



**Universitat
Pompeu Fabra**
Barcelona

Department
of Economics and Business

Economics Working Paper Series

Working Paper No. 1587

Reasoning about other's reasoning

Larbi Alaoui and Antonio Penta

November 2017

Reasoning about Others' Reasoning^{*}

Larbi Alaoui[†]

Antonio Penta[‡]

November 13, 2017

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. From a 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, Nagore Iriberry and Rosemarie Nagel 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.

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

[‡]UW-Madison, 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 rationalizability, 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. Disentangling own understanding from reasoning about others’ incentives to reason, however, 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 for both players (hence, Bob’s opponent is not Ann: in our treatment, Ann is ‘replaced’ by

³Nagel’s (1995) beauty contest game has been one of the main workhorses for level- k reasoning (see also Camerer (2003), Crawford, Costa-Gomes and Iriberrri (2013) and references therein). Other prominent games in the level- k literature are, for instance, the two-persons guessing games of Costa-Gomes and Crawford (2006) and the travelers’ dilemma (e.g., Capra, Goeree, Gomez and Holt (1999)).

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 two experiments. Experiment 1 leverages the existing dataset by applying these methods to the baseline experiments in Alaoui and Penta (2016a). Experiment 2 instead develops simpler variations of those treatments for a new pool of subjects. Both experiments are based on the ‘modified 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* 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 in general level- k behavior 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 for both experiments. Sections 3 and 4 contain, respectively, the treatments and results from Experiment 1. Experiment 2 is discussed in Section 5. Section 6 discusses the results through the lens of the EDR model, and Section 7 discusses more general methodological considerations. Section 8 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.

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 due 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 (2016a). Rampal (2016b) 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. Hence, their cognitive bound is also at least as large as their rationality bound (which is analogous to our behavioral level), 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 find evidence of a significant gap between subjects’ cognitive and rationality bounds.

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, 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 (2016b) 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 (2016c), which extends the EDR model to account for response time, with an application to the experiment by Avoyan and Schotter (2016), and Alaoui and Penta (2017), which generalizes the model to general (non level- k) models of reasoning in coordination

games, and shows how heterogeneity of cognitive abilities between interacting agents favor equilibrium coordination in games of initial response, provided that players agree on their relative sophistication.

Gill and Prowse (2016, 2017) also investigate more explicitly the connection between cognitive abilities and level- k behavior, 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. They also investigate the connections between level- k behavior and non-cognitive abilities.

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 modified 11-20 game

The baseline game remains the *modified 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:⁴

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.

⁴Aside from our use of a similar game, however, our objectives and experimental design are quite different from Arad and Rubinstein's.

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 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 concerns 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 both experiments are undergraduate students from different departments at the Universitat Pompeu Fabra (UPF), in Barcelona. There were 180 subjects in total, with 120 participating in Experiment 1 and 60 in Experiment 2.

All treatments in both experiments are based on the modified 11-20 game above, with one token worth five euro cents. The exact sequences of treatments used in each session and experiment are provided in Appendix A.2. In both experiments, each subject 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 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. 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

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.

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 A.

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 subject's cognitive sophistication. The first criterion, referred to as the exogenous classification, separates subjects by their degree of study. Half 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 exam. 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 A.

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, 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. 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 specifically game theoretic reasoning, and so may induce sharper beliefs among the subjects concerning their relative sophistication compared to their opponents.

In Experiment 2, subjects were not separated by cognitive sophistication. They first took an expanded version of the cognitive sophistication test used in the endogenous classification described above. Only the subjects in the middle half of the distribution participated in the treatments below, but this was not revealed to them. They were never given any information about their performance on the test before the treatments were administered. This 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 Experiment 2 treatments 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 refer to the subjects in Experiment 2 as being in the *'unlabeled' classification*.

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

3 Experiment 1: Treatments

In this Section we discuss the treatments of our experiment. We begin by reviewing the baseline treatments in Alaoui and Penta (2016a), which are necessary to understand our current analysis. In these treatments, discussed in Section 3.1, incentives or beliefs are varied for both agents at the same time. They thus allow to test the basic premise of the EDR model, which is that both beliefs and payoffs may systematically affect the observed ‘level of behavior’. In the 11-20 game, in particular, the EDR model generates the following sets of comparative statics: (i) Holding beliefs about opponents constant, (weakly) higher level- k ’s should be observed when x is increased; (ii) Holding x constant, (weakly) higher level- k ’s should be observed when subjects face opponents that they regard as increasingly sophisticated.

AP’s 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. Our new treatments, designed to disentangle these effects, are introduced in Sections 3.2 and 3.3.

3.1 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 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’.⁶

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

⁶Alaoui and Penta (2016b) provide axiomatic foundations to this assumption of the EDR model.

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.

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.

3.2 Identifying Beliefs: the post-Tutorial Treatments

We explain next the new treatments designed to identify whether subjects play according to their own (binding) cognitive bound in treatments [Hom] and [Het], or whether they are responding to their beliefs about the opponents.

Consistent with the intuitive idea of the *tutorial method* discussed in the Introduction, after having administered the baseline treatments of Section 3.1, 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.

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

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 [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.

3.3 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 3.2), 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 in the low payoffs treatments [Hom] and [HOB], respectively. Hence, the exercise is of a similar spirit to treatment [HOB], in which subjects play against the number chosen by an opponent engaged in treatment [Het]. 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 opponent of the same (other) 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 [HOM] 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.

4 Experiment 1: Results


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. AP's main findings on the baseline treatments can be summarized as follows:⁷

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 reveal 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 towards higher level- k 's when payoffs increase.
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'.

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


⁷Other 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.

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




(a)

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




(b)

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



(c)

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



(d)

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

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.

4.1 Identifying Beliefs – 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*) choose 11, although one could have expected that an even higher fraction would have made that choice.

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). These results are supported by the regressions performed. For all the regressions in this section, we use a standard OLS estimator, and group the endogenous and exogenous treatments together for all the data from the sessions with the tutorial treatments.

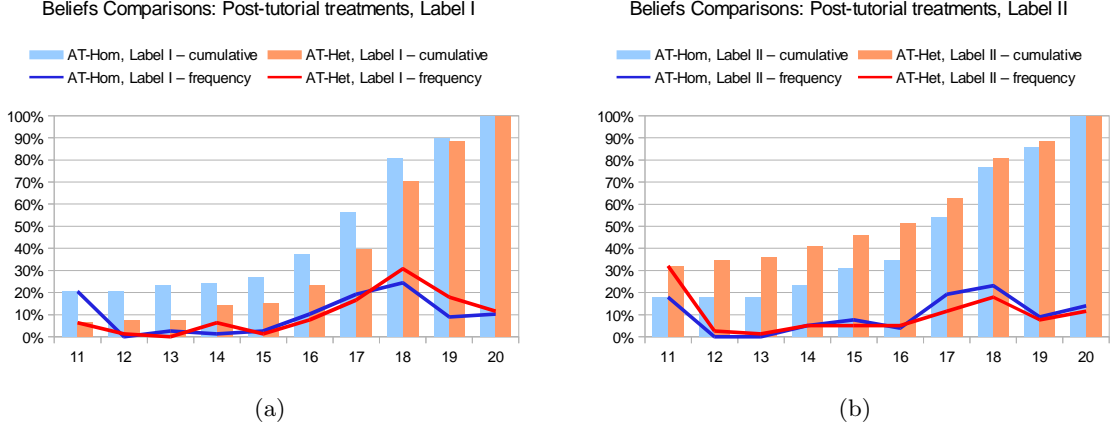


Figure 2: Beliefs comparisons: post tutorial treatments, Label *I* (left) and Label *II* (right)

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 5). Consistent with the patterns observed in Figure 1, the estimated coefficient is negative (-0.71 for label *I*, and -0.85 for label *II*), both statistically significant at the 5% level.⁸ 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.

The shift to lower numbers going from [Hom] to [AT-Hom] for the two labels indicates that the capacity, or cognitive bound, is binding for at least some of the subjects when playing [Hom]. We make this inference because the players face essentially the same opponent in [Hom] and [AT-Hom], and so the tutorial may affect their behavior only if the previous understanding was somehow considered binding. Similarly, the shift to lower numbers when going from [Het] to [AT-Het] for Label *II* also indicates that the cognitive bound is binding for at least some subjects. Interestingly, the same inference cannot be made for Label *I* players, since the coefficient is not significant for the [Het]-[AT-Het] comparison. As we discuss in Section 6, these results are to be expected when taking the perspective of the EDR model.

We then compare [AT-Hom] to [AT-Het] for each label (see Figure 2). In the case of label *I*, there's 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*, 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 expectations. In essence, if both labels

⁸In 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.

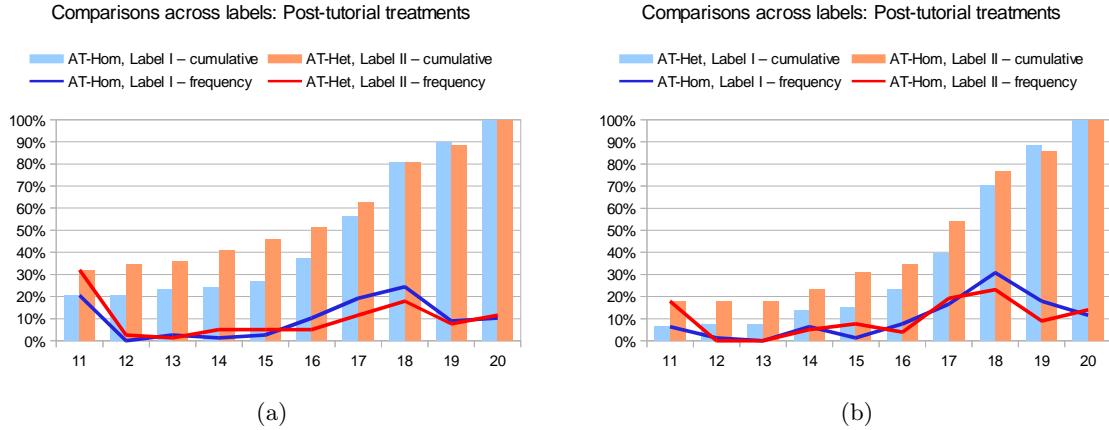



Figure 3: Comparisons across labels, AT-Hom Label *I* to AT-Het Label *II* (left) and AT-Hom Label *II* to AT-Het Label *I* (right).


I and *II* believe that label *II* players are less sophisticated, then they would play according to higher steps of reasoning (meaning lower numbers) when playing against label *I* opponents compared to label *II* opponents. 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. Pushing the analysis further, when comparing [Hom] to [Het] for each label, we find that the effect is smaller than in the comparison between [AT-Hom] to [AT-Het]. As will be discussed in Section 6, also these more subtle findings, which concern the differences on the size of the magnitudes, are to be expected from the viewpoint of the EDR model.

We also find that when comparing [AT-Hom] of Label *I* to [AT-Het] of Label *II*, there is complete FOSD of [AT-Het], and when comparing [AT-Hom] of Label *II* to [AT-Het] of Label *I*, the FOSD relationship is reversed nearly everywhere (Figure 3) The regressions are consistent with the figures, with significance at the 10% and the 5% levels, respectively. Since, for all these treatments, the tutorial has been administered and the subjects' capacity should no longer be a factor for most subjects, we do not expect these differences to be due to cognitive sophistication in understanding the game, but rather to a difference in beliefs between the two labels. In particular, supposing that label *I* and label *II* subjects understand the tutorial equally well, the FOSD relationship for [AT-Hom] of label *I* compared to [AT-Het] of label *II* suggests that label *II* subjects take the label *I* opponents to be more sophisticated than the label *I* subjects do. That is, they overestimate their label *I* opponents' sophistication relative to the way label *I* subjects' estimate opponents from the same group.⁹ Similarly, the FOSD relationship between [AT-Hom] of label *II* and [AT-Het] of label *I* reflects that label *I* subjects underestimate label *II* subjects, relative to label *II*. These effects point to the distinction between understanding of the game, linked to cognitive ability and capacity, and

⁹Equivalently, label *I* subjects underestimate their own group relative to label *II* subjects' expectations of the label *I* group.



(a)



(b)

Figure 4: Asymmetric payoffs treatments, Label *I* (top) and Label *II* (bottom).

understanding of the opponents' capacity.¹⁰ Here, we do not explore whether label *I* or label *II* are more correct in their evaluation of their opponents' sophistication, but it is noteworthy that there is a clear empirical difference in this evaluation.


4.2 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 repeated three times. Given the low number of subjects and the complexity of the instructions for these treatments, the results are overall less informative than those from the previous subsections. But Experiment 2 contains related treatments with a larger sample pool, as we will discuss in Section 5. All the results of the regressions discussed in this subsection are provided in Table 6.

¹⁰This distinction is also discussed in the literature on the curse of knowledge (see, e.g., Camerer, Loewenstein and Weber (1989)).

Asymmetric Payoffs Treatments, Label I


■ HOB, Label I – cumulative
 ■ AP-Het Label I – cumulative
 ■ HOB+, Label I – cumulative
— HOB, Label I – frequency
 — AP-Het Label I – frequency
 — HOB+, Label I – frequency



(a)

Asymmetric Payoffs Treatments, Label II

■ HOB, Label II – cumulative
 ■ AP-Het Label II – cumulative
 ■ HOB+, Label II – cumulative
— HOB, Label II – frequency
 — AP-Het Label II – frequency
 — HOB+, Label II – frequency



(b)

Figure 5: Asymmetric payoffs treatments, with double replacement. Label *I* (top) and Label *II* (bottom).

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 4). This is consistent with the regressions, for which the estimated coefficients when going from [Hom] to [AP-Hom] is -0.74 , and 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 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. An explanation of these findings is provided in Section 6, using the EDR model.




Figure 6: Beliefs comparisons: asymmetric payoffs treatments, Label *I* (left) and Label *II* (right).

When considering the [AP-Het] treatment, the natural comparison is not [Het], but rather [HOB]. 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 are of different labels. 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 5). The relationship between the curves is instead is 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*.

Lastly we compare [AP-Hom] to [AP-Het] for both labels, and find a FOSD relationship of [AP-Het] to [AP-Hom] (see Figure 6). The coefficient estimated in the regressions, however, is not significant for label *II*, while it is significant at the 5% level for label *I*.

5 Experiment 2: Unlabeled Variations

A possible concern with the experimental design of the previous sections 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 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.

Since labels are not provided, we cannot compare how subjects play as opponents' labels

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.

change, but we can analyze how they play against a general distribution of opponents. As we will discuss, we find interesting differences in behavior between the unlabeled and the labeled treatments.

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 modified 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.

Results. We first compare the results for treatments [Un] and [AT-Un]. Graphically, the cumulative distribution for [Un] is higher than that of [AT-Un] everywhere except at 15 and 16 (Figure 7), and the difference between the distributions is slight. We find that the estimated coefficient in the regression is not significant (see Table 7 for the regressions discussed here). Hence, this suggests that subjects' behavior in the [Un] treatment was essentially driven by their beliefs about their opponents, rather than their own (binding) cognitive capacity. The comparison of this result to those of the previous experiment reveals an interesting difference. The comparison of [Hom] to [AT-Hom] and [Het] to [AT-Het] suggests that, in Experiment 1, the cognitive capacity was always binding for label *II* subjects (more clearly so against label *I* subjects – significant at the 1% level, with a -1.92 coefficient – but also against other label *II*'s – coefficient of -0.84 , significant at the 5% level), whereas for label *I* it was only binding when they were playing against other label *I* subjects (coefficient of -0.7 , significant at the 5% level). Hence, subjects' situation when they play against a general distribution of opponents is qualitatively similar to that of label *I* subjects who face label *II*.

In short, unless a subject is prompted to think that they have a relatively lower level of cognitive sophistication (as label *II* subjects are in Experiment 1), or that their opponents have a relatively high cognitive sophistication (as label *I* subjects are in the [Hom] treatment of Experiment 1), subjects in Experiment 2 essentially assume that they are relatively more sophisticated than an opponent who is drawn at random from a general distribution. More broadly, this comparison suggests that capacity is more likely to become binding when beliefs are made sharper (for instance, by attaching labels to specific subgroups of the possible opponents). When, instead, labels are not assigned and beliefs remain diffuse, then beliefs, not capacity, are the main determinants behind choice. Looking ahead to future research, these findings suggest that the design of experiments should account for the precision of the beliefs induced, depending on whether the main focus of the study is on beliefs or cognitive capacity.

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 8). 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+].

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 [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.

6 Through the lens of the EDR model

The experimental results in the main text shed light on the complex interaction between beliefs, incentives and cognitive abilities, even without requiring the EDR framework. But further insights can be gained by looking at the data through the lens of that model.

The basic idea of the EDR model is that a subject '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*),¹¹ his 'behavioral level' may vary with his beliefs about the opponent. Intuitively, 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

¹¹AP use the term 'cognitive bound' to refer to the highest level- k that a subject is able to conceive of, in a given game. The term 'capacity' is due to Georganas, Healy and Weber (2015), essentially with the same meaning.

Pre- and post-tutorial comparisons, Unlabelled Treatments

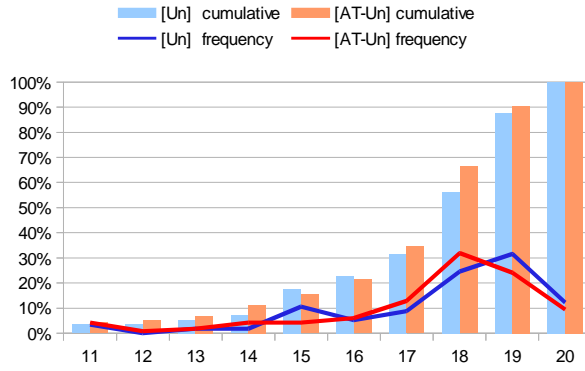


Figure 7: Pre- and post-tutorial, unlabeled treatments.

Asymmetric Payoffs Comparisons, Unlabelled Treatments




Figure 8: Asymmetric payoffs comparisons, unlabeled treatments.

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: $\hat{k}_j^i < \hat{k}_i$.¹²

The second dimension of endogeneity is that the understanding of a game, i.e., capacity, may itself vary with a player’s stakes in the game. For instance, in a version of the 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 x_j .¹³ *In particular, if x_i and x_j are increased symmetrically for both players (as in the low and high payoff treatments in Section 3.1, for each configuration of beliefs), both the cognitive bound and behavioral level should (weakly) increase.* As already discussed, this is one of the main findings in AP’s experiment.

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), an individual’s cognitive bound is binding if he regards his opponent as ‘more sophisticated’. In the EDR model, being of ‘higher sophistication’ is represented by having a 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. Based on this logic, and on the fact that subjects in Experiment 1 are revealed to regard label *I* subjects to be, on average, ‘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.*

Moreover, 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 summarized in Section 3.1.

¹²A 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.

¹³The endogeneity of players’ capacities is modeled as stemming from a cost-benefit analysis: costs represent players’ cognitive abilities; the benefits are related to the games payoffs, such as the x parameter in the 11-20 game. For further details, see Section 2 in Alaoui and Penta (2016a), and Alaoui and Penta (2016b) for an axiomatic foundation of these assumptions and modeling approach.

6.1 The Asymmetric-Payoff Treatments in the EDR model

To illustrate the effects of asymmetric changes of x_i and x_j using the EDR model, suppose that, in the baseline case in which $x_i = x_j = 20$, a particular subject i is such that $\hat{k}_i = k_i = 3$. That is, his behavioral level is 3, and it is equal to his cognitive bound, which therefore is binding. As explained above, this may be because i thinks that j has lower costs of reasoning, and hence he expects him to be (weakly) ‘deeper’ in a game with symmetric incentives. For the sake of the argument, suppose that, if i was able to reason further, he would think that, given j ’s costs of reasoning and $x_j = 20$, j ’s cognitive bound would be $\hat{k}_j^i = 4$ (which i can’t ‘see’ when $x_i = 20$, if he stops reasoning at $\hat{k}_i = 3$). Now, suppose that x_i is increased (say, up to 80), while holding $x_j = 20$ fixed, and that i ’s incentives are now sufficiently high that he performs three more steps of reasoning, so that the new cognitive bound with high incentives is $\hat{k}_i^+ = 6$ (the ‘+’ notation refers to the fact that payoffs are high for i). Now, i can ‘see’ that – given j ’s low incentives – j ’s cognitive and behavioral levels are $\hat{k}_j^i = k_j^i = 4$. Then, i ’s optimal response is to play according to $k_i^+ = 5$, which is higher than the original level of $k_i = 3$, but still lower than the new cognitive bound, $\hat{k}_i^+ = 6$, which therefore is not binding.

Now, holding constant $x_i = 80$ and hence i ’s cognitive bound $\hat{k}_i^+ = 6$, suppose next that x_j is also increased to $x_j = 80$, so that incentives are high for both players, and let $\hat{k}_i^{++} := 6$ denote i ’s cognitive bound now that payoffs are high for both players. Then, if i still regards j as more sophisticated, then he thinks that j has a deeper cognitive bound, $\hat{k}_j^{++} \geq \hat{k}_i^{++} = \hat{k}_i^+$, but not being able to understand more than \hat{k}_i^{++} , then his own cognitive bound $\hat{k}_i^{++} = 6$ becomes binding again, and hence his behavioral bound increases to $k_i^{++} = \hat{k}_i^{++} = 6$.

Note that, in the argument above, the movement from $k_i = 3$ to $k_i^+ = 5$ is thus due to the increase in i ’s cognitive bound alone, and is induced by an increase in x_i holding x_j (as well as i ’s beliefs over j ’s beliefs over i ’s depth of reasoning) constant; the further change from $k_i^+ = 5$ to $k_i^{++} = 6$ instead is determined by i ’s reasoning about the change in his opponent’s incentives.¹⁴ Hence, the overall movement from $k_i = \hat{k}_i = 3$ to $k_i^{++} = \hat{k}_i^{++} = 6$, which results from increasing both x_i and x_j from 20 to 80, would correspond to the comparison between the low-payoff and high-payoff treatments in Experiment 1 (Section 3.1), and between the [Un] and [Un+] treatments in Experiment 2 (Section 5); the ‘asymmetric payoffs’ treatments ([AP-Hom] and [AP-Het] in Section 3.3, and [AP-Un] in Section 5) instead correspond to the intermediate effect explained above.

Predictions of the Model. It follows from the above that the prediction of the EDR model is that *the distribution of actions in treatments [AP-Hom], [AP-Het] and [AP-Un] should be ‘between’ the distributions of the low and high payoff baseline treatments for the corresponding beliefs (which are, [Hom] and [Hom+] for [AP-Hom]; [HOB] and [HOB+] for [AP-Het]; and [Un] and [Un+] for [AP-Un]).* As already discussed in Sections 3.3 and 5, the results in Figures 4, 5 and 8 are roughly in line with these predictions.

Another prediction of the theory is that, *for label I , the increase in k_i from [HOB] to [AP-*

¹⁴In general, in the modified 11-20 game our model implies that \hat{k}_i (weakly) increases whenever x_i is increased. Furthermore, if x_j is held constant, k_j^i increases only if $\hat{k}_j^i \geq \hat{k}_i - 1$ in the first place.

Het] should be at most one. The reason 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. This prediction also seems consistent with the small shift in distribution from [HOB] to [AP-Het], and by 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 in Section 3.3, these results are remarkably consistent with the predictions of the EDR model.

Inferences based on the Model. Moving from the model's predictions to the inferences that can be made when the model is assumed, the difference between [Hom] and [AP-Hom] treatments in Experiment 1 is significant for label *I* (coefficient -0.74, at the 1% level), which suggests that label *I*'s cognitive bound is binding when they play against other label *I* subjects, and the size of the coefficient suggests that they regard other label *I* roughly as sophisticated as they are (the situation in which a subject regards the opponent exactly as sophisticated as him would correspond to a coefficient of exactly 1). Given this, cognitive bounds should also be binding in treatment [Hom+], and hence the fact that essentially no further shift is observed from [AP-Hom] to [Hom+] suggests that the costs of reasoning were such that the higher incentives only justified one extra step of reasoning. For label *II* subjects, instead, neither comparing [Hom] to [AP-Hom] or [AP-Hom] to [Hom+] lead to statistically significant differences, which based on the model implies that the asymmetric payoffs effect are too small to offset the additional costs of reasoning for label *II*. In Experiment 2, the difference between [Un] and [AP-Un] is not significant at the 5% level (although very close to the threshold). This suggests that, by and large, subjects' cognitive bound was not binding in treatment [Un], and hence that unlabeled subjects regarded other unlabeled subjects to be on average less sophisticated than themselves. These inferences on subjects' cognitive bounds entail further predictions for the behavior observed in the post-tutorial treatments. Namely, since the cognitive bounds are not binding in treatment [Un], the tutorial should have no significant effect on the behavior of these subjects, and hence there should be no significant difference between the behavior in treatments [Un] and [AT-Un].¹⁵ As discussed in Section 5, this is precisely what we observed.

The comparison of treatments [AP-Un] and [Un+] instead yields a fairly high effect (the

¹⁵Symmetrically, we could have used the evidence of no shift between [Un] and [AT-Un] to infer that cognitive bounds were not binding in the [Un] treatments, and hence that no further shift should be observed when [Un] is compared with [AP-Un]. In other words, the theory's prediction is that *a shift from [Un] to [AT-Un] should be observed if and only if a shift from [Un] to [AP-Un] is observed, and in the same direction* (magnitudes, however, need not be the same).

OLS coefficient is 0.55, significant at the 5% level).

6.2 The Tutorial Treatments in the EDR model

Through the lens of the EDR model, the Tutorial treatments may be interpreted as setting subjects' costs of reasoning to zero (or very low), so that the cognitive bound is not binding for the subject who received the tutorial. This way, their behavior in the post-tutorial treatments is entirely determined by their beliefs, and therefore can be compared to the baseline treatments [Hom] and [Het] to assess whether the cognitive bound was binding in the latter.

Under this interpretation of the tutorial, if the cognitive bound had not been binding in the pre-tutorial treatment, then the tutorial should have no effect on the post-tutorial treatment (provided that beliefs of all orders are kept the same). Hence, within the EDR model, the choice of a lower number (higher k) in a post-tutorial treatment suggests that – and in fact is only possible if – the cognitive bound was binding in the corresponding pre-tutorial treatment.

Predictions of the Model. Since, as explained above, 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 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 subjects.* This is consistent with the empirical findings in Figure 1 in Section 5, and also in line with the estimated OLS coefficients: the estimated coefficient for label *II* are -1.92 (significant at 1%) going from [Het] to [AT-Het] and -0.85 (significant at 1% level) going from [Hom] to [AT-Hom]; for label *I* subjects instead, the coefficients are -0.7 (significant at 5%) going from [Hom] to [AT-Hom], and -0.32 (significant at 5%) going from [Het] to [AT-Het]. The fact that all these comparisons are statistically significant (albeit with effects that are different in size) suggests that cognitive bounds were binding for at least some subjects in each treatments. Another prediction of the model is that the effect of the tutorial should provide an upper bound to the effect of the asymmetric payoffs, holding everything else constant. So, in Experiment 1, *the shift from [Hom] to [AP-Hom] should be smaller in size than that from [Hom] to [AT-Hom] for both labels.* The data confirm this prediction of the model very clearly. In Experiment 2, we observe no statistically significant difference between [Un] and [AT-Un] treatment. Given this finding (cf. footnote 15), the model implies that *the difference between [Un] and [AP-Un] and the difference between [AP-Un] and [AT-Un] should not be significant either.* As discussed above, the OLS coefficient for the [Un] and [AP-Un] is not statistically significant at the 5%, but the p-value is very close to the cutoff. On the other hand, the comparison between [AP-Un] and [AT-Un] is not significant. Hence, these much more subtle predictions of the theory also find support in the data, albeit with some qualifications and in any case less strongly than the others discussed above.

7 The Tutorial and Replacement Methods: Discussion

In this section we explain how key aspects of our experimental design have broader validity, and provide tools which may be applied to investigate questions that lie outside the typical domain of level- k theories, or of the EDR model in particular.

7.1 The *Replacement Method*

Several of our treatments involve what we call the *replacement method*, which mechanically consists of ‘replacing’ the opponent of a subject’s opponent. The simplest application of this method can be seen in the [HOB] and [HOB+] treatments from Section 3.1. The objective of treatment [HOB] is to test for the degree to which the subjects have a well-formed model of their opponents’ reasoning, and its precise wording is designed to pin down the entire hierarchy of beliefs. For instance, the instructions that a math and sciences student is given concerning his opponent in treatment [HOB] is: “[...] two students from humanities play against each other. You play against the number that one of them has picked.” It is therefore clear that he is playing a humanities playing a humanities subject, who himself is playing a humanities subject, and so forth. Any ambiguity that allowed players to believe that one of them believes (at some high level) that one of them is a student from math and sciences could result in a sophisticated subject behaving as a less sophisticated one, thereby invalidating the unambiguous identification of the beliefs effects.¹⁶ The ‘replacement’ of the opponent’s opponent in treatments [HOB] and [HOB+] is done precisely for this reason, and enables us to clearly isolate the effects of higher-order beliefs on players’ behavior.

The problem of 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. In a recent paper, Kneeland (2015) developed the idea of ‘ring games’, in which players $i = 1, 2, \dots, n$ are such that i ’s actions only have a direct effect on player $i + 1$ ’s payoffs, where we let $n + 1 = 1$. Hence: (i) player 2 only needs to reason about player 1, who has a dominant strategy in Kneeland’s (2015) games, and hence if 2 believes that 1 is rational, then his actions are pinned down by his first-order beliefs; (ii) player 3 only needs to predict 2’s action, but if he understands (i), then he would also form beliefs about 2’s beliefs about 1, and if he thinks that 2 is rational and that 2 believes that 1 is rational, then his behavior would only depend on his second-order beliefs, etc. Our *replacement method* can be seen as a simple way of achieving a similar structure in the lab, starting from a game with an arbitrary payoff structure.¹⁷

¹⁶We note that, in the experiment of Agranov, Potamites, Schotter and Tergiman (2012), the main treatments are analogous to our [Hom] and [HOB] treatments. Hence, the belief effects elicited in that paper identified the combined effects of the beliefs ([Hom] vs [Het]) and higher-order beliefs ([Het] vs. [HOB]) effects identified in our baseline treatments.

¹⁷One 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. In contrast, for general baseline games, our replacement method does not generally ensure a similar truncation of the belief hierarchies: for instance, in the baseline experiments of Alaoui and Penta (2016), all belief of orders higher than 2 are collapsed together in the [HOB] treatment. This advantage of ring games is most

A case in point is provided by the asymmetric payoffs treatments introduced in Sections 3.3. As explained in the introduction, in order to isolate the effects of increasing the stakes on a player's own understanding, it is not enough to vary payoffs asymmetrically, because higher stakes for one subject may affect his beliefs about the opponent's belief about his sophistication. So, for instance, if we observed a subject who chooses 18 in the [Hom] treatment (with low stakes), and 17 in a game in which only his own stakes are increased, we wouldn't know whether this is due to an increase in his own reasoning, or whether his understanding remained the same and he is merely reacting to his belief that the opponent may think that he has become more sophisticated. For this reason, in our [AP-Hom] and [AP-Het] treatments a subject's own payoff are increased, but subjects play against the number chosen by two students playing each other in the game with low stakes. That is, the opponent of a subject's opponent is replaced, to disentangle the subject's *direct* effect of his own reasoning, from the *interactive* effect which also comes into play in the baseline setting without replacement. Also in this case, it is important that the experimental instructions pin down the entire belief hierarchies.

Summing up, in the treatments above, the *replacement method* was applied to the label of the opponent's opponent in treatments [HOB] and [HOB+], to the payoffs of the opponent's opponent in treatment [AP-Hom], and to both the label and payoffs of the opponent's opponent in treatment [AP-Het]. This 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.

As an example, suppose we are interested in studying how sharing norms are affected when individuals interact with people from different cultures. A simple story to anchor ideas could be about tipping norms: the 'normal' tipping rate in the US is much higher than in Europe, and as a consequence waiters in the US and in the EU may react in opposite ways to an intermediate level of tips (for instance, a 12% tip would typically be received with enthusiasm in most European countries, but it would likely generate scorn or disappointment in the US). Coming to intercultural interactions, one question might be to study tipping behavior of a European tourist in the US (possibly with only a partial understanding of local social norms), as well as the waiters' reaction (one might wonder, for instance, whether an American waiter would react differently to a 12% tip if it's given by an American customer or by a European tourist). Even in this stylized illustration, the problem isn't trivial. First, behavior would likely depend on individuals' awareness of their cultural differences. But higher order beliefs come into play as well. For instance, it is reasonable to imagine that the American waiter's willingness to accept a lower tip may depend on his beliefs about how well the European understands the social norm: if he thinks the European knows the norm very well (for instance, if he knows he's been living in the US for a decade), then he might be much less forgiving than if he thought the European had no idea of the cultural differences. But then the European's tip may itself depend on his

crucial as subjects are capable of more sophisticated higher order reasoning. By contrast, the advantage of the replacement treatments is that it allows to investigate higher order beliefs effects (albeit partially identified) in arbitrary games, without altering the underlying payoff structure, and in a way which is very easy to implement in a lab.

beliefs about the waiter’s beliefs about him, and so forth. It would thus be difficult to say which fraction of the difference between inter- and intra-cultural tipping is really due to a subject’s understanding of the other culture (the *direct effect*), and how much instead is due to other higher-order beliefs considerations (the *interaction effect*).

To see how the *replacement method* may help here, consider a simple experiment with two groups of subjects (say, Europeans and Americans), who play an *Ultimatum game* both within groups (as in our [Hom] treatment above), and across groups (as in our [Het] treatment above), with the respondent’s strategy elicited with the minimum acceptable offer (MAO) method, played with feedback. For the sake of the argument, we will interpret offers as ‘tips’, and MAOs as ‘minimum tips which would be accepted without scorn’.¹⁸ Now, suppose that Europeans tip more in the [Het] treatment than in the [Hom] treatment, but still less than Americans’ MAOs in the [Hom] treatment. This could suggest that Europeans have an imperfect understanding of the American social norm, but it may also be that they strategically rely on the possibility that American receivers would be less demanding with them than they would be with fellow American tippers.¹⁹ To disentangle these two effects, a replacement treatment [Rep] could be designed, in which a European faces the MAO selected by an American subject in the [Hom] treatment. Such [Rep] treatment can be used to assess how well subjects understand the social norm of the other culture. That is because it only requires guessing the MAOs of the other group in the [Hom] treatment. For the sake of the argument, suppose that European subjects tip 5% in the [Hom] treatment, 10% in the [Het] treatment, and [18%] in the [Rep] treatment. The difference between [Rep] and [Hom] (13% in this example), would identify the subject’s assessment of the difference in the two social norms (the *direct effect*); the difference between [Rep] and [Het] (8% in this example) instead can be attributed to the *interaction effect*.

Some words of caution are in order, when it comes to applying the replacement method to settings in which social preferences play a major role. For instance, it may be that individuals in the ultimatum game react differently to behavior deemed as unjust if it is targeted to others (which may trigger indignation) or to themselves (which may trigger anger) – see, for instance, Camerer (2001, p.49). Thus, a plain replacement method, in which the opponent’s opponent is replaced by a third party, may turn a setting with a potential for anger into one with a potential for indignation. Simple variations of the method could be devised to address these concerns. For instance, in an experimental setting with multiple rounds, the opponent’s opponent could be replaced by an earlier version of the subject, so that the opponent’s choice would still determine the subject’s overall payoff in the experiment, not that of a third party (so as to maintain the ‘anger’ setup, if deemed desirable).²⁰ Similar variations could be devised for other problems

¹⁸The zero-payoffs for everyone after a rejection in the ultimatum game is perhaps to extreme to match the tipping story exactly, but we keep it nonetheless for illustrative purposes, and to make it clear that the replacement method can be equally applied to settings which arguably have nothing to do with level-k reasoning.

¹⁹In one of the most famous studies on intercultural comparisons of social preferences, Roth, Prasnikar, Okuno-Fujiwara and Zamir (1991) – which compare behavior in the ultimatum games in Israel, Japan, Slovenia and in the US – conclude that “what varies between subject pools is not a property like aggressiveness or toughness, but rather the perception of what constitutes a reasonable offer under the circumstances” (ibid., p.1092).

²⁰Note that this method does not consist of the subject playing against his own past action, as is done in Fragiadakis, Knoepfle and Niederle (2016)).

where identity of players' matters as well. Another variation which may be of particular interest and easy to implement, would be to have a finer distinction of higher order beliefs effects (beyond the second-order beliefs elicited by treatments [HOB] and [HOB+] above, that is). This could be achieved by iterating the replacement method, so as to replace the opponent's opponent (for second order beliefs), as well as the opponent's opponent's opponent (for third order beliefs), and so on. Clearly, this would come at the cost of increasingly complicated instructions, but conceptually it poses no difficulty at all.²¹ Finally, it is worth pointing out that – aside from special payoff structures – the confounding factors which the replacement method is meant to address are likely to be more relevant in settings with a smaller number of players than in larger: In games with many players, the game with replacement may be very similar to the one without replacement.²²

7.2 The *Tutorial Method*

In order to address the cognition-beliefs dichotomy, we designed treatments based on what we call the *tutorial method*, which 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 explained in the introduction, as long as subjects face the same opponents before and after having received this tutorial (an idea related to the replacement method we 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 independent 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 subjects may resort to different, non level- k , reasoning processes – see, e.g., Alaoui and Penta (2017)), 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.

To illustrate how these ideas could also be applied to a setting with feedback, we return to the tipping example from the previous section. Here, the tutorial could provide European subjects information on the normal tipping rate in the US. Then, the post-tutorial treatment would have these subjects play the same kind of opponents as before the tutorial (applying the replacement method, they would be told they play against the MAOs chosen by an American waiter playing with (i) a fellow American tipper (call this [AT-US] treatment), and (ii) an

²¹The same pros and cons apply to 'deeper' ring games as well, maintaining the relative advantages and disadvantages already discussed in footnote 17 above.

²²Suppose that $n \geq 2$ players are engaged in a baseline game G , and let $G(i)$ be the game in which the replacement method is applied to subject $i \in \{1, \dots, n\}$. It should be clear that the dissimilarity between G and $G(i)$ is maximal when $n = 2$, whereas the two games become more and more alike as $n \rightarrow \infty$, barring special payoff structures.

(untrained) European tipper ([AT-EU]). In this case, the [AT-US] treatment would inform to what extent the European subject is willing to adjust to the social norm, clear of any misunderstanding of it, and the comparison between [AT-US] and the [Rep] treatment from Section 7.1 would quantify the extent of the previous misunderstanding. Similarly, and in analogy with the comparison between [Rep] and [Het] treatments Section 7.1 (which identified the interaction effect), the difference between [AT-EU] and [AT-US] would provide a measure of the interaction effect clear of any misunderstanding of the social norm. Hence, as this example shows, these ideas can also be applied to settings that are not necessarily of initial response, and with the tutorial containing information on ‘typical’ distributions of play (such as the ‘normal’ tipping rate in the US). What remains crucial, however, is that such information is not about the behavior observed in earlier experimental rounds. This is to maintain the logical distinction between the role of feedback, and that of the tutorial. In other words, the tutorial method could be super-imposed over a design with feedback, to assess to what extent information about properties of the game, or distribution of behavior in related settings, etc., is considered relevant by the subjects themselves to inform their choices in an experiment.

8 Concluding Remarks

Subjects may reason about their opponents’ reasoning process in a way that is more nuanced than has been typically considered by the current literature. The reason is that it is challenging to identify experimentally 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 the opponents’ depth of reasoning and consequently, their behavior. In this paper we provide two methods, the tutorial and the replacement method, which can be used 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 the sophistication of their opponents but also 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 of cognitive bounds and reasoning about the opponent’s reasoning process given beliefs about their cognitive abilities. 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.

Looking beyond our specific question at hand, the tutorial and replacement methods can be useful when applied to different settings. These methods are to a large extent independent of the features of the underlying game, and are thus portable and easily implemented.

The replacement method, which consists of replacing the opponent’s opponent and controls for higher-order beliefs effects, can be used to disentangle the direct effects from a subject’s preferences and limitations from their beliefs, in a broad sense. For instance, as discussed in Section 7.1, this method can be used to separate subjects’ own preferences to adhere to a social norm from his 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 (2016), to assess for instance to what extent the higher levels of cooperation observed in the high-IQ group are due individuals’ own cognitive abilities, or to the high-IQ environment.

The tutorial method is also portable, 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. As we discussed in Section 7.2, it can also be used in settings with feedback, and in which subjects’ thinking and learning processes may be very different in nature from level- k . Future research can therefore make use of these methods in settings that are very different from the one focused on in this paper.

Appendix

A 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.

Each subject in the first experiment went through a sequence of 18 games. Payoffs are expressed in 'tokens', each worth 5 cents. Subjects were paid randomly, once every six iterations. The order of treatments is randomized (see below). Subjects in the second 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, 7 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.

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. The subjects all took the endogenous classification test first but they were not given any feedback, and remained unlabeled.

A.1 Instructions of the Experiment

We describe in A.1.1 to A.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 A.1.5.

A.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 your 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.

A.1.2 Changing Payoffs: Treatments [HOM+], [Het+] and [HOB+]

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.

A.1.3 Treatments [Tut], [AP-Hom] and [AP-Het]

Before playing treatments [Tut], [AP-Hom] and [AP-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 [AP-Hom]

Your opponent is:

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

Treatment [AP-Het]

Your opponent is:

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

A.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.

A.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.

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.

A.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, AP-Hom, AP-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, AP-Het, Tut, AT-Hom, AT-Het

- **Sequence 4:** Het, Hom, HOB, Hom, Het, HOB, AP-Hom, AP-Hom, AP-Hom, AP-Hom, AP-Hom, AP-Hom, 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 5:** Un, Un+, AP-Un, AP-Un, Un+, Tut-Un, AT-Un, Tut-Un, AT-Un

B Regressions

Table 5: 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** (0.34)	0.47 (0.67)	16.70*** (0.51)	156
From [Hom] to [AT-Hom], Label <i>II</i>	-0.85** (0.34)	0.31 (0.66)	17.09*** (0.50)	156
From [Het] to [AT-Het], Label <i>I</i>	-0.32 (0.26)	0.58 (0.54)	17.29*** (0.41)	156
From [Het] to [AT-Het], Label <i>II</i>	-1.92*** (0.37)	0.30 (0.73)	17.04*** (0.55)	156
From [AT-Hom] to [AT-Het], Label <i>I</i>	1.06*** (0.33)	0.92 (0.65)	15.73*** (0.49)	156
From [AT-Hom] to [AT-Het], Label <i>II</i>	-1.14*** (0.31)	-0.16 (0.89)	16.49*** (0.66)	156
From [AT-Hom] L. <i>I</i> to [AT-Het] L. <i>II</i>	-0.94* (0.52)		16.21*** (0.37)	156
From [AT-Hom] L. <i>II</i> to [AT-Het] L. <i>I</i>	0.86** (0.43)		16.41*** (0.31)	156

Standard errors in parentheses

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

Table 6: 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*** (0.14)	0.88* (0.47)	17.69*** (0.35)	100
From [AP-Hom] to [Hom+], Label <i>I</i>	-0.11* (0.16)	0.88 (0.46)	16.94*** (0.33)	100
From [Hom] to [AP-Hom], Label <i>II</i>	-0.38 (0.41)	-0.18 (0.70)	16.19*** (0.55)	100
From [AP-Hom] to [Hom+], Label <i>II</i>	-0.24 (0.35)	-0.52 (0.83)	15.98*** (0.60)	100
From [HOB] to [AP-Het], Label <i>I</i>	-0.50*** (0.17)	0.67 (0.47)	17.86*** (0.35)	99
From [AP-Het] to [HOB+], Label <i>I</i>	-0.55*** (0.19)	0.60 (0.42)	17.40*** (0.31)	99
From [HOB] to [AP-Het], Label <i>II</i>	-0.28 (0.39)	-0.28 (-.68)	16.49*** (0.54)	100
From [AP-Het] to [HOB+], Label <i>II</i>	-0.32 (0.34)	-0.52 (0.74)	16.33*** (0.54)	100
From [AP-Hom] to [AP-Het], Label <i>I</i>	0.31** (0.13)	0.64 (0.47)	17.06*** (0.34)	119
From [AP-Hom] to [AP-Het], Label <i>II</i>	0.35 (0.29)	-0.28 (0.73)	15.86*** (0.54)	120

Standard errors in parentheses

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

Table 7: Experiment 2, Regressions for all treatments

VARIABLES	Relevant Dummy	Constant	Observations
From [Un] to [AT-Un]	-0.22 (0.31)	17.66*** (0.28)	173
From [Un] to [Un+]	-1.10*** (0.30)	17.68*** (0.30)	174
From [Un] to [AP-Un]	-0.52* (0.27)	17.66*** (0.29)	172
From [AP-Un] to [Un+]	-0.55** (0.24)	17.13*** (0.25)	232

Standard errors in parentheses

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

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. "Cost-Benefit Analysis in Reasoning," *mimeo*.
5. Alaoui, Larbi and Antonio Penta. 2016c. "Endogenous Depth of Reasoning and Response Time, with an application to the Attention-Allocation Task," *mimeo*.
6. Alaoui, Larbi and Antonio Penta. 2017. "Coordination and Sophistication," *mimeo*.
7. Arad, Ayala and Ariel Rubinstein. 2012. "The 11-20 Money Request Game: A Level- k Reasoning Study", *American Economic Review*, 102(7): 3561-3573.
8. Avoyan, Ala and Andrew Schotter (2015), "Attention in Games: An Experimental Study", *mimeo*, New York University.
9. Basu, Kaushik. 1991. "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory." *American Economic Review Papers and Proceedings*, 84(2): 391-395.
10. 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.
11. Camerer, Colin F. 2003. *Behavioral Game Theory*, Princeton University Press, Princeton, NJ, USA.
12. 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.
13. Camerer, Colin F., George Loewenstein; Martin Weber. 1989. "The Curse of Knowledge in Economic Settings", *Journal of Political Economy*, Vol. 97(5), 1232-1254
14. 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.
15. 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.
16. 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.
17. Crawford, Vincent P. 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." *American Economic Review*, 93(1): 133-149.
18. 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.

19. 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.
20. Friedenber, Amanda, Kets, Willemien, and Terri Kneeland. 2017. "Bounded Reasoning: Rationality or Cognition," *mimeo*.
21. Fragiadakis, Daniel E., Knoepfle, Daniel T., and Muriel Niederle. 2016. "Who is Strategic?," *mimeo*.
22. 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.
23. Gill, David., and Victoria Prowse. 2017. "Strategic complexity and the value of thinking" *mimeo*.
24. Georganas, Sotiris., Paul, J. Healy and Roberto A. Weber. 2015. "Cognitive ability and learning to play equilibrium: A level- k analysis." *Journal of Economic Theory*, 159(a): 369-400.
25. 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.
26. Goeree, Jacob K., and Charles A. Holt. 2004. "A Model of Noisy Introspection." *Games and Economic Behavior*, 46(2): 365-382.
27. Goeree, Jacob K., Philippos Louis and Jingjing Zhang. 2017. "Noisy Introspection in the 11-20 Game " *Economic Journal*, forthcoming
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. 2012. "Bounded Reasoning and Higher-Order Uncertainty." *mimeo*.
30. Kneeland, Terri. 2015. "Identifying Higher-order Rationality". *Econometrica*, 2015, Vol. 83, 2065-2079.
31. Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review*, 85(5): 1313-1326.
32. Palacios-Huerta, Ignacio and Oscar Volij. 2009. "Field Centipedes." *American Economic Review*, 99(4): 1619-1635.
33. Rampal, Jeevant. 2016a. "Limited Foresight Equilibrium". *mimeo*
34. Rampal, Jeevant. 2016b. "Opponent's Foresight and Optimal Choice". *mimeo*
35. Roth, Alvin, Vesna Prasnikar, Masahiro Okuno-Fujiwara and Shmuel Zamir, 1991. "Tokyo: An Experimental Study". *American Economic Review*, 1991, vol. 81(5), 1068-95.
36. Sofianos, Andis, Eugenio Proto and Aldo Rustichini. 2016. "Intelligence Personality and Gains from Cooperation in Repeated Interactions", *mimeo*, Univ. of Minnesota
37. Stahl, Dale and Paul Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence". *Games and Economic Behavior*, 1995, vol. 10, issue 1, 218-254.
38. Strzalecki, Tomasz. 2014. "Depth of Reasoning and Higher-Order Beliefs." *Journal of Economic Behavior and Organization*, 108: 108-122.