



Reasoning about Others' Reasoning

Larbi Alaoui
Katharina A. Janezic
Antonio Penta

This version: July 2019
(November 2017)

Barcelona GSE Working Paper Series

Working Paper n° 1003

Reasoning about Others' Reasoning*

Larbi Alaoui[†]

Katharina A. Janezic[‡]

Antonio Penta[§]

June 28, 2019

Abstract

Recent experiments suggest that level- k behavior is often driven by subjects' beliefs, rather than their binding cognitive bounds. But the extent to which this is true in general is not completely understood, mainly because disentangling 'cognitive' and 'behavioral' levels is challenging experimentally and theoretically. In this paper we provide a simple experimental design strategy (the 'tutorial method') to disentangle the two concepts purely based on subjects' choices. We also provide a 'replacement method' to assess whether the increased sophistication observed when stakes are higher is due to an increase in subjects' own understanding or their beliefs over others' increased incentives to reason.

We find evidence that, in some of our treatments, the cognitive bound is indeed binding for a large fraction of subjects. Furthermore, a significant fraction of subjects do take into account others' incentives to reason. Our findings also suggest that in general, level- k behavior should not be taken as driven either by cognitive limits alone or beliefs alone. Rather, there is an interaction between own cognitive bound and reasoning about the opponent's reasoning process. These findings provide support to more subtle implications of the EDR model (Alaoui and Penta, 2016a) than those which were previously tested, and show that the EDR framework is a useful tool for analyzing and understanding the complex interaction of cognitive abilities, incentives, and strategic reasoning.

From a broader methodological viewpoint, the tutorial and replacement methods have broader applicability, and can be used to study the beliefs-cognition dichotomy and higher order beliefs effects in non level- k settings as well.

Keywords: cognitive bound – depth of reasoning – higher-order beliefs – level- k reasoning – replacement method – tutorial method

JEL Codes: C72; C92; D80.

*Some of the material presented in this paper circulated in an old working paper of ours, entitled 'Level- k Reasoning and Incentives' (Alaoui and Penta, 2012). We are indebted to Vincent Crawford, Christian Fons-Rosen, Nagore Iriberry, Christian Michel, Rosemarie Nagel, Marciano Siniscalchi, Alessandro Tarozzi and two anonymous referees for the generous and very helpful comments on the early drafts of this project. We also thank several seminar and conferences audiences where these results were presented. The authors acknowledge financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563). Antonio Penta acknowledges financial support from the ERC Starting Grant #759424.

[†]Universitat Pompeu Fabra and Barcelona GSE. E-mail: larbi.alaoui@upf.edu

[‡]Universitat Pompeu Fabra and Barcelona GSE. E-mail: katharina.janezic@upf.edu

[§]ICREA, Universitat Pompeu Fabra and Barcelona GSE. E-mail: antonio.penta@upf.edu

1 Introduction

Recent experiments have documented that, in games in which individuals behave according to standard models of level- k reasoning, changing subjects' beliefs affects the observed distribution of levels (see, e.g., Agranov, Potamites, Schotter and Tergiman (2012), Georganas, Healy and Weber (2015) and Alaoui and Penta (2016a)). These results suggest that, at least in some settings, level- k patterns of behavior may be driven by individuals' beliefs rather than by their intrinsic cognitive limitations. Whether this is true in general, however, is far less clear, and most work in this area is agnostic on the point. This is largely because disentangling 'behavioral' and 'cognitive' levels in the lab can be difficult, and in fact even theoretical models do not typically distinguish the two.¹

One exception is provided by the Endogenous Depth of Reasoning (EDR) model of Alaoui and Penta (2016a), in which a subject's understanding of a game (his *cognitive bound*, or *capacity*) is formally distinct from his 'behavioral level'.² In the EDR model, holding constant an individual's cognitive bound, the observed level of play may vary with the individual's beliefs about the opponents. For instance, even if a subject understands up to five iterations of the level- k reasoning, he may sometimes play as a level-5, but he may instead play as a level-3 if he thinks the opponent would play as a level-2. The EDR model also allows players' very understanding, or capacity, to vary with the stakes of the game. For example, a subject may understand three iterations of the reasoning process when the stakes are low, but more when the stakes are high enough, depending on his cognitive abilities.

The EDR model provides a formal language with which to ask whether in practice level- k patterns of behavior are mostly driven by subjects' cognitive bounds or by their beliefs, possibly of higher order. But the experimental treatments in Alaoui and Penta (2016a, AP hereafter), which vary subjects' beliefs as well as the stakes for all players at the same time, shed little light on this particular question. AP's treatments also do not disentangle the extent to which the more sophisticated behavior observed in the 'high stakes' treatments is due to agents' deeper understanding or to their beliefs about the increased depth of reasoning of their opponents. Conceptually, the two points are related: if subjects did not reason at all about others' incentives to reason, then the higher sophistication in the 'high stakes' treatments would be exclusively driven by their own increased capacity; conversely, if behavior were purely determined by beliefs, then subjects' behavior should not change, if nothing changes about the opponents. The challenge is to identify these effects in the lab, which is the objective of the present paper.

To assess the extent to which agents' behavior is determined by a binding cognitive bound,

¹Friedenberg, Kets and Kneeland (2017) address a similar question, but from a different perspective. We discuss the differences and similarities with their work in Section 1.1.

²Alaoui and Penta (2016a) use the term 'cognitive bound' to refer to the highest level- k that a subject is able to conceive of, in a given game, which indirectly provides an upper bound to his behavioral level in the game. The term 'capacity' is due to Georganas, Healy and Weber (2015), essentially with the same meaning. Friedenberg, Kets and Kneeland (2017) also use the term 'cognitive bound', but with a different meaning. (See Section 1.1).

as opposed to beliefs, we design treatments based on the following simple idea, which we call the *tutorial method*. Suppose players are engaged in a standard game for level- k reasoning, such as a version of Nagel’s (1995) beauty contest.³ Now entertain the following thought experiment: take a subject, say Ann, whose choice in this game is consistent with level-3 behavior, and provide her with a game theory tutorial which explains the strategic structure of the game (best responses, iterated reasoning, uniqueness of equilibrium, etc.), but without providing any proper factual information (such as information about others’ choices, typical distributions of actions in this game, etc.). Next, ask Ann to play this game again, but against individuals who have *not* received the tutorial. Intuitively – putting aside difficulties in computing the best responses, noise in Ann’s reasoning or choice, and other caveats – if Ann perceives the new pool of opponents as identical to those in her first trial, then her action should change only if the game theory tutorial has made her understand something she deems useful. So, if her level-3 action in the pre-tutorial treatment was purely driven by her beliefs about the opponents (e.g., that they behave as level-2’s), then the tutorial should have no impact on her choice, and her behavior would be level-3 in both rounds. Conversely, if her action shifts (and especially if it shifts towards a higher level- k), then it must be the case that her previous understanding was in some sense ‘binding’, and hence her level-3 choice was not entirely due to her beliefs.

Our second question – understanding whether agents explicitly take into account others’ incentives to reason – is more directly motivated by the central premise of the EDR model, which is that agents’ cognitive bound may itself vary with the payoffs of the game. As explained above, however, the point is inherently related to the broader problem of the cognition-beliefs dichotomy. But disentangling own understanding from reasoning about others’ incentives to reason presents non-trivial conceptual difficulties. For instance, suppose – as assumed in the EDR model – that subjects’ cognitive bounds are increasing in their own stakes in the game. Then, intuitively, one way to disentangle the two effects is to consider ‘asymmetric transformations’ of payoffs in a two-player game, in which stakes are increased for one player (Ann) but not for the other (Bob). This change, however, would not be enough to isolate the effects on Ann’s own understanding, because she may think that Bob could react to her stronger incentives to reason. If this were the case, then a change in Ann’s behavior need not be driven by her own understanding, but by her beliefs about Bob’s reaction to her incentives. In other words, to isolate the effects of Ann’s higher stakes on her own understanding, it is important to hold constant Ann’s beliefs about Bob’s reasoning, of *any* order. For this reason, we design treatments to disentangle own reasoning from reasoning about the opponents’ reasoning. These treatments make use of what we call the *replacement method*. That is, in the asymmetric payoff treatment, Ann does not just play against a subject whose stakes are low; rather, Ann plays against the choice made by a player, Bob, who is engaged in a game in which stakes are low

³This paper will focus on another game, referred to as the acyclical 11-20 game. The logic of the argument would be the same, but we use Nagel’s (1995) beauty contest game here because it has been one of the main workhorses for level- k reasoning (see also Camerer (2003), Crawford, Costa-Gomes and Iriberry (2013) and references therein). Other prominent games in the level- k literature include the two-person guessing games of Costa-Gomes and Crawford (2006) and the travelers’ dilemma (e.g., Capra, Goeree, Gomez and Holt (1999)).

for both players (hence, Bob’s opponent is not Ann: in our treatment, Ann is ‘replaced’ by a low-stakes version of herself). This way, Ann’s beliefs (of any order) about Bob’s reasoning are identical to her beliefs in the low-stakes game, and hence any change in behavior observed when Ann’s stakes are increased can be unambiguously imputed to Ann’s own incentives to reason.

We apply both the *tutorial* and the *replacement* methods to three experiments. Experiment 1 leverages the existing dataset by applying these methods to the baseline experiments in Alaoui and Penta (2016a). Experiments 2 and 3 instead develop simpler variations of those treatments for a new pool of subjects, and address possible concerns on demand effects associated to the specific application of the tutorial method. A fourth experiment addresses possible robustness concerns on social effects which may arise in some of the treatments of Experiment 1. All experiments are based on the ‘acyclical 11-20 game’ from Alaoui and Penta (2016a), but the underlying logic has broader validity and does not rely on any specific feature of the game nor of the EDR model. Hence, our experimental designs suggest a general methodology that can be easily extended to other games and settings: The *tutorial method* can be used to investigate the cognition-beliefs dichotomy in general models of strategic thinking, and the *replacement method* can be used to explore higher-order beliefs effects in general games, with essentially no restrictions on the underlying payoffs.

Methodological considerations aside, our empirical findings show that, for a large fraction of subjects, the cognitive bound is actually binding when they play against opponents who are regarded as more sophisticated. This is a perhaps surprising result for the view that level- k behavior is mainly driven by beliefs: it suggests that, at least in some settings, level- k models are directly applicable to agents’ own understanding. On the other hand, we also find evidence that a large fraction of subjects do reason about others’ incentives to reason, providing support to a much more subtle implication of the EDR model than those that were previously tested. Overall, our results suggest that level- k behavior in general should not be taken as driven either by cognitive limits alone or beliefs alone: it depends on the complex interaction of cognitive bounds, beliefs about opponents’ cognitive abilities, and reasoning about the opponents’ reasoning processes. We also find that the EDR framework is a useful tool for analyzing and understanding this interaction, and that the results are overall consistent with its predictions.

The rest of the paper is organized as follows: Section 1.1 reviews the related literature, Section 2 introduces the baseline game and logistics common to all experiments, and Section 3 presents the specific treatments. Section 4 contains theoretical results on the EDR model which are relevant for the treatments in our experiment, and spells out the identification assumptions used to connect the model to the experimental findings. Section 5 presents the experimental results. Section 6 concludes.

1.1 Related Literature

The classical literature on the level- k and cognitive hierarchy models (e.g. Nagel (1995); Stahl and Wilson (1995) Costa-Gomes, Crawford, and Broseta (2001); Camerer, Ho, and Chong (2004); Costa-Gomes and Crawford (2006)) has analyzed systematic features of observed behavior which suggested that individuals follow distinct patterns of reasoning. This evidence has often being interpreted as being driven by individuals’ limited ability to reason strategically, but models in this literature are typically silent on whether the observed ‘levels of play’ stem from subjects’ cognitive limitations, or perhaps from their beliefs about others’ rationality (of any order) or their ability: most models are consistent with both interpretations.⁴

More recent experiments have focused on how levels of play vary across different games and with different opponents (e.g., Agranov, Potamites, Schotter and Tergiman (2012), Georganas, Healy and Weber (2015) and Alaoui and Penta (2016a)). Their findings suggest that, at least in some settings, level- k patterns of behavior may be driven by individuals’ beliefs rather than by intrinsic cognitive limitations. The distinction between ‘cognitive’ and ‘behavioral’ levels – that is, between the maximum level- k an agent can conceive of, due to his limited ability, and the level of his action, which may be driven by his beliefs – has been made explicit in some recent theoretical models: for instance, Strzalecki’s (2014) notion of level- k type only restricts the support of a type’s beliefs, but level- k behavior may vary as a type’s beliefs are varied; similarly, Alaoui and Penta (2016a) define the ‘cognitive bound’ as the maximum level an agent can conceive of, but that’s distinct from the ‘behavioral level’, which is jointly determined by the cognitive bound and the agent’s beliefs; Georganas, Healy and Weber (2015) also have an analogous distinction, and use the term ‘capacity’ essentially with the same meaning as Alaoui and Penta’s (2016a) ‘cognitive bound’. Similar ideas have been extended to dynamic games by Rampal (2018a). Rampal (2018b) also finds evidence of behavior driven by agents’ beliefs.

Friedenberg, Kets and Kneeland (2017) study a related problem concerning a rationality-cognition dichotomy, but where cognition refers to a distinct concept from ours. More specifically, in Alaoui and Penta (2016a) and Georganas, Healy and Weber (2015), the cognitive bound or capacity refers to a player’s understanding of the game in the sense of the level- k literature, that is as the highest level of iteration of best replies the player is able to conceive of (though, as we discussed, not necessarily the one he plays). In contrast, Friedenberg, Kets and Kneeland (2017) depart from the level- k literature in that they define a player to be ‘cognitive’ if his behavior responds – in any way, rationally or not – to changes in payoffs. This provides a measure of cognitive bound, which in their analysis identifies a lower bound to individuals’ reasoning ability. Similar to ours, their measure of cognitive bound is also at least as large as their rationality bound (which in turn is conceptually analogous to our behavioral level, although formally distinct), but it may be strictly larger than the cognitive level in our sense. Applying this broader notion of cognition to the experimental data from Kneeland (2015), they

⁴This literature has also spurred more sophisticated game theoretic work, which has tackled the challenging question of how to model players who can only conceive of finitely many orders of beliefs, and studied the behavioral implications of those situations (see, e.g., Kets (2017) and Heifetz and Kets (2018)).

find evidence of a significant gap between subjects’ rationality and cognitive bounds, and hence of their reasoning ability.

In AP’s EDR model the cognitive bound is endogenously determined by a player’s cognitive abilities (represented by costs of reasoning) and the incentives to reason (which depend on the game’s payoffs). AP tested the main predictions of the EDR model with the baseline treatments in Section 3.1.1, and showed how it can be used to perform robust predictions across games as well as explain the experimental findings in Goeree and Holt’s (2001) famous ‘little treasures’ experiment. Alaoui and Penta (2018) provide an axiomatic foundation of the model, by characterizing the properties of the reasoning process which justify a cost-benefit approach, as well as particular functional forms for the value of reasoning. Recent extensions of the approach include Alaoui and Penta (2016b, 2018), which extends the EDR model to account for response time, with an application to the experiment by Avoyan and Schotter (2016).

Gill and Prowse (2016, 2017) also investigate more explicitly the connection between level- k behavior and cognitive or non-cognitive abilities, but in a setting with feedback, thereby focusing on learning. They find significant effects of different IQs on the speed of learning, but not on the initial responses.

2 Baseline Game and General Logistics

The experiments are designed not only to test whether individuals play differently when their incentives and beliefs about opponents change, but also to analyze the direction in which their actions change, i.e., towards higher or lower level- k ’s. Moreover, we aim to disentangle whether their action is dictated by their cognitive constraints, given their incentives, or by their beliefs over their opponents’ cognitive constraints. These objectives are reflected in the choice of the baseline game, in the logistics of the experiment and in the subject’s classification criteria. In this section we discuss each of these elements of our design.

2.1 The acyclical 11-20 game

The baseline game remains the *acyclical 11-20 game* throughout:

The subjects are matched in pairs. Each subject enters an (integer) number between 11 and 20, and always receives that amount in tokens. If he chooses *exactly one less* than his opponent, then he receives an extra x tokens, where $x \geq 20$. If they both choose the same number, then they both receive an extra 10 tokens.

This game is a variation of Arad and Rubinstein’s (2012) ‘11-20’ game. The only difference is that the original version does not include the extra reward in case of a tie. As argued by Arad and Rubinstein, the 11-20 game presents a number of advantages in the study of level- k reasoning, which are inherited by our modified version (see also Alaoui and Penta (2016a) for this discussion):

First, using level- k reasoning is natural, as there are no other obvious focal ways of approaching the game. This is desirable because our aim is not to test level- k models against competing models of players' reasoning, but to investigate the effects of changing beliefs of payoffs or the distributions of level- k 's.

Second, the level-0 specification is intuitively appealing and unambiguous, since choosing 20 is a natural anchor for an iterative reasoning process. Moreover, it is the unique best choice for a player who ignores all strategic considerations.

Third, there is robustness to the level-0 specification, in that playing 19 would be the level-1 strategy for a wide range of level-0's, including the uniform distribution over the possible actions, or choosing 20.

Fourth, best-responding to any level- k is simple: level-1 plays 19, level-2 best responds to 19 by playing 18, and so on. Since we do not aim to capture cognitive limitations due to computational complexity, having a simple set of best responses is preferable in this case.

In addition to these points, our modification of the game leads to another useful feature for our objectives. By introducing the extra reward in case of a tie, the best response to 11 is 11, and not 20, as in the version of Arad and Rubinstein. Thus, our modification breaks the cycle in the chain of best responses, which enables us to assign one specific level of reasoning to each possible announcement (with the exception of 11, which corresponds to any level equal to 9 or higher): Action 19 can only be a level-1 strategy, 18 can only be a level-2 strategy, and so forth for every k up to $k = 8$.⁵ In the original 11-20 game, action 19 could have been played by a level-1, but also by a level-11, level-21, or other 'high' levels (levels of form $10n + 1$). Although levels-11 and above appear to be uncommon, it is crucial that these cycles be avoided here. That is because some of the hypotheses that we aim to test concern shifts in the distribution of level- k 's, but these hypotheses could not be falsified in the presence of such cycles.

2.2 General Logistics

The subjects of all experiments are undergraduate students from different departments at the Universitat Pompeu Fabra (UPF), in Barcelona. There were 278 subjects in total, with 120 participating in Experiment 1, 60 in Experiment 2, and 34 in Experiment 3 as well as 64 in robustness experiments (more details on the latter experiment will be provided in Section 3.1.3). Each experimental session took 1.5 hours.

All experimental treatments are based on the acyclical 11-20 game above, with an experimental currency where one token is worth 15 euro cents in experiment 1 and 12.5 in experiments

⁵The fact that every action in this game corresponds to some level- k makes the 11-20 unfit to test level- k models against alternative models of reasoning. That objective would be better attained considering games with large strategy spaces, such as Nagel's (1995) original beauty contest or Costa-Gomes and Crawford's (2006) two-person guessing game. As explained above, however, our objective is to test properties of level- k reasoning when beliefs and payoffs are varied, not to contrast level- k with alternative theories. The latter problem has been the focus of an already extensive literature, which overall has provided strong support for level- k reasoning in a variety of settings. For recent work in this direction, see Kneeland (2015). For a more nuanced view, Goeree, Louis and Zhang (2016) argue that a version of level- k models with noise (namely, the noisy introspection model, see Goeree and Holt (2004)) outperforms the baseline level- k model without noise.

2 and 3. The exact sequences of treatments used in each session and experiment are provided in Appendix B.2. Each subject in the experiments was anonymously paired with a new opponent after every iteration of the game. To focus on initial responses and to avoid learning from taking place, the subjects received no feedback after their play, and they only observed their earnings at the end of the session. As is standard in the literature on initial responses (see, e.g., Costa-Gomes, Crawford and Broseta (2001) and Costa-Gomes and Crawford (2006)), subjects were paid randomly, and therefore did not have any mechanism for hedging against risk by changing their actions. Subjects were informed of the payment method before starting the experiment. Lastly, subjects received no information concerning other subjects' earnings. This serves to avoid that subjects focus on goals other than monetary incentives, such as defeating the opponent or winning for its own sake. The instructions of the experiment were given in Spanish; the English translation and the details on the pool of subjects, the earnings and the logistics of the experiment are in Appendix B.

2.3 Subjects' Classifications

In Experiment 1, we divided the pool of subjects into two groups, according to two criteria (with 3 sessions of 20 subjects each) designed to be indicative of subjects' cognitive sophistication. The first criterion, referred to as the exogenous classification, separates subjects by their degree of study. Half of the students are drawn from humanities, and the other half from math and sciences. Subjects are then made aware of their own classification by being labeled as either 'humanities' or 'math and sciences'. In the endogenous classification, subjects are not separated by degree of study; instead, they are separated by a test that they take at the beginning of the experiment. This test consists of a centipede game, a pirates game and a simplified version of mastermind (see Appendix B for details). The top half of the subjects are labeled as high, and the bottom half as low. Here as well, the subjects are made aware of their own label, 'high' or 'low'. The details are provided in Appendix B.

We use two different classifications for the following reason. In the exogenous classification, the labels are informative of a long-lasting, persistent and salient indicator of the subjects' cognitive sophistication.⁶ The downside, however, is that it is a coarser notion of cognitive sophistication, and is not specifically linked to subjects' ability to reason strategically. In the endogenous classification instead, the labels are made based on a short test and may not necessarily have the persistent strength of the exogenous classification, but the advantage is that the test specifically targets game theoretic reasoning, and so may induce sharper beliefs among the subjects concerning their relative sophistication compared to their opponents.

In Experiments 2 and 3, subjects were not separated by cognitive sophistication. They first took an expanded version of the test used in the endogenous classification described above. This version includes the muddy faces game in addition to the others (see Appendix B). Then,

⁶This is the case especially when considering that at the university from which the subjects are drawn, there is a significant difference in the entry grades between those taking the fields grouped as humanities in our classification and those taking the fields grouped as math and science.

Baseline Treatments	Opponent’s label compared to own	Own payoffs	Opponent’s payoffs	Replacement of opponent’s opponent
Homogeneous [Hom]	same	Low	Low	No
Heterogeneous [Het]	different	Low	Low	No
Higher-Order Beliefs [HOB]	different	Low	Low	Yes
Homogeneous-high [Hom+]	same	High	High	No
Heterogeneous-high [Het+]	different	High	High	No
Higher-Order Beliefs-high [HOB+]	different	High	High	Yes

Table 1: Summary of the baseline treatments

only the subjects in the middle half of the distribution participated in the treatments below. The middle group was identified based on a pre-specified cutoff that was based on the test scores from Experiment 1. Subjects in this group were not given any information about their performance on the test before the treatments were administered. Those subjects who had high or low test results were given tasks to occupy them to prevent disturbances from subjects leaving during the experiment. This particular procedure was carried through for the following reasons. First, taking the test beforehand places the subjects in a similar condition to those of the endogenous classification treatments. Second, the treatments in Experiments 2 and 3 are designed to focus more on the middle part of the distribution, since in Experiment 1 the treatments of the exogenous classification focused on the tails of the distribution and those of the endogenous classification split subjects along the median. We do so to obtain additional information about a different segment of the distribution. We refer to the subjects in Experiments 2 and 3 as being in the ‘unlabeled’ classification.

3 Experimental Design

3.1 Experiment 1: Treatments

We present next the treatments of Experiment 1. In Section 3.1.1 we review the baseline treatments in Alaoui and Penta (2016a), which test the basic premises of the EDR model by varying incentives or beliefs for both agents at the same time. These treatments, however, do not disentangle whether subjects’ change in behavior is due to changing their own incentives or to changing the incentives of the opponents, and whether subjects’ choices are mainly driven by their beliefs about the opponents, or by their own cognitive limitations. The new treatments, designed to disentangle these effects, are introduced in Sections 3.1.2 and 3.1.3.

3.1.1 Baseline Treatments

AP’s baseline treatments, summarized in Table 1, are designed to implement the two sets of comparative statics (on incentives and beliefs) we discussed above.

Varying Incentives. To vary subjects’ incentives to reason, we consider two versions of the game: in the ‘low payoffs’ treatment, we set $x = 20$; in the ‘high payoffs’ treatments, we let $x = 80$. Note that this change does not affect the level- k actions, irrespective of whether the

The Tutorial Treatments	Opponent’s label compared to own	Own payoffs	Opponent’s payoffs	Own Tutorial	Opponent’s Tutorial
Tutorial [Tut]	–	Low	Low	Yes	Yes
Asymm. Tutorial-Homog. [AT-Hom]	same	Low	Low	Yes	No
Asymm. Tutorial-Heterog. [AT-Het]	different	Low	Low	Yes	No

Table 2: Summary of the post-tutorial treatments.

level-0 is specified as 20 or as the uniform distribution. It only increases the rewards for players who stop at the ‘correct’ round of reasoning, and hence the ‘incentives to reason’.⁷

Varying Beliefs. To vary agents’ beliefs, for both specifications of payoffs and for both the classification criteria discussed in Section 2.3, subjects in each treatment are given information concerning their opponent’s label. They play the baseline game against someone from their own label (*homogeneous treatment*) and against someone from the other label (*heterogeneous treatment*). For instance, for the exogenous classification, a student from math and sciences (resp., humanities), is told in the homogeneous treatment that his opponent is a student from math and sciences (humanities). In the heterogeneous treatment, he is told that the opponent is a student from humanities (math and sciences). Identical instructions are used for the endogenous classification, but with ‘high’ and ‘low’ instead of ‘math and sciences’ and ‘humanities’, respectively.

The homogeneous and heterogeneous treatments are designed to test whether the behavior of the subjects varies with the sophistication of the opponent. The next treatment is designed to test whether the subjects believe that the behavior of their opponents also changes when they face opponents of different levels of sophistication. To do so, we consider a *higher order beliefs treatment*: A ‘math and sciences’ subject, for instance, is given the following instructions: “[...] two students from humanities play against each other. You play against the number that one of them has picked.”

In the following, we let [Hom], [Het] and [HOB] denote, respectively, the homogeneous, heterogeneous and higher-order beliefs treatments when payoffs are low, and [Hom+], [Het+] and [HOB+] the corresponding treatments when payoffs are high.

3.1.2 Relaxing cognitive bounds: the post-Tutorial Treatments

We explain next the new treatments designed to identify whether or not subjects play according to their own (binding) cognitive bound in treatments [Hom] and [Het].

Consistent with the intuitive idea of the *tutorial method* discussed in the introduction, after having administered the baseline treatments of Section 3.1.1, we exposed all eighty subjects from four of the six sessions (two for the endogenous and two for the exogenous classifications) to a ‘game theory tutorial’. This tutorial explains how, through the chain of best replies, ‘infinitely sophisticated and rational players’ would play (11, 11). We then proceed with three new (post-tutorial) treatments, each repeated twice, and summarized in Table 2.

⁷Alaoui and Penta (2018) provide axiomatic foundations to this assumption of the EDR model.

Asymmetric Payoffs Treatments	Opponent’s label compared to own	Own payoffs	Opponent’s payoffs	Replacement of opponent’s opponent
Asymm. Payoffs-Homogeneous [AP-Hom]	same	High	Low	Only payoffs
Asymm. Payoffs-Heterogeneous [AP-Het]	different	High	Low	Both label and payoffs

Table 3: Summary of the asymmetric payoff treatments.

In treatment [Tut], we instruct each subject to play the baseline game (with low payoffs) against another subject who has also been given the same tutorial, with no information about his label. In the ‘asymmetric tutorial-homogeneous’ treatment [AT-Hom], we instruct the subjects who had previously received the tutorial to play the baseline game against a player of the same label who had *not* received the tutorial (that is, as in the baseline homogeneous treatment [Hom]). Analogously, the asymmetric tutorial-‘heterogeneous’ treatment [AT-Het] contains the same instructions but with the subjects facing an opponent from a different label (as in baseline treatment [Het]).

Hence, subjects essentially face the same opponents in treatments [AT-Hom] and [Hom] (and in [AT-Het] and [Het]), but the tutorial ensures that their cognitive bounds are not binding in treatments [AT-Hom] and [AT-Het]. Hence, if their cognitive bound was *not* binding in (pre-tutorial) treatments [Hom] and [Het], then the distributions of actions in (post-tutorial) treatments [AT-Hom] and [AT-Het] should be the same as in [Hom] and [Het], respectively.

3.1.3 Reasoning about Others’ Incentives: Asymmetric Payoffs Treatments

In this section we explain the treatments designed to disentangle the effects of increasing payoffs on subjects’ own cognitive bound from their reasoning about others’ incentives to reason. As discussed in the introduction, this question is inherently related to the cognition-belief dichotomy (the object of the treatments in Section 3.1.2), but it is more directly motivated by the basic premise of the EDR.

In the design of treatments [Hom+], [Het+] and [HOB+], relative to [Hom], [Het], [HOB], we increase the payoff for undercutting the opponent for both players in the game. Thus, the shifts in the distributions towards lower numbers observed in Alaoui and Penta (2016a, Section 4) may conflate two distinct effects. The first effect is the possible increase in the cognitive bound of player i , and the second is the change in i ’s beliefs about j ’s cognitive bound due to the change in j ’s incentives. Both effects would determine an increase in the behavioral level, hence a shift of the distribution towards lower actions. The following treatments, summarized in Table 3, are aimed at disentangling the two effects, and testing whether subjects in our experiment reason about their opponents’ incentives independently of their own.

As discussed earlier, the intuitive idea is to increase the stakes for one player without changing the other player’s. To address the problem of higher-order beliefs discussed in the introduction, however, in these treatments we apply the *replacement method* to the game with asymmetric payoffs. That is, in treatments [AP-Hom] and [AP-Het] agents play the high-payoff game against the number chosen by an opponent from their own or the other label respectively

in the low payoffs treatment [Hom]. Hence, the exercise is of a similar spirit to treatment [HOB], in which subjects play against the number chosen by an opponent from a different label who is engaged in treatment [Hom]. Both treatments are administered after the baseline treatments to all forty subjects from two sessions, one exogenous and one endogenous, and each is repeated three times.

Note that these treatments add a further layer of complexity, since the individual is told in treatment [AP-Hom] (resp., [AP-Het]) that he is playing the high-payoff game against the number chosen by an opponent of the same (other) label himself playing the low payoff game against an opponent of their own label. Treatment [AP-Het] is especially complex: for player i , both the payoffs and the label of i 's opponent *and* of the opponent's opponent are different from i 's own payoff and label.

By comparing treatments [AP-Hom] and [AP-Het] with treatments [Hom] and [HOB] and with treatments [Hom+] and [HOB+], we can disentangle the two effects mentioned above. The shift from [Hom] to [AP-Hom] (and from [HOB] to [AP-Het]), due solely to the increase of each subject's own payoffs and not his opponent's, may be attributed to the increase of subjects' own cognitive bound. It should be observed only if the cognitive bound in treatments [Hom] and [HOB] had been binding; the further shift from [AP-Hom] to [Hom+] (and from [AP-Het] to [HOB+]) instead can be imputed to the increase in subjects' beliefs about their opponents' behavior due to the increase of their payoffs.

Robustness experiment. When comparing [Hom] to [AP-Hom], one potential issue is that the replacement method could remove social preferences, if those were present. Specifically, social preferences might impact [Hom] but not [AP-Hom] so that when we compare the behavior across these treatments, we might confound differences in levels with social preference differences. To assess the robustness of our findings, we designed an extra experiment. We conducted a robustness experiment that also contains a new homogeneous treatment with replacement, [Hom-Rep], in order to make [Hom] and [AP-Hom] more comparable.⁸ In [Hom-Rep], the subject played against another subject from their own group who in turn was playing yet another player from their own group. This means that there now is replacement both for the [Hom-Rep] benchmark and the replacement treatment [AP-Hom]. For comparability, the experiment also contained the basic [Hom] treatment. We conducted the experiment with 64 subjects at Universitat Pompeu Fabra, in May 2019.

3.2 Experiment 2: Unlabeled Variations

A possible concern with the design of Experiment 1 is that the joint presence of beliefs and payoff treatments may increase the complexity of the experimental instructions, adding noise to the results. Experiment 2 is designed to simplify the cognitive load of the instructions by replicating the main treatments without label distinctions. For this reason, subjects in this experiment are not given any information on their own or the opponents' performance in the

⁸We thank an anonymous referee for suggesting this treatment.

Baseline Treatments	Own payoffs	Opponent's payoffs	Own Tutorial	Opponent's tutorial	Replacement of opponent's opponent
Unlabeled [Un]	Low	Low	No	No	No
Unlabeled-high [Un+]	High	High	No	No	No
Tutorial-Unlabeled [Tut-Un]	Low	Low	Yes	Yes	No
Asymm.Tut.-Unlab. [AT-Un]	Low	Low	Yes	No	Tutorial only
Asymm.Payoffs-Unlab. [AP-Un]	High	Low	No	No	Payoffs only

Table 4: Summary of the treatments in Experiment 2.

cognitive sophistication test. This change simplifies the instructions, particularly in the post-tutorial treatments and asymmetric payoff treatments. This is because a subject need not keep track at the same time of whether the opponent has taken the tutorial (for the former case, or has different payoffs for the latter) and of his possibly different label as well as the higher order beliefs over the opponent's opponent on both dimensions.

All 60 subjects participate in all treatments. As usual, they do not receive feedback, and they are paid at random based on their behavior on a subset of the treatments (see Appendix A for the exact sequence of treatments, payment details and wording of instructions).

Treatments. The baseline unlabeled treatment, denoted [Un], consists of the subjects playing the baseline acyclical 11-20 game with extra reward $x = 20$ (see Section 2.1), and Unlabeled-high [Un+] consists of the subjects playing the high payoffs game with extra reward $x = 80$.

The tutorial-unlabeled treatment [Tut-Un] is identical to post-tutorial treatment [Tut] except that subjects are not given any label. In Treatment [AT-Un], subjects who have seen the tutorial play against the action chosen by an opponent in pre-tutorial treatment [Un]. This treatment serves to fix both the subject's beliefs that his opponent has not seen the tutorial and his beliefs that his opponent's opponent has not seen the tutorial either (and so on for all higher-order beliefs). Lastly, we have adapted the asymmetric payoffs treatments in the same way, so that Asymmetric Payoffs-Unlabeled [AP-Un] is identical to [AP-Hom] and [AP-Het] without label information. These treatments are summarized in Table 4.

3.3 Experiment 3: Unlabeled Variations with preliminary semi-Tutorial

One possible concern with the previous experiments could be that the tutorial induces level- k thinking. If the subjects were not adopting that kind of reasoning in the pre-tutorial treatments, the difference between the pre- and post-tutorial behavior could in part be due to this form of priming. If that were the case, interpreting the changes in behavior as being due to relaxing a possible cognitive constraint of the subjects may be problematic, as it could be conflating different factors.⁹

To address these concerns, Experiment 3 replicates Experiment 2, but precedes the administration of the treatments with instructions designed to keep the priming as similar as possible between the tutorial and non-tutorial treatments. Specifically, we modified the instructions before the first game to have the following explanation, and kept all else identical:

⁹We are thankful to an anonymous referee for this observation.

Baseline Treatments	Own payoffs	Opponent's payoffs	Own Tutorial	Opponent's tutorial	Replacement of opponent's opponent
Unlabeled [E3-Un]	Low	Low	No	No	No
Unlabeled-high [E3-Un+]	High	High	No	No	No
Tutorial-Unlabeled [E3-Tut-Un]	Low	Low	Yes	Yes	No
Asymm.Tut.-Unlab. [E3-AT-Un]	Low	Low	Yes	No	Tutorial only
Asymm.Payoffs-Unlab. [E3-AP-Un]	High	Low	No	No	Payoffs only

Table 5: Summary of the treatments in Experiment 3.

If you think that your opponent will choose 20, or any number with equal likelihood, then the action that will maximize your earnings would be to choose 19. If you think that following this observation means that your opponent chooses 19, then the action that maximizes your earnings would be to choose 18.

Arguably, these instructions prime level- k reasoning in equal measure as the tutorial treatment, in that they make explicit the start of the chain of reasoning. They are effectively identical to the tutorial except that they do not spell out the *entire* path of reasoning all the way to 11. Hence, if such priming were the main driving force in the results for Experiments 1 and 2, then it would not be the case here. Since, in Experiment 3, all subjects were already primed to think according to level- k when playing the first treatment [E3-Un], changes observed between the pre- ([E3-Un]) and post-tutorial responses ([E3-AT-Un]) should not stem from priming but from changes in their cognitive levels.

In addition, subjects were asked to complete a Theory of Mind test (Stiller and Dunbar 2007) at the end of the experiment to provide a measure of their ability to place themselves in the mind of another person and to thus form higher order beliefs (we further explain this test in Section 5.3).

Treatments. The treatments and logistics are just the same as in Experiment 2, with the only difference being that all 34 subjects were exposed to the semi-tutorial at the very beginning of the session. We thus maintain the corresponding labels, preceded by “E3” to denote they refer to the variation of Experiment 3. These treatments are summarized in Table 5.

4 The EDR model

The basic idea of the EDR model is that a subject’s ‘level of play’, or *behavioral level*, may be endogenous due to two related mechanisms. First, given a subject’s understanding of a game (his *cognitive bound*, or *capacity*), his ‘behavioral level’ may vary with his beliefs about the opponent: e.g, even if a subject understands up to five iterations of the level- k reasoning, he may sometimes play as a level-5 (e.g., choose 15 in the 11-20 game), but sometimes play as a level-3, if he thinks the opponent would play as a level-2. But clearly, it is a matter of definition that one never plays according to a higher level than one’s own capacity. Hence, if \hat{k}_i is the

cognitive bound of subject i , his possible ‘behavioral levels’ are $k_i \leq \hat{k}_i$. And for the same reason, i ’s perception of the opponent’s capacity, \hat{k}_j^i , is also bounded by his own: $\hat{k}_j^i < \hat{k}_i$.¹⁰

The second dimension of endogeneity is that the understanding of a game (i.e., the capacity), may itself vary with a player’s stakes in the game. For instance, in the acyclical 11-20 game, in which $x_i \geq 20$ denotes the extra reward that player i gets for being exactly one below the opponent, the EDR model implies that agent i ’s capacity \hat{k}_i (as well as his perception of the opponent’s capacity, \hat{k}_j^i) is weakly increasing in x_i , and that \hat{k}_j^i is weakly increasing in the opponent’s extra reward, x_j . We introduce next a simple version of the EDR model which formalizes these ideas, as well as the interactions between individuals’ beliefs and incentives to reason, and derive its predictions for the treatments in the experiments above. (All proofs of our results are in Appendix A).

4.1 Baseline Model

Own understanding, Costs and Values: The endogeneity of players’ capacities is modeled as stemming from a cost-benefit analysis: costs represent players’ cognitive abilities; the benefits instead only depend on the game’s payoffs, such as the x parameter in the 11-20 game. Formally, fixing the game payoffs, let $v_i : \mathbb{N} \rightarrow \mathbb{R}_+$ denote the value of reasoning, where $v_i(k)$ represents i ’s value of doing the k -th round of reasoning, given the previous $k-1$ rounds. The cognitive ability of agent i is represented by a cost function $c_i : \mathbb{N} \rightarrow \mathbb{R}_+$, where $c_i(0) = 0$ and $c_i(k)$ denotes i ’s incremental cost of performing the k -th round of reasoning. We say that cost function c' is ‘more (resp. less) sophisticated’ than c'' , if $c'(k) \leq c''(k)$ (resp., $c''(k) \leq c'(k)$) for every k . For any $c_i \in \mathbb{R}_+^{\mathbb{N}}$, we denote by $C^+(c_i)$ and $C^-(c_i)$ the sets of cost functions that are respectively ‘more’ and ‘less’ sophisticated than c_i .

For the 11-20 games of our experiments, the general assumptions of the EDR model imply that: (i) the value of reasoning only depends on i ’s payoffs of the game; (ii) for every i and k , the value $v_i(k)$ is (weakly) increasing in x_i and constant in x_j ; (iii) $v_i = v_j$ if $x_i = x_j$; (iv) the costs of reasoning are constant throughout all variations of the game in all non post-tutorial treatments.¹¹ To obtain testable predictions in our experiments, further assumptions on subjects’ beliefs and the tutorial’s effects are needed, and will be discussed below.

To allow for the case, as in our asymmetric payoff treatments, that j ’s opponent is not i but a low-payoff version of it, besides values v_i and v_j we also introduce $v_{i(j)}$, to denote the value of reasoning of j ’s opponent. This value may coincide with i ’s own value ($v_{i(j)} = v_i$) – as in the standard treatments – or not – as in the asymmetric payoffs treatments ($v_{i(j)} \neq v_i$). More specifically, in general we assume that (v) $v_{i(j)}$ is equal to what would be i ’s value if his extra reward x_i were equal to that of j ’s opponent, $x_{j(i)}$.

¹⁰A different modeling choice would be to assume that the players first consider the sophistication of their opponent, and stop reasoning as soon as they believe they have exceeded it if the opponent is less sophisticated; that is, soon as player i reaches step $\hat{k}_j^i + 1$. This would lead to a different interpretation in that own capacity and beliefs would coincide, but it would be behaviorally equivalent. (cf. Alaoui and Penta, 2016a).

¹¹These assumptions are all implied by the general, ‘detail free’, EDR model (Alaoui and Penta, 2016a), as well as consistent with the axiomatic foundations provided in Alaoui and Penta (2019). We refer to those papers for discussions of the conceptual significance of these assumptions.

For later reference, we define a mapping $\mathcal{K} : \mathbb{R}_+^{\mathbb{N}} \times \mathbb{R}_+^{\mathbb{N}} \rightarrow \mathbb{N}$ such that, $\forall (c, v) \in \mathbb{R}_+^{\mathbb{N}} \times \mathbb{R}_+^{\mathbb{N}}$,

$$\mathcal{K}(c, v) := \min \{k \in \mathbb{N} : c(k) \leq v(k) \text{ and } c(k+1) > v(k+1)\}, \quad (1)$$

where $\mathcal{K}(c, v) = \infty$ if the set in equation (1) is empty. In words, this mapping identifies the first intersection between the value v and the cost c .

Player i 's *cognitive bound* is the value that this function takes at (c_i, v_i) :

$$\hat{k}_i = \mathcal{K}(c_i, v_i). \quad (2)$$

Beliefs and Others' Understanding: To distinguish players' cognitive and behavioral levels, the EDR model also specifies beliefs about the opponent's costs, as well as higher order beliefs, which are then used to derive i 's beliefs about the opponent cognitive bound, his beliefs about j 's beliefs about i 's bound, and so on. In the general EDR model, such beliefs are modeled through *cognitive type spaces*, which can be used to represent arbitrary belief hierarchies over players' costs (cf. Alaoui and Penta, 2016a). Here, however, it suffices to focus on the simpler case of *second-order types with degenerate beliefs*. These beliefs are pinned down by (1) i 's cost function, c_i , (2) i 's beliefs about j 's cost function, c_j^i , and (3) i 's beliefs about j 's beliefs over i 's cost function, c_i^{ij} (which may or may not be such that $c_i^{ij} = c_i$ – the further special case of “common belief ” types).¹² A *type* in the following is thus a triple $t_i = (c_i, c_j^i, c_i^{ij})$. With this notation, we define i 's beliefs about j 's cognitive bound (given his own bound \hat{k}_i , and his beliefs about j 's cost, c_j^i) as:

$$\hat{k}_j^i = \min \left\{ \hat{k}_i - 1, \mathcal{K}(c_j^i, v_j) \right\}. \quad (3)$$

The minimum operator represents the idea that i 's beliefs over j 's capacity are bounded by his own cognitive bound, \hat{k}_i , which effectively only uncovers the understanding of levels $k_i < \hat{k}_i$.

Player i 's beliefs over \hat{k}_j^i , j 's cognitive bound, do not depend on c_i^{ij} , but his beliefs over k_j^i , j 's level of play, do. That is, k_j^i can be below \hat{k}_j^i if i believes that his sophistication is underestimated by j . Put differently, player i attempts to place himself in the mind of j , and views j 's beliefs over his own cognitive bound, \hat{k}_i^{ij} , to be:

$$\hat{k}_i^{ij} = \min \left\{ \hat{k}_j^i - 1, \mathcal{K}(c_i^{ij}, v_{i(j)}) \right\}. \quad (4)$$

The minimum operator in (4) reflects the fact that, just as i 's beliefs over j 's capacity are bounded by \hat{k}_i , so his beliefs about j 's beliefs over i 's capacity are bounded by what i thinks j 's capacity is. The fact that $v_{i(j)}$ is used reflects the idea that i understands that j 's behavior is based on his reasoning about his opponent, which – depending on the treatment – may be i himself ($v_{i(j)} = v_i$) or a replaced version of it with different payoffs ($v_{i(j)} \neq v_i$).

Behavior: For player i 's beliefs over j 's behavior, player i expects j to play according to level

¹²We refer to Alaoui and Penta (2016a) for the general case with non-degenerate beliefs, and an explanation of how second-order types map to the language of cognitive type spaces of the general model.

$\hat{k}_i^{ij} + 1$, provided that i thinks that j is capable of conceiving of such a level, which is the case if $\hat{k}_i^{ij} + 1 \leq \hat{k}_i - 1$. Otherwise, i thinks that j is limited by his own cognitive bound. Hence, for a general second-order type, i 's perception of j 's behavioral bound is:

$$k_j^i = \min \left\{ \hat{k}_i^{ij} + 1, \hat{k}_i - 1 \right\}. \quad (5)$$

Player i then best responds to k_j^i , and hence his *behavioral level* is $k_i = k_j^i + 1$. In the acyclical 11-20 game, the associated actions are $a_i^{k_i} = 20 - k_i$ if $k_i \leq 9$, and $a_i^{k_i} = 11$ otherwise.

Letting x_i, x_j , and $x_{i(j)}$ denote, respectively, the extra reward in the 11-20 game received by player i, j , and j 's opponent, in the following proposition we refer to the 11-20 game with $x_i = x_j = x_{i(j)} = 20$ as the *low payoff game*, to the case with $x_i = 80 \neq x_j = x_{i(j)} = 20$ as the *asymmetric payoff game*, and to the case $x_i = x_j = x_{i(j)} = 80$ as the *high payoff game*.

Proposition 1 *Fix a second-order type $t_i = (c_i, c_j^i, c_i^{ij})$. Then:*

1. *If $c_j^i, c_i^{ij} \in C^+(c_i)$, then the cognitive bound is binding in the low-payoff game, and $k_i = \hat{k}_i$. In the asymmetric payoff game, both k_i and \hat{k}_i are (weakly) higher than in the low payoff game; the cognitive bound may or may not be binding anymore, and $k_i \leq \hat{k}_i$ may also hold with strict inequality. In the high payoff game, \hat{k}_i remains the same as in the asymmetric payoff game, and it is such that $k_i = \hat{k}_i$; k_i may increase or stay the same, but it increases only if the cognitive bound was not binding in the asymmetric payoff game.*
2. *If $c_j^i \in C^-(c_i)$, or if $c_j^i \in C^+(c_i)$ and $c_i^{ij} \in C^-(c_i)$, then the cognitive bound may or not be binding in the low-payoff game, and $k_i \leq \hat{k}_i$ may also hold with strict inequality. In the asymmetric payoff game, \hat{k}_i is (weakly) higher than in the low payoff game; k_i may increase or stay the same, but it increases only if the cognitive bound was binding in the low-payoff game. In the high-payoff game, \hat{k}_i is the same as in the asymmetric payoff game, and may or may not be binding; k_i may increase or stay the same, but it increases only if the cognitive bound was not binding in the asymmetric payoff game.*
3. *For any (c_j^i, c_i^{ij}) and $(x_i, x_j, x_{i(j)})$, replacing c_i with some lower-cost $c_i' \in C^+(c_i)$ always induces a (weakly) higher \hat{k}_i ; k_i may increase or stay the same, but it increases only if the cognitive bound was binding in the first place.*

4.2 Identification Assumptions and Predictions for the Experiments

As briefly mentioned, to obtain testable predictions in our experiments we need to append the EDR's model assumptions (i)-(v) on the properties of the costs and value of reasoning, with identification assumptions on how the treatment variations impact players' beliefs and costs of reasoning (which are the exogenous *types* in the EDR model).

The first identification assumption we introduce formalizes the effects of the tutorial as effectively eliminating the costs of performing any step of level- k reasoning:¹³

¹³While it may well be that the cost is not shifted all the way down to 0 in the post-tutorial treatments, this assumption ties our hands maximally, and allows for clean predictions.

Equivalence classes for c_i in the treatments of Exp.2 (resp., Exp. 3)	Equivalence classes for c_j^i in the treatments of Exp.2 (resp., 3)	Equivalence classes for c_i^{ij} in the treatments of Exp.2 (resp., 3)
all except [Tut-Un] (resp., [E.3-Tut-Un]) [Tut-Un] (resp., [E.3-Tut-Un])	all except [Tut-Un] (resp., [E.3-Tut-Un]) [Tut-Un] (resp., [E.3-Tut-Un])	all except [Tut-Un] (resp., [E.3-Tut-Un]) [Tut-Un] (resp., [E.3-Tut-Un])

Table 6: Equivalence classes for types components in the treatments of Experiments 2 and 3

Equivalence classes for c_i in the treatments of Exp.1	Equivalence classes for c_j^i in the treatments of Exp.1	Equivalence classes for c_i^{ij} in the treatments of Exp.1
[Hom]=[Hom+]=[Het]=[Het+]=[HOB]= [HOB+]=[AP-Hom]=[AP-Het] [Tut]=[AT-Hom]=[AT-Het]	[Hom]=[Hom+]=[AT-Hom] [Het]=[Het+]=[HOB]= =[HOB+]=[AP-Het]=[AT-Het] [Tut]	[Hom]=[Hom+]=[Het]=[Het+]=[AP-Hom] [HOB]=[HOB+]=[AP-Het]=[AT-Het] [Tut]

Table 7: Equivalence classes for types components in the treatments of Experiment 1

IA.1: in the post-tutorial treatments, $c_i(k) = 0$ for all k .

The next assumption restricts the way that subjects' beliefs vary from one treatment to the other. While this may be possible in practice, in order to ensure that the model has bite we impose the most restrictive assumptions on beliefs which is sensible in the present context:

IA.2: For all treatments other than [Tut], [Tut-Un] and [E3-Tut-Un], subject i 's first-order beliefs, c_j^i , only depend on the label of the opponent, and his second order beliefs c_i^{ij} only depend on the label of the opponent's opponent.

For the unlabeled treatments of experiments 2 and 3, IA.2 effectively implies that beliefs c_j^i and c_i^{ij} are constant throughout the experiment, except for treatments [Tut-Un] and [E3-Tut-Un]. The reasons why the latter treatments are treated differently is that in such treatments subjects are informed that the opponent also took the tutorial, and hence – consistent with the spirit of IA.1 – we assume that beliefs $c_j^i(k)$ and $c_i^{ij}(k)$ also get lower for every k . As previously discussed, however, our main interest lies in comparing the pre- and post-tutorial treatments, not in the [Tut] treatment per se. For those pre- and post-tutorial comparisons, in which opponents are effectively the same, IA.2 implies that c_j^i and c_i^{ij} remain unchanged. Table 6 summarizes the implications of identification assumptions IA.1 (first column) and IA.2 (second and third columns) for the treatments in Experiments 2 and 3.

The implications of IA.2 for the labeled treatments in Experiment 1 are more complicated. They imply, for instance, that c_j^i remains the same in the [Het] and [HOB] treatments, but may change moving from [Hom] to [Het]; in contrast, c_i^{ij} is the same in the [Hom] and [Het] treatments, but may change moving from [Het] to [HOB] (and similarly for the high payoff versions of these treatments). Finally, both c_j^i and c_i^{ij} remain the same moving from treatment [X] to [X+], for all $X \in \{\text{Hom, Het, HOB}\}$. The overall implications of assumptions IA.1-2 for the treatments of Experiment 1 are summarized in Table 7.

Under these two basic identification assumptions, the EDR model implies clear predictions on the comparisons for most of our treatments, which we summarize in Proposition 2. In the

following, we let F_X^l denote the cumulative distribution of actions $a \in \{11, \dots, 20\}$ in treatment $[X]$ for label $l \in \{I, II, *\}$ (where “ $l = *$ ” means *unlabeled*, as in Experiments 2 and 3), and denote by \succsim (resp., \succ) the weak (resp., strict) first order stochastic dominance relation.¹⁴

Proposition 2 *For any distribution over subjects which satisfy identification assumptions IA.1-2, under the maintained assumptions of the EDR model (Section 4.1), the following holds:*

1. In Experiment 1, for each $l \in \{I, II\}$: $F_X^l \succsim F_{X+}^l$ for all $X \in \{Hom, Het, HOB\}$.
2. In Experiment 1, for each $l \in \{I, II\}$: (i) $F_{Hom}^l \succsim F_{AP-Hom}^l \succsim F_{Hom+}^l$, with $F_{Hom} \succ F_{AP-Hom}$ only if \hat{k}_i was binding in $[Hom]$ for some i ; and (ii) $F_{HOB}^l \succsim F_{AP-Het}^l \succsim F_{HOB+}^l$, with $F_{HOB} \succ F_{AP-Het}$ only if \hat{k}_i was binding in $[HOB]$ for some i .
3. In Experiment 1, for each $l \in \{I, II\}$: $F_{Hom}^l \succsim F_{AT-Hom}^l$, and $F_{Het}^l \succsim F_{AT-Het}^l$, each strictly only if \hat{k}_i was binding for some i in $[Hom]$ and $[Het]$, respectively.
4. In Experiments 2 and 3: (i) $F_{Un}^* \succsim F_{AP-Un}^* \succsim F_{Un+}^*$, with $F_{Un}^* \succ F_{AP-Un}^*$ only if \hat{k}_i was binding in $[Un]$ for some i ; and (ii) $F_{E3-Un}^* \succsim F_{E3-AP-Un}^* \succsim F_{E3-Un+}^*$, with $F_{Un}^* \succ F_{E3-AP-Un}^*$ only if \hat{k}_i was binding in $[E3-Un]$ for some i .
5. In Experiments 2 and 3: $F_{Un}^* \succsim F_{AT-Un}^*$ and $F_{E3-Un}^* \succsim F_{E3-AT-Un}^*$, each strictly only if \hat{k}_i was binding for some i in $[Un]$ and $[E3-Un]$, respectively.

The remaining identification assumptions only concern the treatments in Experiment 1.

IA.3: For the labeled treatments of Experiment 1, we assume that individuals commonly believe that label I players are more sophisticated than label II . Formally, for label I individuals: if $l_i = I$, $c_j^{i,[Het]} \in C^-(c_j^{i,[Hom]})$, $c_i^{ij,[HOB]} \in C^-(c_i^{ij,[Het]})$; For label II individuals: if $l_i = II$, $c_j^{i,[Hom]} \in C^-(c_j^{i,[Het]})$, $c_i^{ij,[Het]} \in C^-(c_i^{ij,[HOB]})$.

IA.4: For the labeled treatments of Experiment 1, we assume that label II individuals (i) always regard label I 's as more sophisticated than they are, and (ii) they expect label I not to underestimate their sophistication. Formally: (i) $c_j^{i,[Het]} \in C^+(c_i)$ and (ii) $c_i^{ij,[Het]} \in C^+(c_i)$ whenever $l_i = II$.

Proposition 3 *For any distribution over subjects in Experiment 1 that satisfy assumptions IA.2-4, under the maintained assumptions of the EDR model (Section 4.1), the following holds.¹⁵*

1. (i) $F_{HOB}^I \succsim F_{Het}^I \succsim F_{Hom}^I$; (ii) $F_{Hom}^{II} \succsim F_{Het}^{II} \approx F_{HOB}^{II}$; (iii) $F_{HOB+}^I \succsim F_{Het+}^I \succsim F_{Hom+}^I$; and (iv) $F_{Hom+}^{II} \succsim F_{Het+}^{II} \approx F_{HOB+}^{II}$.

¹⁴Given two cumulative distributions $F(x)$ and $G(x)$, we say that F (weakly) first order stochastically dominates G , written $F \succsim G$, if $F(x) \leq G(x)$ for every x . $F \succ G$ if $F \succsim G$ and $F(x) < G(x)$ for some x

¹⁵If part (ii) of IA.4 were dropped, the only change to this Proposition is that parts the \approx relation in parts (ii) and (iv) of point 1 would be weakened to \succsim .

$$2. F_{AT-Het}^I \succsim F_{AT-Hom}^I \text{ and } F_{AT-Hom}^{II} \succsim F_{AT-Het}^{II}.$$

$$3. F_{AP-Het}^I \succsim F_{AP-Hom}^I \text{ and } F_{AP-Hom}^{II} \succsim F_{AP-Het}^{II}.$$

4. For label I , the increase in k_i from [HOB] to [AP-Het] should be at most one.

Intuitively, to understand the effects of changing beliefs in the EDR model, when incentives are symmetric ($x_i = x_j$, as in the baseline treatments in Section 3.1.1), an individual's cognitive bound is binding if he regards his opponent as 'more sophisticated' (i.e., lower cost-of reasoning). Hence, when the incentives to reason are symmetric, individuals with higher costs of reasoning have a lower cognitive bound, which therefore is binding when playing against someone they regard as more sophisticated.

The reason for the asymmetric effect of higher-order beliefs in point 1 is that, for label II subjects, if their cognitive bound is binding in treatment [Het] (respectively, [Het+]) – in which they play against someone they regard as more sophisticated – then it would also be binding in treatment [HOB] (resp., [HOB+]) – in which their opponent may play according to an even deeper behavioral level – and therefore behavior should be the same in these treatments. Label I subjects, instead, would understand that label II subjects play according to a higher behavioral level in the [Het] than in the [Hom] treatment, and hence their behavioral level in treatment [HOB] may be lower than in treatment [Het], which in turn is lower than in [Hom]. *Hence, higher-order beliefs effects (i.e., comparing treatments [Het] and [HOB]) are possible, but they are one-sided: they should be observed, if at all, only for label I subjects.* These are precisely some of AP's main findings, which we summarize in Section 5.1.1.

The reason for point 4 is that, under the assumption that label I subjects regard label II subjects as having a higher cost of reasoning (less sophisticated) than themselves, in the [HOB] treatment the label II cognitive bound can be at most the same as label I 's. If it is strictly less, then increasing label I 's incentives should not affect their behavior, because their cognitive bound was binding in the first place. If instead the cognitive bounds were the same, then label I 's behavior would change, but since the opponents' bound is the same in the two treatments, label I 's behavioral level would only increase by one level.

Based on the logic above, and by IA.3-4 label I subjects are commonly regarded as 'more sophisticated' than label II , *we expect more label II subjects with binding cognitive bounds in treatment [Het] than in treatment [Hom] (and in [Het+] than in [Hom+]), whereas the opposite would be true for label I subjects.* This implies the following proposition:

Proposition 4 *For any distribution over subjects in Experiment 1 that satisfy assumptions IA.1-4, under the maintained assumptions of the EDR model (Section 4.1), the following holds: For label II subjects (resp., label I) shift in behavior from [Het] to [AT-Het] is (weakly) larger (resp., smaller) than from [Hom] to [AT-Hom].*

5 Results

Before examining the results of the individual experiments, we compare our general results to the findings of other level- k papers. Georganas, Healy and Weber (2015) find that it is difficult to compare different types of level- k games as they find that levels of thinking can be uncorrelated or even flip across games. For this reason, we compare our results to Arad and Rubinstein (2012) who use the (cyclical) 11-20 game. They find that the number 20 is played by 6% of subjects, numbers 17, 18 and 19 are played by 74% of subjects and numbers from 11 to 16 are played by 20% of subjects. For our experiments, we find that 8% to 10% play 20, 50% to 60% play numbers 17 to 19 and 32% to 43% play 11 to 16 depending on the experiment. In light of the difference between the cyclical and our acyclical version of the game, these numbers appear comparable with those in Arad and Rubinstein (2012): the higher percentages of subjects who play numbers 11 to 16 are due to the fact that in our 11-20 game, there are no cycles and 11 is the level- ∞ (and level-9 and higher) strategy, while in the 11-20 game in Arad and Rubinstein (2012), 11 is not level- ∞ due to the cyclicity of the game. Our results for the percentages of levels played are also similar to those in Agranov, Potamites, Schotter and Tergiman (2012) where 8% to 10% play level 0 (our outcome 20) and 42% to 77% play levels 1 to 3 (our outcomes 19, 18 and 17). The following sections will present the experiment-specific results.

5.1 Experiment 1

In the following, we will combine the labels for the two classification criteria, and use the term ‘label I ’ to refer indiscriminately to the ‘math and sciences’ or to the ‘high score’ subjects, and the term ‘label II ’ to refer to the ‘humanities’ or ‘low score’ subjects.

5.1.1 Summary of AP’s main results

AP’s main findings on the baseline treatments can be summarized as follows:¹⁶

1. *Beliefs Effects:* For both the low and the high payoff treatments, under both classifications, the distribution of actions for label I subjects is *lower* in the homogeneous than in the heterogeneous treatments. The opposite is true for label II : the distribution of actions for label II subjects is *higher* in the homogeneous than in the heterogeneous treatments. Hence, these patterns are consistent with the assumption that both groups regard label I subjects as ‘more sophisticated’ than label II .
2. *Payoffs Effects:* For all configurations of beliefs, under both classifications, the distribution of actions in the ‘low payoffs’ treatments, $[X]$ – where $X = Hom, Het, HOB$ – first-order stochastically dominates the ‘high payoffs’ treatments, $[X+]$, for both label I and label II subjects. Hence, holding beliefs constant, the distribution of actions shifts

¹⁶Other than the content of the next three bullet points, which summarize the experimental findings in AP, all other experimental results in this paper are new.

towards higher level- k 's when payoffs increase, which is consistent with Proposition 2.1. See Table 8 in the Appendix for the regression results.

3. *Higher-Order Beliefs Effects:* Under both classifications, the distribution of actions for label I subjects is *lower* in the heterogeneous treatment [Het] than in the replacement treatment [HOB]. This suggests that label I subjects expect label II subjects to behave according to higher k 's when they interact with label I , than when they play among themselves, and that label I subjects react to this. For label II subjects instead the distribution of actions in the [Het] and [HOB] treatments are essentially the same. Hence, higher-order beliefs effects are present, but they are 'one-sided', consistent with Proposition 3.1.

Under the assumption, which in fact emerges from the data, that both groups regard label I subjects as 'more sophisticated', the predictions of the EDR model are exactly those observed in the experiment, including the one-sidedness of the higher-order beliefs effects (cf. Alaoui and Penta (2016a)). In the rest of this section and in the following one we discuss our novel experimental results.

5.1.2 Relaxing cognitive bounds – Experimental Results

In this subsection we discuss our findings for the post-tutorial treatments ([Tut], [AT-Hom] and [AT-Het]), which we administered to all eighty subjects from four of the six sessions (two for the endogenous and two for the exogenous classifications), each repeated twice.

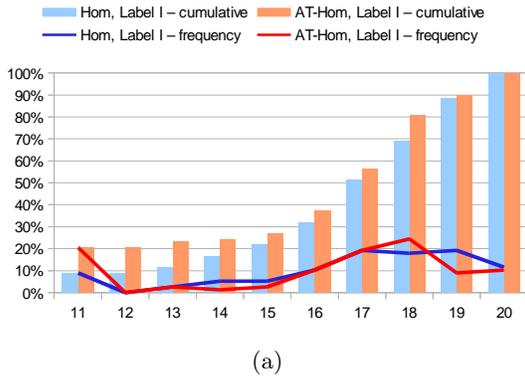
Unsurprisingly, a high fraction of the subjects in treatment [Tut] (48% of label I and 55% of label II) chose 11, although one could have expected that an even higher fraction would have made that choice.

The empirical analysis conducted throughout is as follows. We use panel regressions, clustered at the individual level for all the analyses. We also check all of these comparisons for robustness with a Wilcoxon signed-rank test, and they generally confirm the results of the regressions discussed in this paper. For brevity, we omit the Wilcoxon signed rank tests' p-values from the main text but provide them in Table 14 and discuss those of note.

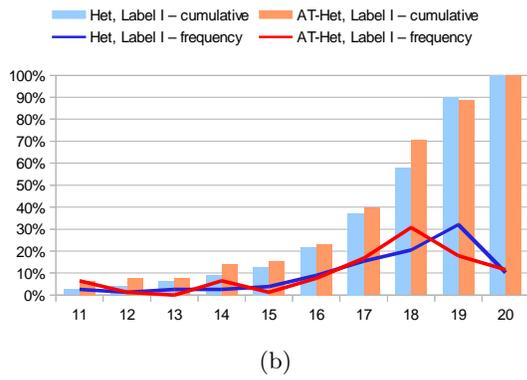
We first consider treatments pre- and post- tutorial. Comparing [Hom] to [AT-Hom], we observe that the distributions of actions shift to the left, with complete first order stochastic dominance (FOSD) of [Hom] to [AT-Hom] for both labels I and II (see Figure 1).¹⁷ These results are supported by the regressions performed. For each label, we regress the chosen action on a dummy which takes value 1 if the treatment is [AT-Hom] and 0 if it is [Hom] (see Table

¹⁷While this is not our focus, we also check whether labels responded differently to treatments by conducting panel regressions with a label I dummy, a treatment dummy and an interaction term. The results are provided in Table 11. The coefficient of the interaction term is significant at the 1% level for the comparisons of the treatments from [Het] to [AT-Het] and [AT-Hom] to [AT-Het], and not the others. Note also that mechanically, if comparing the coefficients of this regression with the regressions splitting the sample into label I and label II , the treatment effect on label II is of course identical to the one in the other regressions for label II . For label I , the sum of the treatment and interaction coefficients is identical to the regression results for label I in the other regressions.

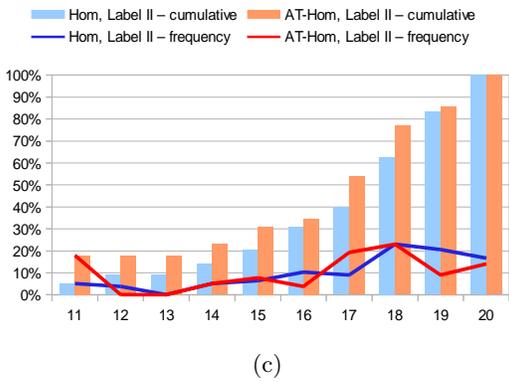
Pre- and post-tutorial comparisons: Homog. treatments, Label I



Pre- and post-tutorial comparisons: Heter. Treatments, Label I



Pre- and post-tutorial comparisons: Homog. treatments, Label II



Pre- and post-tutorial comparisons: Heter. treatments, Label II

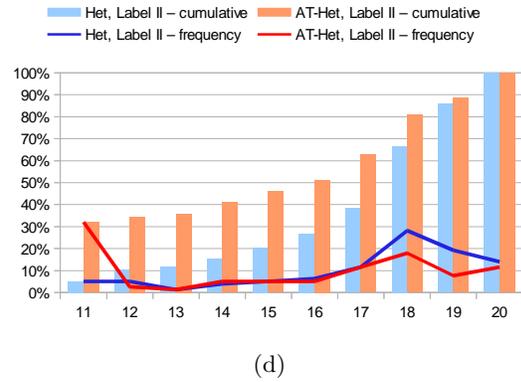


Figure 1: Pre- and post-tutorial comparisons, label *I* (top) and label *II* (bottom)

9). Consistent with the patterns observed in Figure 1, the estimated coefficient is negative (-0.71 for label I , and -0.85 for label II), statistically significant at the 5% level for label II and nearly significant at the 10% level for label I (the Wilcoxon signed-rank test is significant at the 10% level).¹⁸ This means that the tutorial, going from [Hom] to [AT-Hom], induces an average decrease in the number chosen by subjects equal to 0.71 and 0.85 for the two labels. Similarly, when comparing [Het] to [AT-Het], we find relatively weak FOSD everywhere of [Het] over [AT-Het] for label I except at 19, and lack of significance for the estimated coefficient. For label II , there is stronger FOSD everywhere. The estimated coefficient takes value -1.92 , and is statistically significant at the 1% level.

These results are all consistent with Proposition 2.3. The shift to lower numbers for label II going from [Hom] to [AT-Hom] and [Het] to [AT-Het] indicates that the capacity, or cognitive bound, is binding for at least some of the subjects of that label when playing [Hom] and/or [Het]. Recall that we can make this inference because the players face essentially the same opponent in [Hom] and [AT-Hom] (and [Het] to [AT-Het]), and so the tutorial may affect their behavior only if the previous understanding was binding.¹⁹ Similarly, the shift to lower numbers for label I when going from [Hom] to [AT-Hom] also indicates that the cognitive bound is binding in [Hom] when playing against their own type. Moreover, the lack of significance of the coefficient when comparing [Het] to [AT-Het] for label I does not allow us to conclude that their cognitive bound was binding in [Het]. This is also in line with our predictions, given our maintained assumption (cf. IA.3) that label II is commonly believed to be less sophisticated than label I .

Interestingly, Figures 1 and 2 show that players jump to 11 in the post-tutorial treatments much more frequently when their opponent is (weakly) more sophisticated.²⁰ From within the EDR model, this suggests that subjects believe the more sophisticated players to be very sophisticated, and capable of playing according to the highest level of reasoning. Their behavior against the less sophisticated opponent provides a further indication that it is beliefs that drive their behavior against them, and not their cognitive bound. It also suggests that the tutorial served its intended purpose of lowering the costs of cognition significantly. Within the context of the EDR model, it is a direct implication of identification assumption IA.1.²¹

We now analyze differences between homogeneous and heterogeneous post-tutorial treatments. Comparing [AT-Hom] to [AT-Het] for each label (see Figure 2), in the case of label I , there is pronounced FOSD of [AT-Het] over [AT-Hom] everywhere, and the estimated coefficient of the regression is 1.06 and statistically significant at the 1% level. In the case of label II ,

¹⁸In all the regressions that follow, we control for the grouping of the exogenous and endogenous treatments. This control is never statistically significant at the 5% level and only once at the 10% level, and it has a marginal impact on the relevant coefficient compared to when it is omitted. For this reason, we do not discuss it below.

¹⁹We note that the fact that subjects face the same population of opponents is a feature of the experiment. Its formal counterpart within the language of the EDR model is provided by identification assumption IA.2.

²⁰We thank both anonymous referees for bringing our attention to this point.

²¹A small caveat could be the wording of the tutorial as it mentions “game theory” and “rationality”. Subjects might not completely understand these terms. However, the reaction of subjects to the tutorial suggests that this was not an issue.

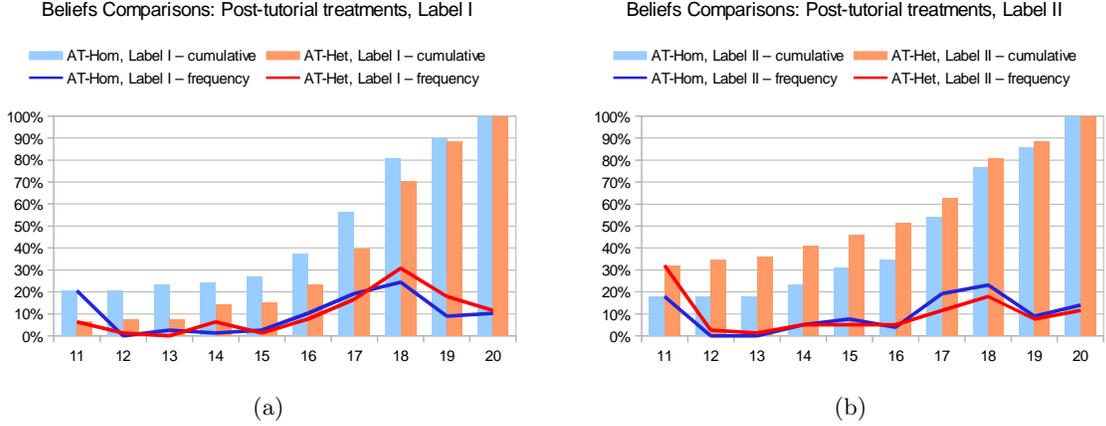


Figure 2: Beliefs comparisons: post tutorial treatments, label *I* (left) and label *II* (right)

instead, the effect is reversed: There is strong FOSD of [AT-Hom] over [AT-Het] everywhere, and the estimated coefficient is -1.14 , also statistically significant at the 1% level. Here as well, the direction of the results is fully consistent with our predictions (Proposition 3.2). Notice that the tutorial should not affect the direction of the comparative statics, because the main relevant factor is not the subjects' own understanding of the game, but rather their beliefs over their opponents' understanding.

Lastly, since we expect more label *II* subjects with binding cognitive bounds in treatment [Het] than in treatment [Hom] (and in [Het+] than in [Hom+]) – and the opposite for label *I* subjects – the model (Proposition 4) predicts that we should observe a (weakly) greater shift in behavior from [Het] to [AT-Het], than from [Hom] to [AT-Hom] for label *II* subjects – and the opposite for label *I*. This is consistent with the empirical findings in Figure 1, and also in line with the estimated OLS coefficients: the estimated coefficients for label *II* are -1.92 (significant at 1%) going from [Het] to [AT-Het] and -0.85 (significant at 5%) going from [Hom] to [AT-Hom]; for label *I* subjects instead, the coefficients are -0.7 (significant at (nearly) 10%) going from [Hom] to [AT-Hom], and -0.32 (not significant) going from [Het] to [AT-Het]. These comparisons are consistent with the model's predictions.

5.1.3 Reasoning about opponents' incentives – Experimental Results

We now discuss the empirical findings for the asymmetric payoff treatments [AP-Hom] and [AP-Het], which were administered after the baseline treatments to all forty subjects from two sessions (one exogenous and one endogenous), each treatment repeated three times. All of the results of the regressions discussed in this subsection are provided in Table 10.

Comparing [Hom], [AP-Hom] and [Hom+] for label *I*, there is a nearly complete FOSD relationship of [Hom] over [AP-Hom] (and of [Hom] over [Hom+]), but the relationship between [AP-Hom] and [Hom+] is less clearly defined (see Figure 3). This is consistent with the regressions, for which the estimated coefficients when going from [Hom] to [AP-Hom] is -0.74 , and

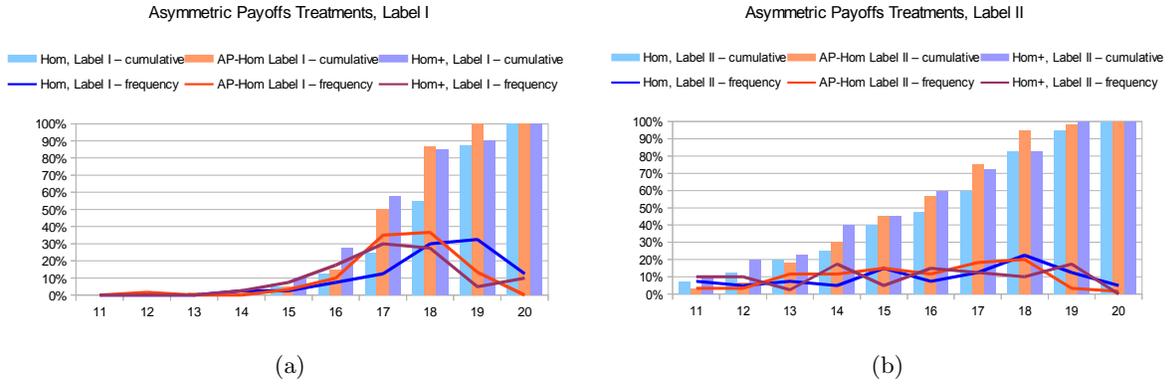


Figure 3: Asymmetric payoffs treatments, label *I* (left) and label *II* (right).

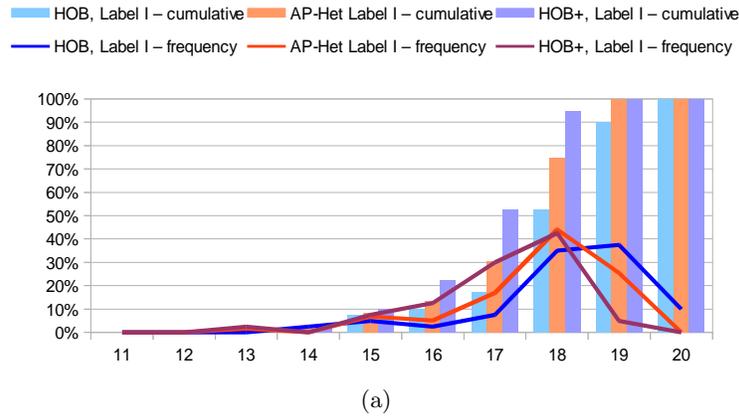
is statistically significant at the 1% level, while it is not significant when going from [AP-Hom] to [Hom+]. For label *II*, the comparisons of [Hom], [AP-Hom] and [Hom+] are ambiguous, and neither of the regressions comparing [Hom] to [AP-Hom] or [AP-Hom] to [Hom+] lead to significance. These results are jointly in line with Proposition 2.2(i) (recall that these predictions are in terms of weak orderings). Moreover, from within the EDR model they indicate that, for label *I*, increasing only the individual's own incentives, without changing either the opponents' incentives or their beliefs over their opponents, leads to them playing according to higher rounds of reasoning. In other words, the change in incentives appear to have led some label *I* subjects to increase their cognitive capacity. For label *II* subjects, instead, the increase in incentives has not had a noticeable impact.

When considering the [AP-Het] treatment, the natural comparison is not [Het], but rather [HOB]; see Proposition 2.2(ii). This is because the only difference between [HOB] and [AP-Het] is in the incentives of the subject, while their opponents are identical in the game they play. With [Het], instead, there is also a difference in the opponents, in that the opponents' opponent is of a different label. Comparing [HOB], [AP-Het] and [HOB+] for labels *I* and *II* separately, we see clear FOSD relationships nearly everywhere of [HOB] to [AP-Het] to [HOB+] for label *I* (see Figure 4). The relationship between the curves is instead more ambiguous for label *II*. The regressions for these comparisons lead to statistical significance at the 1% level for label *I*, but they are not significant for label *II*. These results are consistent with Proposition 2.2(ii).

Lastly we compare [AP-Hom] to [AP-Het] for both labels, and find a FOSD relationship of [AP-Het] to [AP-Hom] (see Figure 5). The coefficient estimated in the regressions, however, is not significant for label *II*, while it is significant at the 5% level for label *I*.

The model, and specifically Proposition 3.4, which predicts a shift from [HOB] to [AP-Het] of at most 1, also seems consistent with the small shift in distribution from [HOB] to [AP-Het], and with the estimates of the OLS coefficient (-0.5, significant at the 5% level). In this case, and consistently with the theory, the movement from [HOB] to [HOB+] for label *I* is mainly due to the increase in the opponents' payoffs, and not solely to the agent's own incentives. In light of the complexity of these treatments and the difficulty of the instructions, both discussed

Asymmetric Payoffs Treatments, Label I



Asymmetric Payoffs Treatments, Label II

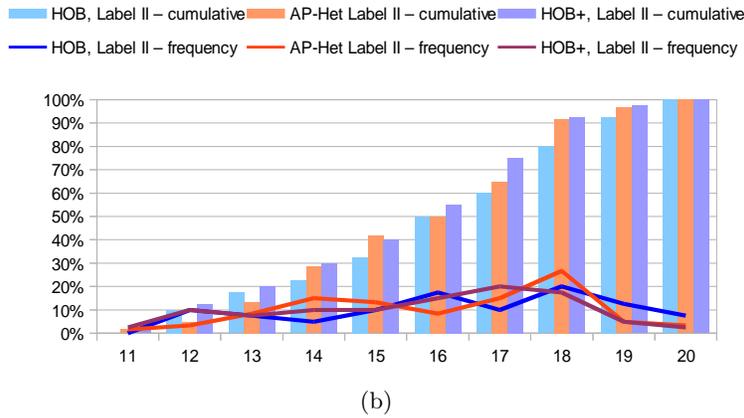
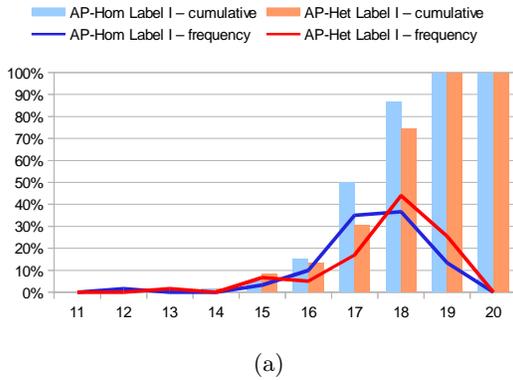


Figure 4: Asymmetric payoffs treatments, with double replacement. Label I (top) and label II (bottom).

Belief Comparisons: Asymmetric Payoffs treatments, Label I



Belief Comparisons: Asymmetric Payoffs treatments, Label II

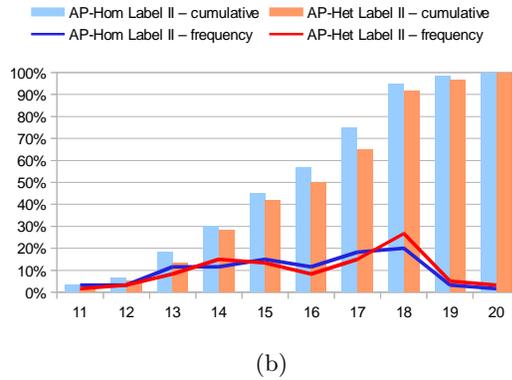


Figure 5: Beliefs comparisons: asymmetric payoffs treatments, label I (left) and label II (right).

in Section 3.1.3, these results are remarkably consistent with the predictions of the EDR model.

Robustness. As discussed in Section 3.1.3, we have conducted an additional robustness experiment that also contains treatments with replacement both for the [Hom-Rep] benchmark and the replacement treatment [AP-Hom]. It also includes the [Hom] treatment, for comparability. This serves to address the possible concern that the replacement method might remove social preferences. A Wilcoxon signed-rank test (selected due to the dependence across the treatments) could not reject the null that behavior in [Hom] and [Hom-Rep] was the same for the whole sample and the label *I* and label *II* subsamples. This suggests that there is no confound of social preferences in the [Hom] treatment compared to a replacement treatment and that we can therefore use the results of Experiment 1 for our analysis.

A second possible concern is that order effects might be the reason for increases in the levels played. We deal with this in three ways (in both the original Experiment 1 and the robustness experiment). First, the order of treatments was randomized across sessions so that order effects should not play a role. Second, all of our reported regression results (for all experiments) were estimated using panel regressions that took into account that the same treatment was asked multiple times. Third, we test for order effects using Wilcoxon signed-rank tests. We find no order effects for repeated treatments other than for [Hom+] in those experiments with tutorial treatments (p-value=0.055) and for [HOB] for those experiments with asymmetric payoffs treatments (p-value=0.064). For those treatments where no order effects were observed, p-values range from 0.15 to 1 with the majority above 0.3. This suggests that order effects are unlikely to drive the results.

5.2 Experiment 2: Results

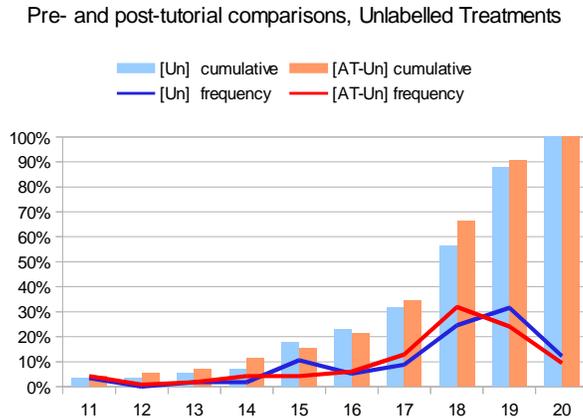


Figure 6: Pre- and post-tutorial, unlabeled treatments.

We first compare the results for treatments [Un] and [AT-Un]. Graphically, the cumulative distribution for [Un] stochastically dominates [AT-Un] everywhere except at 15 and 16 (Figure

6) but the difference between the distributions is slight. We find that the estimated coefficient in the regression is not significant but has the expected sign (see Table 12 for the regressions discussed here). The relationship between [Un] and [AT-Un] is consistent with Proposition 2.5, which predicts weak stochastic dominance. We also note that the Wilcoxon signed-rank test comparing the distributions is significant at the 5% level. Put together, we cannot conclude from these results to which extent subjects' behavior in the [Un] treatment is driven by their own cognitive capacity or by their beliefs about their opponents.

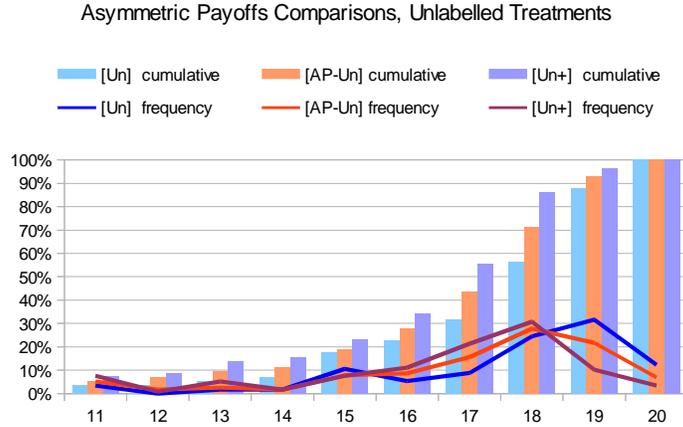


Figure 7: Asymmetric payoffs comparisons, unlabeled treatments.

Comparing the distributions of treatments [Un], [AP-Un] and [Un+], we observe that [Un] first-order stochastically dominates [AP-Un] everywhere, which itself stochastically dominates [Un+] everywhere (see Figure 7). Consistent with these results, the estimated coefficients are statistically significant for the regressions comparing [Un] to [AP-Un], [AP-Un] to [Un+] and [Un] to [Un+] at the 10% level (p-value 0.054), 5% level and 1% level, respectively. These findings indicate that subjects play according to lower sophistication in [Un] than [AP-Un] than [Un+]; this is consistent with Proposition 2.4.

The difference between [Un] and [AP-Un] is in the incentives, holding constant beliefs and higher-order beliefs over the distributions of opponents to change between the two. Hence, playing according to higher sophistication in [AP-Un] than in [Un] is an indication of the cognitive bound increasing. In the comparison between [AP-Un] and [Un+] instead, agents have the same incentives, and hence the difference between the two treatments is due to subjects' beliefs over the opponents. Specifically, since [AP-Un] and [Un+] differ in the opponents' incentives to reason, the fact that behavior is markedly different in these two treatments (the OLS coefficient is -0.55 , with p-value of 0.023) is a clear indication that subjects take into account their opponents' incentives to reason, when they form beliefs over their behavior.

Lastly, Wilcoxon signed-rank tests that examine whether instances of treatments can be differentiated from each other suggest that there are no order effects (p-values range from 0.35 to 0.88).

5.3 Experiment 3: Results

Recall that one of the main objectives of Experiment 3 is to test whether our results comparing pre- and post- tutorial treatments in the earlier experiments could have been driven by the tutorial inducing level- k reasoning rather than by reducing the cognitive costs. Since, in Experiment 3, all subjects were already primed to think according to level- k when playing the first treatment [E3-Un], changes observed between the pre- ([E3-Un]) and post-tutorial responses ([E3-AT-Un]) should not stem from priming but from changes in their cognitive levels.

We observe that [E3-Un] first order stochastically dominates [E3-AT-Un] everywhere (see Figure 8) and the regression coefficient is significant at 10% (see Table 13). This result is consistent with Proposition 2.5, and shows that because subjects played according to a higher level after receiving the tutorial the cognitive constraint in [E3-Un] must have been binding for some subjects. We also find that subjects played according to higher levels going from [E3-Un] to [E3-AP-Un] and from [E3-Un] to [E3-Un+] (significant at 5% and 1%, respectively). This suggests that the higher incentive in payoffs led to more rounds of reasoning (see Figure 9) and is consistent with Proposition 2.4. The results for Experiment 3 are all consistent with the predictions of our model, and with the results of the other experiments. It does not seem, therefore, that the potential level- k priming effect of the tutorial is driving those findings.

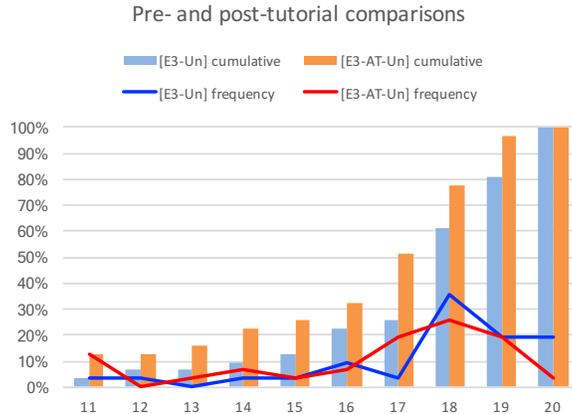


Figure 8: Pre- and post-tutorial, unlabeled treatments, Experiment 3.

The experimental results show that the way that subjects react to changes in the label of the opponent, and in changes of the label of their opponent’s opponent, is entirely consistent with the higher order beliefs effects generated by the EDR model (namely, that they are one-sided, and that they interact in complex ways with other changes in the environments, such as asymmetric changes in the payoff structure).

While our approach is as if, and we do not take a stance on the actual deliberation process of the agent, one could wonder whether subjects are *actually* capable of performing higher order reasoning (see Kets (2017) and Heifetz and Kets (2018) for models which formalize

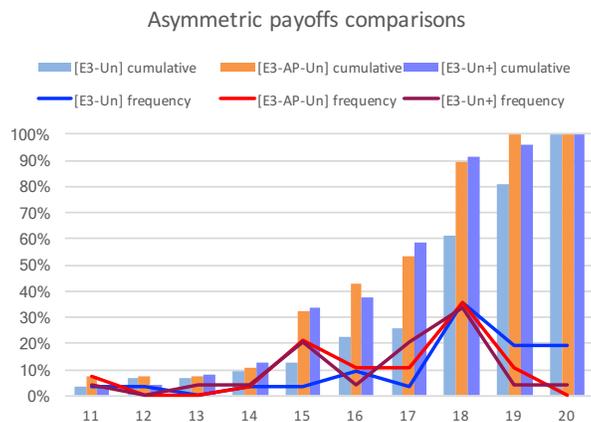


Figure 9: Asymmetric payoffs comparisons, unlabeled treatments, Experiment 3.

the idea that players need not have well-formed higher order beliefs). To gain a more direct answer to this question, we administered a short version of the Theory of Mind test (TOM hereafter; see Stiller & Dunbar 2007, Liddle & Nettle 2006) at the end of the experiment. This test is aimed specifically at testing whether a subject can place themselves in the mind of another person. Subjects are given multiple short stories to read about the interaction of fictional persons with other people. In these stories, the main character is thinking about the motivation behind statements or actions of others. After each story, subjects are asked to complete a series of questions about the story. To rule out that bad performance in the test is based on lack of attention or memory, the test contains a factual part aimed at testing how well subjects remember the story. In the mind part, subjects have to answer questions about others' reasoning process (and others' reasoning process about others' reasoning up to several levels). Bad performance in both parts suggests that the subject did not remember the story correctly while good performance in the factual part coupled with bad performance in the mind part suggests that the subject is able to remember the story but unable to place themselves in the mind of someone else. Test results show that more than 70% of subjects answered more than 50% of the TOM questions correctly (see the distribution of TOM scores in Figure 10). The results are not driven by the results of the factual part which can be seen in Figure 11 in the Appendix. Most of the TOM questions require higher order beliefs of multiple levels. Only one subject was unable to correctly answer any question that requires one level of thinking. This suggests that the majority of test subjects is capable of reasoning about others' reasoning.

In addition to the TOM, subjects had to complete a test of reasoning before starting the experiment. This test contained the Muddy Faces (also known as Dirty Faces) game which examines the 'level of iterated rationality' (Weber 2001). On average, subjects scored nearly 85% in this test, suggesting that they are capable of reasoning iteratively which is a prerequisite of level- k reasoning. Together, these results suggest that the experimental subjects are indeed capable of reasoning about others' reasoning and that we can thus use the results from the replacement method.

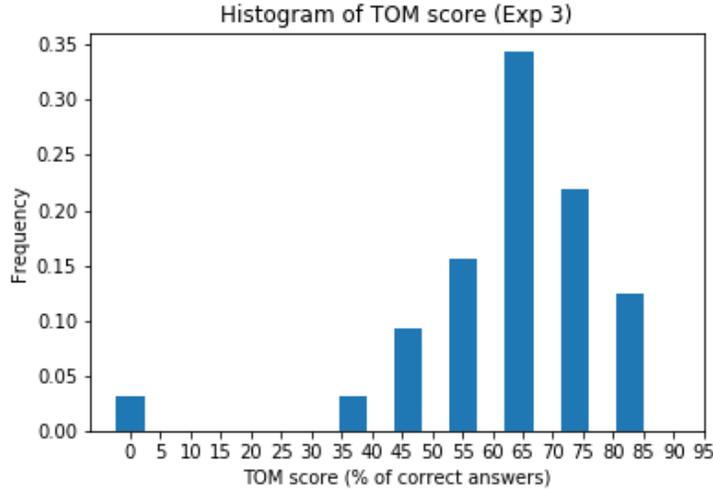


Figure 10: Distribution of total TOM scores.

To test for the potentiality that test results are driven by order effects, we conducted Wilcoxon signed-rank tests that examine whether instances of treatments can be differentiated from each other. The resulting p-values suggest that there is only one order effect present in the experiment; [E3-AT-Un] with a p-value equal to 0.057 (p-values for the other treatments range from 0.41 to 0.49). The results are thus unlikely to have been driven by order effects and, as in the other experiments, we use panel regressions to take into account that treatments were administered multiple times.

5.4 Individual Effects

The experimental results sections above analyze behavior at the aggregate level. One possible concern with this approach is that behavior within an individual might not be consistent and that this might be averaged out. In this section, we examine individual-level results which can shed light on this concern. We find that most subjects are highly consistent in their behavior across treatments.

To analyze individual level behavior across treatments, we created a score of violations against the theory. For example, the theory predicts that the level played in [Hom] is weakly lower than that in [Hom+]. A violation would then be if the subject played a strictly higher level in [Hom] (similarly for [Un] and [Un+] in experiments 2 and 3). For all individuals across all experiments, we count the number of violations across all comparisons of treatments where the theory makes a clear prediction. Figures in Section D.2 of the Appendix show that for each experiment, the number of violations is relatively low. For example, for Experiment 2, more than 60% of subjects had zero violations of the theory and another 15% had only one violation

of the theory in the averaged treatment violations score. These results suggest that subjects are highly consistent in their behavior. The results for each experiment individually can be seen in the Appendix.²²

A second potential issue with looking at aggregate results is that some individuals might shift their behavior a lot, driving the observed average effects. In order to examine this possibility, we calculated the number of levels that each subject shifts for all comparisons in the regressions. We then plot the distribution of these shifts in levels. In the figures in section D.3 of the appendix, it can be seen that most of the subjects shift levels moderately and for most comparisons a very low percentage of subjects shifts by many levels.²³ Due to their very low frequency, it is unlikely that the results are driven by these individuals.

These figures also show that the shift in levels of reasoning is larger for the tutorial than the payoff treatments. This suggests that the tutorial is successful at removing the cognitive constraint so that levels of play are determined by beliefs only.

Finally, these figures allow us to examine by how many levels label I subjects shifted from [HOB] to [AP-Het]. Figure 17 shows that 80% of label I subjects shifted by one level or less which supports Proposition 3.4.

6 Concluding Remarks

Individuals may reason about others' strategic reasoning in a way that is more nuanced than has been typically considered by the current literature. In particular, it is not clear whether subjects' choices are constrained by their own cognitive limitations or rather by their beliefs about their opponents' limitations. It is also unclear whether subjects would take into account how changing the opponents' stakes may change their depth of reasoning, and hence their behavior. In this paper we provided several experiments designed to address these questions.

Our findings indicate that subjects do indeed reason about their opponents' reasoning process, and that they form beliefs not only about their opponents' sophistication, but also account for the change of this sophistication with the incentives to reason. We also find that, while beliefs play a clear role in subjects' behavior, the cognitive bounds of a significant fraction of subjects are binding and determine their behavior when facing opponents they view as more sophisticated. These results suggest that in general, level- k behavior should not be taken as driven either by cognitive limits alone or beliefs alone. In some settings it is a function of both, and depends on the complex interaction between cognitive bounds, beliefs about the opponent's cognitive abilities, and reasoning about others' reasoning. We also find that the EDR framework of Alaoui and Penta (2016a) is a useful tool for analyzing and understanding this interaction, and that the results are overall consistent with its predictions.

²²Violations can also be calculated based on instances of each treatment i.e. comparing the level played in the first instance of [Hom] to the first instance of [Hom+] and the second instance of [Hom] to the second instance of [Hom+] and so forth. Figures for these violation scores are available upon request.

²³Due to the large number of figures, we only show the results for Experiment 3 (and for [HOB] to [AP-Het] for label I for Experiment 3). All other figures are available upon request.

While we have focused our analysis on a specific game and setting, our experimental design can be applied to a number of other contexts. It is based on two key ideas, which we have referred to as the *replacement* and the *tutorial* methods. These methods are to a large extent independent of the features of the underlying game, and are thus portable and easily implemented to other settings as well.

The *replacement method* consists of replacing the opponent’s opponent and controls for higher-order beliefs effects. Controlling subjects’ beliefs hierarchies is a well-known difficulty in designing game theoretic experiments, particularly if they are aimed at isolating the effects of beliefs manipulations or at identifying subjects’ higher order reasoning.²⁴ The precise wording of the treatments based on the replacement method is designed to pin down the entire hierarchy of beliefs, thereby addressing this challenge.²⁵

The replacement method was used in several of our treatments. The basic idea, however, has much broader applicability. It can be applied to essentially any setting (level- k or not) to disentangle *direct* and *interactive* effects involved in general comparative statics exercises. For instance, it can be used to separate subjects’ own preferences to adhere to a social norm from their beliefs about the consequences faced when deviating from it. The method can also be further adapted to more complex settings, such as the interesting experiment by Sofianos, Proto and Rustichini (2019), to assess for instance to what extent the higher levels of cooperation observed in the high-IQ group are due to individuals’ own cognitive abilities, or to the high-IQ environment.

The *tutorial method* instead was designed to address the cognition-beliefs dichotomy. It consists in studying how subjects’ behavior is affected by receiving a tutorial which contains non-factual information about the strategic structure of the game. As long as subjects face the same opponents before and after having received this tutorial (an idea related to the replacement method discussed above), then their behavior should change if and only if the tutorial has made them understand something they deem useful. This has enabled us to assess whether subjects’ understanding was somehow binding, relative to the information provided by the tutorial, or whether their action was rather driven by their beliefs, whatever understanding of the game they had before or after the tutorial. This idea can be applied independently of whether the underlying reasoning takes the form of level- k thinking. For instance, one could imagine a game with multiple equilibria (e.g., stag hunt, or other coordination games in which

²⁴We stress that the importance of the replacement method is to disentangle various higher order beliefs effects by progressively changing the orders of beliefs one-by-one, keeping all higher order beliefs constant from one design to the next. In this sense, while the experiment in Agranov et al. (2012) contains treatments analogous to our [Hom] and [HOB] treatments, it still cannot be considered as a full implementation of the method we are describing, since comparing [Hom] and [HOB] directly does not enable us to disentangle first-order effects from higher-orders. We also note, as we discuss next, that the replacement method can be applied to beliefs as well as payoffs, as we did in the AP-treatments.

²⁵In a recent paper, Kneeland (2015) develops the idea of ‘ring games’, which has a similar objective. One important difference between ring games and the replacement method, is that by having player 1 have a dominant strategy, Kneeland’s ring games allow for an exact identification of different orders of beliefs in rationality, which our replacement method does not. In contrast, the advantage of replacement treatments is that it allows to investigate higher order beliefs effects (albeit partially identified) in arbitrary games, without altering the underlying payoff structure, in a way that is easy to implement in the lab.

subjects may resort to different, non level- k , reasoning processes), and have the tutorial simply explain the properties of the different equilibria (such as risk-dominance and efficiency in stag-hunt). The exercise would go through in that case essentially unchanged, to assess the extent to which subjects' understanding before the tutorial was binding or not.

Similar to the replacement method, the tutorial method is portable to other settings, and is especially suited to understanding the cognition-beliefs dichotomy for different forms of reasoning. It need not only apply to level- k reasoning or to games of initial response. Future research can therefore make use of both these methods in settings that are very different from the one focused on in this paper.

Appendix

A Proofs

Proof of Proposition 1. For point 1, suppose that $c_j^i, c_i^{ij} \in C^+(c_i)$:

- Let v_i denote i 's value of reasoning in the low payoff game, and notice that assumptions (i)-(v) in p.15 implies that $v_i = v_j^i = v_{i(j)}$ in that game. Hence, if $c_j^i, c_i^{ij} \in C^+(c_i)$, we have $\mathcal{K}(c_j^i, v_j) \geq \mathcal{K}(c_i, v_i) = \hat{k}_i$ and $\mathcal{K}(c_i^{ij}, v_{i(j)}) \geq \mathcal{K}(c_i, v_i) = \hat{k}_i$. The former inequality implies (by eq. (3)) that $\hat{k}_j^i = \hat{k}_i - 1$, which together with the latter inequality and (4) implies $\hat{k}_i^{ij} = \hat{k}_i - 2$. By (5), it follows that $k_j^i = \hat{k}_i - 1$, and hence $k_i = k_j^i + 1 = \hat{k}_i$.
- Let v'_i denote i 's value of reasoning for the high x_i in the asymmetric payoff game. By assumptions (i)-(v) in p.15, v_j^i and $v_{i(j)}$ remain the same as in the low payoff case, whereas $v'_i(k) \geq v_i(k)$ for every k . Letting $\hat{k}'_i := \mathcal{K}(c_i, v'_i)$ and k'_i denote, respectively, i 's cognitive bound and behavioral level in the asymmetric payoff game, it follows that $\hat{k}'_i \geq \hat{k}_i$. Equations (3)-(5) then immediately imply that \hat{k}_i^{ij} and k_j^i weakly increase from the low-payoff to the asymmetric payoff game, and hence $k'_i \geq k_i$.
- Let \tilde{v}_i, \tilde{v}_j and $\tilde{v}_{j(i)}$ denote the value of reasoning in the high payoff game for i, j and j 's opponent, respectively. By assumptions (i)-(v) in p.15, $\tilde{v}_i = v'_i$ and $\tilde{v}_i = \tilde{v}_j^i = \tilde{v}_{i(j)}$. Hence, if $c_j^i, c_i^{ij} \in C^+(c_i)$, we have $\mathcal{K}(c_j^i, \tilde{v}_j) \geq \mathcal{K}(c_i, \tilde{v}_i) = \mathcal{K}(c_i, v'_i) = \hat{k}'_i$ and $\mathcal{K}(c_i^{ij}, \tilde{v}_{i(j)}) \geq \mathcal{K}(c_i, \tilde{v}_i) = \mathcal{K}(c_i, v'_i) = \hat{k}'_i$. Letting k''_i and \hat{k}''_i denote i 's behavioral level and cognitive bound in the high payoff game, the same arguments as in the low payoff game at this point imply that $k''_i = \hat{k}''_i = \hat{k}'_i$, which in turn implies that $k''_i > k'_i$ only if $k'_i < \hat{k}'_i$.

For point 2, first consider the case $c_j^i \in C^-(c_i)$: Let v_i denote i 's value of reasoning in the low payoff game, and notice that assumption (iii)-(v) in p.15 implies that $v_i = v_j^i = v_{i(j)}$ in that game. Hence, if $c_j^i \in C^-(c_i)$, we have $\mathcal{K}(c_j^i, v_j) \leq \mathcal{K}(c_i, v_i) = \hat{k}_i$.

- First note that, by eq. (3)-(5) and the definition of k_i , it is always the case that $k_i \leq \hat{k}_i$. The inequality can be strict only if $\mathcal{K}(c_j^i, v_j) < \hat{k}_i - 1$, which is possible (though not necessary) if $c_j^i \in C^-(c_i)$.
- Let v'_i denote i 's value of reasoning for the high x_i in the asymmetric payoff game. By assumptions (i)-(v) in p.15, v_j^i and $v_{i(j)}$ remain the same as in the low payoff case, whereas $v'_i(k) \geq v_i(k)$ for every k . Letting $\hat{k}'_i := \mathcal{K}(c_i, v'_i)$ and k'_i denote, respectively, i 's cognitive bound and behavioral

level in the asymmetric payoff game, it follows that $\hat{k}'_i \geq \hat{k}_i$. Since in the low payoff game we already had $\hat{k}_j^i \leq \hat{k}_i - 1$, and hence (by 4) $\hat{k}_i^{ij} \leq \hat{k}_i - 2$, it follows that the (weak) increase in i 's cognitive bound does not affect either \hat{k}_j^i or \hat{k}_i^{ij} . If in the low payoff game it was the case that $\hat{k}_j^i < \hat{k}_i - 1$ and $\hat{k}_i^{ij} < \hat{k}_i - 2$, and hence $k_j^i < \hat{k}_i - 1$, and hence $k_i = k_j^i + 1 < \hat{k}_i$, then we would have $k'_i = k_i$. It follows that if $k'_i > k_i$ (which could be the case if $\hat{k}_i^{ij} = \hat{k}_i - 2$, and $k_j^i = \hat{k}_i - 1$), then it must have been the case that k_i and \hat{k}_i were equal in the low payoff game.

- The argument for the high payoff game is the same as the corresponding one for point 1 above.

Now consider the case, also in point 2, in which $c_j^i \in C^+(c_i)$ and $c_i^{ij} \in C^-(c_i)$:

- Again, note that, by eq. (3)-(5) and the definition of k_i , it is always the case that $k_i \leq \hat{k}_i$. Now, let v_i denote i 's value of reasoning in the low payoff game, and notice that assumption (iii)-(v) in p.15 implies that $v_i = v_j^i = v_{i(j)}$ in that game. Hence, if $c_j^i \in C^+(c_i)$, we have $\mathcal{K}(c_j^i, v_j) \geq \mathcal{K}(c_i, v_i) = \hat{k}_i$. The former inequality implies (by (3)) that $\hat{k}_j^i = \hat{k}_i - 1$. Given this, $k_i < \hat{k}_i$ is possible only if $\mathcal{K}(c_i^{ij}, v_{i(j)}) < \hat{k}_i - 1$, which is possible (though not necessary) if $c_i^{ij} \in C^-(c_i)$.

Given this, the arguments for the asymmetric and high payoff games are the same as those we considered above for the other case of point 2.

For point 3, let v_i, v'_i and \tilde{v}_i denote, respectively, i 's value of reasoning in the low, asymmetric, and high payoff games. As in point 1, the maintained assumptions imply that $v'_i = \tilde{v}_i$ and $v'_i(k) \geq v_i(k)$ for all k . Also, for any $w_i \in \{v_i, v'_i\}$ $\hat{k}'_i := \mathcal{K}(c'_i, w_i) \geq \mathcal{K}(c_i, w_i) =: \hat{k}_i$ if $c'_i \in C^+(c_i)$ – that is, for any payoff configuration, i 's cognitive bound always weakly increases as cost c_i is replaced with a lower cost c'_i . Since equations (by 3)-(5) are all (weakly) increasing in i 's cognitive bound, and i 's behavioral level is increasing in k_j^i , it follows that also the behavioral level (weakly) increases if c_i is replaced with a lower cost c'_i . ■

Proof of Proposition 2. All results follow directly from Proposition 1 and IA1-2, noting that a higher behavioral level translates to lower numbers chosen in the acyclic 11-20 game:

1. Points 1 and 2 of Proposition 1 imply that, for any c_i, c_j^i and c_i^{ij} – which by IA-2 remain constant between X and $X+$ for any $X \in \{\text{Hom}, \text{Het}, \text{HOB}\}$ – the behavioral level weakly increases for any i from the low payoff to the high payoff game. Hence, for each $l \in \{I, II\}$: $F_X^l \succsim F_{X+}^l$ for all $X \in \{\text{Hom}, \text{Het}, \text{HOB}\}$.
2. Points 1 and 2 of Proposition 1 imply that, for any c_i, c_j^i and c_i^{ij} – which, by IA-2, remain constant between [Hom], [AP-Hom] and [Hom+], and between [HOB], [AP-Het] and [HOB+] – the behavioral level weakly increases for any i from the low payoff to the asymmetric payoff game, and from the asymmetric to the high payoff game. The former increase is strict only if $\hat{k}_i = k_i$ in the low payoff game for some i . Hence, for each $l \in \{I, II\}$: (i) $F_{Hom}^l \succsim F_{AP-Hom}^l \succsim F_{Hom+}^l$, with $F_{Hom}^l \succ F_{AP-Hom}^l$ only if \hat{k}_i was binding in [Hom] for some i ; and (ii) $F_{HOB}^l \succsim F_{AP-Het}^l \succsim F_{HOB+}^l$, with $F_{HOB}^l \succ F_{AP-Het}^l$ only if \hat{k}_i was binding in [HOB] for some i .
3. Point 3 of Proposition 1 implies that, for any c_j^i and c_i^{ij} – which by IA-2 remain constant between X and [AT- X] for any $X \in \{\text{Hom}, \text{Het}\}$ – the behavioral level weakly increases for any i if their costs are lowered from c_i to some $c'_i \in C^+(c_i)$, a condition which is satisfied in the post-tutorial treatment under assumptions IA.1, and strictly so only if $\hat{k}_i = k_i$ in the first place. It follows that for each $l \in \{I, II\}$: $F_{Hom}^l \succsim F_{AT-Hom}^l$, and $F_{Het}^l \succsim F_{AT-Het}^l$, each strictly only if \hat{k}_i was binding for some i in [Hom] and [Het], respectively.

4. Points 1 and 2 of Proposition 1 imply that, for any c_i, c_j^i and c_i^{ij} – which, by IA-2, remain constant between [Un], [AP-Un] and [Un+] in both experiments 2 and 3 – the behavioral level weakly increases for any i from the low payoff to the asymmetric payoff game, and from the asymmetric to the high payoff game. The former increase is strict only if $\hat{k}_i = k_i$ in the low payoff game for some i . Hence: (i) $F_{Un}^* \succsim F_{AP-Un}^* \succsim F_{Un+}^*$, with $F_{Un}^* \succ F_{AP-Un}^*$ only if \hat{k}_i was binding in [Un] for some i ; and (ii) $F_{E3-Un}^* \succsim F_{E3-AP-Un}^* \succsim F_{E3-Un+}^*$, with $F_{Un}^* \succ F_{E3-AP-Un}^*$ only if \hat{k}_i was binding in [E3-Un] for some i .
5. Point 3 of Proposition 1 implies that, for any c_j^i and c_i^{ij} – which by IA-2 remain constant between Un and AT-Un in both experiments 2 and 3 – the behavioral level weakly increases for any i if their costs are lowered from c_i to some $c_i' \in C^+(c_i)$, a condition which is satisfied in the post-tutorial treatment under assumptions IA.1, and strictly so only if $\hat{k}_i = k_i$ in the first place. It follows that $F_{Un}^* \succ F_{AT-Un}^*$ and $F_{E3-Un}^* \succ F_{E3-AT-Un}^*$, each strictly only if \hat{k}_i was binding for some i in [Un] and [E3-Un], respectively.

■

Proof of Proposition 3. Under the maintained assumptions (i)-(v) of the EDR model, plus identification assumptions IA.1-4, the cost of reasoning c_i and the value of reasoning v_i, v_j^i and $v_{j(i)}^i$ remain constant across all treatments within the same point of the proposition. Hence, also \hat{k}_i does not change within each point of the proposition. The only things that change, within each point, are thus beliefs c_j^i and higher order beliefs c_i^{ij} . For each point we describe what these changes are and how they impact behavior across treatments:

1. IA.2 implies that $c_j^{i,[Het]} = c_j^{i,[HOB]}$ and $c_i^{ij,[Het]} = c_i^{ij,[Hom]}$ for any i of any label.
 - (a) For parts (i) and (iii), IA.3 implies that for any i with label I , we have $c_j^{i,[Het]} \in C^-(c_j^{i,[Hom]})$, $c_i^{ij,[HOB]} \in C^-(c_i^{ij,[Het]})$. For each $X \in \{\text{Hom}, \text{Het}, \text{HOB}\}$, let $\hat{k}_i^{ij,[X]}$, $\hat{k}_j^{i,[X]}$ and $k_j^{i,[X]}$ denote, respectively, the values taken by equations (4)-(5) when $c_j^i = c_j^{i,[X]}$ and $c_i^{ij} = c_i^{ij,[X]}$, and $k_i^{[X]}$ denote i 's corresponding behavioral level. first compare treatment [Hom] and [Het]: since $c_i^{ij,[Het]} = c_i^{ij,[Hom]}$ and $c_j^{i,[Het]} \in C^-(c_j^{i,[Hom]})$, it follows that $\hat{k}_j^{i,[Het]} \leq \hat{k}_j^{i,[Hom]}$ and $\hat{k}_i^{ij,[Het]} \leq \hat{k}_i^{ij,[Hom]}$, which in turn implies $k_j^{i,[Het]} \leq k_j^{i,[Hom]}$ and hence $k_i^{[Het]} \leq k_i^{[Hom]}$; then compare [Het] and [HOB]: since $c_j^{i,[Het]} = c_j^{i,[HOB]}$, $\hat{k}_j^{i,[Het]} = \hat{k}_j^{i,[HOB]}$, but $c_i^{ij,[HOB]} \in C^-(c_i^{ij,[Het]})$ implies $\hat{k}_i^{ij,[HOB]} \leq \hat{k}_i^{ij,[Het]}$, which in turn implies $k_j^{i,[HOB]} \leq k_j^{i,[Het]}$ and hence $k_i^{[HOB]} \leq k_i^{[Het]}$ part (i) follows. The same argument also applies to $X+$, which implies part (iii).
 - (b) For parts (ii) and (iv), note that IA.3 implies, for any i with label II , we have $c_j^{i,[Hom]} \in C^-(c_j^{i,[Het]})$, $c_i^{ij,[Het]} \in C^-(c_i^{ij,[HOB]})$, and by IA.4 we have $c_j^{i,[Het]} \in C^+(c_i)$. Maintaining the same notation as above, first compare treatment [Hom] and [Het]: since $c_i^{ij,[Het]} = c_i^{ij,[Hom]}$ and $c_j^{i,[Het]} \in C^-(c_j^{i,[Hom]})$, it follows that $\hat{k}_j^{i,[Het]} \geq \hat{k}_j^{i,[Hom]}$ and $\hat{k}_i^{ij,[Het]} \geq \hat{k}_i^{ij,[Hom]}$, which in turn implies $k_j^{i,[Het]} \geq k_j^{i,[Hom]}$ and hence $k_i^{[Het]} \geq k_i^{[Hom]}$; under part (ii) of IA.4, we also have $c_j^{i,[Het]} \in C^+(c_i)$, and hence $k_i^{[Het]} = \hat{k}_i^{[Het]}$. Next, compare [Het] and [HOB]: since $c_j^{i,[Het]} = c_j^{i,[HOB]}$, $\hat{k}_j^{i,[Het]} = \hat{k}_j^{i,[HOB]}$, but $c_i^{ij,[HOB]} \in C^-(c_i^{ij,[Het]})$ implies $\hat{k}_i^{ij,[HOB]} \leq \hat{k}_i^{ij,[Het]}$, which in turn implies $k_j^{i,[HOB]} \leq k_j^{i,[Het]}$ and hence $k_i^{[HOB]} \leq k_i^{[Het]}$. Since, by IA.3, $c_i^{ij,[HOB]} \in C^+(c_i^{ij,[Het]})$ and $c_j^{i,[Het]} \in C^+(c_i)$, we also have $k_i^{[HOB]} = \hat{k}_i^{[HOB]} = \hat{k}_i^{[Het]}$, and hence also $k_i^{[HOB]} = k_i^{[Het]}$ if part (ii) of IA.4 also holds. Part (ii) follows. The same argument also applies to $X+$, which implies part (iv).

2. IA.2 implies that $c_i^{ij,[AT-Het]} = c_i^{ij,[AT-Hom]}$ for any i of any label. IA.3 implies that for any i with label I , we have $c_j^{i,[AT-Het]} \in C^-(c_j^{i,[AT-Hom]})$. Maintaining the same notation as above, since $c_i^{ij,[AT-Het]} = c_i^{ij,[AT-Hom]}$ and $c_j^{i,[AT-Het]} \in C^-(c_j^{i,[AT-Hom]})$, it follows that $\hat{k}_j^{i,[AT-Het]} \leq \hat{k}_j^{i,[AT-Hom]}$ and $\hat{k}_i^{ij,[AT-Het]} \leq \hat{k}_i^{ij,[AT-Hom]}$, which in turn implies $k_j^{i,[AT-Het]} \leq k_j^{i,[AT-Hom]}$ and hence $k_i^{[AT-Het]} \leq k_i^{[AT-Hom]}$. It follows that $F_{AT-Het}^I \succsim F_{AT-Hom}^I$. The argument for $F_{AT-Hom}^{II} \succsim F_{AT-Het}^{II}$ is symmetric, swapping the inequalities, since IA.3 implies that for any i with label II , $c_j^{i,[AT-Hom]} \in C^-(c_j^{i,[AT-Het]})$, $c_i^{ij,[AT-Het]} \in C^-(c_i^{ij,[AT-HOB]})$.
3. The result follows from the same argument as in the previous point, just replacing AT-X with AP-X, for $X \in \{\text{Hom}, \text{Het}\}$.

■

Proof of Proposition 4. The result follows directly from the argument in the main text. ■

B Logistics of the Experiments

The experiments were conducted at the Laboratori d'Economia Experimental (LEEX) at Universitat Pompeu Fabra (UPF), Barcelona. Subjects were students of UPF, recruited using the LEEX system. No subject took part in more than one session. Subjects for the first experiment were paid 3 euros for showing up (students coming from a campus that was farther away received 4 euros instead). Subjects' earnings averaged 15.8. Subjects had a showup fee of 4 euros in the second experiment, and their earnings averaged 18 euros. The payments of subjects in Experiment 3 averaged at 14 euros.

Each subject in the first experiment went through a sequence of 18 games. Payoffs are expressed in 'tokens', each worth 15 cents. Subjects were paid randomly, once every six iterations. The order of treatments is randomized (see below). Subjects in the second and third experiment each went through a sequence of 9 games, and were paid randomly based on three iterations. In those, to compensate for there being fewer games from which the payments were drawn, 8 tokens were worth 1 euro. For all experiments, subjects only observed their own overall earnings at the end, and received no information concerning their opponents' results. The setup was identical for Experiment 3.

Our subjects for the first experiment were divided in 6 sessions of 20 subjects, for a total of 120 subjects. Three sessions were based on the exogenous classification, and each contained 10 students from the field of humanities (humanities, human resources, and translation), and 10 from math and sciences (math, computer science, electrical engineering, biology and economics). Three sessions were based on the endogenous classification, and students were labeled based on their performance on a test of our design (see Alaoui and Penta (2016a)). In these sessions, half of the students were labeled as 'high' and half as 'low'.

There were 60 subjects for the second experiment and 34 for the third experiment. The subjects all took the endogenous classification test first but they were not given any feedback, and remained unlabeled.

B.1 Instructions of the Experiment

We describe in B.1.1 to B.1.4 the instructions as worded for a student from math and sciences in the first experiment. The instructions for students from humanities would be obtained replacing these labels everywhere. Similarly, labels high and low would be used for the endogenous classification. The related instructions for the second experiment are in B.1.5.

B.1.1 Baseline Game and Treatments [Hom], [Het] and [HOB]

Pick a number between 11 and 20. You will always receive the amount that you announce, in tokens.

In addition:

- if you give the same number as your opponent, you receive an extra 10 tokens.
- if you give a number that's exactly one less than your opponent, you receive an extra 20 tokens.

Example:

- If you say 17 and your opponent says 19, then you receive 17 and he receives 19.
- If you say 12 and your opponent says 13, then you receive 32 and he receives 13.
- If you say 16 and you opponent says 16, then you receive 26 and he receives 26.

Treatments [Hom] and [Het]:

Your opponent is:

- a student from maths and sciences (treatment [Hom]) / humanities (treatment [Het])
- he is given the same rules as you.

Treatment [HOB]:

In this case, the number you play against is chosen by:

- a student from humanities facing another student from humanities. In other words, two students from humanities play against each other. You play against the number that one of them has picked.

Treatment [Hom-Rep]

In this case, the number you play against is chosen by:

- a student from maths and sciences facing a student from maths and sciences. In other words, two students from maths and sciences play against each other. You play against the number that one of them has picked.

B.1.2 Changing Payoffs: Treatments [HOM+], [Het+], [HOB+] and [Hom-Rep]

You are now playing a high-payoff game. Pick a number between 11 and 20. You will always receive the amount that you announce, in tokens.

In addition:

- if you give the same number as your opponent, you receive an extra 10 tokens.
- if you give a number that's exactly one less than your opponent, you receive an extra 80 tokens.

Example:

- If you say 17 and your opponent says 19, then you receive 17 and he receives 19.
- If you say 12 and your opponent says 13, then you receive 92 and he receives 13.
- If you say 16 and you opponent says 16, then you receive 26 and he receives 26.

Treatments [HOM+] and [Het+]

Your opponent is:

- a student from maths and sciences playing the high-payoff game (treatment [Het+]) / humanities (treatment [Het+])
- he is given the same rules as you.

Treatment [HOB+]

In this case, the number you play against is chosen by:

- a student from humanities playing the high payoff game with another student from humanities. In other words, two students from humanities play the high payoff game with each other (extra 10 if they tie, 80 if exactly one less than opponent). You play against the number that one of them has picked.

B.1.3 Treatments [Tut], [AT-Hom] and [AT-Het]

Before playing treatments [Tut], [AT-Hom] and [AT-Het], the subjects were given the following ‘tutorial’:

Game Theory Tutorial: According to game theory, if the players are infinitely rational, then the game should be played in the following way. Both players should say 11.

Explanation: Suppose the two players are named Ana and Beatriz. If Ana thinks Beatriz plays 20, then Ana would play 19. But then Beatriz knows that Ana would play 19, so she would play 18. Ana realizes this, and so she would play 17.... they both follow this reasoning until both would play 11. Notice that if Beatriz says 11, then the best thing for Ana is to also say 11.

Treatment [Tut]

Your opponent is:

- a student who has also been given the game theory tutorial.

Treatment [AT-Hom]

Your opponent is:

- a student from maths and sciences.
- he has not been given the game theory tutorial.

Treatment [AT-Het]

Your opponent is:

- a student from humanities.
- he has not been given the game theory tutorial.

B.1.4 Treatments [AP-Hom] and [AP-Het]

Treatment [AP-Hom]

In this case, the number you play against is chosen by:

- a student from maths and sciences playing the low payoff game with another student from maths and sciences. In other words, two students from maths and sciences play the low payoff game with each other (extra 10 if they tie, 20 if exactly one less than opponent). You play against the number that one of them has picked.

Treatment [AP-Het]

In this case, the number you play against is chosen by:

- a student from humanities playing the low payoff game with another student from humanities. In other words, two students from humanities play the low payoff game with each other (extra 10 if they tie, 20 if exactly one less than opponent). You play against the number that one of them has picked.

B.1.5 Treatments [Un], [Un+], [AP-Un], [Tut-Un], [AT-Un]

The treatments for the second experiment contain no information concerning own or opponents’ label, and are adjusted accordingly. The third experiment contains identical treatments.

Treatments [Un], [Un+] and [AP-Un] are identical to [Hom] (and Het), [Hom+] (and Het+) and [AP-Hom] (and AP-Het), respectively, of the first experiment, but with the following information concerning the opponent:

Your opponent is given the same rules as you.

Treatments **Tut-Un** is also preceded by the same game theory tutorial as **Tut** and the same game, followed by:

Your opponent has also seen the game theory tutorial.

Treatment [AT-Un] is identical to treatment [AT-Hom] (and AT-Het), with the following information concerning the opponent:

In this case, you are playing against a subject who has not seen the game theory tutorial, and who himself (or herself) plays against a subject who hasn't seen the tutorial either. In other words, the two subjects have played one another. You play against the number that one of them has chosen.

B.2 Sequences

In the first experiment, our 6 groups (3 for the endogenous and 3 for the exogenous classification) went through four different sequences of treatments. Two of the groups in the exogenous treatment followed Sequence 1, and one followed Sequence 2. The three groups of the endogenous classification each took a different sequence: respectively sequence 1, 3 and 4. All the sequences contain the treatments [Hom], [Het], [HOB], [Hom+], [Het+], [HOB+]. The order of the main treatments is different in each sequence, both in terms of changing the beliefs and the payoffs. Some sequences include treatments [AP-Hom], [AP-Het] while others included [Tut], [AT-Hom] and [AT-Het].

- **Sequence 1:** Hom, Het, HOB, Het, Hom, HOB, Hom+, Het+, HOB+, Het+, Hom+, HOB+, Tut, AT-Hom, AT-Het, Tut, AT-Hom, AT-Het

- **Sequence 2:** Hom, Het, Het, Hom, HOB, HOB, AP-Hom, AP-Hom, AP-Hom, AP-Hom, AP-Hom, AP-Hom, Hom+, Het+, Het+, Hom+, HOB+, HOB+

- **Sequence 3:** Hom+, Het+, HOB+, Het+, Hom+, HOB+, Hom, Het, HOB, Het, Hom, HOB, Tut, AT-Hom, AT-Het, Tut, AT-Hom, AT-Het

- **Sequence 4:** Het, Hom, HOB, Hom, Het, HOB, AP-Hom, AP-Het, AP-Hom, AP-Het, AP-Hom, AP-Het, Het+, Hom+, HOB+, Hom+, Het+, HOB+

- **Sequence Robustness:** Hom, Het, HOB, Hom-Rep, Het, Hom, HOB, Hom-Rep, AP-Hom, AP-Het, AP-Hom, AP-Het, Het+, Hom+, HOB+, Hom+, Het+, HOB+

The second experiment had 60 subjects in total with unlabeled groups, and they went through the following sequence:

- **Sequence Unlabeled:** Un, Un+, AP-Un, AP-Un, Un+, Tut-Un, AT-Un, Tut-Un, AT-Un

B.3 Details of the Cognitive Test

The cognitive test that subjects played in Experiment 1 takes roughly 30 minutes to complete, and consists of three questions. First, subjects are asked to play a computerised version of the board game Mastermind. Second, subjects are given a typical centipede game of seven rounds, and are asked what an infinitely sophisticated and rational agent would do. Third, subjects are given a lesser known pirates game, which is a four player game that can be solved by backward induction. Subjects are asked what the outcome of this game would be, if players were infinitely sophisticated and rational. Each question was given a score, and then a weighted average was taken. Subjects whose score was higher (lower) than the median score were labeled as high (low). Subjects in Experiments 2 and 3 saw an additional question before the other questions. This question was a 'muddy faces' game where subjects had to perform iterated reasoning to answer the three sub-questions correctly. We report next the instructions of the test, as administered to the students (see the online appendix for the original version in Spanish).

Instructions of the Test. This test consists of three questions. You must answer all three within the time limit stated.

Question 1:

In this question, you have to guess four numbers in the correct order. Each number is between 1 and 7. No two numbers are the same. You have nine attempts to guess the four numbers. After each attempt, you will be told the number of correct answers in the correct place, and the number of correct numbers in the wrong place.

Example: Suppose that the correct number is: 1 4 6 2.

If you guess : 3 5 4 6, then you will be told that you have 0 correct answers in the correct place and 2 in the wrong place.

If you guess : 3 5 6 4, then you will be told that you have 1 correct answer in the correct place and 1 in the wrong place.

If you guess : 3 4 7 2, then you will be told that you have 2 correct answers in the correct place and 0 in the wrong place.

If you guess : 1 4 6 2, then you will be told that you have 4 correct answers, and you have reached the objective.

Notice that the correct number could not be (for instance) 1 4 4 2, as 4 is repeated twice. You are, however, allowed to guess 1 4 4 2, in any round.

You have a total of 90 seconds per round: 30 seconds to introduce the numbers and 60 seconds to view the results.

Question 2:

Consider the following game. Two people, Antonio and Beatriz, are moving sequentially. The game starts with 1 euro on the table. There are at most 6 rounds in this game:

Round 1) Antonio is given the choice whether to take this 1 euro, or pass, in which case the game has another round. If he takes the euro, the game ends. He gets 1 euro, Beatriz gets 0 euros. If Antonio passes, they move to round 2. Round 2) 1 more euro is put on the table. Beatriz now decides whether to take 2 euros, or pass. If she takes the 2 euros, the game ends. She receives 2 euros, and Antonio receives 0 euros. If Beatriz passes, they move to round 3.

Round 3) 1 more euro is put on the table. Antonio is asked again: he can either take 3 euros and leave 0 to Beatriz, or pass. If Antonio passes, they move to round 4.

Round 4) 1 more euro is put on the table. Beatriz can either take 3 euros and leave 1 euro to Antonio, or pass. If Beatriz passes, they move to round 5.

Round 5) 1 more euro is put on the table. Antonio can either take 3 euros and leave 2 to Beatriz, or pass. If Antonio passes, they move to round 6.

Round 6) Beatriz can either take 4 euros and leaves 2 to Antonio, or she passes, and they both get 3.

Assume Antonio and Beatriz are infinitely sophisticated and rational and they each want to get as much money as possible. What will be the outcome of the game?

- a) Game stops at Round 1, with payoffs: (Antonio: 1 euro Beatriz: 0 euros)
 - b) Game stops at Round 2, with payoffs: (Antonio: 0 euro Beatriz: 2 euros)
 - c) Game stops at Round 3, with payoffs: (Antonio: 2 euros Beatriz: 1 euro)
 - d) Game stops at Round 4, with payoffs: (Antonio: 1 euro Beatriz: 3 euros)
 - e) Game stops at Round 5, with payoffs: (Antonio: 3 euros Beatriz: 2 euros)
 - f) Game stops at Round 6, with payoffs: (Antonio: 2 euros Beatriz: 4 euros)
 - g) Game stops at Round 6, with payoffs: (Antonio: 3 euros Beatriz: 3 euros)
- You have 8 minutes in total for this question.

Question 3:

Four pirates (Antonio, Beatriz, Carla and David) have obtained 10 gold dobl ones and have to divide up the loot. Antonio proposes a distribution of the loot. All pirates vote on the proposal. If half the crew or more agree, the loot is divided as proposed by Antonio.

If Antonio fails to obtain support of at least half his crew (including himself), then he will be killed. The pirates start over again with Beatriz as the proposer. If she gets half the crew (including herself) to agree, then the loot is divided as proposed. If not, then she is killed, and Carla then makes the proposal. Finally, if her proposal is not agreed on by half the people left, including herself, then she is killed, and David takes everything.

In other words:

Antonio needs 2 people (including himself) to agree on his proposal, and if not he is killed.

If Antonio is killed, Beatriz needs 2 people (including herself) to agree on her proposal, if not she is killed.

If Beatriz is killed, Carla needs 1 person to agree (including herself) to agree on her proposal, and if not she is killed.

If Carla is killed, David takes everything.

The pirates are infinitely sophisticated and rational, and they each want to get as much money as possible. What is the maximum number of coins Antonio can keep without being killed?

Notice that *the proposer* can also vote, and that exactly half the votes is enough for the proposal to pass.

You have 8 minutes in total for this question.

Scoring in Experiment 1. In the mastermind question, subjects were given 100 points if correct, otherwise they received 15 points for each correct answer in the correct place and 5 for each correct answer in the wrong place in their last answer. In the centipede game, subjects were given 100 points if they answered that the game would end at round 1, otherwise points were equal to min 0, (6 round) 15. In the pirates game, subjects obtain 100 if they answer 100, 60 if they answer 10, and min 0, (80 x) 10 otherwise. The overall score was given by the average of the three.

Question 0 (Experiments 2 and 3):

There are three people, A, B and C, each with a circle on their forehead. The circle can be white or black. Every person can see the circle on the others' forehead but not the one on their own. In reality, A and C have a white circle and B has a black circle:

They are given the following instructions, in this order, and can observe the reaction of the others:
If you know that your circle is black, take a step forward. Who will take a step forward?

Now, they are informed that at least one of them has a black circle. They are then asked: If you know the color of your circle, take a step forward. Who will take a step forward?

They observe the reaction to the previous question (in other words, they see who took a step forward). They are asked: Now that you have seen who stepped forward, if you know the color of your circle, take a step forward. Who will take a step forward? (Include only those new persons who take a step forward, don't include anyone who already took a step forward in the previous questions.)

Scoring in Experiments 2 and 3. Scoring for the three common games was the same as in Experiment 1. The muddy faces game gave a total of 120 points if the correct answer was given in each sub-question. Partial points were given depending on how close the answer was to the correct iterative reasoning. In order to calculate the overall score, the scores for each of these sub-questions were added to those from the three other questions and the resulting sum was divided by 4.2.

C Regressions

For all following tables, the number of observation refers to the number of treatment observations. These observations are clustered at id level into groups that have size equal to the number of subjects. For example, for Table 9, there are 156 treatment observations clustered at the id level to form 39 groups (i.e. 39 subjects played these treatments). All standard errors are thus clustered at the individual level, taking into account the dependence of the treatment outcomes.

Table 8: Experiment 1, Regressions on Payoffs Effects (joint for all sequences)

VARIABLES	Relevant Dummy	Classification Dummy	Constant	Observations
From [Hom] to [Hom+], Label <i>I</i>	-0.50*** (0.19)	0.22 (0.52)	17.21*** (0.33)	235
From [Hom] to [Hom+], Label <i>II</i>	-0.64** (0.32)	-0.10 (0.46)	16.92*** (0.34)	236
From [Het] to [Het+], Label <i>I</i>	-0.62*** (0.21)	0.37 (0.38)	17.50*** (0.33)	233
From [Het] to [Het+], Label <i>II</i>	-0.74*** (0.28)	0.39 (0.48)	16.58*** (0.39)	236
From [HOB] to [HOB+], Label <i>I</i>	-1.15*** (0.24)	0.34 (0.38)	17.77*** (0.30)	236
From [HOB] to [HOB+], Label <i>II</i>	-0.97*** (0.26)	-0.07 (0.45)	16.97*** (0.37)	234

Clustered standard errors in parentheses

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

Table 9: Experiment 1, Regressions from Post-tutorial treatments

VARIABLES	Relevant Dummy	Classification Dummy	Constant	Observations
From [Hom] to [AT-Hom], Label <i>I</i>	-0.71 (0.47)	0.47 (0.67)	16.67*** (0.49)	156
From [Hom] to [AT-Hom], Label <i>II</i>	-0.85** (0.42)	0.32 (0.66)	17.09*** (0.46)	156
From [Het] to [AT-Het], Label <i>I</i>	-0.32 (0.31)	0.58 (0.54)	17.29*** (0.48)	156
From [Het] to [AT-Het], Label <i>II</i>	-1.92*** (0.49)	0.30 (0.72)	17.04*** (0.54)	156
From [AT-Hom] to [AT-Het], Label <i>I</i>	1.06*** (0.36)	0.92 (0.65)	15.73*** (0.57)	156
From [AT-Hom] to [AT-Het], Label <i>II</i>	-1.14*** (0.40)	-0.16 (0.88)	16.49*** (0.56)	156

Clustered standard errors in parentheses

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

Table 10: Experiment 1, Regressions for asymmetric payoff treatments

VARIABLES	Relevant Dummy	Classification Dummy	Constant	Observations
From [Hom] to [AP-Hom], Label <i>I</i>	-0.74*** (0.11)	0.88* (0.47)	17.69*** (0.42)	100
From [AP-Hom] to [Hom+], Label <i>I</i>	-0.11 (0.17)	0.88 (0.45)	16.94*** (0.35)	100
From [Hom] to [AP-Hom], Label <i>II</i>	-0.38 (0.55)	-0.18 (0.69)	16.19*** (0.50)	100
From [AP-Hom] to [Hom+], Label <i>II</i>	-0.24 (0.37)	-0.52 (0.81)	15.98*** (0.48)	100
From [HOB] to [AP-Het], Label <i>I</i>	-0.50*** (0.14)	0.68 (0.46)	17.86*** (0.37)	99
From [AP-Het] to [HOB+], Label <i>I</i>	-0.55* (0.29)	0.60 (0.41)	17.40*** (0.39)	99
From [HOB] to [AP-Het], Label <i>II</i>	-0.28 (0.36)	-0.28 (0.67)	16.49*** (0.54)	100
From [AP-Het] to [HOB+], Label <i>II</i>	-0.32 (0.25)	-0.52 (0.73)	16.33*** (0.47)	100
From [AP-Hom] to [AP-Het], Label <i>I</i>	0.31** (0.13)	0.64 (0.45)	17.06*** (0.38)	119
From [AP-Hom] to [AP-Het], Label <i>II</i>	0.35 (0.24)	-0.28 (0.72)	15.86*** (0.48)	120

Clustered standard errors in parentheses

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

Table 11: Experiment 1, Regressions to examine effect of labels for going "From treatment x to y" (using a dummy for Label I, treatment dummy, and a label-treatment interaction term)

VARIABLES	Label I Dummy	Treatment Dummy 0: if treatment x 1: if treatment y	Interaction	Constant	Obs.
From [Hom] to [AP-Hom]	1.67*** (0.45)	-0.38 (0.54)	-0.36 (0.55)	16.10*** (0.49)	200
From [AP-Hom] to [Hom+]	1.8*** (0.59)	-0.24 (0.36)	0.13 (0.40)	15.72*** (0.38)	200
From [HOB] to [AP-Het]	1.63*** (0.45)	-0.28 (0.35)	-0.22 (0.38)	16.35*** (0.38)	199
From [AP-Het] to [HOB+]	1.4*** (0.48)	-0.32 (0.24)	-0.23 (0.38)	16.07*** (0.37)	199
From [AP-Hom] to [AP-Het]	1.63*** (0.45)	0.35 (0.24)	-0.04 (0.27)	15.72*** (0.38)	239
From [Hom] to [AT-Hom]	-0.21 (0.61)	-0.85** (0.42)	0.14 (0.61)	17.26*** (0.31)	312
From [Het] to [AT-Het]	2.00*** (0.61)	-1.92*** (0.49)	1.60*** (0.57)	17.19*** (0.34)	312
From [AT-Hom] to [AT-Het]	2.00*** (0.61)	-1.14*** (0.40)	2.21*** (0.53)	16.41*** (0.45)	312

Clustered standard errors in parentheses

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

Table 12: Experiment 2, Regressions for all treatments

VARIABLES	Relevant Dummy	Constant	Observations
From [Un] to [AT-Un]	-0.22 (0.32)	17.66*** (0.28)	173
From [Un] to [Un+]	-1.10*** (0.32)	17.68*** (0.28)	174
From [Un] to [AP-Un]	-0.52* (0.28)	17.66*** (0.28)	172
From [AP-Un] to [Un+]	-0.55** (0.28)	17.13*** (0.25)	232

Clustered standard errors in parentheses

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

Table 13: Experiment 3, Regressions for all treatments

VARIABLES	Relevant Dummy	Constant	Observations
From [E3-Un] to [E3-AT-Un]	-1.03* (0.60)	17.75*** (0.42)	90
From [E3-Un] to [E3-Un+]	-1.01*** (0.37)	17.72*** (0.41)	87
From [E3-Un] to [E3-AP-Un]	-1.02** (0.41)	17.70*** (0.42)	83
From [E3-AP-Un] to [E3-Un+]	-0.04 (0.40)	16.72*** (0.38)	112

Clustered standard errors in parentheses

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

Table 14: Experiments 1-3, Wilcoxon signed rank test results for all treatments

Experiment	Treatment	P-value
1	[Hom] to [AT-Hom], Label I	0.074*
	[Hom] to [AT-Hom], Label II	0.052**
	[Het] to [AT-Het], Label I	0.269
	[Het] to [AT-Het], Label II	0.000***
	[AT-Hom] to [AT-Het], Label I	0.015**
	[AT-Hom] to [AT-Het], Label II	0.010***
	[Hom] to [AP-Hom], Label I	0.000***
	[Hom] to [AP-Hom], Label II	0.331
	[AP-Hom] to [Hom+], Label I	0.601
	[AP-Hom] to [Hom+], Label II	0.435
	[HOB] to [AP-Het], Label I	0.002***
	[HOB] to [AP-Het], Label II	0.438
	[AP-Het] to [HOB+], Label I	0.021**
	[AP-Het] to [HOB+], Label II	0.055*
	[AP-Hom] to [AP-Het], Label I	0.029**
	[AP-Hom] to [AP-Het], Label II	0.211
2	[Un] to [AT-Un]	0.033**
	[Un] to [Un+]	0.000***
	[Un] to [AP-Un]	0.003***
	[AP-Un] to [Un+]	0.003***
3	[E3-Un] to [E3-AT-Un]	0.042**
	[E3-Un] to [E3-Un+]	0.016**
	[E3-Un] to [E3-AP-Un]	0.000***
	[E3-AP-Un] to [E3-Un+]	0.900

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

D Additional Figures

D.1 TOM scores

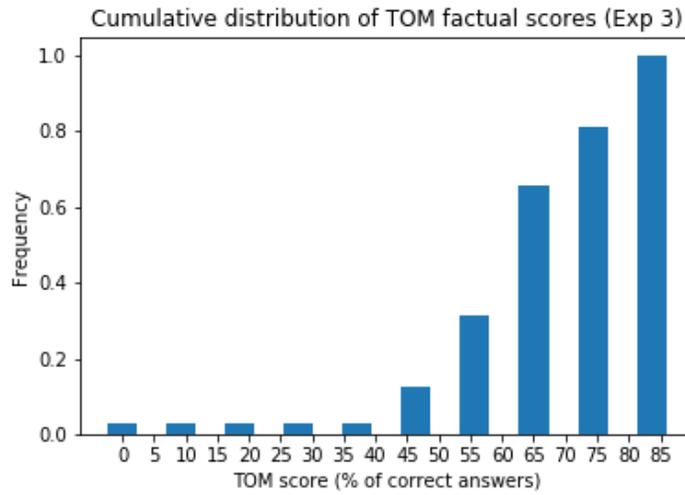


Figure 11: Distribution of factual TOM scores.

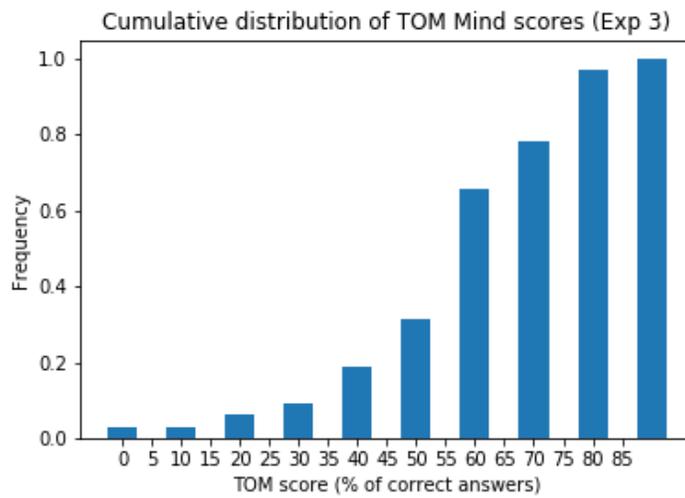


Figure 12: Distribution of TOM mind scores.

D.2 Individual behavior: violations of theory

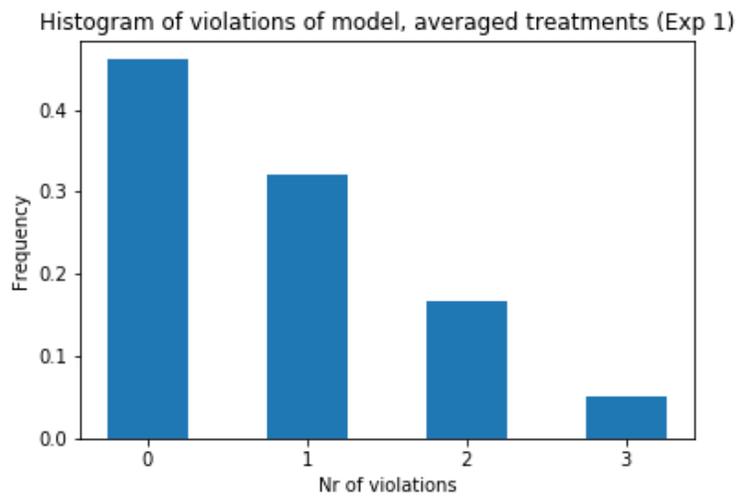


Figure 13: Experiment 1 (tutorial sessions): Number of violations of theory for averaged treatments.

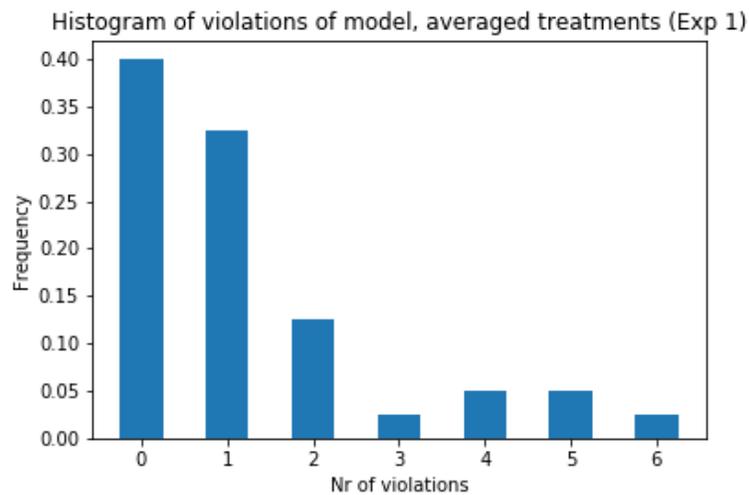


Figure 14: Experiment 1 (payoff sessions): Number of violations of theory for averaged treatments.

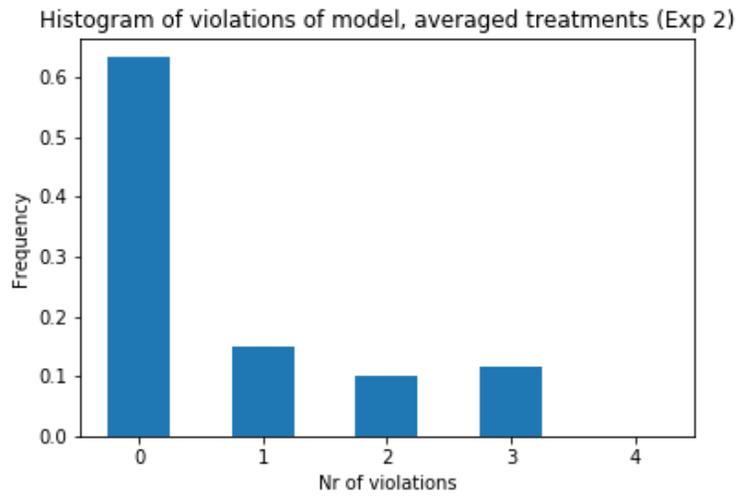


Figure 15: Experiment 2: Number of violations of theory for averaged treatments.

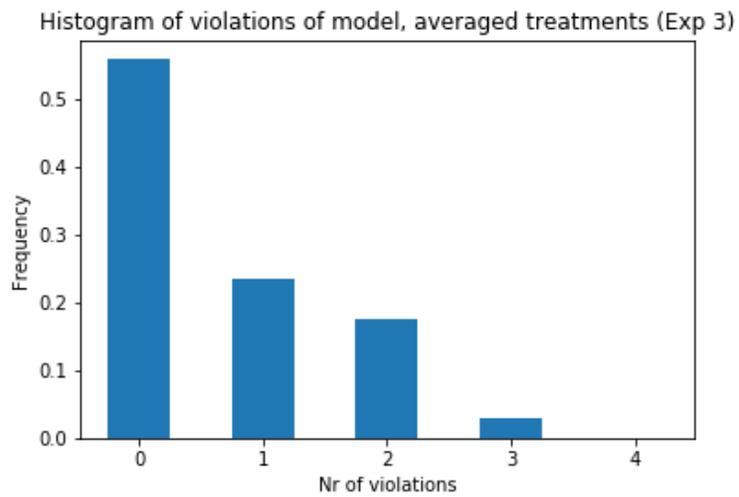


Figure 16: Experiment 3: Number of violations of theory for averaged treatments.

D.3 Individual behavior: shifts in behavior

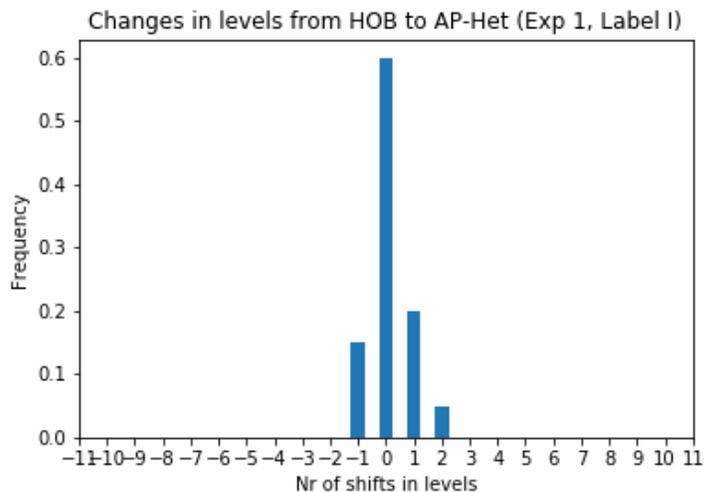


Figure 17: Experiment 1: Frequency of shifts in level played from [HOB] to [AP-Het].

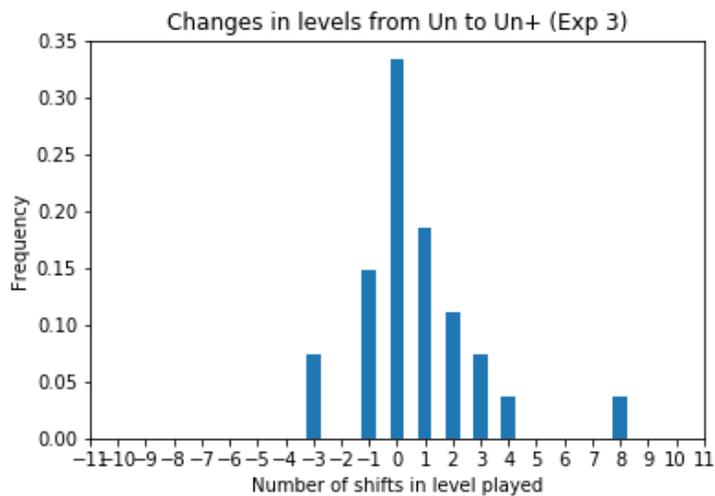


Figure 18: Experiment 3: Frequency of shifts in level played from [Un] to [Un+].

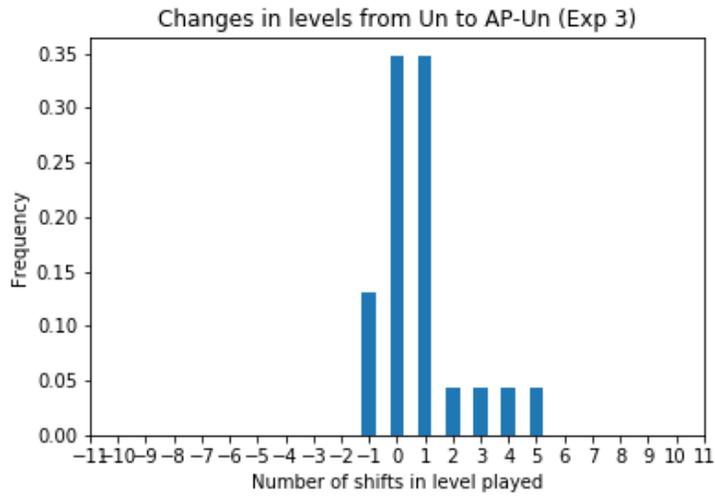


Figure 19: Experiment 3: Frequency of shifts in level played from [Un] to [AP-Un].

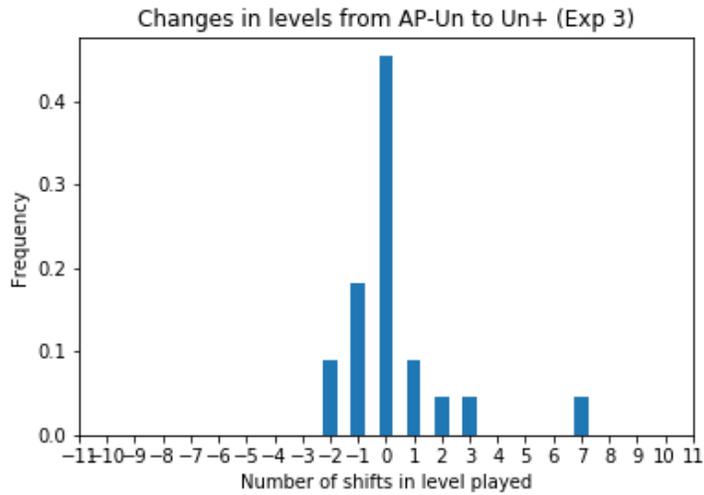


Figure 20: Experiment 3: Frequency of shifts in level played from [AP-Un] to [Un+].

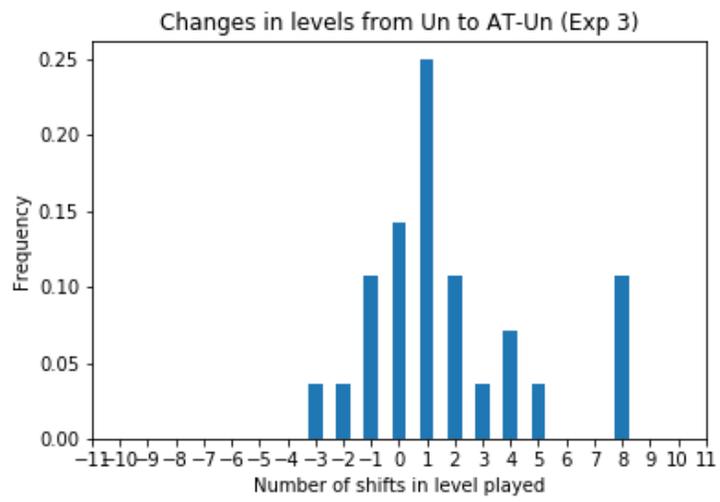


Figure 21: Experiment 3: Frequency of shifts in level played from [Un] to [AT-Un].

References

1. Agranov, Marina, Andrew Caplin and Chloe Tergiman. 2015. "Naive Play and the Process of Choice in Guessing Games" *Journal of Economic Science Association*, 1(2): 146-157.
2. Agranov, Marina, Elizabeth Potamites, Andrew Schotter and Chloe Tergiman. 2012. "Beliefs and Endogenous Cognitive Levels: An Experimental Study" *Games and Economic Behavior*, 75(2): 449-63.
3. Alaoui, Larbi and Antonio Penta. 2016a. "Endogenous Depth of Reasoning," *Review of Economic Studies*, 83, Issue 4: 1297-1333.
4. Alaoui, Larbi and Antonio Penta. 2016b. "Endogenous Depth of Reasoning and Response Time, with an application to the Attention-Allocation Task," *mimeo*.
5. Alaoui, Larbi and Antonio Penta. 2018. "Cost-Benefit Analysis in Reasoning," *mimeo*.
6. Arad, Ayala and Ariel Rubinstein. 2012. "The 11-20 Money Request Game: A Level- k Reasoning Study", *American Economic Review*, 102(7): 3561-3573.
7. Avoyan, Ala and Andrew Schotter (2015), "Attention in Games: An Experimental Study", *mimeo*, New York University.
8. Basu, Kaushik. 1991. "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory." *American Economic Review Papers and Proceedings*, 84(2): 391-395.
9. Bosch-Domènech, Antoni, Jose García-Montalvo, Rosemarie Nagel, and Albert Satorra. 2002. "One, Two, (Three), Infinity...: Newspaper and Lab Beauty-Contest Experiments." *American Economic Review*, 92(5), 1687-1701.
10. Camerer, Colin F. 2003. *Behavioral Game Theory*, Princeton University Press, Princeton, NJ, USA.
11. Camerer, Colin F., Teck-Hua Ho, and Juin Kuan Chong. 2004. "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics*, 119(3): 861-898.
12. Camerer, Colin F., George Loewenstein; Martin Weber. 1989. "The Curse of Knowledge in Economic Settings", *Journal of Political Economy*, Vol. 97(5), 1232-1254
13. Capra, C. Monica, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt. 1999. "Anomalous Behavior in a Traveler's Dilemma?" *American Economic Review*, 89(3): 678-690.
14. Costa-Gomes, Miguel A., and Vincent P. Crawford. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review*, 96(5): 1737-1768.
15. Costa-Gomes, Miguel A., Vincent P. Crawford, and Bruno Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica*, 69(5): 1193-1235.
16. Crawford, Vincent P. 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." *American Economic Review*, 93(1): 133-149.
17. Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." *Journal of Economic Literature*, 51.

18. Crawford, Vincent P., and Nagore Iriberri. 2007. "Level- k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica*, 75(6): 1721-1770.
19. Friedenber, Amanda, Kets, Willemien, and Terri Kneeland. 2017. "Bounded Reasoning: Rationality or Cognition," *mimeo*.
20. Fragiadakis, Daniel E., Knoepfle, Daniel T., and Muriel Niederle. 2016. "Who is Strategic?," *mimeo*.
21. Gill, David., and Victoria Prowse. 2016. "Cognitive ability and learning to play equilibrium: A level- k analysis." *Journal of Political Economy*. 126(4), 1619-1676.
22. Gill, David., and Victoria Prowse. 2017. "Strategic complexity and the value of thinking" *mimeo*.
23. Georganas, Sotiris., Paul, J. Healy and Roberto A. Weber. 2015. "On the persistence of strategic sophistication. " *Journal of Economic Theory*, 159(a): 369-400.
24. Goeree, Jacob K., and Charles A. Holt. 2001. "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions." *American Economic Review*, 91(5): 1402-1422.
25. Goeree, Jacob K., and Charles A. Holt. 2004. "A Model of Noisy Introspection." *Games and Economic Behavior*, 46(2): 365-382.
26. Goeree, Jacob K., Philippos Louis and Jingjing Zhang. 2017. "Noisy Introspection in the 11-20 Game " *Economic Journal*, forthcoming
27. Heifetz, A., and W. Kets. 2018. "Robust multiplicity with a grain of naiveté." *Theoretical Economics*, 2018, Vol. 13, 415-465.
28. Ho, Teck-Hua, Colin Camerer, and Keith Weigelt. 1998. "Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'." *American Economic Review*, 88(4): 947-969.
29. Kets, Willemien. 2017. "Bounded Reasoning and Higher-Order Uncertainty." *mimeo*.
30. Kneeland, Terri. 2015. "'Identifying Higher-order Rationality". *Econometrica*, 2015, Vol. 83, 2065-2079.
31. Liddle, Bethany, and Daniel Nettle. 2006. "Higher-order Theory of Mind and Social Competence in School-age Children. *Journal of Cultural and Evolutionary Psychology*, 2006, Vol. 4, 3-4, 231-246.
32. Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review*, 85(5): 1313-1326.
33. Palacios-Huerta, Ignacio and Oscar Volij. 2009. "Field Centipedes." *American Economic Review*, 99(4): 1619-1635.
34. Rampal, Jeevant. 2018a. "Limited Foresight Equilibrium". *mimeo*
35. Rampal, Jeevant. 2018b. "Opponent's Foresight and Optimal Choice". *mimeo*
36. Roth, Alvin, Vesna Prasnikar, Masahiro Okuno-Fujiwara and Shmuel Zamir, 1991. "Tokyo: An Experimental Study". *American Economic Review*, 1991, vol. 81(5), 1068-95.
37. Sofianos, Andis, Eugenio Proto and Aldo Rustichini. 2019. "Intelligence Personality and Gains from Cooperation in Repeated Interactions", *Journal of Political Economy*, Forthcoming.

38. Stahl, Dale and Paul Wison. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence". *Games and Economic Behavior*, 1995, vol. 10, issue 1, 218-254.
39. Stiller, James and R. I. M. Dunbar. 2007. "Perspective-taking and memory capacity predict social network size". *Social Networks*, 2007, vol. 29, 93-104.
40. Strzalecki, Tomasz. 2014. "Depth of Reasoning and Higher-Order Beliefs." *Journal of Economic Behavior and Organization*, 108: 108–122.
41. Weber, Roberto A. 2001. "Behavior and Learning in the "Dirty Faces" Game." *Experimental Economics*, 2001, vol. 4, issue 3, 229-242.