DEPARTMENT OF ECONOMICS,
MANAGEMENT AND STATISTICS

UNIVERSITY OF MILAN – BICOCCA

DEMS WORKING PAPER SERIES

Limited backward induction: foresight and behavior in sequential games

Marco Mantovani

No. 289 – December 2014

Dipartimento di Economia, Metodi Quantitativi e Strategie di Impresa
Università degli Studi di Milano - Bicocca
http://dems.unimib.it/

Limited backward induction: foresight and behavior in sequential games

Marco Mantovani*

December 2014

Abstract

The paper tests experimentally for limited foresight in sequential games. We develop a general out-of-equilibrium framework of strategic thinking based on limited foresight. It assumes the players take decisions focusing on close-by nodes, following backward induction – what we call limited backward induction (LBI). The main prediction of the model is tested in the context of a modified Game of 21. In line with the theoretical hypotheses, our results show most players think strategically only on close-by nodes without reasoning backwards from the end of the game. A small fraction of subjects play close to equilibrium, while few others try to exploit the limited foresight of their opponent. The results provide strong support for LBI, and cannot be accounted for using the most popular models of strategic thinking, let alone equilibrium analysis.

JEL classification: D03, C72, C91

Keywords: Behavioral game theory, sequential games, strategic thinking, level-k, limited foresight.

*Department of Economics, University of Milan Bicocca. E-mail: `marco.mantovani@unimib.it`

Acknowledgements: the author would like to thank Martin Duwfenberg, Antonio Filippin, Georg Kirchsteiger, Astrid Gamba, Frank Heinemann, Nagore Iriberri, Arno Riedl, Bram De Rock and Alex Roomets as well as participants to presentations in Bari, Berlin, Brussels, Maastricht, Milan, and New York for valuable comments and suggestions to improve the paper. Financial support from the FSR-Marie Curie incoming postdoc program of the Académie de Louvain is gratefully acknowledged.

1 Introduction

How do you figure out moves in a chess game? Presumably, most people think of what the other is going to do next. Some think of their own next move as well, and maybe of the opponent's choice that follows. Deep consideration of further stages characterizes chess lovers and professionals.[1] When, within this horizon, a player conceives his move as a best-response to the ones that follow, and those are deduced in the same way, he is performing backward induction on a limited number of stages. In a nutshell, this represents what we call limited backward induction (LBI). It is reasonable that one looks further ahead when one of the kings is menaced, or as the action space shrinks throughout the game. This paper presents a framework that catches such features of strategic thinking in sequential games of perfect information, and a novel experiment that tests for it.

The benchmark for strategic ability set by game theory is hardly matched by human beings, as widely documented by the experimental literature.[2] Backward induction is no exception.[3] Different studies have suggested that the players use limited look-ahead in extensive form games [e.g., Binmore et al., 2002; Johnson et al., 2002]. For instance, in Johnson et al. [2002], the subjects need to uncover the payoffs at different stages by clicking on them. The authors show that the players focus on the current and the following stage, paying little attention to more distant ones. Despite these results and the underlying intuition, until recently, no model of strategic thinking addressed the specific challenge that dynamic strategic environments pose to individual reasoning in the presence of limited foresight, 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 of stages on which it is performed.

Under LBI, a player faces what we call a limited-foresight game (LF-game), which is a section of a subgame of the original game that only encompasses the future nodes that are closest to the current decision node. The nodes that are included in the LF-game depend on the player's foresight. This is endogenously determined after considering how relevant the stakes are – i.e., how beneficial strategic thinking is – and how complex the game is – i.e., how costly strategic

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

[2] See, among other excellent surveys, the one in Selten [1998].

[3] On the theoretical side, backward induction has been the object of a lively debate [cf., Aumann, 1995; Battigalli, 1997; Ben-Porath, 1997; Binmore, 1996]. See Brandenburger and Friedenberg [2014a] for an axiomatic approach to backward induction and Bonanno [2001] for a foundation based on temporal logic.

thinking is – in the original game. The foresight determines the nodes that are object of strategic consideration.

The LF-game is then completed by assigning intermediate payoffs to its terminal nodes. This is done by projecting the payoffs that arise beyond one's foresight on the terminal node of the LF-game with which they are consistent. Finally, actions consistent with a subgame perfect equilibrium (SPE) of the LF-game are taken. The higher the foresight, the more nodes are included in the LF-game so that, in the limit, this coincides with the original game, and the actions taken are consistent with subgame perfection.

We use a variant of the Game of 21 [Dufwenberg et al., 2010; Gneezy et al., 2010] to identify LBI. In this simple game players alternate to choose numbers within a range. Those are added up, until a certain target number is reached. The player who reaches it wins a prize, the other gets nothing. By picking the correct numbers, one of the players can secure the victory from the first move. This advantage transfers to his opponent in case of error. By backward induction, one can identify a set of dominant strategies where these opportunities are always exploited. Any SPE is in weakly dominant strategies. Thus, level-k players also play consistent with equilibrium. This game is also known as the race game, and we use this more generic label, as our target number is not 21. Consistently with LBI, previous results [Dufwenberg et al., 2010; Gneezy et al., 2010; Levitt et al., 2011] show little compliance with SPE. The subjects find it hard to substitute a subgame with its outcome, and discover the solution only as they approach the end of the game.[4]

We design an experiment with two treatments. In a baseline, the subjects play a standard race game. In a second treatment, we introduce an intermediate small prize. In both treatments, subjects have the possibility (incentivized) to claim victory of any prize at any time in the game. The design allows us to identify reasoning based on limited foresight. Evidence of limited foresight comes from individuals solving for the intermediate prize before they do for the final one. However, this seemingly optimal behavior on a LF-game may come not only from limited foresight, but also from attempts to exploit the limited ability of the opponent. Claims and the comparison with the baseline treatment are used to test that the observed actions on the path to the intermediate prize are due to limited foresight and not to this type of sophisticated behavior.

[4]Dufwenberg et al. [2010] show that solving a shorter game helps finding the solution to a longer one and that learning is gradual. Gneezy et al. [2010] find that subjects discover the positions on the path to the prize sequentially, starting from the last one, by backward analysis. Levitt et al. [2011] report an interesting correlation between a subject's performance in the race game and his ranking as chess player, suggesting the game can capture the ability at backward induction.

Our results show most subjects solve for the intermediate prize before they do for the final one. That is, they reason backwards only on close-by nodes. The majority of the population proves able to run no more than two or three iterations of backward induction in our race game. A smaller fraction plays consistent with equilibrium, and few players use sophisticated strategies of the type described above.[5] Consistently, reasoning efforts, as recorded by the timing of decisions, are concentrated in the nodes that are closest to a prize – i.e., when this enters the LF-game. On aggregate, we provide clear evidence of limited foresight, and in support of LBI.

In the last twenty years, different models have addressed limited strategic ability.[6] Most of them target simultaneous-move games, but some have been adapted to extensive-form ones. For instance, McKelvey and Palfrey [1998] propose the agent quantal response equilibrium (AQRE) as the QRE counterpart for extensive-form games. Ho and Su [2013] and Kawagoe and Takizawa [2012] adapt level-k models to dynamic games.[7] In these models the key ingredients mirror their static analogue, and the dynamic context plays no specific role on how the players reason.

Closer to our spirit is the work of Jehiel [1995] on limited forecast models. As we do, they assume the players do not make a complete plan of actions, and only reason on a limited number of stages. However, on the one hand, Jehiel proposes an equilibrium notion, contrary to our out-of-equilibrium one. On the other, the way in which the players produce their forecasts is our main object of investigation, while it is not specified in the limited forecast equilibrium. Thus, the two approaches differ in their objective. So will, in general, their predictions.[8]

An independent attempt, similar to ours, is being carried out by Roomets [2010]. Contrary to LBI, in his Horizon-Based Limited Foresight model the foresight is exogenous and the way in which intermediate payoffs are derived is largely unspecified. Most notably, in Roomets [2010] there is no experimental test of the model.

Recently, levels of reasoning have been endogenized in a level-k model in a way similar to ours by Alaoui and Penta [2014]. LBI also shares many other features with standard level-k models [Costa-Gomes et al., 2001; Stahl and Wilson, 1995]. Both are based on a hierarchy of decision rules, where each level best

[5] Experiments on level-k thinking have suggested a similar distribution of levels. Note, however, that the *type* of reasoning investigated is different.

[6] See Crawford et al. [2013] for a survey of those models.

[7] See also Brandenburger and Friedenberg [2014b] for a more general discussion of levels of reasoning not based on limited foresight in extensive-form games.

[8] See also Jehiel [1998a] where a learning justification is provided for the limited forecast equilibrium, and Jehiel [1998b, 2001] where the concept is applied to repeated games.

responds to the rule that is one step lower in the hierarchy. The chain of best replies is anchored to the behavior of a non-strategic level, which, in our case, is represented by the player choosing before the foresight bound. As level-k, LBI constitutes an out-of-equilibrium model of behavior, and should be understood to capture initial responses to a game. However, the underlying type of reasoning is different, and LBI predictions generally diverge from the level-k ones, as made clear by our experiment.

The paper proceeds as follows. Section 2 lays out the theoretical framework. Section 3 presents the experimental design and procedures. Results follow in Section 4. Section 5 shows through examples how LBI can be applied to most studied sequential games, and discusses the relation with level-k and the limited forecast equilibrium. Section 6 concludes.

2 Theoretical framework

2.1 A sketch of Limited Backward Induction

Figure 1 shows a four-stage game. Each outcome $\mathbf{a}, ..., \mathbf{p}$ is a vector in $\mathbb{R}^2$, identifying von Neumann-Morgenstern utilities for each player. Player 1 knows what Player 2 is choosing at every node in stage four. Reasoning by backward induction, he substitutes each subgame with its solution and iterates this procedure backwards, finally choosing his best reply in stage one.

Suppose Player 1 has only limited foresight, and is not able to run backward induction from the terminal nodes of the game. The dashed line after stage two represents his foresight bound. He best replies to what he believes the next mover is choosing, but cannot substitute the subgames that follow with their solution. His level of foresight is two, and solves an LF-game in two stages. We call pseudo-terminal nodes (histories) the terminal nodes (histories) of the LF-game.

FIGURE 1: A FOUR-STAGE SEQUENTIAL GAME WITH A FORESIGHT BOUND

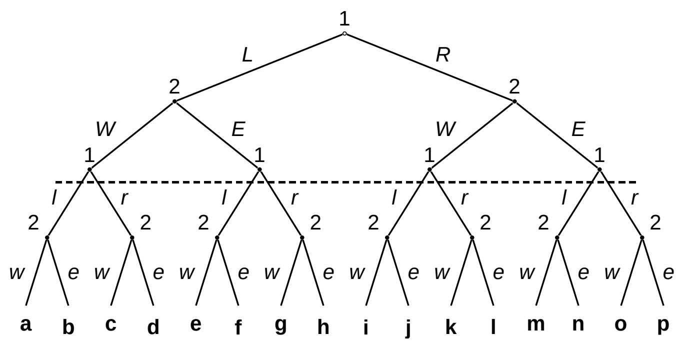

To figure out what Player 2 is choosing in stage 2, Player 1 needs some intermediate payoffs for the pseudo-terminal nodes of the LF-game. We assume the players know the terminal payoffs of the game, and use this information to retrieve the intermediate payoffs.[9] Each payoff of the LF-game is derived as a projection of the payoffs of the original game that are consistent with each pseudo-terminal history. For instance, the payoffs of the LF-game at the pseudo-terminal history (L, W) are a projection 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 had a foresight of three stages, the dashed line in figure 1 would move one stage downwards. The LF-game would be larger, and its payoffs would be derived from smaller sets of final outcomes. In general, we propose as projection function the median point in the range of the available payoffs.[10]

We label F_κ the decision rule of a player with a foresight of κ stages. In reasoning backward from the κ-th stage, he implicitly believes that the agent choosing there acts according to F_1. This rule is non-strategic, in the sense that it does not consider the decisions of any other player, and serves as an anchor of the LBI reasoning. The agents controlling the next-to-last decision nodes are believed to act according to F_2, and so on. Thus, F_κ best responds to $F_{\kappa-1}$, who best responds to $F_{\kappa-2}, \ldots$, who best responds to F_1.

Provided that the foresight bounds of a sequence of players do not coincide, the beliefs about the next players' moves are generally incorrect. Moreover, the moves of one single player in different nodes need not be consistent one with the other.[11]

2.2 General notation

Consider a finite game of perfect information $\Gamma = \langle \mathcal{I}, (N, \preccurlyeq), l, U \rangle$, where $\mathcal{I} = \{1, 2, ..., I\}$ is the set of players, and N is the set of nodes, partially ordered by the successor relation $\preccurlyeq$, so that $n' \preccurlyeq n$ if and only if node n' is on the unique path from the initial node ϕ to n. The terminal (resp. non-terminal) nodes are Z (resp. H); perfect information implies information sets are singletons $\{h\}$, where $h \in H$. For simplicity, we will generally write h for $\{h\}$. The mapping $l : H \to \mathcal{I}$ specifies the player who moves at each information set. Denote H_i the set of information

[9] The assumption that players know the terminal payoffs, but ignore the portion of the game tree that leads to them may seem troublesome at first glance. It is, however, natural in most contexts: players know the possible payoffs of a chess game, the size of the cake in a bargaining problem, the possible outcomes of a sequential voting mechanism, etc. Their knowledge of these payoffs is unrelated to their knowledge and consideration of the game tree that leads to them. Nevertheless we relax this assumption in Section 5.2.

[10] That is: the simple average between the minimum and the maximum of the available payoffs. The issues related to the projection function are discussed in Section 2.4.

[11] See Section 5.3 for a discussion of the issue.

sets where i moves. At each information set in H_i, player i's set of moves is $A_i[h]$, where each move $a_i[h] \in N$ is an immediate successor of h.

The successor relation induces a temporal rank on the nodes: ϕ is of rank 1; n is of rank t if the maximum of the rank of n' such that $n' \prec n$ is $t-1$. Let $r(n)$ denote the rank of n. The extensive-form payoff functions are given by $u_i : Z \to \mathbb{R}$, and $U = (u_1, u_2, \ldots, u_I)$. Let $\mathcal{Z}$ be the set of all subsets of Z. Function $\zeta : H \to \mathcal{Z}$ is such that $\zeta(h)$ maps h to $\{z \in Z : h \prec z\}$ – the terminal nodes that are successors of h. Finally denote with $l_{\tilde{H}}$ the restriction of l to $\tilde{H} \subseteq H$.

2.3 Foresight

A player's *foresight* represents his depth of strategic thinking. It is a positive integer that identifies the LF-game tree, and is endogenously derived from a local optimization problem.[12] We define the functions $b_i : H \times \mathbb{N} \to \mathbb{R}_+$ and $c_i : H \times \mathbb{N} \to \mathbb{R}_+$, where $b_i(h,s)$ (resp. $c_i(h,s)$) represents the incremental benefit (resp. cost) at node h from increasing one's foresight from $s-1$ to s stages. We assume that the benefit from reasoning depends on the payoffs that are achievable at node h. The cognitive cost depends on the complexity of the subgame with initial node h as well as on the player's cognitive ability. We do not impose restrictions on the shape of these functions, apart from the following.

Assumption 1. *The benefit and cost of reasoning respect the following conditions:*

i) if, $\forall z, z' \in \zeta(h)$, $u_i(z) = u_i(z') \Rightarrow b_i(h, \cdot) = 0$;

ii) consider two subgames Δ and Δ' rooted in h and h'. If these are identical up to a rescaling of the payoffs, meaning that, for every $z \in \zeta(h)$ and its corresponding $z' \in \zeta(h')$, $u_i(z) = \alpha u_i(z')$, then $\alpha > 1 \Rightarrow b_i(h, \cdot) \geq b_i(h', \cdot)$;

iii) $c_i(\cdot, 0) = 0$;

iv) if $\#\zeta(h) > \#\zeta(h') \Rightarrow c_i(h, \cdot) \geq c_i(h', \cdot)$.

A player increases his foresight as far as the benefit of doing so exceeds its cost. Let the foresight function $k : \mathbb{R}_+^{\mathbb{N}} \times \mathbb{R}_+^{\mathbb{N}} \to \mathbb{N}$ be such that, fixed any node h:

$$k(b_i, c_i) = \min \{ s \in \mathbb{N} : b_i(h,s) \geq c_i(h,s) \wedge b_i(h,s+1) < c_i(h,s+1) \}$$

[12] In a preliminary version (available online) we endogenized the foresight in a similar, though formally different way. As the work of Alaoui and Penta [2014], endogenizing levels of reasoning in a level-k model, was made available in the meantime, we decided to adapt our presentation to theirs when convenient. See also Diasakos [2008], where a similar exercise is made for dynamic individual decision problems.

.

The foresight of player i at h is the value taken by this function at $(b_i(h,\cdot), c_i(h,\cdot))$. We generally denote it with the Greek letter κ. The corresponding decision rule under LBI is F_κ. At node h, κ determines the nodes of the LF-game, $\tilde{N} = \{h' \in N : h \preccurlyeq h' \wedge r(h') - r(h) \leq \kappa\}$, and the corresponding pseudo-terminal nodes $\tilde{Z} = \{z \in \tilde{N} : z \in Z \vee r(z) - r(h) = \kappa\}$.

The intuition that dynamic strategic thinking is affected in the way assumed by exogenous variations in the stakes of the game and its complexity is backed by a number of experimental results. For instance, in the centipede game, most players are consistent with SPE if the stakes are high [Rapoport et al., 2003], and a marginal increase in complexity shifts behavior away from SPE [Crosetto and Mantovani, 2012].

2.4 Intermediate payoffs and the LF-game

The LF-game is completed by associating payoffs to the pseudo-terminal nodes $\tilde{Z}$. We assume intermediate payoffs are projections of the set of consequences that are consistent with the pseudo-terminal node they refer to. Let $v_i : \mathcal{Z} \rightarrow \mathbb{R}$ be a function such that $v_i(\zeta(h))$ maps the payoffs that are achievable at successors of h to the intermediate payoff of h. A generic LF-game is then $\tilde{\Gamma} = \langle \mathcal{I}, (\tilde{N}, \preccurlyeq), l_{\tilde{H}}, V \rangle$, where $\tilde{H} = \tilde{N} \setminus \tilde{Z}$, and $V = (v_1, \ldots, v_n)$.

In applications one wants to make intermediate payoffs operational. As a general rule, we propose v_i to map the viable terminal payoffs into the median value within their range, or, equivalently, the average between the minimum and the maximum in the set:[13]

$$v_i(\zeta(h)) = \frac{\max_{z \in \zeta(h)}(u_i(z)) + \min_{z \in \zeta(h)}(u_i(z))}{2}.$$

This projection function gives a rough idea of what one can gain after a certain node is reached with minimal computation, and a parsimonious use of information arising at distant nodes. A one-fits-all solution is unlikely to be optimal for every possible game. The proposed rule is not sensitive to any manipulation of the payoffs that leaves their range unaffected. This may be questionable, for example, if the payoff structure is particularly simple. However, fixing a projection function restricts the degrees of freedom of the model and we thus discourage

[13] The question of how the subjects project the payoffs is an empirical one. However, any experimental investigation over it must take into account that foresight and projection function are generally jointly tested. In our experiment we manage to disentangle the two by making the projection function irrelevant and focusing on the foresight.

ad-hoc changes to it.[14]

The decision rule F_κ prescribes to take decisions following backward induction on the LF-game. The player choosing at h finds the optimal moves of the LF-subgames whose roots have as immediate successors only nodes in $\tilde{Z}$. These subgames are discarded, leaving behind only their solution. He then iterates this procedure on the immediate predecessors of these subgames, on their predecessors, and so on, until he reaches h. His moves are then consistent with a SPE of the LF-game. The following statements are true.

Proposition 1. i. *F_κ always prescribes at least one move.* **ii.** *If the maximal rank of a node in Γ is $T+1$ and $t+\kappa \geq T+1$, then the moves prescribed by F_κ at h, $r(h)=t$, are all and only those that are part of a SPE of Γ.*

The proofs are self-evident and are omitted. When the foresight of a player reaches the terminal nodes, his moves are consistent with a SPE of the subgame he is deciding in. This entails, in particular, that for $\kappa \to \infty$, F_κ always prescribes all and only the moves that are consistent with a SPE of the game. In general, higher foresight implies earlier consistency with SPE.

3 Experimental design

3.1 The Race Game: parameters and treatments

To test for LBI, we design a novel experiment based on a perfect-information game, known as the race game. In our race game, two players start at position 1 and take turns choosing an integer between 1 and 6. The chosen numbers are added up, so that the position at stage s is given by the the sum of the numbers chosen in stages $1,...,s$ plus one. When a player reaches 66, he wins a prize $P=100$, and the other gets nothing. Any race game can be solved by backward induction. A player wins when choosing at positions $65,...,60$. Thus, a player choosing at 59 is meant to lose. This position can be reached from $58,...,53$, implying that a player choosing at 52 is meant to lose. Iterating this reasoning unveils a sequence of losing positions.[15] A player that reaches any of these positions, can secure the victory of the prize by reaching all subsequent losing positions.

Denote $\mathcal{T}$ the ordered set of all positions; the set of losing positions is $\mathcal{L}=\{t \in \mathcal{T} : t = 66 - i(6+1),\text{ for } i = 1,2,...\}$. The set of winning positions is $\mathcal{W} = \mathcal{T} \setminus \mathcal{L}$.

[14] The issue is similar to that of defining the level zero player in level-k models, or the noise element in QRE. See Hargreaves Heap et al. [2014] for a discussion of the former.

[15] We follow the terminology of Gneezy et al. [2010]. Actually, the player who chooses at a losing position is meant to lose the game, but the player who reaches it is winning.

Thus, the set of losing positions in our race game is

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

A weakly dominant strategy prescribes to reach the closest losing position at winning positions, and choose whatever number at losing positions. Because $1 \notin \mathcal{L}$, Player 1 has an initial advantage, in the sense that he wins the game by playing any dominant strategy. Whenever a player chooses a move from a dominated strategy, the advantage transfers to the opponent.

We introduce two modifications to the base game to identify limited foresight. **Intermediate prize**: the player who reaches the intermediate position $40 \notin \mathcal{L}$ wins a small prize $p = 30$. Denote $\mathcal{L}_p$ and $\mathcal{W}_p$ the set of losing and winning positions on the path to p. As

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

Player 1 has the initial advantage also with respect to prize p.
Claims: players can claim they are going to win p and P at any time, and independently of who is moving. One can claim one or both prizes, eventually at the same time. Claims are irrevocable and non-strategic: a player gets no information on his opponent's claims, and those are not affecting his payoff. However, one's own claims are payoff-relevant: after claiming P (resp. p) at position t, a player gets $66 - t$ (resp. $40 - t$) on top of the prize, in case the claim is realized. In case not, he pays a fine, valued 15.

Our design features two treatments. In $T0$ the subjects play the standard race game with payoff-relevant claims and no intermediate prize. In $T1$ they play the same race game with the intermediate small prize p, which is the only treatment variable.

3.2 Equilibrium and out-of-equilibrium hypothesis

The standard race game has nice features for studying limited foresight. It is a zero-sum game with only two possible outcomes, allowing us to abstract from preference-related issues. The game features a weakly dominant strategy, essentially unique, for each player, which makes beliefs about the others' strategies irrelevant. Finally, the set of payoffs that is achievable after some node is the same across all nodes within a stage. For instance, the same number of terminal nodes where a player wins P follows position 31 and 32, though the former belongs to $\mathcal{L}$ and the latter to $\mathcal{W}$. It follows that the same intermediate payoffs are projected on every pseudo-terminal node of any LF-game. This holds irrespectively of the projection function, and its specification is then irrelevant. We hereafter summarize the predictions of alternative models of behavior, and draw the experimental hypotheses.

Because Player 1 has the initial advantage, he wins P in any SPE. Player 2 wins p in $T1$. They both play a weakly dominant strategy and pass only through positions $\mathcal{L}$ and $\mathcal{L}_p$. A similar statement holds for dynamic level-k models. L_0 identifies a fictitious random player, existing only as a belief in the mind of higher-level players. For $k \geq 1$, each L_k believes the others are L_{k-1} and best replies to this belief. Given that L_0 plays a dominant strategy with positive probability, L_k never plays a dominated strategy. Thus, every level mimics SPE in $T0$. In $T1$, L_1 may try to win both prizes. Higher levels play as in SPE, regardless of the parameters. The players only pass through the losing positions $\mathcal{L}_p$ and $\mathcal{L}$. With a possible exception for the first move, aggregate play is identical to SPE.[16]

According to LBI, new positions enter the LF-game as the game proceeds. In $T0$, if no terminal node is included in the LF-game, the same intermediate payoffs are associated to all pseudo-terminal nodes, and a player is indifferent between his moves. As soon as a terminal node is included in his LF-game, he reaches the positions in $\mathcal{L}$.[17] Previous results indeed indicate that individuals are unable to figure out their dominant strategy from the beginning. They rather discover it as they gain experience, starting from the losing positions that are close to the end. These results can certainly be attributed to limited strategic ability, and are consistent with LBI. However they are not conclusive. In particular, the standard game does not allow to distinguish reasoning on close-by stages – i.e., limited foresight – from reasoning on limited subgames. We call this latter type of reasoning limited subgame perfection (LSP).

According to LSP, backward reasoning starts at the terminal nodes, but is iterated backward for a limited number of stages. A player's strategy switches to equilibrium as soon as he gets to a subgame he can solve.[18] LSP is not distinguishable from LBI in $T0$, but departs from it in $T1$. An LSP player acts consistent with equilibrium over any subgame he can solve. As any subgame represents a smaller race game, he reaches $\mathcal{L}$, and wins P, if choosing at a position in $\mathcal{W}$. He

[16] We here use the dynamic level-k model of Kawagoe and Takizawa [2012]. The authors discuss different types of randomness associated to L_0, which may correspond to different behavior of L_1 in $T1$. The aggregate prediction does not change if we consider more sophisticated models, as the one by Ho and Su [2013], or allowed the players to revise their beliefs after observing a move inconsistent with them.

[17] F_1 realizes how to win only when the distance to the prize is lower than 6, F_2 when it is lower than 12, and so on.

[18] Gneezy et al. [2010] appeal to a sort of LSP in their learning argument.

only reaches $\mathcal{L}_p$ if he is choosing at a position in $\mathcal{L}$. Thus, under LSP, the players reach $\mathcal{L}$ before $\mathcal{L}_p$. Under LBI, in $T1$ prize p is included in the LF-game before P. A player is able to distinguish $\mathcal{L}_p$ from $\mathcal{W}_p$ before $\mathcal{L}$ from $\mathcal{W}$. For instance, he recognizes the difference between $32 \in \mathcal{W}_p$ and $33 \in \mathcal{L}_p$ before that between $31 \in \mathcal{L}$ and $32 \in \mathcal{W}$. Thus, the following hypothesis holds under LBI.

Hypothesis 1. *The positions in $\mathcal{L}_p$ are reached before the positions in $\mathcal{L}$.*

The intermediate prize generates LF-games where optimal behavior departs from the optimal behavior of players reasoning on limited subgames. This is necessary to identify limited foresight. However, in the presence of the intermediate prize there are no longer dominant strategies. Seemingly optimal behavior on an LF-game may result not only from limited foresight, but also from beliefs about the limited ability of the opponent. That is, sophisticated players that try to win both prizes. A sophisticated player postpones reaching $\mathcal{L}$ in $T1$ with respect to what he would do in $T0$, in order to win both prizes. We exploit the fact that sophisticated players play differently in $T1$ with respect to $T0$ to test for this alternative explanation. The following hypothesis holds under LBI.

Hypothesis 2. *The positions in $\mathcal{L}$ are reached at the same time in $T1$ and in $T0$.*

Because other factors may induce different behavior in the two treatments, the test is conservative, but not very informative on sophisticated behavior.[19] An alternative view on the issue is provided by claims. Under general assumptions, claiming decisions are informative on what prize a player is targeting, and allow us to isolate potentially sophisticated behavior at the individual level.

Consider a player that is able to solve a subgame that starts at some position $\bar{t}$, with $\bar{t} < m$ and $\bar{t} \in \mathcal{W} \cap \mathcal{W}_p$. In $T0$ he reaches a position in $\mathcal{L}$, claims and wins P. In $T1$, he either targets P, or *p and then P*, depending on his belief about the strategic ability of his opponent. Including claiming decisions, the player has three undominated strategies: 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 belief about the strategic ability of the opponent is represented by the probability that he does not solve the game within position 40. Let q be this probability. We can prove that in $T1$, for $\bar{t} < 37$ there exists no q such that $S_2 \succeq S_3$ and $S_2 \succeq S_1$.

Thus, given our parameters, a *sophisticated* player claims P as soon as he knows how to win it, and never chooses S_2. Intuitively, in order to choose S_2, a player must believe q is high enough to try to win both prizes, but not so high

[19] For instance, one could think that the presence of p makes the game more complex, resulting in a lower strategic performance of the subjects.

to claim P immediately. In our game, the above interval does not exist at any relevant $\bar{t}$.[20] A derivation of this result can be found in Appendix A. This also includes the analysis of both constant absolute and constant relative risk aversion, where we show that risk aversion is not a concern in our context.[21] Thus, *sophisticated* players may reach $\mathcal{L}_p$ before $\mathcal{L}$, but should not claim p before P. On the other hand, under LBI the players claim a prize as soon as they intentionally reach a position in $\mathcal{L}$ or $\mathcal{L}_p$. As positions in $\mathcal{L}_p$ are reached before positions in $\mathcal{L}$, the following hypothesis hold under LBI.[22]

Hypothesis 3. *Prize p is claimed before prize P.*

Only under LBI none of the previous hypotheses would be rejected.

3.3 Procedures

The experiment took place at the EELAB of the University of Milan-Bicocca in June, 2013. The computerized program was developed using *Z-tree* [Fischbacher, 2007]. The subjects' interface was similar to the one used by Gneezy et al. (2010). For each treatment we run 4 sessions with 12 subjects per session, for a total of 96 participants. Participants were undergraduate students from various disciplines,[23] recruited through an announcement on the EELAB website.

Each subject participated only in one treatment. Subjects played 8 race games. In each of these rounds they were randomly assigned a new partner ('strangers' matching) and a role (Player 1 or 2). We did not impose any time constraint on individual decisions, so that the subjects' strategic performance was not biased by time pressure.

Instructions were read aloud. An English translation of the instructions can be found in Appendix B. Before starting the experiment, participants filled in a control test to ensure everybody understood the instructions. At the end, they were asked to fill in a questionnaire. We gathered qualitative information about the expectations from the game, the opponent, and the strategy followed. We recorded their perception on how easy they found the experiment. Finally, we

[20] In the extremely optimistic scenario of being able to claim P at position 41, an interval of beliefs that sustain S_2 exists if $\bar{t} = 37$ or $\bar{t} = 39$, and is such that $q \in [0.77, 0.79]$.

[21] The intuition for this is is that with risk aversion S_2 becomes more attractive relatively to S_3, but less attractive relatively to S_1. Thus, as before, until very close to m, there is no reasonable parameter of risk aversion that sustains S_2.

[22] For obvious reasons, in any SPE the players claim the prize they are going to win at the initial position, and the same holds for level-k. Under LSP, the players claim P before p.

[23] Economics and business, law, medicine, psychology, mathematics and natural sciences, statistics, pedagogy.

elicited self-reported quantitative measures of risk preferences, using the questions of the SOEP German Panel.[24] Sessions took on average 50 minutes, including instructions, control test and final questionnaire.

During the experiment subjects earned Experimental Currency Units (ECU). At the end, one game was selected at random for each subject, and the ECU earned in the corresponding game were paid to him privately, according to the exchange rate $10ECU = 1€$. The average payment was 11.10€, with a minimum of 3.50€, and a maximum of 25.40€. Subjects received an initial endowment of 5€ that could be partially spent to pay fines.

4 Results

We first analyze actions and claims to test for Hypothesis 1-2. We then report on the timing of decisions. 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. There are ten intervals on the path to P, and six intervals on the path to p.[25]

4.1 Errors

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.

In both treatments, around 80 percent of the subjects choosing at a winning position makes a P-error in the first interval. In intervals 2-7, this percentage declines of about 20 percentage points in $T0$, while remains stable in $T1$. In both treatments the error rate drops sharply in the last three intervals.[26] The rate of p-errors is above 90 percent in the first interval. It then decreases significantly, and is below 40 percent in interval 5.

The moment a subject understands the solution of the game is captured by the last interval where he makes an error. When a subject stops making errors, we can no longer register errors made by his opponent. As a consequence the last errors we observe are a lower bound for the actual ones.[27] Figure 2 shows

[24] See Dohmen et al. [2011]

[25] The intervals for P and p do not perfectly overlap, because $\mathcal{L}$ and $\mathcal{L}_p$ are disjoint. The sets of intervals are: $\{[1,\ldots,3],[4,\ldots,10],\ldots,[60,\ldots,66]\}$ and $\{[1,\ldots,5],[6,\ldots,12],\ldots,[34,\ldots,40]\}$.

[26] As a player choosing at random makes an error 83 percent of the time, the choices of many players are indistinguishable from random play until the last three intervals.

[27] This implies that we are not allowed to extract from them the distribution of foresight levels in the population.

FIGURE 2: AVERAGE LAST ERROR OVER ROUNDS

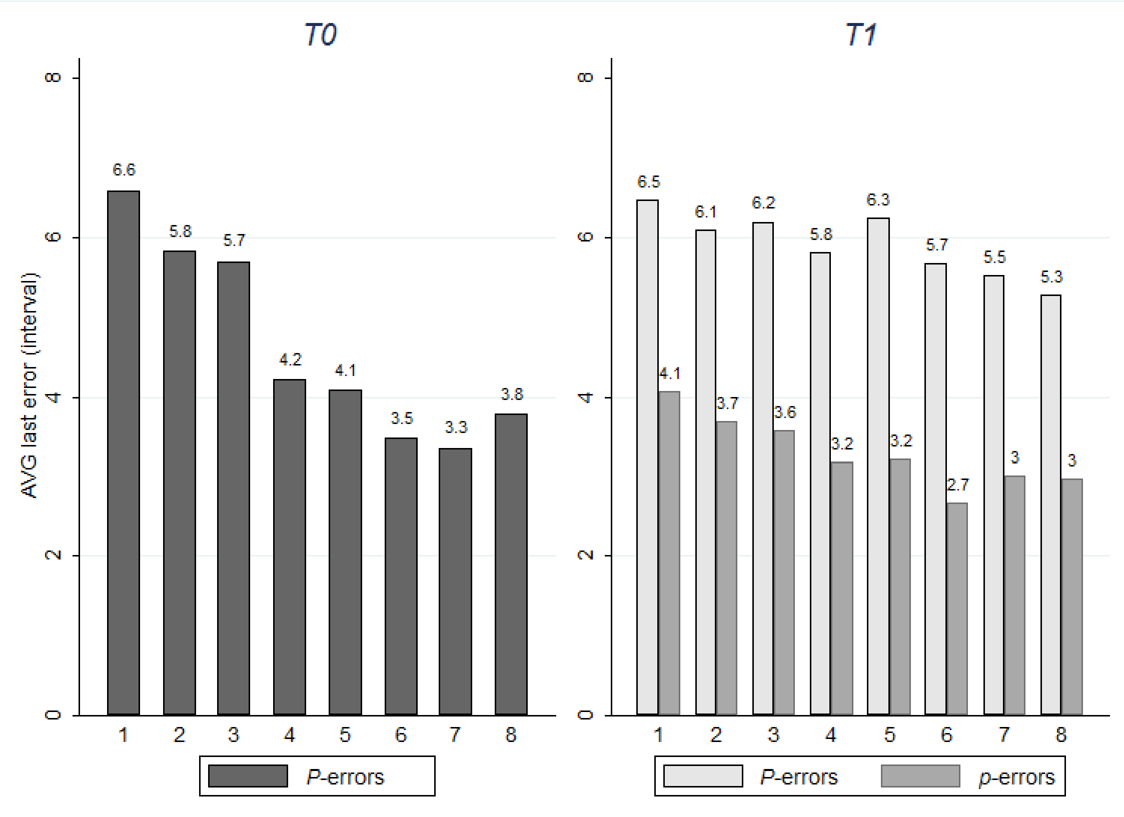

the average interval of the last error for each round. In $T0$, the figure decreases from above the sixth to below the fourth interval. In $T1$ the average interval is similar in the first rounds, but the reduction over rounds is smaller. The last p-error occurs before the last P-error, on average, in all rounds. The magnitude of the difference is up to three intervals.

We estimate the GLS panel regression model with random effects:

$$Last_{it} = \gamma_0 + \gamma_{Treat} Treat_i + \gamma_X X_{it} + \gamma_t t + u_i + \epsilon_{it},$$

where $Last_{it}$ is the interval where individual i makes his last P-error in round t, $Treat_i$ is a treatment dummy, X_{it} is a set of controls, including individual characteristics, such as gender, age, field of study, self-reported attitudes toward risk aversion and assessment of how easy the experiment was. The individual-specific random effect is u_i, and ϵ_{it} is the error term.[28] Standard errors are clustered at the session level.

Table 1 presents the results from various specifications. Model (1) tests that the average last error occurs at the same time across treatments, without consid-

[28]Obviously the analysis is restricted to subjects that won prize P.

Table 1: Last errors: across-treatment GLS estimates

		Last P-err			Adj. last P-err	
	(1)	(2)	(3)	(4)	(5)	(6)
T1	0.94	1.13	0.62	0.46	0.56	0.08
	(0.73)	(0.74)	(0.68)	(0.74)	(0.73)	(0.68)
round		-0.30*			-0.28*	
		(0.07)			(0.08)	
female		2.16*	2.23*		2.26*	2.32*
		(0.47)	(0.45)		(0.41)	(0.40)
age		0.03	0.03		0.02	0.03
		(0.03)	(0.03)		(0.03)	(0.03)
risk		0.02	0.03		-0.07	-0.06
		(0.12)	(0.12)		(0.13)	(0.13)
easy_exp		-0.35	-0.36		-0.37	-0.38
		(0.37)	(0.36)		(0.33)	(0.33)
2.faculty		0.36	0.45		0.13	0.22
		(0.50)	(0.48)		(0.37)	(0.37)
3.faculty		1.53*	1.44*		1.77*	1.68*
		(0.70)	(0.66)		(0.69)	(0.67)
laterounds			-1.71*			-1.70*
			(0.36)			(0.36)
laterounds*T1			1.12*			1.07*
			(0.39)			(0.40)
constant	5.36*	5.83*	5.22*	5.34*	6.55*	5.98*
	(0.65)	(1.79)	(1.65)	(0.65)	(1.17)	(1.03)
N	384	384	384	384	384	384
R-squared	0.03	0.24	0.24	0.01	0.23	0.23

Standard errors in parentheses (clustered at the session level)

* $p < 0.05$

ering any other covariate. The last P-error occurs around one interval later in $T1$, but the difference is not significant. Model (2) adds a number of covariates. As before, the coefficient for the treatment dummy is positive, but not significant. The interval of the last P-error significantly decreases over rounds. We also find winning women make their last P-error more than two intervals later than men. This result is consistent with other recent studies that find men can perform on average more steps of strategic reasoning than women [Dittrich and Leipold, 2014]. Students of economics (1.faculty), statistics, maths and natural sciences (2.faculty) make their last error significantly earlier than students of other majors (3.faculty).

In model (3) we include the interaction between treatment and a dummy for early (1-4) versus late (5-8) rounds. The treatment has no effect in early rounds. In late rounds, $T1$ is associated with a significant delay in last P-errors, estimated in 1.6 intervals.[29,30] This result could be due to the emergence of *sophisticated* behavior or to subjects experiencing greater learning hurdles in $T1$. To check for the relevance of *sophisticated* behavior, if a subject claims both prizes, wins prize p, and does not make any P-error thereafter, we attribute to him a last P-error equal to his last p-error, and estimate the previous models on the adjusted variable (models 4-6). Though the difference across treatments reduces in magnitude, all of the results carry over to the adjusted variable.

On aggregate, behavior regarding prize P does not differ substantially in $T0$ and $T1$. Moreover, the differences that emerge over rounds are only marginally explained by *sophisticated* behavior.

Result 1. *On aggregate, and in particular in early rounds, the positions in $\mathcal{L}$ are reached at the same time in $T1$ and in $T0$, supporting Hypothesis 2.*

Figure 3 reports the distributions of the last errors over intervals, by treatment. The distributions for prize P are bimodal. A fraction of subjects, higher in $T0$, does not make any error. The majority of the population stops making errors only when close to the final position, and in particular around the intervals 7 and 8. In $T1$, a majority of the subjects also makes the last p-error less than three intervals before the prize. The supports of the distributions of the last P- and p-error are different, as prize p is reached before P. However, recall from Section 3.2 that the last P-error should happen before the last p-error under LSP behavior, and that no error should arise under equilibrium and level-k.

In Figure 4, we plot the distribution of the individual differences between the last P- and p-error. Both the original distribution of the last P-error, and the one adjusted as described above are represented. The difference between the two provides a rough measure of the relevance of *sophisticated* behavior. Both distributions are biased toward positive values. That is, last P occurs later than

[29]The reduction in the interval of last P-errors in late vs early rounds is significant also in $T1$ (Chi-squared test: $z = 13.9$; P-val < 0.01).

[30]Similar results are obtained interacting the treatment with a continuous variable for the rounds - i.e., same intercept across treatments, but different slopes.

FIGURE 3: DISTRIBUTION OF LAST ERRORS

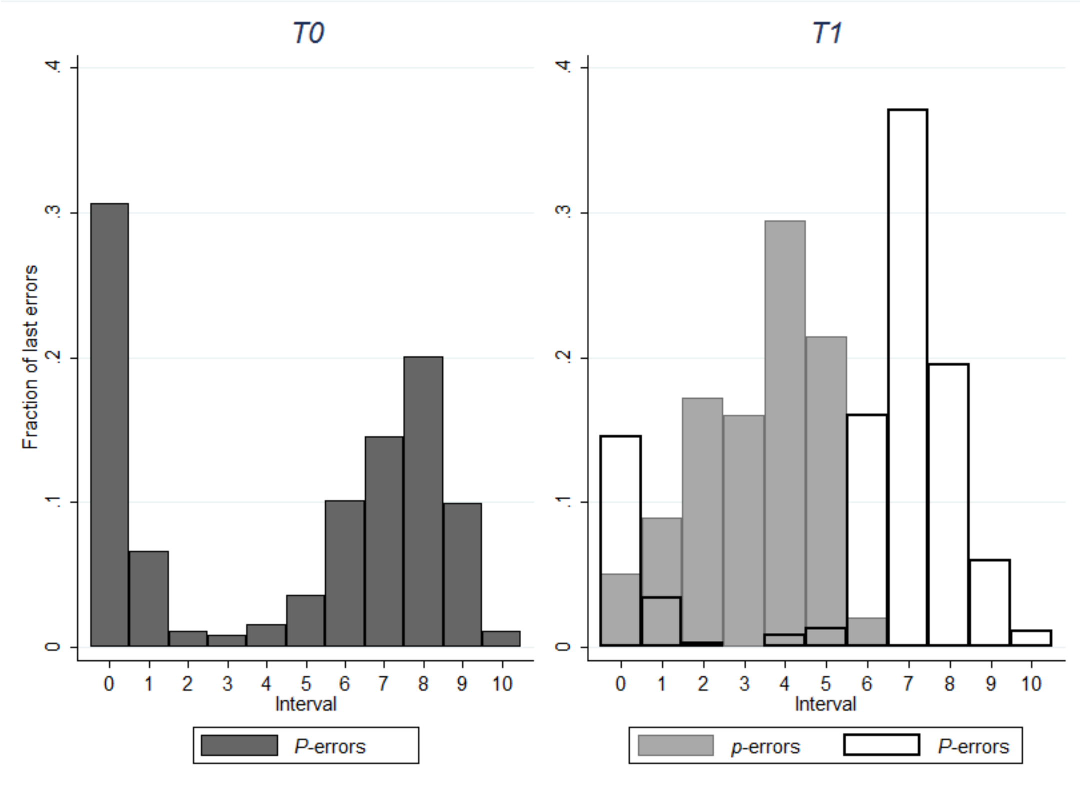

FIGURE 4: DIFFERENCES BETWEEN THE LAST P-ERROR AND LAST p-ERROR

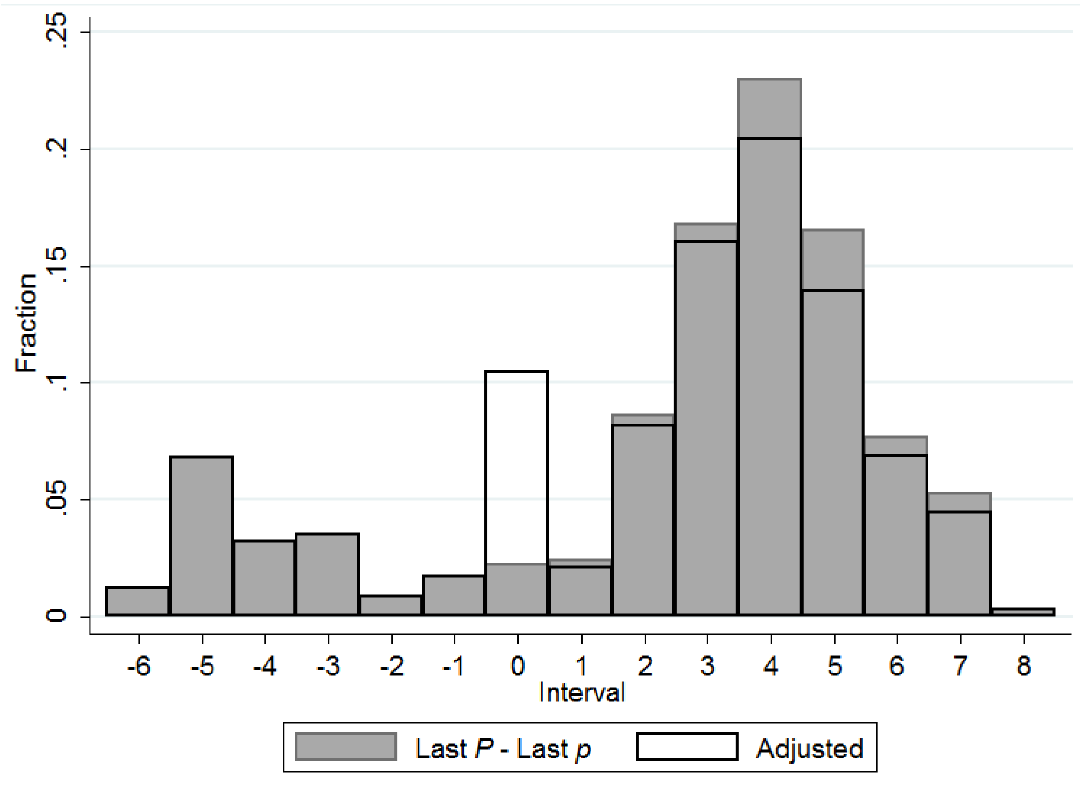

TABLE 2: TESTS ON DIFFERENCE BETWEEN LAST P- AND p-ERRORS

	Obs.	Est. diff.	Last P (Std. err)	F stat.	P-val.	Est. diff.	Adj. Last P (Std. err)	F stat	P-val.
$\mathbf{H_0}$: Last P = Last p	192	**2.75**	(0.35)	61.00	.004	**2.38**	(0.38)	38.31	.008

Standard errors clustered at the session level. Bold indicates significance at the .05 level.

FIGURE 5: DISTRIBUTION OF CLAIMS

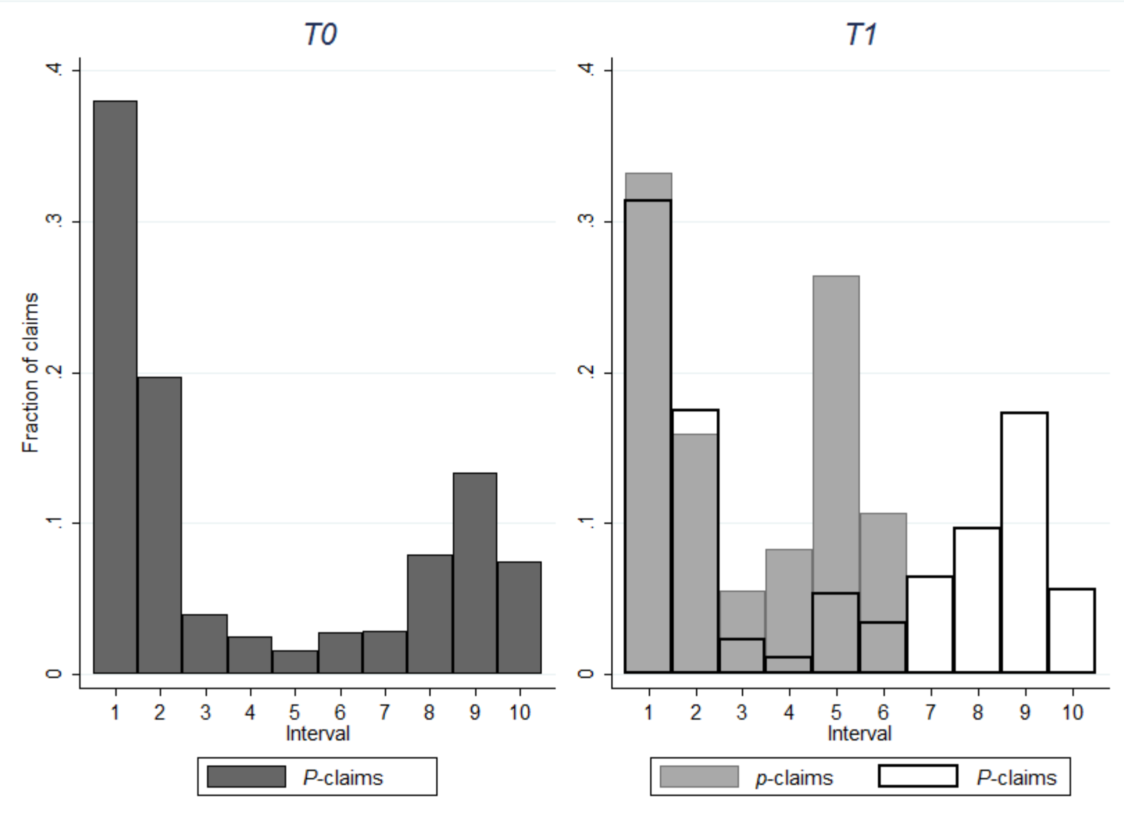

last p-error. Table 2 reports a formal test of this hypothesis, based on one independent observation per session. The average difference is estimated in 2.75 intervals and is statistically different from zero, even when considering adjusted last P-errors.

Result 2. *The positions in $\mathcal{L}_p$ are reached before the positions in $\mathcal{L}$, supporting Hypothesis 1.*

4.2 Claims

In both treatments, around 65 percent of the subjects claims he would have won prize P; 60 percent claims prize p in $T1$. Most claims end with the claiming player winning the prize. However, some claims are unwarranted, and a fine is imposed on 29 percent of both the p-claims and the P-claims. Figure 5 displays

TABLE 3: P-CLAIMS: ACROSS-TREATMENT GLS ESTIMATES

	P-claim			Just. P-claim			Adj. Just. P-claim		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
T1	0.36	0.51	0.18	0.75	0.96	0.42	0.30	0.64	0.20
	(0.64)	(0.49)	(0.50)	(0.96)	(0.69)	(0.67)	(0.97)	(0.70)	(0.68)
win_P		1.10*	1.11*		1.79*	1.85*		1.48*	1.53*
		(0.36)	(0.36)		(0.42)	(0.36)		(0.37)	(0.31)
round		-0.24*			-0.36*			-0.36*	
		(0.04)			(0.05)			(0.05)	
female		2.61*	2.63*		2.74*	2.77*		2.87*	2.91*
		(0.50)	(0.52)		(0.45)	(0.48)		(0.42)	(0.41)
age		-0.05	-0.05		-0.03	-0.03		-0.03	-0.03
		(0.05)	(0.05)		(0.06)	(0.06)		(0.06)	(0.06)
risk		-0.39*	-0.39*		-0.27	-0.27		-0.31*	-0.31*
		(0.14)	(0.14)		(0.15)	(0.14)		(0.16)	(0.15)
2.faculty		-0.72*	-0.73*		-0.26	-0.23		-0.37	-0.36
		(0.35)	(0.36)		(0.55)	(0.60)		(0.54)	(0.58)
3.faculty		1.34	1.31		1.64*	1.62*		1.54*	1.51*
		(0.78)	(0.76)		(0.65)	(0.60)		(0.72)	(0.68)
easy_exp		-0.08	-0.06		-0.06	-0.07		-0.08	-0.09
		(0.27)	(0.27)		(0.32)	(0.30)		(0.30)	(0.28)
laterounds			-1.33*			-1.79*			-1.73*
			(0.16)			(0.10)			(0.10)
laterounds*T1			0.68*			1.06*			0.81*
			(0.18)			(0.12)			(0.16)
constant	4.86*	7.89*	7.38*	5.68*	7.01*	6.19*	5.68*	7.52*	6.72*
	(0.41)	(2.29)	(2.25)	(0.82)	(1.96)	(1.85)	(0.82)	(1.98)	(1.88)
N	504	504	504	307	307	307	331	331	331
R-squared	0.00	0.32	0.32	0.02	0.38	0.38	0.01	0.39	0.38

Standard errors in parentheses (clustered at the session level)

* $p < 0.05$

the distribution of claims over intervals for $T0$ and $T1$. The distribution of P-claims traces that of last errors, separated between early (intervals 1-2) and late claimers (intervals 7-9).

We estimate a GLS panel model with random effects, similar to the one presented above, on the interval where prize P is claimed. Table 3 presents the results. Standard errors are always clustered at the session level. Results mimic

TABLE 4: WITHIN-TREATMENT TESTS ON CLAIMS

	Treat	Claims	Obs.	Est. diff.	(Std. err)	F stat.	P-val.
$\mathbf{H_0}$: Claim P = Claim p	$T1$	All	192	**1.98**	(0.30)	22.83	.017
	$T1$	Justif.	96	**3.58**	(0.41)	139.20	.001
$\mathbf{H_0}$: Claim P = Last P	$T0$	All	247	0.50	(0.80)	0.39	.576
	$T1$	All	257	1.03	(0.65)	2.56	.208
$\mathbf{H_0}$: Claim p = Last p	$T1$	All	231	-0.07	(0.27)	0.07	.80

Standard errors clustered at the session level. Bold indicates significance at the .05 level.

those obtained for the last errors. In particular, on aggregate claims are not made earlier in $T0$ with respect to $T1$. The interaction of late rounds and treatment is again significant, but the marginal linear effects of treatment in early and late rounds are not significant (Chi-squared test: early, $\chi^2 = 0.13$; P-val = .716; late, $\chi^2 = 3.12$, P-val = .077). Self-reported risk preferences have a significant impact: among those who claim P, the more a subject is risk-loving, the earlier he claims on average. The analysis of the treatment effect does not change if we restrict to claims that are *justified*, in the sense that the player does not make any error after he claims, nor if, in addition, we consider adjusted last P-errors to account for the possibility of *sophisticated* behavior (models 4-6 and 7-9, respectively).[31]

Around two thirds of the P-claims are *justified* in $T0$. The figure is identical in T1 if we consider adjusted last P-errors. Otherwise it is around 60 percent, which is also the percentage of *justified* p-claims.[32] Thus, in both treatments, some subjects claim before they make their last error. In Figure 6 we represent the distributions of individual differences between one's last errors and the corresponding claims. This difference is smaller than one for more than 70 percent of the subjects. The value of minus one corresponds to claiming in the interval that follows that of the last error, and is modal in both treatments. Overall we cannot reject that the mean of these distributions is zero in any treatment, and for both p and P, as reported in Table 4.

In $T1$ the distribution of p-claims is first order stochastically dominant with respect to that of P-claims. Recalling that *sophisticated* players should not claim p before P, we can test whether this is the case for subjects that claim both prizes. Results show P-claims occur on average 2 intervals later than p-claims and that this difference is statistically significant (see Table 4). Notably, the subjects for which the last P-error is different from the adjusted one - i.e., those that we iden-

[31]If we restrict to justified claims, risk preferences cease to play a role: indeed, there is no risk in claiming once you know how to win. Consistent with the idea that more risk-loving players are more likely to try to win both prizes, the coefficients are again significant if we consider claims that are justified against adjusted last errors.

[32]All claims are *justified* in the three intervals that precede a prize.

FIGURE 6: DIFFERENCES BETWEEN THE LAST ERROR AND THAT OF THE CLAIM

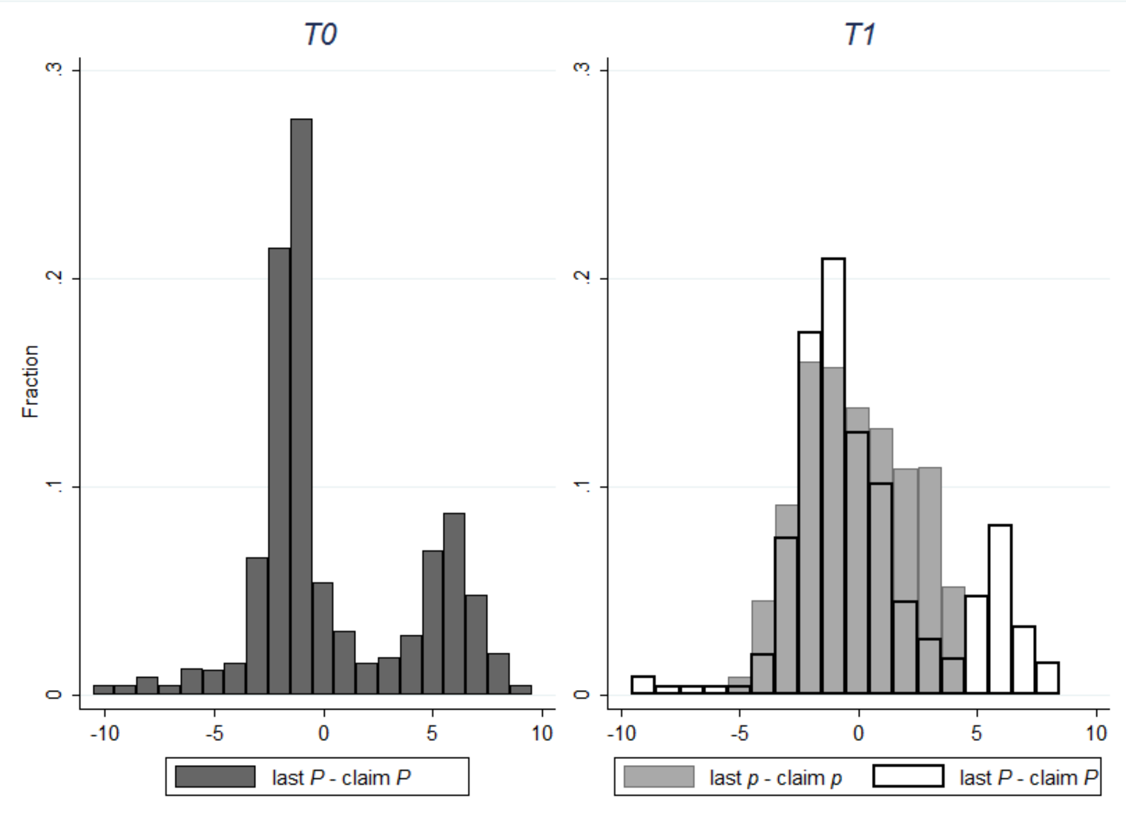

tify as sophisticated from their errors - claim P in the same interval as p on average (1.40 vs 1.37).

Result 3. *On aggregate, prize p is claimed before prize P supporting Hypothesis 3.*

4.3 Timing

We consider the time allocated to a decision task as a proxy of the effort allocated to it.[33] Equilibrium reasoning seems to imply that effort is concentrated at the beginning of the game, after which the players follow the planned strategy. This is also true for level-k and for any theory where an action plan is specified for the whole game. Under LSP the effort in taking a decision drops as soon as the subgame that includes the decision node is solved. Before that, the subjects are reasoning on distant subgames. It is not clear what this implies for reasoning efforts. Nevertheless, no special effort should be exerted around the small prize.

Under LBI it is not worth reasoning until a prize enters the LF-game. Before, all intermediate payoffs are the same, and the LF-game provides no information

[33]Recall the players were not time constrained.

FIGURE 7: AVERAGE TIME FOR EACH DECISION OVER POSITIONS

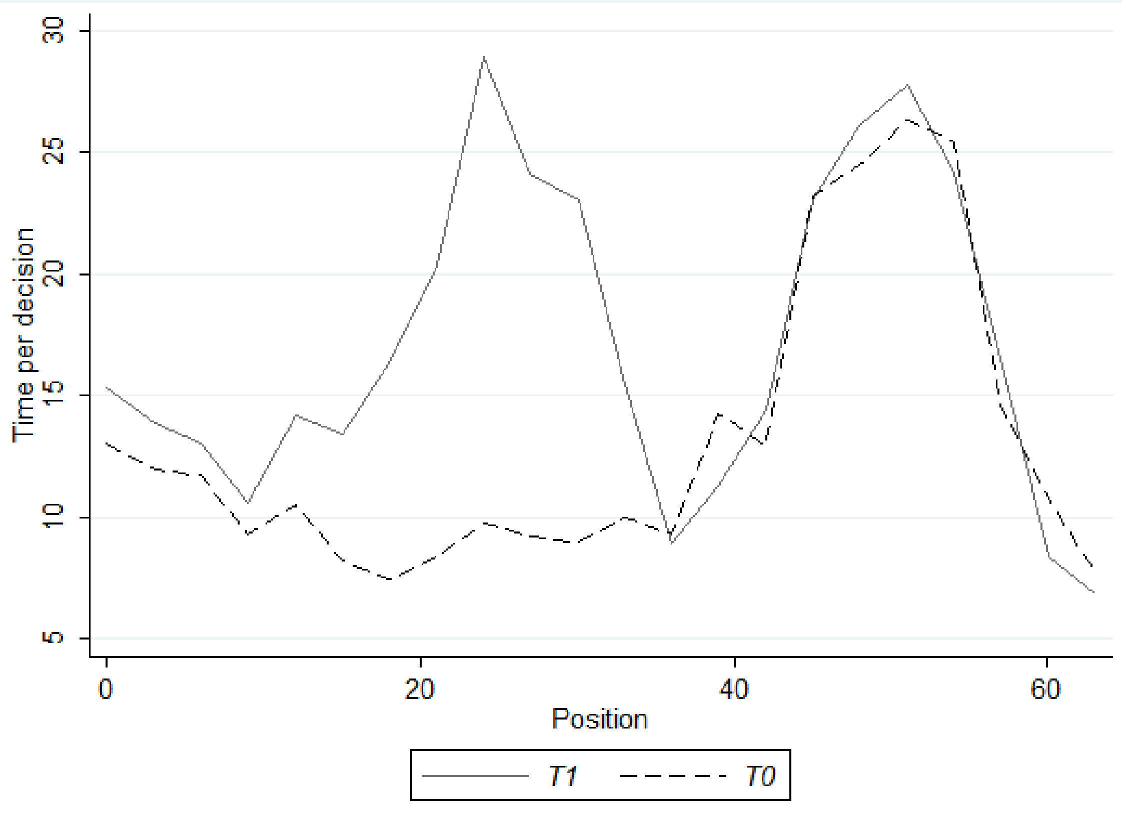

on how to choose meaningfully. Absent in-depth consideration of different alternatives, choices are taken quickly.[34] The players exert reasoning effort as soon as a prize is included in the LF-game. The time needed to reach a decision increases as the players get close to p and P, and then decreases once it is clear who is going to win the prize.

In Figure 7, the average number of seconds per decision is plotted against the set of positions, for both treatments.[35] In $T0$, the average decision time remains flat around ten seconds until position 40. It then increases, peaking around position 50, and drops back to ten seconds by position 60. From position 40 on the graph is similar in $T1$. It shows another symmetric steep increase and fall between positions 20 and 40. This behavior perfectly matches the pace of reasoning efforts under LBI, and is at odds with the other models. Consistent with the results on errors and claims, decision times suggest that prizes are included in the LF-game at two or three intervals of distance from each prize, on average.

Overall we find evidence in favor of every single hypothesis derived from LBI. At

[34]Indeed, the error rate in the first intervals are consistent with random play.

[35]We plot the three-position averages in order to eliminate idiosyncratic fluctuations between adjacent positions.

the individual level, we also find evidence of some equilibrium and *sophisticated* players. There seems to exist a huge divide between those latter subjects and the rest. This is also proven by the answers to the final questionnaire, which included an open question about the chosen strategy. Around 15 percent of the subjects identified the full set of losing positions as the guide for their strategy. Half of the subjects mentioned as "strategic" a subset of the last three losing positions. In $T1$ those players mention the last losing positions relative to both p and P. Around one third of our sample stated they were trying to move as quickly as possible to the "hot-spots", close to the prizes.

5 Applications and discussion

In this section we first show how LBI can be applied to some classical examples, drawing predicted behavior, and comparing it to the experimental evidence.[36] We then lay out an extension to the theoretical model. Finally we compare LBI to level-k models and to the limited forecast equilibrium.

5.1 Examples

Centipede game. The centipede game (Figure 8) is a two-player finite sequential game. The players alternate choosing whether to 'take' and end the game (T) or to 'pass' to the other player (P). The payoff from T in the current decision node is greater than that received if the other plays T in the next one, but smaller than the payoffs earned if the other plays P as well. The player active at the last decision node gets more from T than from P, and is therefore expected to take. By backward induction, in the unique SPE the game ends at the first terminal node, implying a huge efficiency loss.

Experimental evidence has found little compliance with SPE in the laboratory.[37] The typical results feature a bell-shaped distribution of endnodes. Table 5 presents the results from the four and six-leg centipede in the seminal paper of McKelvey and Palfrey [1992]. Beyond the failure of SPE, it is a robust finding that longer games are associated, ceteris paribus, with higher endnodes.

A player whose foresight is limited does not reason strategically on the whole game. For example, assume Player 1 follows decision rule F_4. At the beginning of the game in Figure 8, he reasons backward from the second of the decision nodes

[36] Roomets [2010] also applies his Horizon-Based Limited Foresight model to these games, and his results are in line with ours.

[37] See, for example, Kawagoe and Takizawa [2012]; Levitt et al. [2011]; McKelvey and Palfrey [1992]; Rapoport et al. [2003].

FIGURE 8: THE SIX-LEGS CENTIPEDE GAME IN MCKELVEY AND PALFREY [1992]

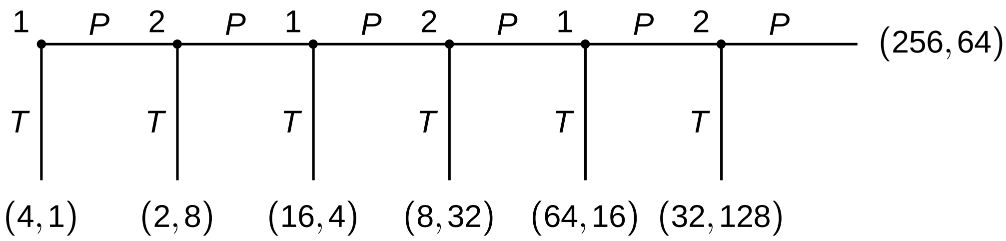

TABLE 5: RESULTS FROM MCKELVEY AND PALFREY [1992]

	Proportion of observations in each terminal node						
	1	**2**	**3**	**4**	**5**	**6**	**7**
Four legs	.07	.36	.37	.15	.05		
Six legs	.01	.07	.20	.38	.25	08	.01

where his opponent moves. The payoffs that follow action P at that node are projected on the pseudo-terminal history (P, P, P, P). The resulting intermediate payoffs, computed as the median within the range of available payoffs, are (144, 72). Thus, Player 2 is expected to pass at that node ($72 > 32$). The same holds at previous nodes, leading Player 1 to choose P. A similar reasoning can be applied, for every decision rule (F_1, F_2,...), to every decision node. In this example the intermediate payoff from P is always higher than the payoff from T. Thus, a player does not play T unless all terminal nodes are included in his LF-game.[38] Ceteris paribus, the longer the centipede, the later this happens.

In the six-leg (four-leg) centipede, for Player 1, F_1 prescribes to always pass, F_2 and F_3 to take at the next to last node, F_4 and F_5 to take from the third (first) node on, and higher levels to always take. For Player 2, F_1 and F_2 prescribe to take at the last node; F_3 and F_4 to take from the fourth (second) node on, and higher levels to always take. The results in Table 5 can be explained by a population composed by F_2, F_3 and F_4. Those results are based on averages over ten repetitions, and the initial-response results display a shift toward lower foresight levels.[39]

In a generic centipede game, it is not always necessary for a player to choose T that the terminal nodes of the centipede game are included in the LF-game. The point at which F_κ prescribes to play T depends on the progression of the

[38] This is true not only for the proposed projection function. For example, if a player earns x by taking at some node, the average of the n payoffs that follow is in the form $(1/2 + 4 + 2 + 16...)x/n > x$.

[39] For the first game: a majority of F_2 and F_3, plus some F_1 and F_4.

FIGURE 9: THE SEQUENTIAL BARGAINING GAME IN JOHNSON ET AL. [2002]

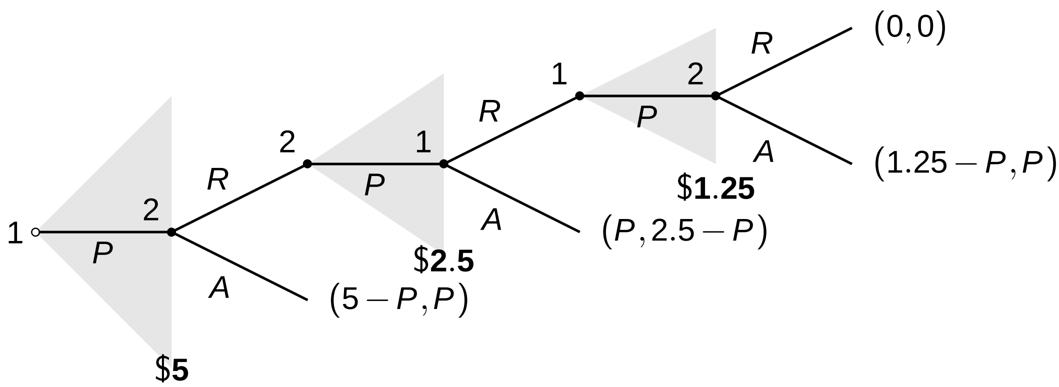

payoffs, because this determines the relation between the intermediate payoffs and the last visible terminal ones.[40]

Sequential bargaining. In sequential bargaining (Figure 9) two players must agree on the division of a cake. The players alternate in proposing a split. The player receiving the proposal either accepts, in which case the game ends and the proposed split is implemented, or rejects, in which case they move to the next round. The cake shrinks every time an offer is rejected. In finite bargaining, the last round of bargaining is an ultimatum game, and its solution provides the minimal acceptable offer in the previous one. This reasoning can be iterated backward to the initial node. In the unique SPE, in every round, the players submit as proposers the minimal acceptable offer, and accept as responders any offer weakly higher than the minimal acceptable one. Thus, the first offer is implemented, and the game ends. Experimental evidence shows that offers are, on average, more equitable than in the SPE, and those that are close to equilibrium are often rejected.[41] The first offers are also independent of the number of rounds of bargaining, even if this determines the roles in the last round and, thus, who retains more bargaining power.

Figure 9 displays the game used in Johnson et al. [2002]. The cake is initially worth $5, and it is halved at every new round. The first SPE offer is $1.26. The average observed one is $2.10. Half of the offers below $2 are rejected. In a treatment where other-regarding preferences are disabled, using robots as opponents, the average first offer decreases to $1.84.

Under LBI, a player does not realize the last round is an ultimatum game if this is not included in the LF-game. If the LF-game of Player 1 includes the first

[40] Kawagoe and Takizawa [2012] present a survey of classical initial-response results in the centipede game. Behavior is consistent with the population shares outlined above.

[41] See the review in Roth [1995].

two rounds of bargaining, he foresees intermediate payoffs of \$0.625 each following a rejection in round two. Reasoning backward, he understates, with respect to SPE, his minimal acceptable offer in the second round, and expects Player 2 to offer $P \cong \$0.63$ and keep \$1.87. He then overstates the minimal acceptable offer of his responder in the first round, and offers $P \cong \$1.87$. If Player 2 has a similar foresight, he rejects any offer below this value.[42] This prediction matches the average offer in the robot treatment of Johnson et al. [2002].

Given a distribution of foresight levels in the population, the length of the game does not affect the behavior of all players whose foresight does not reach the terminal nodes. Thus, despite its role in allocating bargaining power, under LBI the length of the game has no impact on the distribution of the first offers.

5.2 An extension to infinite games and further developments

Stage payoffs and sight. We here relax, for a broad class of games, the assumption that the players know the terminal payoffs of the game. In infinite games, and where retrieving information about the payoffs is costly, the players may disregard the payoffs that arise beyond a certain stage.[43] This extension will also prove useful in Section 5.3 when comparing LBI to the limited forecast equilibrium.

To discriminate between close and distant payoffs, it must be possible to represent terminal payoffs as the sum of the payoffs gained along the game. That is, the payoff function must be additively separable. Let $\pi_i : N \to \mathbb{R}$ be a function identifying single-stage payoffs. A game satisfies Additive Separability (AS) if:

$$u_i(z) = \sum_{h:h \preccurlyeq z} \pi_i(a_j[h]),$$

where $a_j[h]$ is the move of the player who active at h.

An agent with a *sight* $\lambda \in \mathbb{N}$ at node h, with $r(h) = t$, sees all successors of h with rank smaller or equal to $t + \lambda$ and the corresponding stage payoffs. If $t + \lambda \geq \max(r(z))$, the player sees all relevant information about the game from the first stage, as in our base model.

The *sight* is exogenous. If the payoffs satisfy AS, the *sight* of a player always

[42]If Player 1 only sees the first round of bargaining, he offers $P \cong \$1.26$. Low foresight levels do not match subgame perfect first offers in general. The initial offer oscillates as foresight varies, and is lower (higher) when an even (odd) number of rounds are not included in the LF-game.

[43]Disregarding distant payoffs, and considering the close ones is an instance of present-bias, though not in the standard sense. Present-bias and, more generally, discounting refer to preferences. As we assume the terminal payoffs are von-Neumann-Morgenstern utilities, we exclude preference-related aspects from our analysis.

provides information about payoffs to players, and identifies a game. The foresight of the players is determined on this truncated game. In particular, the cost and benefit of further foresight levels, and the intermediate payoffs are assessed using only the information provided by one's *sight*.

The decision rule $F_{\kappa,\lambda}$ follows backward induction on a LF-game with κ stages, where the payoffs obtained within λ stages are projected on the pseudo-terminal nodes. This extended version of LBI applies to infinite games, implicitly assuming that players treat them as finite ones.[44] If the *sight* and the *foresight* coincide LBI admits as special cases myopic behavior, and beliefs of myopic behavior on the part of others.

Developments. A note is needed on what our account of foresight does not cover. First, in our model there is no separation between one's foresight and his beliefs about others' foresight. This is what excludes, for example, the possibility of strategic sophistication. Nor there is a separation between one's foresight and the 'type' of reasoning that is implemented within its bounds. We decided to focus on this simple framework for the sake of isolating the predictions of limited foresight, in a context where these restrictions seem plausible. These generalizations of the model, that both our intuition and experimental results indicate as well deserving, are left for future research.

5.3 Comparison with related models

Strategic thinking in LBI and level-k. LBI and level-k are first-response models. They explain out-of-equilibrium behavior of untrained subjects. They are based on a hierarchy of decision rules, such that each best replies to the next lower one. This chain of best replies is anchored to a rule that is non-strategic, in the sense of not being based on any belief about others. Level zero (L_0) is usually modeled as a randomizer; F_1 optimizes over the intermediate payoffs that follow his decision node.

A level-k player believes the opponents are of level $k-1$, and that they believe the others are of level $k-2$. When applied to dynamic contexts the players specify an action plan for the whole game that is consistent with their beliefs [cf., Ho and Su, 2013; Kawagoe and Takizawa, 2012]. Under LBI, F_κ implicitly imposes decreasing levels of foresight on the agents that are active at the following nodes: the belief on the next active player is consistent with $F_{\kappa-1}$, that on the following with $F_{\kappa-2}$, and so on. Following this decision rule, a player's moves

[44] Relating infinite-horizon games to finite ones is an old topic in game theory [e.g., Fudenberg and Levine, 1983].

need not form a consistent action plan.[45]

Thus, level-k and LBI imply different inconsistencies across levels of higher order beliefs. Within the LF-game, the player assumes the others' beliefs to be consistent with his owns: one's belief about the next player's belief about the decision rule of the third player in the row is consistent with his own belief about this latter player's decision rule. However, beliefs regarding the choices of the same player, including himself, at any two different nodes, are inconsistent, as a lower foresight is imposed both on the others' and on one's own future self.[46] In level-k it is the other way around.

The different features of LBI with respect to level-k models make the two models best suited for different situations. We argue that LBI is better suited to address games where the dynamic aspects are salient. Our experiment provides such an example, showing a case where LBI succeeds where level-k fails.

LBI and the limited forecast equilibrium. LBI and the limited forecast equilibrium (LFE) [Jehiel, 1995] share the same motivation: studying extensive-form games where players have limited foresight. In every stage an action is chosen based on predictions about a limited number of the forthcoming moves. In the LFE, the strategies are constrained to be justified, in the sense that they maximize the payoff obtained within one's horizon, given the forecast. The forecasts are constrained to be correct, in the sense of being consistent with the equilibrium strategies.

Under LBI, actions also maximize the payoff, given one's forecast, but forecasts need not be correct. Correct forecasts are justified in Jehiel [1998a] through a learning argument. The equilibrium captures the limiting outcome after the players have gained sufficient experience. LBI targets, instead, out-of-equilibrium, first-response behavior.

Example 1 in Jehiel [1995, p.504] shows the difference. The players choose sequentially from a binary action space, with $A_1 = \{U, D\}$, and $A_2 = \{L, R\}$. Their action remains valid for the current stage and the following one. The stage payoffs of the game are displayed in Table 6. They depend on the actions that are valid in each stage. The game is infinitely repeated. The players maximize their average payoff within their horizon, given their forecasts, and do not consider what happens beyond it. In the example, the players have a foresight of two

[45] The fact that players do not form a complete action plan implies differences in behavior when a game is played in the normal and in the extensive form, a result backed by a number of studies [e.g., Cooper and Van Huyck, 2003].

[46] If beliefs did not feature this inconsistency, a player's horizon would immediately extend to that of his future self, turning down the assumption of limited foresight. One can see this as if, under LBI, the players were split into agents, each choosing at a different node.

TABLE 6: STAGE PAYOFFS IN EXAMPLE 1 OF JEHIEL [1995]

UL	**UR**	**DL**	**DR**
(3,2)	(0,3)	(2,1)	(2,0)

stages, so that they only forecast the next player's action. For comparability, we liken them to LBI players with a sight of two stages, and an identical foresight, so that they forecast the same moves, and optimize over an identical set of payoffs.[47]

The strategy profile where Player 1 starts by playing U, plays U after L, D after R thereafter, and Player 2 always plays L, is a LFE [Jehiel, 1995, p. 507]. Under LBI, both players believe their opponent to act myopically. Player 1 expects Player 2 to choose R after U, and L after D. He takes action D. Player 2 expects Player 1 to choose U after L, and D after R. He takes action L. Thus, the players play differently than in the above LFE. Moreover, their choices are inconsistent with any LFE, because they are based on incorrect forecasts.[48]

6 Conclusion

The paper studies whether agents reason according to limited foresight. We present a general framework of out-of-equilibrium behavior in sequential games, Limited Backward Induction, according to which players take decisions following backward induction over the close-by nodes that fall within their foresight. LBI formalizes the common intuition that players forecast decisions only within a limited horizon, and act consistently with these imprecise forecasts.

We develop an experimental design, based on the race game, which allows to test for limited foresight. In a baseline treatment, the subjects compete to reach a prize which is achieved at the end of the game. In a second treatment, they can also obtain a small prize which is achieved before the end of the game, by giving the opponent the opportunity to secure the victory of the final prize. Results show the subjects solve for the small prize before they do for the final one. As sophisticated strategies that aim at winning both prizes can be excluded as an overall explanation, the data provide clear evidence of limited foresight. Backward reasoning is the main cognitive procedure used by the subjects. As assumed

[47]In a related model, Jehiel [2001] allows players to take into account also the expected payoffs beyond the horizon of their forecasts, so that payoffs are more similar to our standard intermediate ones.

[48]The players may realize that their forecasts are not correct, and modify them. To capture how behavior would evolve requires a learning model based on LBI, which goes beyond the scope of this paper. Under which conditions LBI would converge to LFE in such a model is an interesting question that we leave for future research.

under LBI, it does not proceed backwards from the terminal nodes. Rather, it is routinely performed on the few stages that follow the current decision node. Alternative models, including equilibrium and dynamic level-k models, fail to rationalize the evidence.

We believe our results represent a significant contribution to the understanding of initial-response behavior in sequential games. Nevertheless, quantifying precisely the heterogeneity of foresight types, and understanding how portable these are across games, call for further experimental investigation. A full classification will likely include also equilibrium and sophisticated players, which also show up in our sample.

On the theoretical side, LBI is a flexible framework and applies to all sequential games with perfect information, including infinite ones. As the foresight grows, in the limit LBI mimic subgame perfection. At the other end, it encompasses perfect myopia as a special case. The framework is suitable for further developments. A promising direction seems to be the merger of models of foresight, levels of reasoning and sophistication in a unified framework that would then apply to any dynamic game. Simultaneously, one could study how limited foresight evolves with experience, as similarly done in Mengel [2014]. Our results provide some insights for this latter line of research, since we observe diverging learning patterns in the two treatments for reasons that seem related to the partial sophistication of the subjects.

The possible presence of limited foresight bares important consequences and insights for many contexts. For instance, this paper shows how LBI can rationalize observed behavior in widely studied games, such as centipede and sequential bargaining games. By ignoring limited foresight, one might also estimate incorrectly individual discount factors in all settings that include planning and anticipating distant decisions. More generally, LBI is a promising tool for all situations where end-of-game effects seem relevant. For instance, a relation between limited foresight and bubble formation in experimental asset markets is already being hypothesized by Bosch-Rosa and Meisner [in progress].

References

ALAOUI, L. AND A. PENTA, "Endogenous depth of reasoning," Barcelona GSE Working Paper No. 653 (2014).

AUMANN, R. J., "Backward induction and common knowledge of rationality," *Games and Economic Behavior* 8 (1995), 6–19.

BATTIGALLI, P., "On rationalizability in extensive games," *Journal of Economic Theory* 74 (1997), 40–61.

BEN-PORATH, E., "Rationality, Nash equilibrium and backwards induction in perfect-information games," *The Review of Economic Studies* 64 (1997), 23–46.

BINMORE, K. G., "A note on backward induction," *Games and Economic Behavior* 17 (1996), 135–137.

BINMORE, K. J., J. MCCARTHY, G. PONTI, L. SAMUELSON AND A. SHAKED, "A backward induction experiment," *Journal of Economic Theory* 104 (2002), 48–88.

BONANNO, G., "Branching time, perfect information games, and backward induction," *Games and Economic Behavior* 36 (2001), 57–73.

BOSCH-ROSA, C. AND T. MEISNER, "Cognitive bubbles," mimeo (in progress).

BRANDENBURGER, A. AND A. FRIEDENBERG, "Are admissibility and backward induction consistent?," mimeo (2014a).

———, "How many levels do players reason? An observational challenge and solution," mimeo (2014b).

COOPER, D. J. AND J. B. VAN HUYCK, "Evidence on the equivalence of the strategic and extensive form representation of games," *Journal of Economic Theory* 110 (2003), 290–308.

COSTA-GOMES, M., V. CRAWFORD AND B. BROSETA, "Cognition and behavior in normal-form games: an experimental study," *Econometrica* 69 (2001), 1193–1235.

CRAWFORD, V. P., M. A. COSTA-GOMES AND N. IRIBERRI, "Structural models of nonequilibrium strategic thinking: theory, evidence, and applications," *Journal of Economic Literature* 51 (2013), 5–62.

CROSETTO, P. AND M. MANTOVANI, "Availability of information and representation effects in the centipede game," Jena Economic Research Paper No. 51 (2012).

DIASAKOS, T. M., "Complexity and bounded rationality in individual decision problems," School of Economics & Finance Discussion Paper No. 1314 , University of St Andrews (2008).

DITTRICH, M. AND K. LEIPOLD, "Gender differences in strategic reasoning," CESifo Working Paper Series No. 4763 (2014).

DOHMEN, T., A. FALK, D. HUFFMAN, U. SUNDE, J. SCHUPP AND G. G. WAGNER, "Individual risk attitudes: measurement, determinants, and behavioral consequences," *Journal of the European Economic Association* 9 (2011), 522–550.

DUFWENBERG, M., R. SUNDARAM AND D. J. BUTLER, “Epiphany in the Game of 21,” *Journal of Economic Behavior & Organization* 75 (2010), 132–143.

FISCHBACHER, U., “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics* 10 (2007), 171–178.

FUDENBERG, D. AND D. LEVINE, “Subgame-perfect equilibria of finite and infinite-horizon games,” *Journal of Economic Theory* 31 (1983), 251–268.

GNEEZY, U., A. RUSTICHINI AND A. VOSTROKNUTOV, “Experience and insight in the Race game,” *Journal of Economic Behavior & Organization* 75 (2010), 144–155.

HARGREAVES HEAP, S., D. ROJO ARJONA AND R. SUGDEN, “How portable is Level-0 behavior? A test of Level-k theory in games with non-neutral frames,” *Econometrica* 82 (2014), 1133–1151.

HO, T.-H. AND X. SU, “A dynamic level-k model in sequential games,” *Management Science* 59 (2013), 452–469.

JEHIEL, P., “Limited horizon forecast in repeated alternate games,” *Journal of Economic Theory* 67 (1995), 497–519.

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

———, “Repeated games and limited forecasting,” *European Economic Review* 42 (1998b), 543–551.

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

JOHNSON, E. J., C. CAMERER, S. SEN AND T. RYMON, “Detecting failures of backward induction: monitoring information search in sequential bargaining,” *Journal of Economic Theory* 104 (2002), 16–47.

KAWAGOE, T. AND H. TAKIZAWA, “Level-k analysis of experimental centipede games,” *Journal of Economic Behavior & Organization* 8 (2012), 548–566.

LEVITT, S. D., J. A. LIST AND S. E. SADOFF, “Checkmate: exploring backward induction among chess players,” *The American Economic Review* 101 (2011), 975–990.

MCKELVEY, R. D. AND T. R. PALFREY, “An experimental study of the centipede game,” *Econometrica* (1992), 803–836.

———, “Quantal response equilibria for extensive form games,” *Experimental economics* 1 (1998), 9–41.

MENGEL, F., "Learning by (limited) forward looking players," *Journal of Economic Behavior and Organization* 108 (2014), 59–77.

RAPOPORT, A., W. E. STEIN, J. E. PARCO AND T. E. NICHOLAS, "Equilibrium play and adaptive learning in a three-person centipede game," *Games and Economic Behavior* 43 (2003), 239–265.

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

ROOMETS, A., "On limited foresight in games," mimeo (2010).

ROTH, A. E., "Bargaining experiments," in J. H. Kagel and A. E. Roth, eds., *The handbook of experimental economics* (Princeton University Press, Princeton, 1995), 253–248.

SELTEN, R., "Features of experimentally observed bounded rationality," *European Economic Review* 42 (1998), 413–436.

STAHL, D. O. AND P. W. WILSON, "On players' models of other players: theory and experimental evidence," *Games and Economic Behavior* 10 (1995), 218–254.

A Sophisticated behavior and risk aversion

Denote M the position where P is obtained, and m the position where p is obtained. Let $F < 0$ be a generic fine. The current position is $\bar{t}$. The expected payoff from S_1 is $P + (M - \bar{t})$. The expected payoff from S_2 is $p + (m - \bar{t}) + q(P + (M - m - a_{opp,m}))$, where q is the probability that the opponent does not solve the game within m, and $a_{opp,m}$ is the action of the opponent at m. The expected payoff from S_3 is $p + (m - \bar{t}) + q(P + (M - \bar{t})) + (1 - q)F$. Using simple algebra one obtains that, for a risk neutral player, $S_2 \succeq S_3$ and $S_2 \succeq S_1$ if and only if:

$$\frac{\Delta_P + \Delta_M}{P + \Delta_M - a_{opp,m}} \leq q \leq \frac{-F}{m + a_{opp,m} - \bar{t} - F} \tag{1}$$

where $\Delta_P = P - p$ and $\Delta_M = M - m$. It follows that, in the game played in $T1$, for $\bar{t} < 37$ there exists no q such that $S_2 \succeq S_3$ and $S_2 \succeq S_1$.

A similar reasoning applies to risk averse players. Consider a player, whose utility function is $u_i(x) = -e^{-\alpha x}$, $\alpha > 0$, featuring constant absolute risk aversion. Then $S_2 \succeq S_3$ and $S_2 \succeq S_1$ if and only if:

$$\frac{1 - e^{-\alpha(\Delta_P + \Delta_M)}}{1 - e^{-\alpha(P + \Delta_M - a_{opp,m})}} \;\leq\; q \;\leq\; \frac{1 - e^{-\alpha F}}{1 + e^{-\alpha(P + M - \bar{t})} - e^{-\alpha(P + \Delta_M - a_{opp,m})} - e^{-\alpha F}} \tag{2}$$

FIGURE 10: BELIEFS SUSTAINING S_2: CARA (STARS), CRRA (CIRCLES)

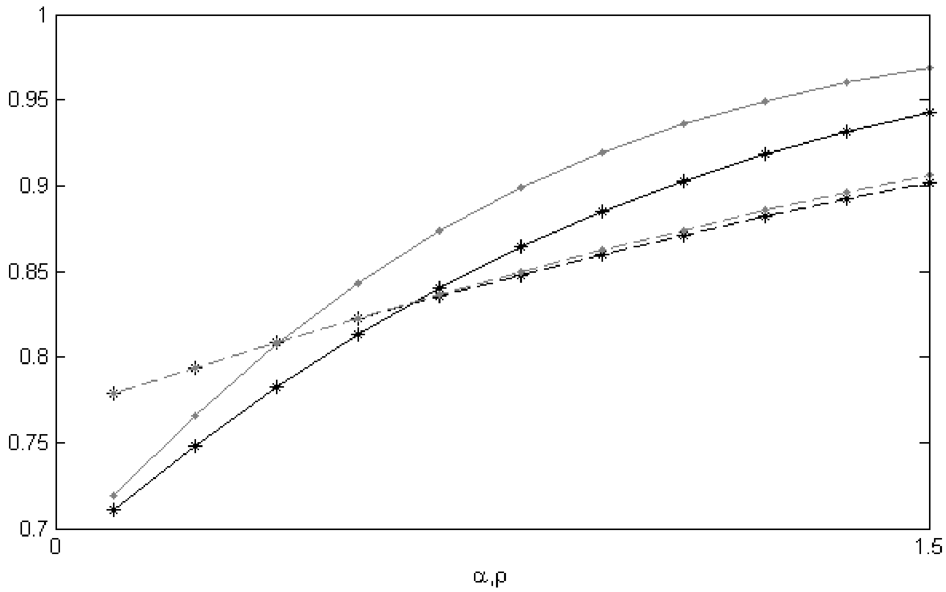

Consider now a player whose utility function is $u_i(x) = \frac{x^{1-\rho}}{1-\rho}$, $\rho > 0$ featuring constant relative risk aversion. Slightly abusing notation for simplicity, let $u^{j,\omega}$, $j \in \{1,2,3\}$, $\omega \in \{g,b\}$, be the (rescaled) utility of the player when his strategy is S_j, conditional on state ω. In case the opponent does not solve the game within m, $\omega = g$; otherwise, $\omega = b$.

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

Then $S_2 \succeq S_3$ and $S_2 \succeq S_1$ if and only if:

$$
\frac{u^{1,\omega} - u^{2,b}}{u^{2,g} - u^{2,b}} \le q \le \frac{u^{3,b} - u^{2,b}}{(u^{2,g} - u^{2,b}) - (u^{3,g} - u^{3,b})} \tag{3}
$$

If we apply our parameters to the previous inequalities we obtain that, in G_1, if a player 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 $S_2 \succeq S_1$, and (ii) under CRRA, for $\rho < 0.5$ there exists no q such that $S_2 \succeq S_3$ and $S_2 \succeq S_1$.

In Figure 10 the bounds of the interval for q, in the case of CARA (stars) and CRRA (circles), are plotted against the coefficient of (constant or relative) risk aversion. To sustain S_2, q must be higher than the dashed line, and lower than the solid line. It is assumed that $\bar{t} = 34$ 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 of beliefs appears. The interval never grows beyond a magnitude of 0.06, corresponding to a belief q between .83 and .97.

B Experimental Instructions

Welcome to this experiment in decision-making. Please, read these instructions carefully. 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). You will receive 50 ECU as initial endowment. 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

- The game you will play has two players, *P* and *Q*.
- The players decide **sequentially**: they take turns, one after the other. Each decision consists in a number of steps, between 1 and 6 (included).
- You 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.
- You 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 15 ECU.
- When a player declares he is going to win [a prize], his opponent is not informed and can himself declare he is winning [the same prize].
- The number of ECU earned are the sum of the initial endowment, 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 be randomly assigned to role P or Q.
- In every new round you will play against a **new** partner, chosen at random between the other participants.
- You will never play twice with the same partner.

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 Euros.

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.