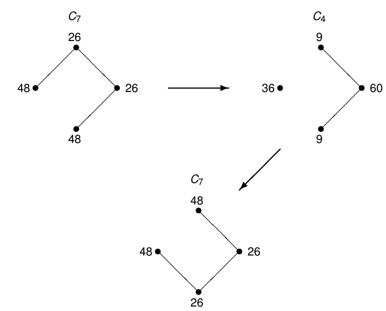

Table 4.2: Farsighted deviations from VNMFS sets

Subjects were randomly assigned to individual terminals and were not allowed to communicate during the experiment. Instructions were read aloud (see Appendix D for an English translation of the instructions). Participants were asked to fill in a control questionnaire; the experiment started only when all the subjects had correctly completed the task.

Sessions took on average 90 minutes, including instructions, control and final questionnaire phases. Average payment was 16.10 Euro (no show up fee was paid) with a minimum of 4.70 and a maximum of 32.40 Euro.

	Date	Participants	Groups (Ind. Obs)	Treatment
1	Jun 2010	24	6	T1
2	Jun 2010	24	6	T1
3	Jun 2010	24	6	T1
4	Jun 2010	24	6	T1
5	Jun 2010	24	6	T1
6	Jun 2010	24	6	T1
7	Apr 2012	16	4	T2
8	Apr 2012	16	4	T2
9	Apr 2012	16	4	T2
10	May 2012	16	4	T3
11	May 2012	16	4	T3
12	May 2012	16	4	T3
13	May 2012	16	4	T3
14	May 2012	16	4	T2
15*	May 2012	8	2	T2
16*	May 2012	8	2	T3

* Sessions 15 and 16 were run at the same time.

Table 4.3: Sessions

4.4 Results

In this section we first show how both perfect myopia and farsightedness are inconsistent with the networks formed, whereas limited farsightedness can reconcile the different results in all treatments. We then investigate this hypothesis using individual data, finding clear evidence of the relevance of limited levels of farsightedness.

We start by considering groups' final networks. Figure 4.4 classifies groups with respect to their final network. Figure 4.5 provides the same information for each repetition (period). In every treatment, around three out of four groups reach a PWS network[25]. This percentage increases consistently across repetitions within each single treatment, except between the second and third repetition of T3.

The distribution within PWS networks shows different patterns across treatments. In T1 and T2 the fraction of groups ending up in the VNMFS set is consistently higher than that ending up in $g^{\varnothing}$. This difference increases across repetitions with the farsighted and the myopic outcome accounting for around 70 and below 20 percent of the final networks, respectively, in the last repetition.

This pattern is almost reversed in T3. The final network is $g^{\varnothing}$ for half of the groups,

[25]This high percentage is reassuring on the subjects' ability to understand the game, as it would hardly result from generalized non-meaningful play.

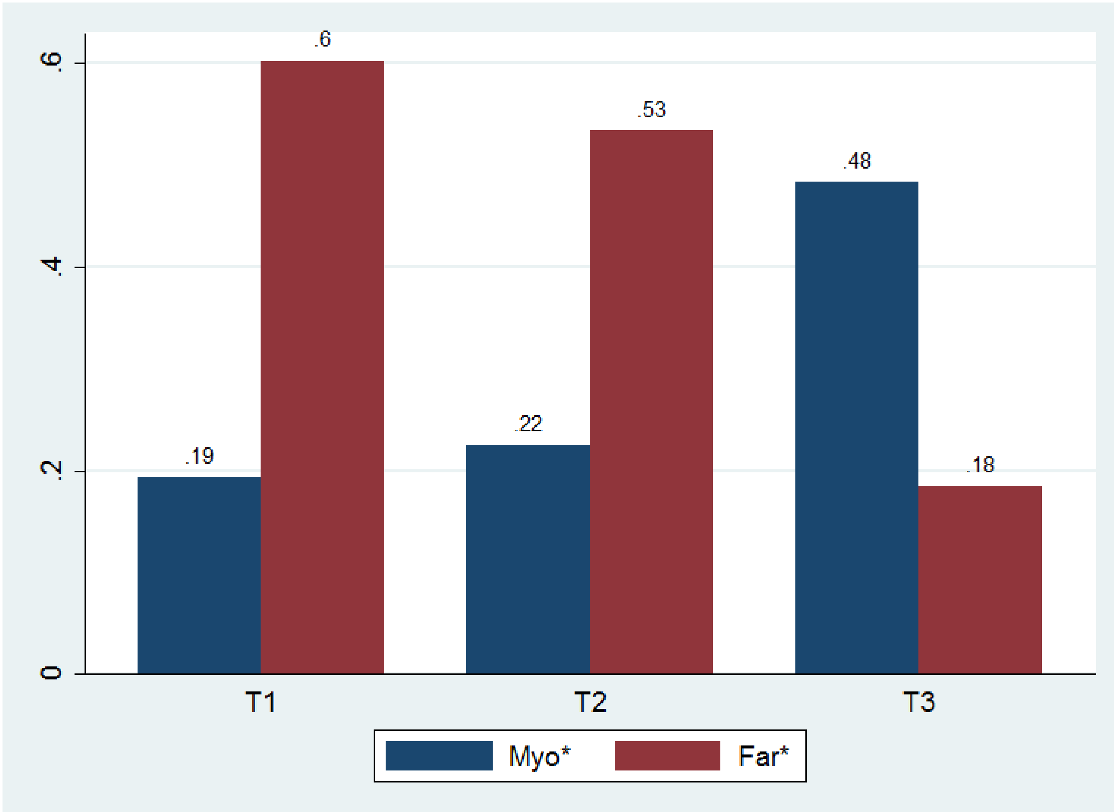

Figure 4.4: Group final network, by type of outcome

with this percentage peaking at 60 percent in the second repetition. A VNMFS set is reached by about 20 percent of the groups in all repetitions.[26]

We use the Pearson's chi-square and the Likelihood Ratio test to determine whether the relative frequencies of the myopic and the farsighted outcome differ or not within treatments and conclude that those differences are statistically significant at the 0.05 level in each single treatment.[27] Running the tests for each single repetition leads to significant differences in repetitions two and three of T1 and repetition three of T2.[28]. Given that those differences go in opposite direction in T3, with respect to T1 and T2, those results imply a rejection of both Hypothesis 1 and 2.

[26] Around 90 percent of the groups move from the empty network. As a consequence we gather indirect evidence about the behavior of groups that do not start from a pairwise stable network.

[27] We run the tests on the distributions obtained for outcomes - i.e. myopic, farsighted, other -, for network classes and for single networks. We run them against different assumptions for the frequencies that are not being tested under the null hypothesis (H_0: equality of frequencies for myopic and fardighted outcome): uniform distribution, uniform given the actual cumulative frequency of myopic and farsighted, actual frequencies. The results are identical across all specifications.

[28] Repetition two of T2 is significant at the 0.1 level. Note that we collected fewer observations for T2 and T3 than for T1.

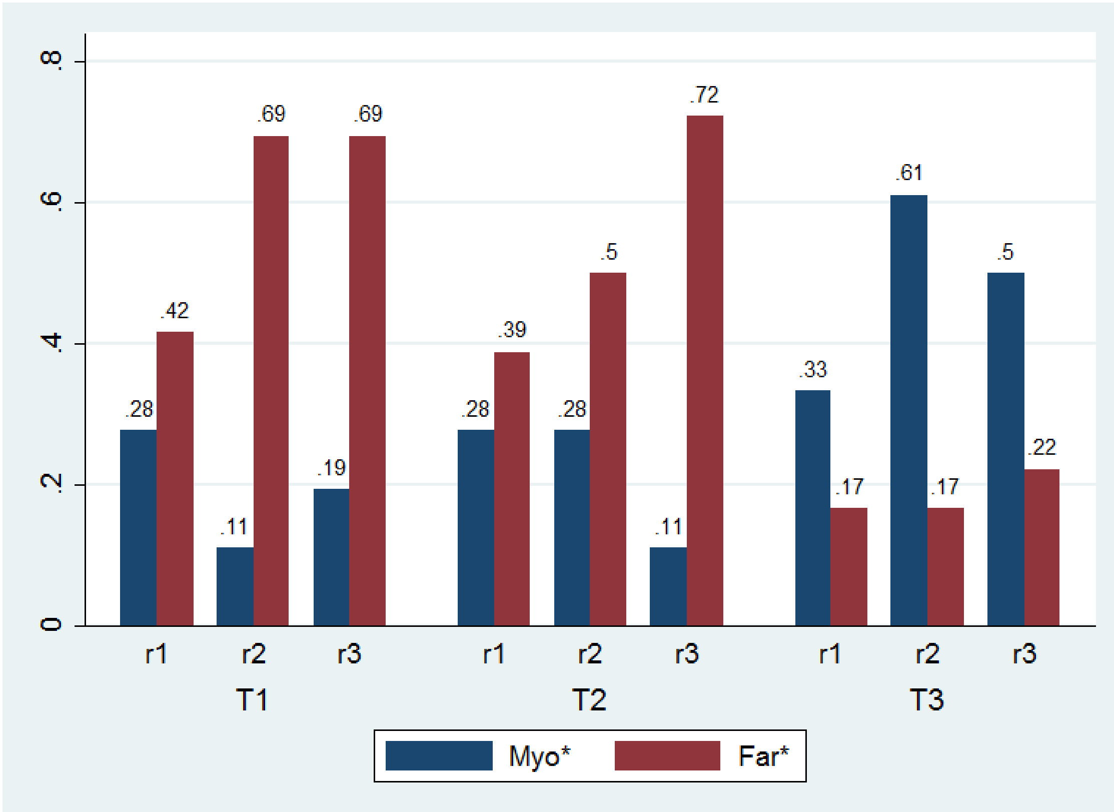

Figure 4.5: Group final network, by type of outcome and repetition

Result 1: *The predicted stable networks (PWS) account, on aggregate, for 75 percent of our groups' final outcomes. The VNMFS sets account for most of those observations in T1 and T2, but not in T3. The reverse holds for the myopic prediction, which shows some success only in T3. Thus, both perfect myopia (H.1) and farsightedness (H.2) fail to rationalize our results.*

We use a two-sample Kolmogorov-Smirnov test to compare the distribution of aggregate outcomes - i.e. myopic, farsighted, other - across treatments. As expected, we find that the distribution of outcomes in T1 and T2 are significantly different from that in T3 at the 0.05 level. When comparing T1 and T2 we do not find their distributions to be significantly different. This leads us to reject both Hypothesis 3 and 4, as we do not find the inaquality of the payoffs nor strong stability to affect systematically the stability of the VNMFS sets. The results are, instead, consistent with Hypothesis 5, supporting limited farsightedness.

Between one fifth and one third of the groups did not end up neither in the myopic nor in the farsighted prediction; we generally refer to this category as "other". Remarkably, a vast majority of those, between 72 and 77 percent, ended the game in networks that are direct neighbors of either of the two. The specific figures are as follows: in T1, 50 and 23 percent of those ended up at one step from the empty and the complete network, and thus in C_{10} and C_2, respectively; in T2, 61 percent resulted in C_9, at one step from the VNMFS set, 16 percent in C_2; in T3, 39 percent were in C_2, while 33 percent in C_4. We note that in T2 and T3, the groups that were close to a VNMFS set happened to be precisely ont the first step of the farsighted deviations outlined above.

Result 2: *The asymmetric performance of the VNMFS sets in T3 with respect to T1 and T2 can not be explained by payoff inequality or coalitional stability, leading to a rejection of both Hypothesis 3 and 4. The results are, instead, consistent with limited farsightedness (H.5).*

Table 4.4 reports the change in the outcome of individual groups from Period 1 to 2 and from Period 2 to 3, for all treatments. For example, the row "Farsighted" from the upper-left panel (T1, Period 1 - Period 2) shows that in T1, among those groups who reached the complete network in period 1, only 7 percent switched to the empty, myopic network in period 2, whereas 93 percent of the groups also reached the complete network in period 2. But among those groups who ended up in the empty network in period 1 (row "Myopic"), only 20 percent stayed at the empty network in period 2, whereas 50 percent switched to the complete network, and 30 percent to an unstable network. Similarly, among the groups who ended up in some other network in period 1 (row "Other"), 55 percent of them switched to the complete network in period 2, while only 18 percent of them switched to the empty network.

Table 4.4 shows that groups that reached a VNMFS set in a previous period are able to replicate the result in T1 and T2: the Farsighted-Farsighted cell displays a fraction close or above 80 percent in the corresponding panels. The other categories display greater mobility across time. Some of them reach a VNMFS set, others fluctuate among the empty network and the Other category. Again, a striking difference appears comparing those results with the right-hand side panels, corresponding to T3. Around two thirds of the groups that end in the empty network in one repetition, replicate this outcome in the subsequent one. This is the only outcome showing some persistence; the farsighted outcome, in particular is much less stable across repetitions.

The columns labeled "Destinations" report the major receivers of the outflows from each class of network, and their share of those outflows. We are particularly interested in the results for C_7 in T3. It turns out that two thirds of the groups that left a VNMFS set in T3 did so consistently with the short farsighted deviation described above (destinations C_3 and C_4, see Table 4.2) [29].

[29] Note also that of the groups that left the VNMFS set in T2 (C_5), more than 90 percent did so consistently with a farsighted deviation (destination C_9).

		T1			T2			T3		
			Period 2			**Period 2**			**Period 2**	
		Myopic	**Farsighted**	**Other**	**Myopic**	**Farsighted**	**Other**	**Myopic**	**Farsighted**	**Other**
	Myopic	0.20	0.50	0.30	0.60	0.20	0.20	0.67	0.17	0.17
Period 1	**Farsighted**	0.07	0.93	0.00	0.00	0.86	0.14	0.33	0.33	0.33
	Other	0.18	0.55	0.27	0.33	0.33	0.33	0.67	0.11	0.22
			Period 3			**Period 3**			**Period 3**	
		Myopic	**Farsighted**	**Other**	**Myopic**	**Farsighted**	**Other**	**Myopic**	**Farsighted**	**Other**
	Myopic	0.60	0.00	0.40	0.20	0.40	0.40	0.64	0.09	0.27
Period 2	**Farsighted**	0.00	0.92	0.08	0.11	0.78	0.11	0.00	0.33	0.67
	Other	0.67	0.34	0.00	0.00	1.00	0.00	0.50	0.50	0.00

Table 4.4: Group flows by treatment and period

	Out, In (Ratio)			Destinations (Share)			Average stay		
	T1	**T2**	**T3**	**T1**	**T2**	**T3**	**T1**	**T2**	**T3**
$\mathbf{C_1}$	115, 136 (.85)	60, 72 (.83)	65, 91 (.71)	C_2(100%)	C_2(100%)	C_2(100%)	2.28	3.00	3.67
$\mathbf{C_2}$	135, 146 (.92)	75, 77 (.97)	96, 103 (.93)	C_1(21%), C_3(12%), C_4(67%)	C_1(24%), C_3(7%), C_4(69%)	C_1(39%), C_3(10%), C_4(51%)	2.60	2.90	3.24
$\mathbf{C_3}$	18, 18 (1)	10, 10 (1)	21, 21 (1)	C_2(18%), C_7(82%)	C_2(20%), C_7(80%)	C_2(76%), C_7(24%)	1.61	1.10	1.43
$\mathbf{C_4}$	129, 135 (.96)	66, 66 (1)	70, 76 (.92)	C_2(26%), C_5(30%), C_7(24%)	C_2(23%), C_5(43%), C_7(26%)	C_2(31%), C_5(30%), C_7(33%)	2.12	1.99	1.84
$\mathbf{C_5}$	46, 47 (.98)	26, 55 (.47)	24, 25 (.96)	C_4(50%), C_9(50%)	C_4(8%), C_9(92%)	C_4(58%), C_9(42%)	4.49	4.42	3.44
$\mathbf{C_6}$	35, 35 (1)	8, 8 (1)	10, 10 (1)	C_4(43%), C_9(57%)	C_4(37%), C_9(62%)	C_4(90%), C_9(10%)	2.89	1.00	3.00
$\mathbf{C_7}$	49, 50 (.98)	34, 36 (.94)	32, 42 (.76)	C_4(10%), C_8(22%), C_9(61%)	C_3(15%), C_4(26%), C_9(50%)	C_3(41%), C_4(25%), C_9(31%)	1.54	1.28	5.05
$\mathbf{C_8}$	16, 16 (1)	11, 11 (1)	13, 13 (1)	C_7(19%), C_{10}(81%)	C_7(82%), C_{10}(18%)	C_7(69%), C_{10}(31%)	2.87	1.54	1.38
$\mathbf{C_9}$	96, 96 (1)	51, 59 (.86)	31, 31 (1)	C_5(12%), C_6(12%), C_{10}(67%)	C_5(53%), C_{10}(39%)	C_6(19%), C_7(16%), C_{10}(51%)	1.94	4.03	1.32
$\mathbf{C_{10}}$	72, 77 (.93)	24, 24 (1)	19, 21 (.90)	C_9(9%), C_{11}(90%)	C_8(33%), C_9(54%), C_{11}(12%)	C_8(47%), C_9(37%), C_{11}(17%)	4.06	2.96	3.05
$\mathbf{C_{11}}$	0, 65 (0)	2, 3 (.66)	2, 3 (.66)	–	C_{10}(100%)	C_{10}(100%)	1.52	1.67	1.00

Table 4.5: Descriptive statistics, by treatment and class of networks

The last three columns display the average number of consecutive stages the groups stayed in a network, which we consider as another marker of the absorbing power of a network. In T1, when groups reach the complete network, they immediately decide to end the game.[30] In T2 and T3 the players cannot decide to stop the game when they reach a VNMFS set, probably due to the asymmetries in their payoffs. Nevertheless they spend more consecutive stages there than in any other class. In T2 this results in a high percentage of groups ending the game in C_5. In T3 the players leave C_7 more often before the end of the game, despite staying there for more than five rounds, on average [31]. Consistently, on average a game lasted longer in T3 (22.93 stages), followed by T2 (21.5) and T1 (17.73).

All the presented results are in line with with Hypothesis 5 as a way to reconcile the aggreagte outcomes. The latters, as said, can not be rationalized using traditional theoretical arguments. In T3, the VNMFS are Pareto efficient, Pareto dominant among the PWS networks, and strongly stable (a condition not met in T2). Our interpretation is that the VNMFS sets are less robust to limitedly farsighted deviations in T3. As discussed in Section 4.3.2 there are farsighted improving paths in two steps leaving any $g \in C_7$ and reaching another network in the same class. Both steps imply a strict improvement in the final network with respect to the current one (see table 4.2). Deviations leaving the VNMFS set in T2 are longer and less feasible as they include some players adding a link only to be as well off in the final network as they are in the current one.

An alternative interpretation would be that the multiplicity of networks that are in a VNMFS set generate coordination problems among the players. This problem is not present in T1 and has an obvious solution in T2, given the sequential nature of the game[32]. In T3, agents with two links are worse off than the agents with one link, in network class in C_7. Hence, agents have a strategic incentive to build only one link, and let the others build two. However, this interpretation is confuted by our data. According to it, we would observe the agents having problems in reaching C_7, and not moving away once they are there. We observe almost the opposite. As shown in Table 4.5, in T3 the groups ended the game in C_7 only in ten out of the forty-two times they accessed it. The same ratio (for C_5) is twenty-nine out of fifty-five in T2. Thus the groups have more problems

[30] This fact explains why g^N displays a low average stay, despite it is the final network for a majority of the groups.

[31] Note the relatively high numbers for C_5 in T1 and T3; those networks feature relatively low payoffs and are not Nash stable (the connected players can be better off by cutting two of their existing links), though they are PWS. Note also the high number for C_9 in T2. Those networks are often reached when an unsatisfied player in a VNMFS set takes a non-myopic move. As expected, this deviation is generally unsuccessful, in the sense that the group is stuck in C_9 until a backward move is taken by the same player.

[32] As the connected agents in a VNMFS set are better off, the first agents that are proposed a link on a path to C_5 should build them.

staying in C_7 than accessing it.

We thus explore the relevance of limited farsightedness, analyzing individual behavior. Before doing so, we should stress that limited farsightedness, as its extreme counterparts, is meant to be a tool for assessing the stability of a certain state. As such, it should not be interpreted as a proper model of individual strategic behavior, and the following analysis should be understood accordingly.

We build the vectors of choices of virtual players endowed with different levels of limited farsightedness, including myopia, according to Definition 3. Those are vectors of dummies, $f^{ij}_{i_k g_t}$, for $k = 1, 2, \ldots$, containing the ideal actions of an individual i, with level of farsightedness k, choosing over link ij in netwrok g_t.

Recall that an action is consistent with farsightedness of level k if it lies on a farsighted improving path of length (weakly) shorter than k; $k = 1$ is identical to myopia. To lie on a farsighted improving path, an action must imply moving from the current network. Categorizing choices that imply inaction - i.e. staying in the current network - is more problematic. According to definition 3, those actions are consistent with farsightedness of level k if moving would not be farsighted of level k, which equals assuming that a farsighted agent should always take any farsighted improving path, imposing a strong restriction on farsighted behavior.[33] As a throughout theoretical analysis of limited farsightedness goes beyond the scope of the present paper, we will tackle this issue by running the analysis twice, on two set of decisions: the full set of choices, and its restriction to the actions that imply moving from the current network - i.e. excluding those that result in inaction. We will refer to the former set as *choices*, and to the latter as *paths*[34].

In Figure 4.6 we represent the fraction of choices that are consistent with myopia and progressive levels of limited farsightedness, over stages. Starting relatively low, the fraction of choices that are consistent with myopia remains approximately stable, above 60 percent, in the central part of the game, and is somewhat higher in the last stages. Including farsightedness of level two increases the fraction of consistent choices by about 15 percentage points. Another 10 percent is added by farsightedness of level three, whilst higher level of farsightedness result in improvements that are only marginal[35]. This picture suggests, once more, that myopic incentives were a main guide for decision making in our framework; however agents often departed from those, following short farsighted deviations, with relevant consequences on the final outcomes.

We perform a regression analysis to explore the relation between the actual choices, $a^{ij}_{ig_t}$, and the ideal benchmarks, $f^{ij}_{i_k g_t}$, up to a level of farsightedness of four. This exercise

[33]We note that this restriction is not problematic for myopic behavior.

[34]This set actually identifies the paths - i.e. sequences of different networks - the groups walk through

[35]The picture is qualitatively similar across treatments

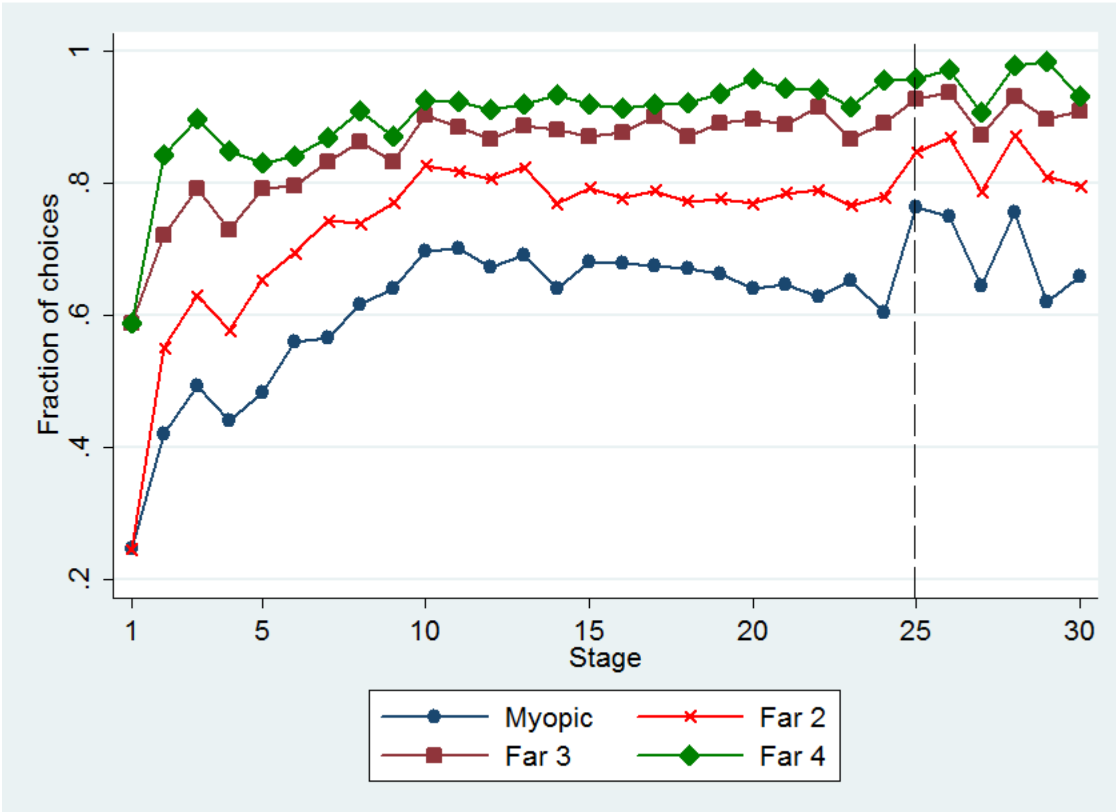

Figure 4.6: Fraction of choices consistent with each behavioral benchmark

suffer from many statistical limitations. In particular, the number of choices each agent takes is endogenous, as groups can decide when to stop a game. We apply a two-steps Heckman selection model (Heckman, 1979) to address this issue.[36]

We estimate a (panel) linear probability model (LPM) with random effects, where the actions $a^{ij}_{ig_t}$ are regressed, conditional on being observed, over the benchmark choices, $F^{ij}_{ig_t} = \{f^{ij}_{ikg_t}\}^4_{k=1}$, and a set of controls, $X^{ij}_{ig_t}$, including characteristics of the choice problem and of the individual. The unobservable characteristics of the individual i, assumed independent from the attributes of the decision problem, are captured by ν_i, resulting in the LPM specification:

$$P(a^{ij}_{ig_t} = 1 \mid z^*_i > 0, F^{ij}_{ig_t}, X^{ij}_{ig_t}) = \beta_0 + F^{ij}_{ig_t}\beta + \hat{\lambda}_{it}\beta_\lambda + X^{ij}_{ig_t}\gamma + \nu_i \qquad (4.1)$$

[36]We are aware of the limitations of this approach in the case of a binary independent variable; Nicoletti and Peracchi (2001) show that the bias of two-stage methods might not be severe when the correlation of unobservables is low.

where $\hat{\lambda}_{it}$ is the estimate of the inverse Mills ratio from the selection equation:

$$z^*_{it} = \delta W_{it} + u_i$$

$$z_{it} = \begin{cases} 1 & \text{if } z^*_{it} > 0 \\ 0 & \text{if } z^*_{it} \leq 0 \end{cases} \tag{4.2}$$

where z^*_{it} is the latent variable capturing the propensity of a choice to be selected, and z_{it} is a dummy variable indicating whether we observe the choice or not. We use as restrictions[37] in the selection equation dummies for each treatment and for each type of final outcome (myopic, farsighted, other). We do not include the treatments in the main regression because we have no reason to think that they have any effect on single decisions, but for the effect of the different payoffs, which are already accounted for through our main regressors. A similar reasoning holds for the final outcomes of the groups. The restrictions are justified as both the treatments and the group final outcomes are relevant determinants of the time when the agents stop the game.

We run this specification on both *choices* and *paths*, with and without group fixed effects[38], giving the four specifications shown in Table 4.6. There is a major shift between the left-hand side and the right-hand side specifications. When considering choices, myopic behavior and farsightedness of level two have a positive and significant coefficient. For higher levels of farsightedness the coefficient eventually turns negative (though not significantly different form zero). The picture is reverted with paths. Myopia has a negative and significant coefficient, and those for all farsightedness levels are positive and significant. Level two is the only variable to show a stable explanatory power across specifications.

Combining the results, we see that subjects often refused to move from the current network, because of myopic incentives. When they move, they do so also against their immediate interest, following farsighted improving paths (even of relatively high length); nevertheless, they regularly do not move even if a lengthy farsighted improving path is available. The coefficients for farsightedness of level two (and myopia) suggest that agents generally take myopic and farsighted improving paths of length two when those are available, and refuse to move when neither of the two is.

This interpretation is consistent with the aggregate results, and in particular with the observation that PWS networks express a high absorbing power, even in those cases

[37] That is, we include those variables in W_{it}, but not in $X^{ij}_{ig_t}$.

[38] Errors are always clustered at the group level.

	Choices Group effects: No	Choices Group effects: Yes	Paths Group effects: No	Paths Group effects: Yes
Myopic	.150***	.142***	−.144***	−.178***
	(.027)	(.027)	(.013)	(.014)
Farsighted 2	.048**	.040*	.051***	.052***
	(.022)	(.023)	(.014)	(.014)
Farsighted 3	.014	.031	.059***	.122***
	(.028)	(.029)	(.019)	(.020)
Farsighted 4	-.020	−.035*	.379***	.364***
	(.020)	(.019)	(.036)	(.035)
Inv. Mills	.116***	.013	.079***	.004
	(.023)	(.034)	(.011)	(.018)
Constant	.671***	.888***	.534***	.640***
	(.119)	(.143)	(.069)	(.069)
N. obs	6166	6166	3003	3003
N. subjects		288		
N. groups		72		

*,**,*** statistical significant at the 10%, 5% and 1% level, respectively.
Controls include the stage, the repetition and a set of individual characteristics.

Table 4.6: Estimates results for the main regressions of a two-steps Heckmen selection model (Robust Std errors in parentheses)

where they are eventually left by the subjects. The results for farsightedness of level two are suggestive, as it is exactly the level that would explain the differences between T3 and the other treatments. Overall, low levels of farsightedness look like important determinants of individual behavior.

Result 3: *Individual behavior is best explained by low level of farsightedness (nesting myopia). Despite the observed impact of myopic incentives, the subjects often disregard them and take farsighted deviations. This limitedly farsighted behavior consistently explains the differences across treatments, supporting Hypothesis 5 as a rationale for Result 1.*

We are aware that the statistical approach suffers from important limitations. We do not properly take into account the effect of the past choices of the same individual and of the group, though it is likely that the path of a group has a huge influence on the behavior of the subjects. Moreover, the different results for paths and choices are, at least partially, an artifact of the way in which the regressors are constructed. In particular, the ex ante probability that modifying a network is consistent with some level of farsightedness is increasing in the level itself.

Despite those limitations, the interpretation of the aggregate results as determined by the behavior of limitedly farsighted agents is supported by our analysis.

4.5 Conclusion

This paper reports an experimental test of the most used stability notions for network formation. In particular, by studying the performance of pairwise stability and of von Neumann-Morgenstern farsighted stability, we test whether subjects behave according to myopia or farsightedness when forming a network, allowing for limited levels of farsightedness. As far as we know this is the first experimental investigation into this issue.

The results show that both of the extreme theories, perfect myopia and farsightedness, are inconsistent with our data, and suggest that the subjects are only limitedly farsighted. Agents reach a stable network in 75 percent of the cases, and more so as the game is repeated. In two of the treatments, a vast majority reach a von Neumann-Morgenstern farsightedly stable set. In the third treatment, where the farsighted prediction is not robust to limitedly farsighted deviations, they fail to do so, and 50 percent of them end up in the myopic prediction.

The properties of the treatments enable us to attribute this asymmetry to a form of limited farsightedness, and individual behavior analysis confirms this interpretation: low levels of farsightedness, nesting myopia as the lowest level, best explain our data.

Our results opens the way to new interesting research questions. While the literature has concentrated on the extreme cases of perfect myopia and perfect farsightedness, our experimental results suggests that an intermediate approach could provide a valuable alternative and a promising refinement of pairwise stability.

Appendix C

Proofs

Proof of Proposition 1. To avoid reporting the farsighted improving path for each single network, let g^i be a generic network in class C_i and $c_i \subset C_i$ a generic proper subset of the corresponding class. We will write $g^i \to g$ with $g \in C_j$, and $g^i \to g$ with $g \in c_j$, when the generic network g^i in class C_i reaches with a farsighted improving path all the networks in class C_j or only a proper subset c_j of C_j, respectively.

i In T1 the list of farsighted improving paths among the networks in $\mathbb{G}$ is the following:

$F(g^{\varnothing}) = \{g \mid g \in C_{10} \cup C_{11}\}$
$F(g^2) = \{g \mid g \in C_1 \cup C_{10} \cup C_{11}\}$
$F(g^3) = \{g \mid g \in C_1 \cup c_2 \cup c_5 \cup C_{10} \cup C_{11}\}$
$F(g^4) = \{g \mid g \in C_1 \cup c_2 \cup c_5 \cup C_{10} \cup C_{11}\}$
$F(g^5) = \{g \mid g \in C_1 \cup c_2 \cup C_{10} \cup C_{11}\}$
$F(g^6) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup c_5 \cup C_{10} \cup C_{11}\}$
$F(g^7) = \{g \mid g \in C_1 \cup c_2 \cup c_3 \cup c_4 \cup C_5 \cup C_{10} \cup C_{11}\}$
$F(g^8) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup C_5 \cup c_7 \cup C_{10} \cup C_{11}\}$
$F(g^9) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup c_5 \cup c_6 \cup c_7 \cup C_{10} \cup C_{11}\}$
$F(g^{10}) = \{g \mid g \in c_2 \cup c_4 \cup c_5 \cup c_6 \cup C_{11}\}$
$F(g^N) = \varnothing$.

It follows that $g^N \in F(g)$, for all g in $\mathbb{G} \setminus g^N$ and $F(g^N) = \varnothing$. Thus $\{g^N\}$ is the unique VNMFS set.

ii In T2 the list of farsighted improving paths among the networks in $\mathbb{G}$ is the following:

$F(g^{\varnothing}) = \{g \mid g \in C_5\}$
$F(g^2) = \{g \mid g \in C_1 \cup C_5 \cup c_9\}$
$F(g^3) = \{g \mid g \in C_1 \cup c_2 \cup C_5 \cup C_9\}$
$F(g^4) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup C_5 \cup c_9\}$
$F(g^5) = \{g \mid g \in C_9 \cap \bar{A}_{g^5}\}$

$F(g^6) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup C_5 \cup C_9\}$
$F(g^7) = \{g \mid g \in C_1 \cup c_2 \cup c_3 \cup c_4 \cup C_5 \cup C_9\}$
$F(g^8) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup C_5 \cup c_7 \cup C_9\}$
$F(g^9) = \{g \mid g \in c_4 \cup (C_5 \cap A_{g^9}) \cup (C_9 \setminus g^9)\}$
$F(g^{10}) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup C_5 \cup c_7 \cup c_8 \cup C_9\}$
$F(g^N) = \{g \mid g \in C_5 \cup C_9 \cup C_{10}\}$
The set $\{g \mid g \in C_5\}$ is a VNMFS set. It is reached by any network outside the set and there are no paths between any two networks in the set. Let us check that it is unique.

Consider first a candidate set that does not include any network in C_5. It must then be reached by each single network in C_5, which implies this set should include at least two networks that belong to C_9. Given that $\{g' \mid g' \in C_9 \setminus g\} \subset F(g)$ for every $g \in C_9$, a set including two networks in C_9 is not internally stable.

Now consider a candidate that includes at least one network $g \in C_5$. Then it should include at least one network $g' \in C_9$, such that $g' \notin F(g)$ and $g \notin F(g')$. This condition is impossible as all networks in C_9 that are not adjacent to a network in C_5 are reached by a farsighted improving path form this network, and all networks in C_9 that are adjacent to a network in C_5 reach this network with a farsighted improving path. We conclude that $\{g \mid g \in C_5\}$ is the unique VNMFS set.

iii In T3 the list of farsighted improving paths among the networks in $\mathbb{G}$ is the following:
$F(g^\varnothing) = \{g \mid g \in C_7\}$
$F(g^2) = \{g \mid g \in C_1 \cup C_7 \cup c_{10}\}$
$F(g^3) = \{g \mid g \in C_1 \cup c_2 \cup C_7 \cup C_{10}\}$
$F(g^4) = \{g \mid g \in C_1 \cup c_2 \cup c_5 \cup C_7 \cup c_{10}\}$
$F(g^5) = \{g \mid g \in C_1 \cup c_2 \cup C_7 \cup c_{10}\}$
$F(g^6) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup c_5 \cup C_7 \cup C_{10}\}$
$F(g^7) = \{g \mid (g \in C_7 \wedge \exists i \text{ s.t. } d_i(g) \neq d_i(g')) \vee g \in C_{10}\}$
$F(g^8) = \{g \mid g \in C_1 \cup c_4 \cup C_7 \cup C_{10}\}$
$F(g^9) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup c_5 \cup c_6) \cup C_7 \cup C_{10}\}$
$F(g^{10}) = \{g \mid g \in C_1 \cup c_2 \cup c_4 \cup c_5 \cup c_6 \cup C_7 \cup c_8 \cup C_{10} \setminus g^{10}\}$
$F(g^N) = \{g \mid g \in C_7 \cup C_{10}\}$
A network $g \in C_7$ is reached with a farsighted improving path from any other network except for the network $g' \in C_7$, where each single agent has the same degree as in g. By definition each dyad $\{g, g'\}$ is a VNMFS set. Let us check there is no other VNMFS set.

Given the previous argument, any set containing $g \in C_7$ and any other network $g'' \neq g'$ (as defined above) does not satisfy internal stability. Consider now a candidate set that does not include any network in C_7. As it must be reached by networks in C_7, it will necessarily include one and only one (for internal stability) network in C_{10}. Then it must necessarily include $g^{\varnothing}$, which violates internal stability.

APPENDIX D

EXPERIMENTAL INSTRUCTIONS

Welcome to this experiment in decision-making. In this experiment you can earn money. The amount of money you earn depends on the decisions you and other participants make. Please read these instructions carefully. In the experiment you will earn points. At the end of the experiment we will convert the points you have earned into euros according to the rate: 6 points equal 1 Euro. You will be paid your earnings privately and confidentially after the experiment. **Throughout the experiment you are not allowed to communicate with other participants in any way. If you have a question please raise your hand. One of us will come to your desk to answer it.**

Groups

- At the beginning of the experiment the computer will randomly assign you - and all other participants - to a group of 4 participants. Group compositions do not change during the experiment. Hence, you will be in the same group with the same people throughout the experiment.

- The composition of your group is anonymous. You will not get to know the identities of the other people in your group, neither during the experiment nor after the experiment. The other people in your group will also not get to know your identity.

- Each participant in the group will be assigned a letter, A, B, C, or D, that will identify him. On your computer screen, you will be marked 'YOU' as well as with your identifying letter (A, B, C or D). You will be marked with your identifying letter (A, B, C or D) on the computer screens of the other people in your group.

- Those identifying letters will be kept fixed within the same round, but will be randomly reassigned at the beginning of every new round.

Length and articulation of the experiment

- The experiment consists of 3 rounds, each divided into stages.
- The number of stages in each round will depend on the decisions you and the other people in your group make.
- After a round ends, the following will start, with the same rules as the previous: actions taken in one round do not affect the subsequent rounds.

General rules: rounds, stages, formation and break of links

- In each round the task is to form and break links with other members of the group.
- You will have the possibility to link with any other participant in your group. That is, you can end up with any number of links (0, 1, 2 or 3).
- Thus, the number of links that can be formed in your group will be a number between 0 and 6 (0, 1, 2, 3, 4, 5, 6). The set of links that exist in your group at the same time is called a network.
- Your group starts the first stage of every round with zero links.
- In every stage a network of links is formed, based on your and the other group participants decisions. This network is called the current network.
- Your group will enter a new stage with the links that exist in the network that is formed in the previous stage, according to the following linking rules

Stage rules

- In each stage the **computer** will select for each group a single link among the six possible at random. A link cannot be selected twice in two consecutive stages.
- The participants involved in that link will be asked to take a decision in that stage, the others will be informed about the selected link and will be asked to wait for others' decisions.
- If this link does not exist at the beginning of the stage, the decision will be whether to form that link or not. If this link exists at the beginning of the stage, the decision will be whether to keep or to break that link.
- Thus, in each stage at most **one** link can be formed or broken.

Stopping rules

- After every stage you and the other people in your group will be asked if you are willing to modify the current network. You can answer YES or NO.
- If ALL the people in your group answer NO the round ends and the points associated to the current network are considered to compute your earnings.
- If at least one person in your group answers YES, the group moves to the next stage.
- After stage 25 a random stopping rule is added. In this case, even if you or any of the other people in your group are willing to modify the current network, the round will end with probability 0.2.

Earnings

- To every participant in every network is associated a number of points.
- You will receive points according to the network that exists in your group at the end of each round.
- Your total earnings will be the sum of the earnings in each of the 3 rounds.
- Thus, the points associated to the networks you and the other people in your group form at every stage, except for the last of each round, are not considered for the computation of your earnings.
- You are always informed about the points associated to the current network on screen. On the top of your screen, you are always informed of the points you earned in the previous rounds.
- You can learn about the points associated to every other network through the points sheet you find attached to the instructions. It displays the points associated to every class of networks:
  - In every network, the black dots are the participants in the group; the lines are the existing links.
  - Every class of network is characterized by the number of links each participant has.
  - The numbers close to every black dots indicate the number of points a person with that number of links is earning in that specific class of networks.

- An example will clarify the relation between network and points and the developing of the experiment. You will also practice through a training stage.

Concluding remarks

You have reached the end of the instructions. It is important that you understand them. If anything is unclear to you or if you have questions, please raise your hand. To ensure that you understood the instructions we ask you to answer a few control questions. After everyone has answered these control questions correctly the experiment will start.

Bibliography

Atiker, Evren, William S Neilson, and Michael K Price, "Activity Bias and Focal Points in Extensive Form Games," February 2011.

Aumann, Robert J., "Reply to binmore," *Games and Economic Behavior*, 1996, *17* (1), 138–146.

Bala, Venkatesh and Sanjeev Goyal, "A noncooperative model of network formation," *Econometrica*, 2003, *68* (5), 1181–1229.

Bernheim, B Douglas, "Rationalizable strategic behavior," *Econometrica: Journal of the Econometric Society*, 1984, pp. 1007–1028.

Berninghaus, Siegfried K, Karl-Martin Ehrhart, and Marion Ott, "Forward-looking behavior in Hawk–Dove games in endogenous networks: Experimental evidence," *Games and Economic Behavior*, 2011.

Binmore, Kenneth G., "A Note on Backward Induction," *Games and Economic Behavior*, 1996, *17* (1), 135–137.

Binmore, Kenneth J., John McCarthy, Giovanni Ponti, Larry Samuelson, and Avner Shaked, "A backward induction experiment," *Journal of Economic Theory*, 2002, *104* (1), 48–88.

Bloch, Francis and Matthew O Jackson, "Definitions of equilibrium in network formation games," *International Journal of Game Theory*, 2006, *34* (3), 305–318.

Bolton, Gary E., "A comparative model of bargaining: Theory and evidence," *The American Economic Review*, December 1991, *81* (5), 1096–1136.

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong, "A cognitive hierarchy model of games," *The Quarterly Journal of Economics*, 2004, *119* (3), 861–898.

Capablanca, Jose R., *Chess Fundamentals, Revised*, Random House Puzzles & Games, 2006.

Cappelletti, Dominique, Werner Güth, and Matteo Ploner, "Being of two minds: Ultimatum offers under cognitive constraints," *Journal of Economic Psychology*, 2011, *32* (6), 940 – 950.

Carrillo, Juan and Arya Gaduh, "The Strategic Formation of Networks: Experimental Evidence," 2012.

Chwe, Michael Suk Young, "Farsighted coalitional stability," *Journal of Economic Theory*, 1994, *63* (2), 299–325.

Costa-Gomes, M., V.P. Crawford, and B. Broseta, "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, 2003, *69* (5), 1193–1235.

Cox, J.C. and D. James, "Clocks and Trees: Isomorphic Dutch Auctions and Centipede Games," *Econometrica: Journal of the Econometric Society*, March 2012, *80* (2), 883–903.

Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri, "Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications," *Journal of Economic Literature. Forthcoming*, 2012.

Crosetto, Paolo and Marco Mantovani, "Availability of Information and Representation Effects in the Centipede Game," *Jena Economic Research Paper*, 2012, (2012-051).

Crumbaugh, Charles M. Evans Gary W.;, "Effects of prisoner's dilemma format on cooperative behavior," *Journal of Personality and Social Psychology*, April 1966, *3* (4), 486–488.

Devetag, Giovanna and Massimo Warglien, "Games and phone numbers: Do short-term memory bounds affect strategic behavior?," *Journal of Economic Psychology*, 2003, *24* (2), 189 – 202. <ce:title>The Economic Psychology of Herbert A. Simon</ce:title>.

Diasakos, Theodoros M, "Complexity and Bounded Rationality in Individual Decision Problems," *Carlo Alberto Notebooks*, 2008.

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner, "Individual Risk Attitudes: Measurement, Determinants, And Behavioral Consequences," *Journal of the European Economic Association*, 2011, *9* (3), 522–550.

Duffy, Sean and John Smith, "Cognitive Load in the Multi-Player Prisoner's Dilemma Game: Are There Brains in Games?," July 2012. Available at SSRN: http://ssrn.com/abstract=1841523.

Dufwenberg, Martin and Georg Kirchsteiger, "A theory of sequential reciprocity," *Games and Economic Behavior*, 2004, *47* (2), 268–298.

_ , Ramya Sundaram, and David J. Butler, "Epiphany in the Game of 21," *Journal of Economic Behavior & Organization*, 2010, *75* (2), 132–143.

Dutta, Bhaskar, Sayantan Ghosal, and Debraj Ray, "Farsighted network formation," *Journal of Economic Theory*, 2005, *122* (2), 143–164.

Ellingsen, Tore and Robert Ostling, "When does communication improve coordination?," *The American Economic Review*, 2010, *100* (4), 1695–1724.

Engelmann, Dirk and Martin Strobel, "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments," *The American Economic Review*, 2004, *94* (4), 857–869.

Eyster, Erik and Matthew Rabin, "Cursed Equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

Falk, Armin and Michael Kosfeld, "It's all about Connections: Evidence on Network Formation," *Review of Network Economics*, 2012, *11* (3).

Fehr, Ernst and Klaus M Schmidt, "A theory of fairness, competition, and cooperation," *The quarterly journal of economics*, 1999, *114* (3), 817–868.

Fey, M., R.D. McKelvey, and T.R. Palfrey, "An experimental study of constant-sum centipede games," *International Journal of Game Theory*, 1996, *25* (3), 269–287.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 2007, *10* (2), 171–178.

Gneezy, Uri, Aldo Rustichini, and Alexander Vostroknutov, "Experience and insight in the Race game," *Journal of Economic Behavior & Organization*, 2010, *75* (2), 144–155.

Goeree, Jacob K, Arno Riedl, and Aljaž Ule, "In search of stars: Network formation among heterogeneous agents," *Games and Economic Behavior*, 2009, *67* (2), 445–466.

Groot, Adriaan D. De, Fernand Gobet, and Riekent W. Jongman, *Perception and memory in chess: Studies in the heuristics of the professional eye.*, Van Gorcum & Co, 1996.

Harrison, Glenn W. and E. Elisabet Rutström, "Risk Aversion in the Laboratory," in James C. Cox and Glenn W. Harrison, eds., *Risk Aversion in Experiments*, Vol. 12 of *Research in Experimental Economics*, Emerald Group Publishing Limited, 2008, pp. 41–196.

_ and Kevin A. McCabe, "Expectations and fairness in a simple bargaining experiment," *International Journal of Game Theory*, 1996, *25* (3), 303–327.

Herings, P, Ana Mauleon, and Vincent J Vannetelbosch, "Rationalizability for social environments," *Games and Economic Behavior*, 2004, *49* (1), 135–156.

_ , _ , and Vincent Vannetelbosch, "Farsightedly stable networks," *Games and Economic Behavior*, 2009, *67* (2), 526–541.

Hixon, J. Gregory; Stein-Seroussi Alan; Gilbert Daniel T Swann William B.;, "The fleeting gleam of praise: Cognitive processes underlying behavioral reactions to self-relevant feedback," *Journal of Personality and Social Psychology*, Jul 1990, *59* (1), 17–26.

Ho, Teck-Hua and Xuanming Su, "A Dynamic Level-k Model in Centipede Games," *Management Science*, 2013.

Jackson, Matthew O and Alison Watts, "The evolution of social and economic networks," *Journal of Economic Theory*, 2002, *106* (2), 265–295.

— and **Anne Van Den Nouweland**, "Strongly stable networks," *Games and Economic Behavior*, 2005, *51* (2), 420–444.

— and **Asher Wolinsky**, "A strategic model of social and economic networks," *Journal of economic theory*, 1996, *71* (1), 44–74.

Jehiel, Philippe, "Limited horizon forecast in repeated alternate games," *Journal of Economic Theory*, 1995, *67* (2), 497–519.

— , "Learning to play limited forecast equilibria," *Games and Economic Behavior*, 1998, *22* (2), 274–298.

— , "Repeated games and limited forecasting," *European Economic Review*, 1998, *42* (3-5), 543–551.

— , "Limited foresight may force cooperation," *Review of Economic Studies*, 2001, *68* (2), 369–391.

— , "Analogy-based expectation equilibrium," *Journal of Economic theory*, 2005, *123* (2), 81–104.

Johnson, Eric J., Colin Camerer, Sankar Sen, and Talia Rymon, "Detecting failures of backward induction: Monitoring information search in sequential bargaining," *Journal of Economic Theory*, 2002, *104* (1), 16–47.

Kagel, John H. and Alvin E. Roth, *The handbook of experimental economics*, Vol. 1, Princeton University Press Princeton, 1995.

Kawagoe, Toshiji and Hirokazu Takizawa, "Level-k analysis of experimental centipede games," *Journal of Economic Behavior & Organization*, May 2012, *8* (2-3), 548–566.

Konishi, Hideo and Debraj Ray, "Coalition formation as a dynamic process," *Journal of Economic Theory*, 2003, *110* (1), 1–41.

Kosfeld, Michael, "Economic networks in the laboratory: A survey," *Review of Network Economics*, 2004, *3* (1), 20–41.

Kreps, D., *Game Theory and Economic Modeling*, Oxford University Press, 1990.

Levitt, Steven D., John A. List, and Sally E. Sadoff, "Checkmate: Exploring Backward Induction among Chess Players," *The American Economic Review*, 2011, *101* (2), 975–990.

Levy, David N. L. and Monty Newborn, *How computers play chess*, Computer Science Press, Inc., 1991.

Maniadis, Zacharias, "Aggregate Information and the Centipede Game: a Theoretical and Experimental Study," Nov 2011.

Mauleon, Ana and Vincent Vannetelbosch, "Farsightedness and cautiousness in coalition formation games with positive spillovers," *Theory and Decision*, 2004, *56* (3), 291–324.

McKelvey, Richard D. and Thomas R. Palfrey, "An experimental study of the centipede game," *Econometrica*, 1992, pp. 803–836.

McKelvey, Richard D and Thomas R Palfrey, "Quantal response equilibria for normal form games," *Games and economic behavior*, 1995, *10* (1), 6–38.

McKelvey, Richard D. and Thomas R. Palfrey, "Quantal response equilibria for extensive form games," *Experimental economics*, 1998, *1* (1), 9–41.

Nagel, Rosemarie and Fang Fang Tang, "Experimental results on the centipede game in normal form: an investigation on learning," *Journal of Mathematical Psychology*, 1998, *42* (2), 356–384.

Ochs, Jack and Alvin E. Roth, "An Experimental Study of Sequential Bargaining," *The American Economic Review*, 1989, *79* (3), 355–384.

Page, Frank H and Myrna Wooders, "Strategic basins of attraction, the path dominance core, and network formation games," *Games and Economic Behavior*, 2009, *66* (1), 462–487.

__ **, Myrna H Wooders, and Samir Kamat**, "Networks and farsighted stability," *Journal of Economic Theory*, 2005, *120* (2), 257–269.

Palacios-Huerta, Ignacio and Oscar Volij, "Field Centipedes," *The American Economic Review*, 2009, *99* (4), 1619–1635.

Ponti, Giovanni, "Cycles of Learning in the Centipede Game," June 1996, (96-22 ISSN 1350-6722).

Pruitt, Dean G, "Reward Structure and Cooperation: the Decomposed Prisoner's Dilemma Game," *Journal of Personality and Social Psychology*, Sep 1967, *7* (1), 21–27.

Rapoport, Amnon, William E. Stein, James E. Parco, and Thomas E. Nicholas, "Equilibrium play and adaptive learning in a three-person centipede game," *Games and Economic Behavior*, 2003, *43* (2), 239–265.

Reynolds, Robert I., "Search heuristics of chess players of different calibers," *The American journal of psychology*, 1982, *95* (3), 383–392.

Roomets, Alex, "On Limited Foresight in Games," 2010.

Rosenthal, Robert W., "Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox," *Journal of Economic Theory*, 1981, *25* (1), 92–100.

Selten, Reinhard, "Features of experimentally observed bounded rationality," *European Economic Review*, 1998, *42* (3-5), 413–436.

Shiv, Baba and Alexander Fedorikhin, "Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making," *Journal of Consumer Research*, December 1999, *26* (3), 278–292.

Smith, Vernon L., "Economics in the Laboratory," *The Journal of Economic Perspectives*, 1994, *8* (1), 113–131.

Sobel, Joel, "Interdependent preferences and reciprocity," *Journal of Economic Literature*, 2005, *43* (2), 392–436.

Stahl, Dale O and Paul W Wilson, "On Players? Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 1995, *10* (1), 218–254.

Tercieux, Olivier and Vincent Vannetelbosch, "A characterization of stochastically stable networks," *International Journal of Game Theory*, 2006, *34* (3), 351–369.

Tversky, Amos and Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice," *Science*, 1981, *211*, 453–458.

Watts, Alison, "A dynamic model of network formation," *Games and Economic Behavior*, 2001, *34* (2), 331–341.

__ , "Non-myopic formation of circle networks," *Economics Letters*, 2002, *74* (2), 277–282.

Xue, Licun, "Coalitional stability under perfect foresight," *Economic Theory*, 1998, *11* (3), 603–627.

Zauner, Klaus G, "A Payoff Uncertainty Explanation of Results in Experimental Centipede Games," *Games and Economic Behavior*, 1999, *26* (1), 157–185.

Ziegelmeyer, Anthony and Katinka Pantz, "Collaborative networks in experimental triopolies," Technical Report, Max Planck Institute of Economics, Strategic Interaction Group 2005.