The Effect of Incentives on Choices and Beliefs in Games
An Experiment *

Teresa Esteban-Casanelles[†] Duarte Gonçalves[‡]

First version: February 27, 2020.

This version: October 1, 2020.

Abstract

In this paper, we present experimental evidence establishing that the level of incentives affects both gameplay and beliefs. Holding fixed the actions of the other player, we find that, in the context of dominance-solvable games, higher incentives make subjects more likely to best-respond to their beliefs. Moreover, higher incentives result in more responsive beliefs but not necessarily less biased. We provide evidence that incentives affect effort and that it is effort, and not incentives directly, that accounts for the changes in belief formation. The results support models where, in addition to choice mistakes, players exhibit costly attention.

Keywords: Belief Formation; Incentives; Level-k; Strategic Complexity; Quantal Response.

JEL Classifications: C72, C92, D83.

*We are very grateful to Mark Dean, Yeon-Koo Che and Jacopo Perego for many insightful comments and conversations. We also thank Larbi Alaoui, Alessandra Casella, Kfir Eliaz, Navin Kartik, Rosemarie Nagel, Antonio Penta, Evan Sadler, Karl Schlag, Michael Woodford, Evan Friedman, Guy Aridor and the participants at Columbia University's Micro Theory Colloquium and Experimental Lunch for valuable feedback.
This material is based upon work supported by the National Science Foundation under Grant Number 1949395 as well as by Columbia University's Program for Economic Research Experimental Grant and Columbia Experimental Laboratory in the Social Sciences.

[†]Department of Economics, Columbia University; teresa.estebancasanelles@columbia.edu.

[‡] Department of Economics, Columbia University and BRU-IUL; duarte.goncalves@columbia.edu.

1. Introduction

Making choices in strategic settings is challenging: people need to form an understanding of the environment, how others may act and what to do. One would expect that with larger stakes people would commit more effort to thinking and make better choices. However, there is surprisingly little evidence on how scaling up the incentive level affects gameplay — the distribution of actions of the players —, which can prove crucial in understanding the role of cognitive limitations in explaining behavior in strategic environments.

Existing models disagree in whether and how varying the incentive level, that is, scaling up payoffs, affects behavior in games. Some existing models, such as Nash equilibrium (Nash 1951) or level-k (Nagel 1995; Stahl and Wilson 1994; Camerer et al. 2004), posit that changing the incentive level does not affect behavior.[1] In other models, such as quantal response equilibrium (McKelvey and Palfrey 1995), stronger incentives affect behavior because mistakes become more costly. In line with recent theoretical work in individual decision making (Caplin and Dean 2015; Alaoui and Penta 2016, 2018), another possibility is that higher incentives make it worthwhile to commit more effort to reasoning and forming beliefs about others' choices.

This poses an important question that we address in this paper: does the incentive level affect behavior? Naturally, if the incentive level does not affect behavior, then this simplifies the search for models that can accurately describe behavior, making any incentive considerations unnecessary. If the incentive level does affect behavior, it is also important to know why. On the one hand, incentives may affect choice implementation mistakes, leading to better responses as mistakes become more costly as in quantal response. On the other hand, incentives may lead players to think harder and form different beliefs. In individual decision-making, existing experimental literature supporting models of costly reasoning and rational inattention (Caplin and Dean 2013; Dean and Neligh 2019) suggests that beliefs play a major role in how incentives and the environment's complexity affect choices. However, it remains unclear how costly reasoning would express itself in strategic environments where, instead of an exogenous state, uncertainty regards others' choices: Do beliefs about

[1] In level-k models, players have fixed types determining their beliefs about their opponents' gameplay which translates in invariance with respect to the incentive level.

others become less biased when incentives increase? Do they become more concentrated around a Nash equilibrium of the game?

In strategic settings, as opposed to individual choice settings, experimental evidence on whether and how incentive levels affect behavior is lacking. Aside from specific games with strong other-regarding concerns (e.g. the ultimatum game), the experimental literature has so far mostly focused on the effects of *distorting* the relative incentives to take any given action, rather than the effect of *scaling the incentive level*, leaving relative incentives unchanged.[2] In the former case, different choices and beliefs may arise even if the same thought process or heuristic is pursued, making it ambiguous why different outcomes emerge. As such, how costly reasoning expresses itself in strategic settings, where the object of uncertainty — others' gameplay — is not exogenous but endogenous, fundamentally concerns how individuals' behavior and beliefs are affected by the *scale* of incentives.

This paper presents experimental results testing whether and how varying the incentive level affects behavior, combining choice and belief elicitation as well as decision times. Within the context of a new class of dominance-solvable games, this paper presents findings based on data of an online experiment where we manipulate subjects' incentive levels. We show that incentives have a significant effect on both gameplay and beliefs through two distinct channels: choice implementation mistakes and belief formation. The first underlies models in which players face implementation costs and have noisy best-responses to beliefs.[3] The second constitutes the main feature in belief-driven models of costly cognition or reasoning. We find support for both channels and that their relative importance depends on the complexity of the strategic environment.

Our experimental design has two main features. First, it devises a methodology to identify the effect of incentives not only on gameplay but also on beliefs. If incentives affect gameplay, players can report different beliefs not because incentives affect belief formation directly, but because they anticipate that their opponents are going to play differently. Relying on a modification of the role of an observer closely related to Alaoui et al.'s (2020)

[2]Examples of papers focusing on incentive distortion are Goeree and Holt (2001), Heinemann et al. (2009), Alaoui et al. (2020) and Friedman and Ward (2019). The few exceptions that look at *scaling* efforts – namely McKelvey et al. (2000), Brocas et al. (2014) and Pulford et al. (2018) — find no significant difference. As we argue below, this null result is likely due to the relatively small incentive differences in these experiments.

[3]These mistakes could be due to individuals failing to realize a given action is best given their beliefs or to "trembling hands".

replacement method in a concurrent paper, we vary a player's incentives while holding fixed the opponent's gameplay (and incentive level). In this way, changes in beliefs are due to changing the player's incentive level only, as the event about which the player forms beliefs remains the same.

Second, we use a new class of games, termed diagonal games (Gonçalves 2020). This class of games provides a disciplined way to generate dominance-solvable games that vary in the number of iterations of deletion of strictly dominated strategies while keeping the number of strategies fixed and with minor changes in the payoff structure. Within this class of games, the number of steps to reach the dominance solution can be plausibly interpreted as a measure of the strategic complexity of the game. As such, we can contrast the effects of incentives in simpler and more complex games, that is, games that take more iterations to reach the dominance solution. Furthemore, this structure allows us to identify the strategies played by level-k players (Nagel 1995; Stahl and Wilson 1994) and provides testable predictions for quantal response models (Goeree et al. 2005, 2019).

We have a 2×2×2 between-subject design with random assignment. We manipulate the subject's incentive level (high vs. low), her opponent's incentive level (high vs. low) and the strategic complexity of the game, specifically, the number of steps to reach the dominance solution (2 vs. 3). This design is not only able to distinguish between the effects of one's own-incentive level and how this depends on the strategic complexity of the game but, as a by-product, it also provides evidence on how the opponent's incentive level affects beliefs and gameplay. Moreover, to avoid any learning effects or other confounders, each subject plays a single time. We elicit the subject's choice and beliefs about their opponent's choices — the probability they assign to each of the opponent's action being chosen — in an incentive-compatible manner and we measure the subject's response time, the number of seconds taken to state their beliefs and choose an action.

We first establish that gameplay and beliefs are significantly different depending on whether own incentives are high or low. One finding is that higher incentives lead to a decrease in the frequency with which dominated strategies are played as well as in the belief that opponents play dominated strategies. A second finding is that the action corresponding to the dominance solution is not necessarily played more often: In the simpler game, subjects do choose it more frequently and they also believe their opponent chooses the dominance-solution action more often. In the more complex game, however, both the frequency with which the dominance-solution action and the belief that their opponent will choose it decreases.

We then focus on examining two channels through which incentives may affect gameplay: choice implementation mistakes and belief formation. With respect to choice implementation mistakes, we show that with higher own incentives individuals choose better actions. Not only are best-response rates to own beliefs higher, the distribution of choices shifts towards subjectively better actions, that is actions with higher expected payoffs according to each player's beliefs. This, together with the fact that average gameplay is significantly less uniform, supports the conclusion that higher incentives result in lower implementation mistakes and can account for the decrease in frequency of the dominance-solution action in the more complex game.

In what concerns belief formation, we report three main findings. The first is that beliefs are more responsive — less uniform — under higher incentives, indicating that players put more effort into forming beliefs about their opponents. The second is that the bias in beliefs decreases from low to high own incentives, but only when the opponent has high incentives too. When instead the opponent has low incentives, increasing own incentives produces no significant change in the simpler game and results in slightly more biased beliefs in the more complex game.

The third and perhaps most important finding is that incentives seem to affect the amount of effort players commit and that it is effort in turn that accounts for these patterns in beliefs. We find that higher own incentives lead to longer response times and that these, and not incentives directly, account for the changes in beliefs. Moreover, greater effort is associated with more accurate beliefs when the opponent has high incentives or the environment is simple; when otherwise, effort turns out to be detrimental to belief accuracy. This is suggestive evidence that weak incentives result in underthinking in simple environments but stronger incentives can lead to overthinking in more complex settings.

Lastly, we explore which of these two channels — choice implementation mistakes or belief formation — is the most relevant in accounting for the changes in gameplay induced by the different incentive levels. We perform a structural estimation exercise based on logit quantal response to understand how much the variation in beliefs contributes to the variation in gameplay due to different incentive levels. The analysis of the counterfactual evidence suggests that while in the simpler game the induced changes in beliefs are relevant in accounting for gameplay and changes in gameplay, in the more complex game the variation in gameplay is mostly due to choice implementation mistakes.

Studying the effects of players' incentive level on both beliefs and gameplay yields important insights for developing models that deliver better predictions. In contrast with the limited existing work, we show that incentive levels do affect choices and beliefs. Our results further suggest that models of strategic behavior should incorporate two main components. First, that players make payoff-dependent mistakes in a manner much similar to control costs or quantal response. Second, that players' belief formation takes time and follows a costly reasoning process, but also that higher incentives may lead beliefs to be more biased, not less. We hope that this evidence may help the development of a new generation of game-theoretic models that incorporate these elements.

Related Literature

Our paper lies in the intersection of different literatures: the effects of incentives on choices and beliefs in strategic settings — experimentally and theoretically —, implementation mistakes and strategic sophistication.

The most relevant literature to our paper is the experimental literature looking at how incentives affect choices. In strategic environments, most of the literature has looked at settings where other-regarding concerns are central such as the ultimatum game. Most papers, in both laboratory (Hoffman et al. 1996; Carpenter et al. 2005) or field settings (Slonim and Roth 1998) find no differences in offers, only in rejection rates. Andersen et al. (2011) increase incentives by different factors, going as high as 1,000, and show that offer proportions are lower in high stakes treatments compared to the lowest treatment and unconditional rejection rates decrease as the stakes increase. In games with significant other-regarding concerns, monetary payoffs may not be good indicators of payoff functions describing players' behavior. Our paper differs from these by focusing on games with a neutral frame and a payoff structure that at least mitigates potential other-regarding concerns.

There are three papers that vary the incentive level in a strategic setting with less salient other-regarding concerns but that find no effects: McKelvey et al. (2000), Brocas et al.

(2014) and Pulford et al. (2018). McKelvey et al. (2000) focus on asymmetric matching pennies with both payoff distortion and scaling, where incentives were scaled by a factor of 4 corresponding to a difference of $0.40 per game. The authors show that the differences are small (between 2-5 pp.), inconsistent across sessions and not significant. Brocas et al. (2014) look at two-player, incomplete information betting games to study reasoning with private information and include a treatment with a 5-fold increase in payoffs (<$5 difference per game). Pulford et al. (2018) examine behavior in 3×3 and 4×4 normal-form games without strictly or weakly dominated strategies and scale incentive levels also by a factor of 5 (<$5 difference per game). None of these studies finds any statistically significant difference in gameplay. In contrast to these, our incentive treatment has a scaling factor of 40 with an implied difference of almost $20 between the high and low incentive cases per game and subjects play a single round. Other significant differences are that our compensation method precludes confounding variations in gameplay with risk aversion[4] and that we focus on dominance-solvable games and elicit beliefs.

More broadly, Camerer and Hogarth (1999) review the effects of incentives in a series of experiments with decision-making and game-theoretic environments, comparing across studies. The authors also find that higher incentives reduce framing effects and anchoring, though these effects are smaller in environments with other-regarding preferences. Other related experimental papers are Caplin and Dean (2013) and Dean and Neligh (2019), who show that scaling up incentives by a factor short of 20 improves accuracy in a perception task.

This paper also contributes to the theoretical literature on game theoretic models by testing in a controlled environment whether and how the incentive level affects choices and beliefs in games. Structural quantal response equilibrium (McKelvey and Palfrey 1995) implies that higher incentive levels scale down the influence of the variance in choices, which can be interpreted as costly implementation of choices (Mattsson and Weibull 2002).[5] The formulation of regular quantal response equilibrium (Goeree et al. 2005), instead, is omis-

[4]Aside from using an incentive-compatible compensation scheme that is robust to risk attitudes, the games in our experiment are themselves robust to risk attitudes in different senses. We rely on dominance-solvability and each step of deletion is identified based only on ordinal preferences. Moreover, the identification of level-k gameplay in diagonal games is also only based on ordinal ranking of outcomes.

[5]The equivalence between additive random utility and control costs was shown by Fudenberg et al. (2015), in an individual decision-making context.

sive regarding if and how equilibrium varies with scaling payoffs. The predictions of other game-theoretic models such as level-k (Nagel 1995; Stahl and Wilson 1994, 1995) and cognitive hierarchies (Camerer et al. 2004) are invariant to changes in the incentive level. The rational inattention framework (Sims 2003; Caplin and Dean 2015; Matějka and McKay 2015) or the endogenous depth of reasoning model Alaoui and Penta (2016; 2018) would naturally imply that incentives affect choices by leading individuals to adopt different information acquisition or reasoning strategies and ultimately reach different beliefs about the state of the world.

Another strand of the literature of particular relevance to this paper concerns the experimental literature focusing on implementation mistakes in games. The closest papers to ours are Costa-Gomes and Weizsäcker (2008), Rey-Biel (2009) and Friedman and Ward (2019). Costa-Gomes and Weizsäcker (2008) and Rey-Biel (2009) examine choices and belief reports in 3×3 two-player games and report best-response rates to stated beliefs of about 54% and 67%, respectively. Costa-Gomes and Weizsäcker (2008) find that failure to best-respond to stated beliefs results in objective payoff losses. Friedman and Ward (2019) report average best-response rates to stated beliefs of around 75% in asymmetric matching pennies games. These studies suggest that while belief reports are predictive of choices, implementation mistakes are also a relevant element. Differently from these studies, we study how the incentive level affects implementation mistakes.

Finally, this paper broadly contributes to the literature on strategic sophistication. Based on level-k, many papers classify subjects into different sophistication levels by using their choices (e.g. Costa-Gomes and Crawford 2006; Arad and Rubinstein 2012; Kneeland 2015) and find substantial level-1 and level-2 gameplay. Georganas et al. (2015) show that this classification is fragile, leading to inconsistencies across different games. Alaoui et al. (2020) compare how gameplay changes by increasing the payoff associated with undercutting in a dominance-solvable variant of the 11-20 game and provide evidence that strategic sophistication is also affected by payoff distortion, even when it should not. Our paper examines how choices in dominance-solvable games are affected by scaling incentives and emphasizes the role played by individuals' beliefs.

Overview

The remainder of the paper is organized as follows: Section 2 presents the methodological framework. Section 2.1 discusses the need for a specific experimental methodology to identify the effect of incentive levels on beliefs and choices in games and explains our approach. Section 2.2 justifies the choice of games, introduces the class of diagonal games and the specific instances of this class that the subjects faced. In Section 3, we outline the experimental design, discussing the methods, treatments and other details. Section 4 presents the results, starting with an overview of the data and a discussion of the general patterns observed in gameplay (Section 4.1), then focusing on (i) how incentives affect implementation mistakes (Section 4.2), (ii) the patterns in beliefs and their relation to effort (Section 4.3), and concluding with (iii) the examination of the relative contribution to the variation of gameplay of choice implementation mistakes and of belief formation (Section 4.4). Finally, Section 5 concludes.

2. Methodological Framework

2.1. Identifying the Effect of Incentives on Beliefs and Gameplay

In this section we formally outline the main question we seek to test in this paper: whether beliefs and gameplay are affected by changes in the incentive level. We then present the methodology that allows to disentangle the direct incentive effects from other potential confounding factors.

Most game-theoretic models can be seen as defining conditions (1) on players' beliefs about others' behavior, and (2) on how players' choices depend on these beliefs. Models such as Nash equilibrium can be interpreted as imposing equilibrium conditions such that beliefs are consistent with gameplay and players best-respond to these beliefs. Similarly, quantal response equilibrium imposes that beliefs are correct, but that players make payoff-dependent mistakes. The level-k model and other related models (e.g. cognitive hierarchies) can be thought of as restricting players' to have specific beliefs about their opponents and simply best respond to those beliefs, with no constraints posed by equilibrium considerations.

This paper seeks to understand whether the beliefs and the choices are affected by varying player i's incentive level, that is, by scaling payoffs by a positive factor. However, in order to be able to study the effects of incentives on beliefs and choices we need to keep fixed the event about which the player forms beliefs. Changes in a player's payoffs may affect her beliefs about her opponents' gameplay not because own incentives directly motivate players to put in more effort in forming better beliefs about others' gameplay but instead because a player considers others would think she would play differently because she has stronger incentives.

An example illustrates why simply comparing beliefs and gameplay under high or low incentives — for all players or a subset of these — may not work. Suppose two players, A and B, play a given game and we increase both players' incentives. If increasing B's incentives affects B's distribution of gameplay, we are then changing the event about which A forms beliefs and therefore cannot assert whether A's incentives affect her belief formation or not. Suppose instead that we only increase A's incentives. Then it may be the case that B thinks that A is going to play differently, leading B to adjust her gameplay. Again this changes the event about which A forms beliefs, introducing a confounding factor due to higher-order beliefs. In sum, comparing gameplay and beliefs under different incentive levels will not identify whether the effect is due to changes in beliefs.

A way to avoid such issues is to adapt to strategic settings the position of an observer (e.g. Huck and Weizsäcker 2002) for whom others' choices are held fixed. Concurrent work by Alaoui et al. (2020) also follows this approach to avoid these potential confounders and coined it the replacement method. Differently from their case, which examines the effects of distorting payoffs on actions, we focus on understanding the implications of scaling up incentives on both actions and beliefs.

The methodology is as follows. In the context of two player games, suppose players A and B play a given game. Then, another player, $\tilde{A}$, is presented with another game where $\tilde{A}$'s choices and the B's choices determine $\tilde{A}$'s payoff, but where $\tilde{A}$'s choices do not affect player B's payoff. It is then as if Player A is "replaced" by player $\tilde{A}$. We focus on the situation where $\tilde{A}$ is presented with the exact same game as A up to scaling payoffs.

A natural concern arises from this method: while $\tilde{A}$'s choices have no effect on her opponent's payoffs, A's choices do have an effect on B's payoffs. This may induce different considerations for A and $\tilde{A}$ when making choices, namely in terms of other-regarding considerations. Therefore, we propose the following modification that essentially entails that neither A's nor $\tilde{A}$'s choices affect B's payoffs.

In this paper we will focus on the case of two-player symmetric games with available actions S, a base payoff function $u : S \times S \to \mathbb{R}$ and an agent-specific payoff $u_i = \alpha_i \cdot u$, which depends on i's choices, her opponent's and i's incentive level $\alpha_i \in \{\alpha_L, \alpha_H\}$. Each agent i is assigned to one of four groups: $\{HH, LL, HL, LH\}$, where the first letter denotes the agent's own incentive level, and the second that of the agent's opponent. Each agent i will be matched to an opponent j such that both i's and j's actions determine i's payoffs. The matching function $m : I \to I$ satisfies the following conditions:

1. An agent cannot be matched with herself: $m(i) \neq i$.
2. An agent's opponent cannot be matched with the agent: $m(m(i)) \neq i$.
3. Agents in groups HH and LL are matched with other agents within the same group: if $\theta_i \in \{HH, LL\}$, then $\theta_{m(i)} = \theta_i$.
4. Agents in groups HL and LH are matched with agents in groups LL and HH, respectively: if $\theta_i = HL$ $(= LH)$, then $\theta_{m(i)} = LL$ $(= HH)$.

This matching procedure deserves some comments. The first condition ensures the player is not playing against herself. The second tackles a potential pitfall that would emerge in the case of having agents whose actions affect each others' payoffs and other agents in the role of observers, whose actions do not affect their opponents' payoffs. Instead, condition 2 makes agents symmetric: regardless of whether their opponent faces the same incentives, their choices do not affect their opponent's payoffs. The third and fourth conditions serve to fix higher-order beliefs: as B's opponent is not A, even if both A and B have the same incentive level, it could be that B's opponent does not, in which case this would be relevant information for A to understand how B is to play. Conditions 3 and 4 fix expectations: A knows that B, her opponent, her opponent's opponent and so on all have the same incentives and therefore makes knowledge of the specific matching redundant. The cognitive burden is also not significant: the agent only has to form beliefs about gameplay either of agents in the HH or in the LL groups in order to assess which action to take.

The effect of incentive levels is immediately identified when agents are randomly assigned to groups and the matching — complying with the procedure above — is also random. Holding the opponent's incentive level fixed at H (L), the effect of own incentives is obtained by comparing beliefs and choices in groups LH and HH $(LL$ and $HL)$. Moreover, this method provides the interesting by-product identifying the effect of the opponent's incentive level on the agents' beliefs and choices by comparing instead HL to HH, and LL to LH.

2.2. Games and Testable Implications

Having discussed how to disentangle the effect of incentives on beliefs and choices in games from other potential confounding factors, we now discuss the second major design question: the choice of the games themselves. We first go over the properties that we sought the games to have and then present the class of games — introduced and developed in Gonçalves (2020) — that disciplined our design decisions, highlighting some of their theoretically interesting and, we argue, experimentally desirable properties.

In order to decide on the games used in this experiment we strove to design games that allowed for simple manipulations of the degree of strategic complexity — in a well-defined sense — and that did not exhibit some of the features which could play the role of confounding factors when focusing on testing the effects of incentives on beliefs and choices. We searched for games with the following properties:

– First, we focused on dominance-solvable games as arguably the number of steps to reach the dominance solution could constitute a measure of strategic complexity. In particular, our goals was that the payoff structure and the procedure of iterated elimination of strictly dominated strategies would produce a clear, linear, ranking of strategies. Moreover, we wished that the changes in the payoff structure required to vary the number of steps were minor so that the games would be as similar as possible. While strategic complexity may not be necessary related with the number of steps to reach the dominance solution,[6] we sought a class of games where this ranking was plausibly associated with game complexity.

– Second, we wanted for the games to entail testable predictions by the predominant models in explaining experimental evidence in games, namely quantal response, level-k, and their generalizations.

– Third, according to existing experimental evidence, (e.g. Costa-Gomes and Weizsäcker 2008), subjects seem to coordinate on Pareto efficient outcomes. For this reason, we wished that the dominance solution was not neither Pareto dominated, nor Pareto dominant, nor salient, in the sense that the associated pair of payoffs occurs more than once.

– Finally, for the sole purpose of minimizing data requirements, we searched for games satisfying these properties within the class of symmetric, two-player games. Specifically, we wanted that the underlying strategic reasoning was symmetric.

Studied in Gonçalves (2020), diagonal games constitute a flexible class of games that accommodates these and other appealing properties. Moreover, it provides a disciplined way to generate two-player dominance-solvable games with varying steps of iterated deletion of strictly dominated strategies for arbitrary finite number of strategies with minimal payoff changes. This is desirable as it constrains the degrees of freedom in generating such games for experimental purposes. Furthermore, it provides a way to compare games within this class in terms of their strategic complexity, a ranking with respect to which different models of strategic behavior generate meaningful comparative statics as discussed in Gonçalves (2020). We now present the two instances of this class that we use in this experiment and highlight the properties that are of interest for this paper. These properties are general within this class of games.[7]

As the number of actions is a free parameter in the definition of games of this class, we were free to generate games of different sizes. Striving for the simplest environment for which we could meaningfully elicit beliefs and vary the number of iterations to reach the

[6] For instance, the reasoning support choosing a weakly dominant strategies may prove more difficult in some games than that to iteratedly eliminate dominated strategies. Indeed some experimental evidence suggests this: see, e.g. Grosskopf and Nagel (2008).

[7] A discussion of this class of games can be found in the online appendix.

Actions		Player 2 s_1	s_2	s_3	s_4
Player 1	s_1	40,40	70,30	80,20	10,10
	s_2	30,70	40,40	70,30	80,20
	s_3	20,80	30,70	40,40	70,30
	s_4	10,10	20,80	30,70	40,40

(a) **4×4, IDS: 2**

Actions		Player 2 s_1	s_2	s_3	s_4
Player 1	s_1	40,40	70,30	10,20	10,10
	s_2	30,70	40,40	70,30	10,20
	s_3	20,10	30,70	40,40	70,30
	s_4	10,10	20,10	30,70	40,40

(b) **4×4, IDS: 3**

Figure 1. **Games**

dominance solution (IDS),[8] we chose games where players have 4 actions. This constitutes the minimum number of actions needed to be able to have meaningful strategic interaction and be able to manipulate the number of IDS: 2 and 3. We focused on having simple, round payoffs, multiples of 10, with the same number of digits and, in particular, the same number of non-zero digits, in order to prevent salience considerations.

Specifically, the two instances of diagonal games that we are studying in this paper are those in Figure 1. Panels 1a and 1b show two symmetric dominance-solvable games with 4 strategies for each player and that require 2 and 3 IDS, respectively. The dominance solution, (s_1,s_1), is readily reached by iteratedly deleting strategies with higher indices. Note that two forces may drive players away from just coordinating on any of the actions and move towards the dominance solution: (i) seeking a higher payoff for themselves and (ii) taking preventive action in the expectation that the opponent will herself play actions with lower indices.

This class of games leads to a clear ranking of the players' actions $\triangleright$, coinciding with the actions' indices. Note that if an action is a higher index than another — e.g. s_n,s_m, with $n > m$ — then either (i) the one with the higher index is deleted at earlier iterations or (ii) they are both deleted at the same iteration but the one with lower index iteratedly[9] strictly dominates the one with higher index. This implies that even two actions that are found to be (iteratedly) strictly dominated at a given round of deletion, they are ranked among themselves. An example of this can be found in panel 1a, where actions s_3 and s_4 are both strictly dominated by s_2 and deleted in the first round and yet are ranked: s_3 strictly dominates s_4.

[8] By the number of iterations to reach the dominance solution we mean the number of iterations of simultaneous deletion of strictly dominated strategies.

[9] Considering only actions of the opponent that survived up to then.

This strict dominance relation leads to a testable implication by quantal response equilibrium as it predicts that the gameplay frequencies of actions s_2, s_3 and s_4 (s_3 and s_4) are ordered and decreasing in the 2 IDS (3 IDS) game. This is because one of the postulates of quantal response equilibrium (Goeree et al. 2005) is that actions associated with higher expected payoffs are played with greater frequency. As these actions are ranked according to strict dominance, for any beliefs the players might hold, their expected payoffs are equivalently ranked.

The level-k model, assuming that level 0's uniformly randomize, makes a different set of predictions. First, level 1 players would play s_2 in the 2 IDS game and s_3 in the 3 IDS one, as these are the best-responses to uniformly randomizing opponents. Second, higher levels play lower indices: in the 2 IDS game, level 2 or higher play s_1, while in the 3 IDS game level 2 plays s_2 and level 3 or higher play s_1. Third, in the 2 IDS game the proportion of level 0 players can be estimated by adding together the frequency of with which s_3 and s_4 are played — that is, the total strictly dominated gameplay — and scaling it up by 2 to adjust for the fact that these represent half of the total strategies.

The uniform randomization assumption also implies that the frequencies of s_3 and s_4 should be the same. We can then adjust the proportion of frequencies of the remaining strategies.[10] An analogous procedure is feasible for the 3-IDS game (and for games in this class). Moreover, dominance-k[11] gameplay is identical to level-$(k-1)$ gameplay. Finally, note that this does not rely on a cardinal interpretation of payoffs: applying any strictly increasing transformation to the payoffs results in the same prediction.

The third property we desired is also met: the payoff vector associated with the dominance solution, (40,40), is neither Pareto dominant nor Pareto dominated by any other outcome. Moreover, as the corresponding payoff pair is not salient in the sense that it is repeated along the main diagonal. Other individual payoffs are also repeated parallel to the main diagonal, further mitigating issues of particular payoffs being salient in the bimatrix. Both these properties were sought to help avoiding potential other-regarding concerns as these may

[10]Fraction of level 1 gameplay = Frequency of s_2 – Fraction level 0 gameplay / (Total # strategies – # strictly dominated strategies) and analogously for level 2 or higher, by doing the same adjustment to the frequency of s_1.

[11]A dominance-k player does k steps of deletion of strictly dominated strategies and then best-responds to a uniform prior over the remaining strategies. Cf. Costa-Gomes and Crawford (2006).

act as confounding factors when analyzing the effect of incentives and strategic complexity on gameplay.

A possible worry is that the payoff structure itself would render the payoffs in the diagonal salient and induce a particular form of reasoning. In order to prevent salience of a particular pattern, the games displayed to subjects had the rows and columns randomly permuted. Consequently, the probability a given arrangement of the payoff structure is exhibited is extremely small, $1/(24)^2 \approx 0.0017$, and so is the probability that the diagonal structure of the payoffs is apparent to the subjects.

3. Experimental Design

In this section we detail the remaining logistics of the experimental design.

3.1. Treatments and Compensation

The experiment implemented a 2×2×2 design, corresponding to own and opponent's assignment (or, alternatively, to the 4 incentive group treatments as specified in Section 2.1) and to the specific games subjects played, with 2 or 3 IDS as in Figure 1. Subjects were sorted into one of the 8 treatments uniformly at random and *play only once*, being this a one-shot game implementation. Beliefs and choices were simultaneously elicited.

The order of rows and columns was randomized and subjects were informed of this. They received information on the possible incentive treatments but no information on games that they were to play except that these were 4×4 two-player games. They faced two practice rounds before knowing the treatment they were assigned to: in the first, the game had a strictly dominant action, the other had no pure-strategy Nash equilibrium; both were 4×4 two-player games and no feedback was provided. The subjects faced no time constraint but to finish the experiment under two hours. The subjects took on average about 22 minutes to complete the experiment.

From their own choices and their opponent's subjects obtained "Action Points". These correspond to the payoffs obtained in the game they faced and a randomly selected opponent complying with the matching conditions laid out in Section 2.1. From their belief reports, subjects got "Guess Points" $= 100 \times \left(1 - \frac{1}{2}||\hat{\sigma}^i_{-i} - \sigma_{-i}||^2\right)$, where $||\bullet||$ denotes the L2 norm,

$\hat{\sigma}^i_{-i}$ is the belief reported by the subject and σ_{-i} the gameplay of all subjects in the same treatment as the opponent other than the subject herself.[12] Only either the Action Points or the Guess Points were chosen, with equal probability, to implement the bonus, precluding hedging. The bonus they received was equal to the prize of their assigned own incentive treatment with probability (%) equal to the points they got. The choice to have points translated into probability points sought to preclude (or at least mitigate) risk-aversion concerns. This compensation mechanism induces incentive-compatible belief elicitation, robust to risk-aversion.

The high (own) incentive treatment corresponded to a prize of $20.00 and the low to a prize of $0.50. Points gained counted towards the probability of earning the prize. The subjects received $2.00 just for completing the experiment. Average earnings were of $8.94 ($24.97 per hour) and subjects were paid within 48 hours after the conclusion of the session.

All sessions were conducted on the platform Amazon Mechanical Turk, constraining potential workers to be adults based in the United States to control for the value of incentives. One of the main reasons why we chose to run the experiment on this online platform are that the this pool of subjects faces a more homogeneous opportunity cost of time as other tasks are readily available to take up and time here is very clearly money. The fact that these tasks are taken up primarily for the purpose of monetary compensation makes incentives more salient and mitigates concerns with false negatives due to subjects being intrinsically motivated to perform well regardless of monetary compensation. Moreover, in such a setting and in contrast to what is often the case of an on-site laboratory setting, subjects can exit the experiment at any moment without any inconvenience for themselves or others.

Although a potential concern is that the incentives in the low incentive treatment are too meager and could demotivate subjects, there is no clear evidence of that; much the opposite. We see that, in the low treatment, subjects do spend far more time than necessary to complete the task. While it is true that the difference of time spent between the high and low incentive treatments in completing the task is on average less than a minute, subjects in the low treatment still spend at least one minute on average. This is a significant amount of time for a task that can be completed within less than 3 seconds. If subjects were demoti-

[12] Our belief-elicitation incentive scheme corresponds to a special case of Hossain and Okui's (2013) binarized scoring rule.

vated, they would want to spend the least amount of time in this task since the opportunity cost — time — is high to subjects in our sample.

We also highlight that $0.50 for a single, one-shot decision is at least in line if not above the standard pay per answer in experiments,[13] and that the average hourly pay with the low-incentive bonus exceeds the U.S. minimum Federal hourly wage. Additionally, even low incentives are better than no incentives at all: As Gächter and Renner (2010) report in the context of a repeated game setting, belief accuracy is significantly higher when belief elicitation is paid on average €0.45 per report than when it is not paid for at all.

3.2. Experimental Sessions and Other Procedures

We conducted the experimental sessions during the days 9-10 and 15-17 of January 2020 with a set target of 100 subjects per treatment. Each individual could only take the experiment once. We recruited a total of 834 subjects with the treatment with fewer subjects having exactly 100 subjects.

The sequence of pages in the experiment proceeded as follows: (i) instructions were provided together with comprehension questions, (ii) a captcha check was performed to control for bots, (iii) two practice rounds were presented, (iv) own and opponent's incentives were revealed, (v) the choice and beliefs were elicited, (vi) a brief questionnaire on sociodemographic variables was presented and finally (vii) debriefing information on payment was shown.[14] The experiment was programmed using oTree.[15] Screenshots of the interface and the instructions are provided in the online appendix.

The questionnaire asked basic information: age, sex, education and prior exposure to game theory. In our sample, 44.36% of the subjects identified themselves as women and 55.64% as men. Subjects' age ranged from 20 to 73 years olds, with an average of about 37 years and standard deviation of about 11 years with the distribution being right-skewed.

[13] Pay per round/choice (beyond fixed fees) is similar or above, for instance, to that in Costa-Gomes and Weizsäcker (2008), Rey-Biel (2009) and Holt and Smith (2016) and strictly higher than in Agranov and Ortoleva (2017), Fudenberg and Liang (2019) and Dewan and Neligh (2020). All of these papers refer to on-site laboratory experiments except for the last two, which run experimental sessions on Amazon Mechanical Turk, paying a *maximum* of $0.06 and $0.015 per choice/round, respectively.

[14] Subjects spend on average about 6 minutes on the instructions and 5 on the ensuing questions, 4 minutes on the practice rounds and 1.5 minutes on the main task. Debriefing — including information on realized payoffs, completion code and questionnaire on socio-demographic information — took on average 2-3 minutes.

[15] oTree (Chen et al. 2016) is an adaptation of Django, a Python-based web framework.

Table 1. **Subjects per ISCED Level**

≤2	3	4	5	6	7	8
0.96%	12.83%	5.16%	27.34%	42.21%	11.15%	0.36%

Note: ISCED levels are as follows: ≤2: Incomplete high school or less; 3: High school; 4: Business, technical, or vocational school after high school; 5: Some college or university qualification, but not a bachelor, 6: Bachelor or equivalent; 7: Master or post-graduate training; 8: Ph.D.

Reported education ISCED levels are presented in Table 1. Moreover, 81% of the subjects claimed to have no prior exposure to game theory, 11.15% had been exposed to game theory but outside an academic environment and 7.89% had had formal training in game theory. Finally, we conducted tests and verified that indeed the samples across treatments were balanced by relying on t-tests for means (age) and Fished-exact tests of independence (sex, education and exposure to game theory).

4. Results

In this section we discuss how the data supports the main ideas put forth: (1) that incentives affect subjects' effort, and (2) that this in turn leads not only to changes in belief reports but also in subjects better responding to their reported beliefs. We present an overview of the data before discussing each of these topics. In the last subsection we argue that results are consistent with incentives affecting *both* belief formation and noise in actions.

4.1. Overview of the Data

We start by presenting the main statistics of the data. An immediate first observation is that the different treatments do have an effect on the subjects' behavior.

When players face significantly higher incentives their choices and beliefs are markedly different. Figures 2 and 3 exhibit the action frequencies and average beliefs.[16] In these figures it is possible to see that comparing high and low own incentives leads, in general, to statistically significant differences in choices and beliefs, which contrasts with the absence of significant differences in the existing literature (McKelvey et al. 2000; Brocas et al. 2014; Pulford et al. 2018). A reason behind this may be that in our experiment not only is the ratio

[16]We included in the appendix the choice frequencies and average beliefs for each of the 8 treatments — Table 11. Tests for differences in the distributions of choices and beliefs can be found in Table 18 in the online appendix.

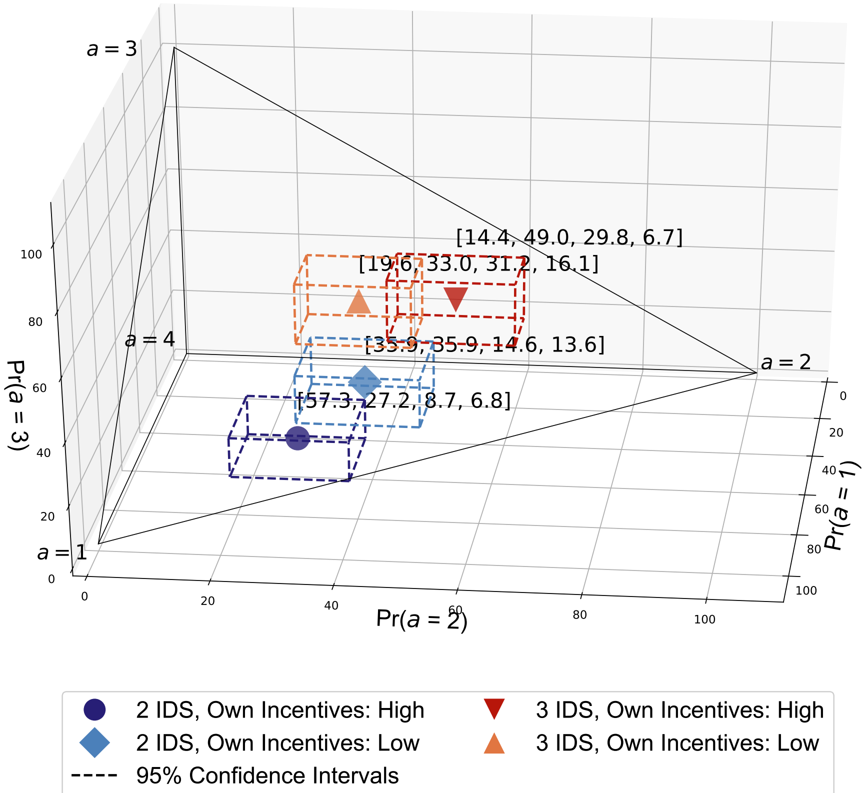

(a) **Opponent's Incentives: High**

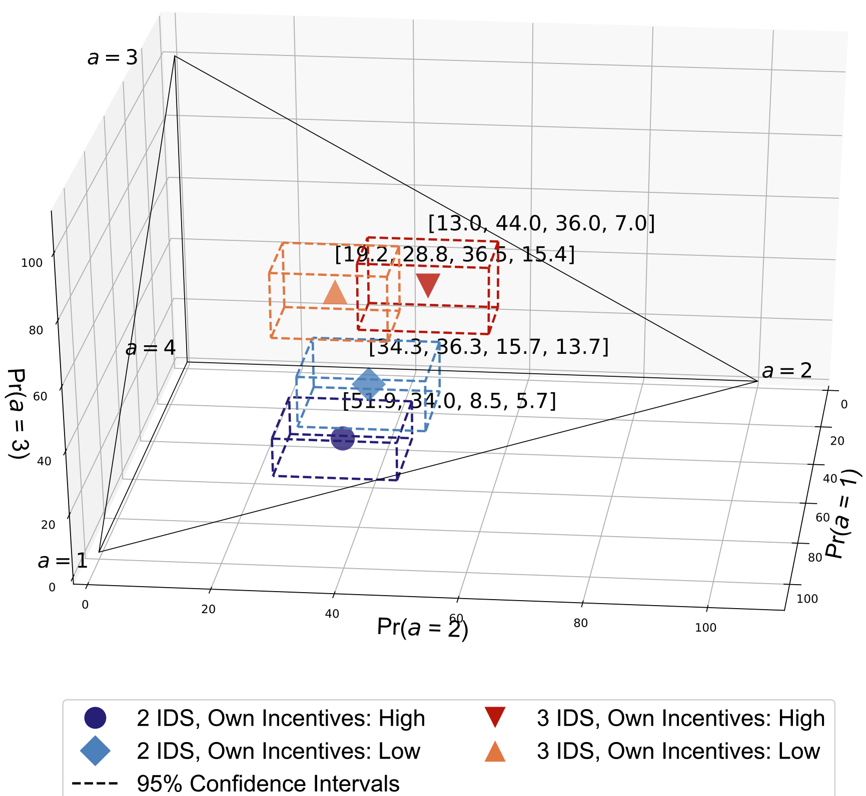

(b) **Opponent's Incentives: Low**

Figure 2. **Choice Frequencies**

Note: This figure shows the estimated probability of each action being played. Confidence intervals are provided for the vector as a whole implementing sup-t bootstrapped confidence intervals (Montiel Olea and Plagborg-Møller 2018). Labels correspond to the choice frequencies: [Action 1, Action 2, Action 3, Action 4] (%).

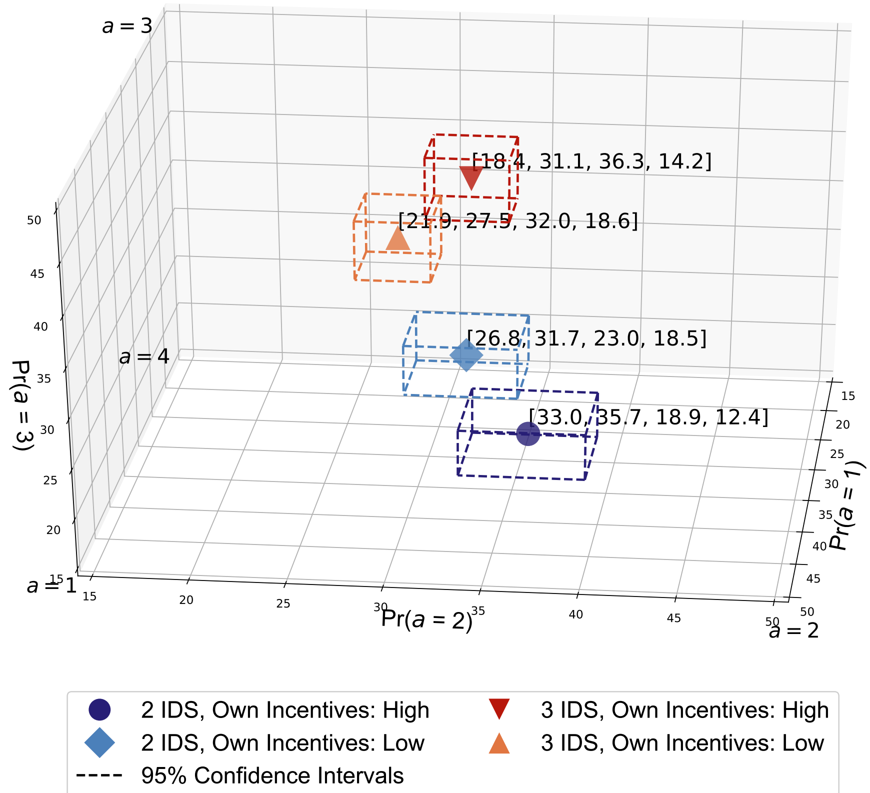

(a) **Opponent's Incentives: High**

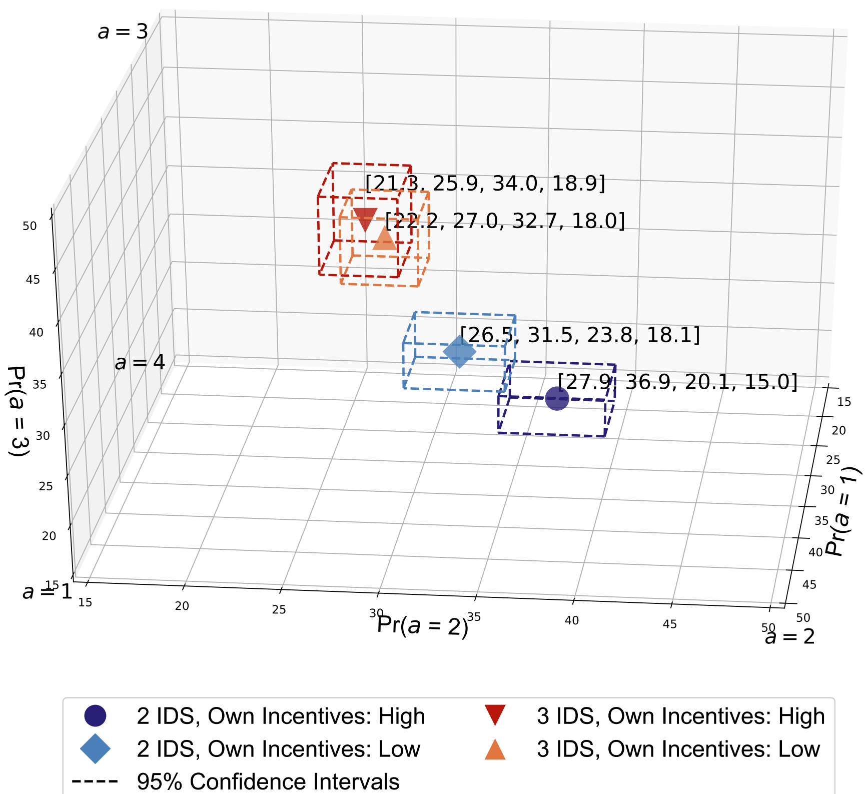

(b) **Opponent's Incentives: Low**

Figure 3. **Average Beliefs**

Note: This figure represents the average belief reports. Confidence intervals are provided for the vector as a whole implementing sup-t bootstrapped confidence intervals (Montiel Olea and Plagborg-Møller 2018). Labels correspond to the average probability assigned to each action: [Action 1, Action 2, Action 3, Action 4] (%).

of the incentive treatments substantial, with incentives being scaled by a factor of 40, when going from low to high own incentives, but also the incentive level is substantially different: for a single choice taking on average less than 2 minutes, the difference in potential reward is of $19.50. In contrast, in the mentioned literature the highest scale factor for a single choice is of 5 and the highest payoff difference between low and high incentives is $5.

Figures 2 and 3 also show a clear pattern of how incentives affect gameplay and beliefs: strictly dominated actions[17] are chosen less often and the probability assigned to these also decreases. In the 2-IDS game, increasing own incentive level decreases strictly dominated play from 28% to 15% and in the 3-IDS game from 16% to 7%. This is further confirmed in Table 2, where we present the differences in actions and beliefs when increasing own incentives. In particular, we show that when own incentives increase, subjects make fewer mistakes and they believe the opponents are to make fewer mistakes. These differences are statistically significant at $p<.05$, except for the change in beliefs about opponents when opponents have low incentives and play the more complex game. Moreover, when the *opponent's* incentives increase, subjects believe they are less likely to choose strictly dominated actions but only when the subjects themselves have high incentives. Otherwise, their beliefs about opponent's dominated gameplay are not significantly different. In other words, subjects are only responsive to the effect of their opponent's incentives on their choice of action when they have high incentives.

A naïve hypothesis would be that as incentives increase, the dominance solution strategy (action 1) is played more often. This is indeed true in the simpler 2-IDS game: not only is the difference in gameplay of action 1 higher, it is so by almost 20 pp. Moreover, subjects also assign a higher probability to opponents playing the dominance solution, although this is not statistically significant when opponents have low incentives. In the more complex 3-IDS game, higher incentives lead to a *decrease* — by more than 5 pp. — in the frequency of play of the dominance solution and to subjects assigning a *lower* probability of this action being played by their opponents.

[17]As described in section Section 2.2, actions are ordered by their index in ascending order. Action 4 is always strictly dominated by action 3, which is iteratedly dominated by acrtion 2, and so on. Therefore, figures that are closer to the lower left corner are closer to the dominance solution (which is always action 1). In the case of the 2-IDS game, actions 3 and 4 are strictly dominated; in the 3-IDS game, only action 4 is strictly dominated.

Table 2. **Difference in Dominated Gameplay**

(a) **Own Incentives: High vs. Low**

IDS	Opponent's Incentives	Difference in Choices		Difference in Beliefs	
2	High	-12.62	(5.69)	-10.20	(2.56)
2	Low	-15.26	(5.61)	-6.77	(2.48)
3	High	-9.34	(4.31)	-4.36	(1.76)
3	Low	-8.38	(4.39)	0.82	(1.74)

(b) **Opp. Incentives: High vs. Low**

IDS	Own Incentives	Difference in Beliefs	
2	High	-3.82	(2.29)
2	Low	-0.39	(2.74)
3	High	-4.67	(1.79)
3	Low	0.51	(1.70)

Note: Standard errors in parentheses.

(a) **IDS: 2** (b) **IDS: 3**

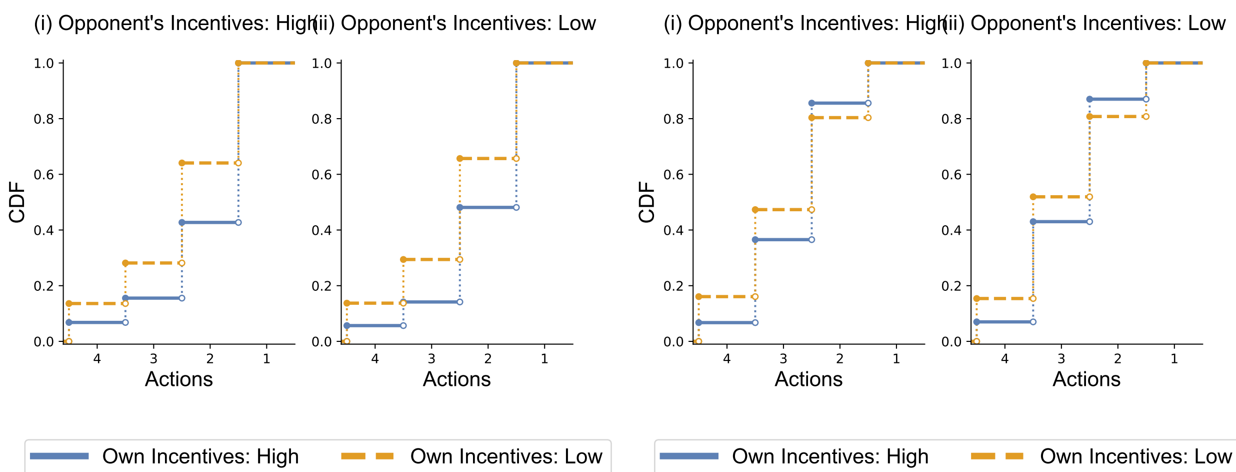

Figure 4. **Distribution of Gameplay according to ⊳**

Note: The actions are ranked according to ⊳.

As shown in Figure 4, gameplay shifts towards actions associated with higher levels of rationality in the simpler game but not in the more complex one. In 2-IDS game there is a first-order stochastic dominance shift towards actions with lower indices, that is, actions ranked higher according to ⊳. However, in the 3-IDS game, this distributional shift is not verified due to the mentioned decrease in the dominance solution. In the 3-IDS game, action 2 — and not the dominance solution, action 1 — is played more often.

This result is especially surprising in the light of recent evidence on similar games, stressing the difference between manipulating the overall incentives, the incentive *level*, and *distorting* the incentive structure, that is, scaling up payoffs as opposed to changing the relative incentives of one alternative over another. A recent paper by Alaoui et al. (2020) finds that, in a dominance-solvable game with a similar structure to diagonal games, where players

Table 3. **Difference in Dominance Solution Gameplay**

(a) **Own Incentives: High vs. Low**

IDS	Opponent's Incentives	Difference in Choices		Difference in Beliefs	
2	Pooled	19.42	(4.80)	3.78	(1.57)
3	Pooled	-5.72	(3.63)	-2.24	(1.33)
2	High	21.36	(6.79)	6.24	(2.33)
3	High	-5.22	(5.12)	-3.53	(1.72)
2	Low	17.57	(6.76)	1.39	(2.08)
3	Low	-6.23	(5.14)	-0.91	(2.03)

(b) **Opp. Incentives: High vs. Low**

IDS	Own Incentives	Difference in Beliefs	
2	High	5.08	(2.34)
2	Low	0.23	(2.06)
3	High	-2.88	(1.92)
3	Low	-0.25	(1.84)

Note: Standard errors in parentheses.

are motivated to choose actions with lower indices than the opponents,[18] *distorting* the incentives associated with undercutting the opponent's choice does lead to a distributional shift in gameplay towards actions associated with higher levels of rationality. The surprising element is that, in a simpler game, increasing the incentive *level* does not lead to such shift. This shift does occur in our 2-IDS game, as the distributional shift towards actions ranked higher according to $\triangleright$ implies a shift towards actions associated with higher levels of rationality. In contrast, in the 3-IDS game, this distributional shift is not verified due to the mentioned decrease in the dominance solution. This is surprising, especially considering that in the 3-IDS game, players require fewer iterations to reach the dominance solution (3 vs. 9) and have fewer actions (4 vs. 10) when compared to the game in Alaoui et al. (2020).

A simple explanation for why higher incentives lead to the dominance-solution action being chosen less frequently is that higher incentives lead to gameplay being less noisy. This hypothesis is supported by the data.[19] As we explore in the following section, this is closely related to subjects tending to choose actions with higher subjective expected payoffs more often.

Of interest is also the fact that, on average, higher own incentives lead to higher payoffs. That is, subjects do objectively better in terms of game payoffs — which translate to the probability with which they obtain the prize — in the high own-incentive treatment. The

[18]The game is a variant of the 11-20 game mentioned where players get payoffs equal to the chosen number, from 11 to 20, and get a small bonus if they choose the same number and a large bonus if they undercut their opponent by one. In our games, the equivalent distortion would be to increase the payoffs associated with (s_n, s_{n+1}) for $n = 1, 2, 3$.

[19]Table 12 in the appendix shows that indeed under higher own incentives gameplay is less uniform.

difference is statistically significant ($p<.05$) and reflects an increase in average payoffs of about 10% in the simpler game and 6% in the more complex game.[20]

Finally, we note that belief reports are significantly correlated to actions chosen and that the average beliefs track the change in gameplay induced by changes in the incentive level.[21] But, if gameplay is clearly associated with belief reports, the incentive level that subjects face significantly affects the relation between choices and beliefs, namely whether people best-respond to their beliefs or not.

4.2. Implementation Mistakes

While beliefs are, as expected, correlated with choices, not only are both affected by incentives but also choices are better responses to beliefs under higher own incentives. This manifests itself in a variety of indicators that we discuss in this section, from the increase in best-response rate to belief reports to changes in the probability of choosing not the best action — according to the subject's beliefs — but also the second best, third best and worst one. Furthermore, we find that a prominent element in determining subjective best-response rate is how much effort subjects put in as measured by response times. Finally, we discuss evidence suggesting that changes in beliefs track changes in gameplay induced by the higher opponent's incentives.

In line with the basic intuition of control costs and quantal response, higher incentives lead to fewer implementation mistakes and, therefore, to higher subjective best-response rates. This is exactly what one can observe in Table 4: increasing the incentive level leads to an increase in the rate at which players choose the action which yields the highest expected payoff according to their reported beliefs. The increase is not only statistically significant ($p<.05$) as it is pronounced, exceeding 10 pp. across all treatments. This implies that the incentive level affects the subjective best-response rate and, as such, existing estimates of best-response rates (e.g. Costa-Gomes and Weizsäcker 2008; Rey-Biel 2009) have to be interpreted in light of the incentives that subjects faced. Note that under low incentives, the rate of subjective best responses (≈ 45%) is lower than the one in Costa-Gomes and Weizsäcker (2008) and Rey-Biel (2009) (≈ 60-70%), which could be due to the fact that

[20]The statistical test can be found in the online appendix, Table 19.
[21]See Tables 20 and 21 in the online appendix.

Table 4. **Subjective Best-Response (BR) Rates**

IDS	Opponent's Incentives	Own Incentives High	Own Incentives Low	Difference	
2	High	66.02	42.72	23.30	(6.75)
2	Low	64.15	46.08	18.07	(6.79)
3	High	63.46	46.43	17.03	(6.67)
3	Low	59.00	45.19	13.81	(6.93)

Note: Standard errors in parentheses.

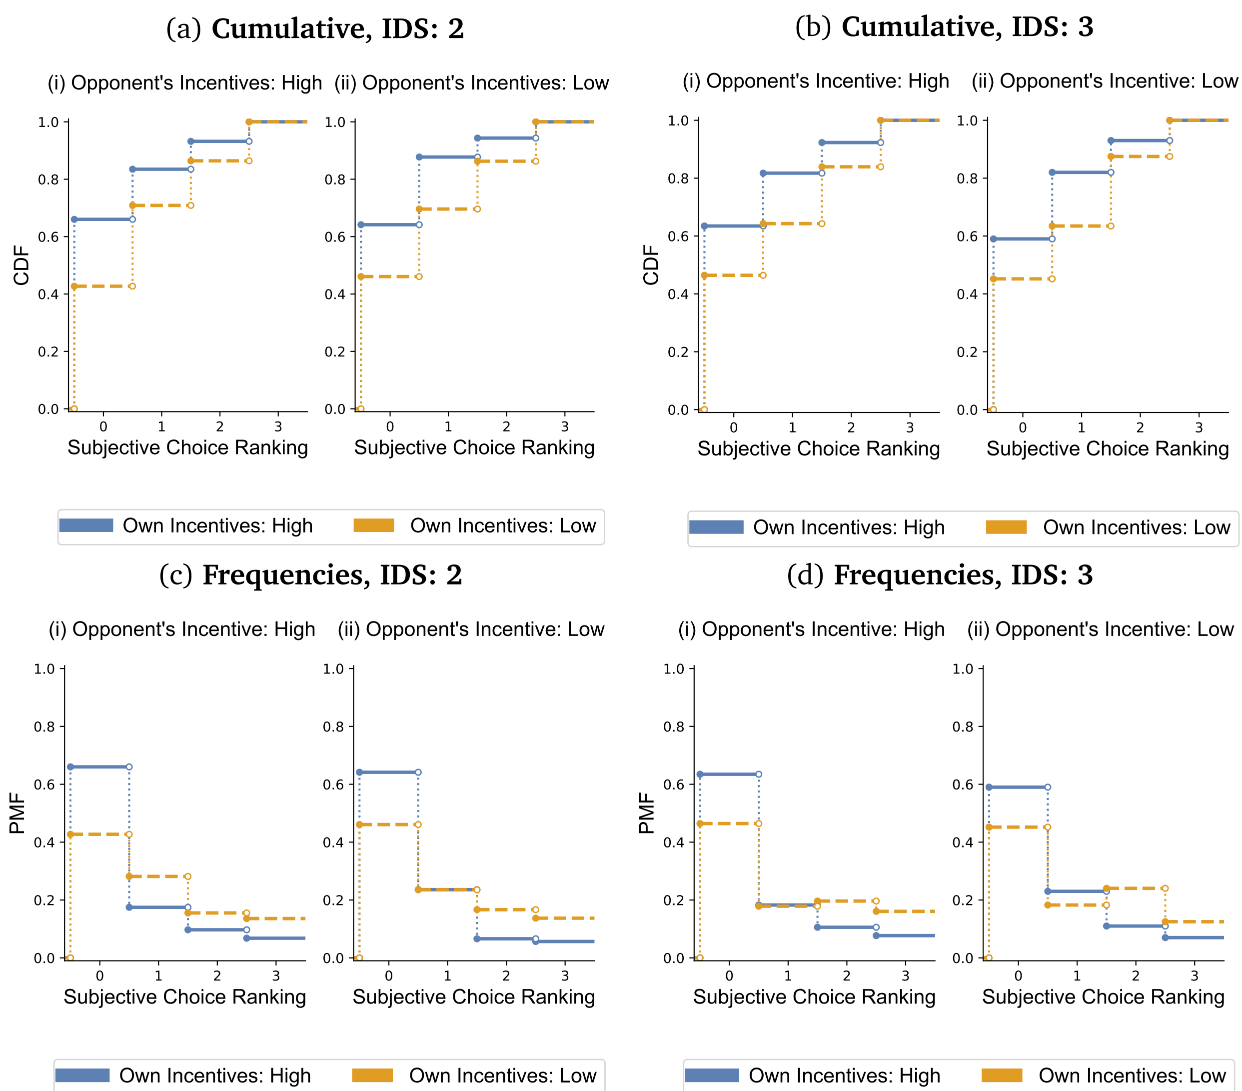

Figure 5. **Distribution of Subjective Choice Ranking**

Note: The choices are ranked according to subjective expected payoff given the individual belief reports, with 0 denoting the event that the subject chose the action with highest subjective expected payoff and 3 that with the lowest.

this papers use 3×3 games instead of 4×4. In contrast, the best-response rate under high own incentives exhibits similar values, which implies subjects better respond to their beliefs, once one accounts for the fact they face a greater number of alternative actions. There is no statistically significant effect of the opponent's incentive level or game complexity in subjective best-response rates.

The effect of own incentives on implementation mistakes goes beyond the fact that subjects choose their (subjectively) best alternative more often with higher incentives: in fact, there is a distributional shift towards actions with higher subjective payoff. In Figures 5a and 5b it is possible to observe a clear first-order stochastic dominance shift towards subjectively better actions — statistically significant at 5% for the 3-IDS case by means of a Wald test. Moreover, action frequencies are ranked as in rank-dependent choice equilibrium (Goeree et al. 2019), a generalization of regular quantal response equilibrium (Goeree et al. 2005): the action with the highest expected payoff is played the most often, the one with second highest expected payoff is played the second most often and so on. A significant difference with respect to this solution concept is that this holds when considering the *subjective* expected payoff and not the *objective* one — where the expectation is taken with respect to actual gameplay frequencies. In fact, actions frequencies are not ranked according to the objective choice ranking — see Table 11 in the appendix — and the postulate of rank-dependence of action frequencies is rejected[22] when considering the objective expected payoff but not the subjective one.

A related finding is that reported beliefs are better predictors of gameplay than the observed gameplay frequencies. On average, across all treatments, the subjective best-response rate is 12.35 pp. higher than the objective best-response rate. This is also robust to comparing each treatment separately.[23] These findings suggest that if the effect of implementation mistakes is in line with costly implementation models as quantal response, it is crucial to take into account that implementation mistakes are with respect to one's own beliefs and do not refer to the actual observed distribution of actions.

That higher incentives lead to fewer implementation mistakes is also visible when considering the distribution of payoff loss, that is, difference between the (subjective) expected

[22] The frequency of the objectively third best action in the 3-IDS game is statistically significantly higher than the frequency of the second best at a 5% significance level; see panel 11b(i) of Figure 11 in the appendix.
[23] See Table 13 in the appendix.

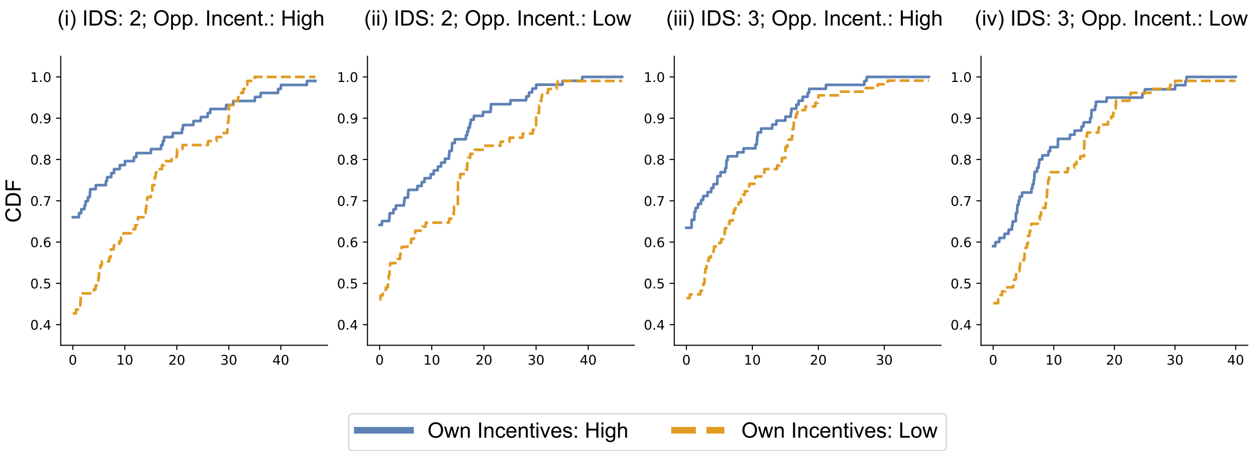

Figure 6. **Distribution of Implementation Loss**

Note: Implementation loss denotes the difference between the maximum subjective expected payoff attainable and the subjective expected payoff associated with the action chosen by the subject.

Table 5. **Determinants of Subjective Best-Response Rates and Implementation Loss**

Dep. Variable:	Subj. BR		Implementation Loss	
	(1)	(2)	(3)	(4)
Log RT	10.33	10.40	-1.56	-1.59
	(2.21)	(2.24)	(0.38)	(0.38)
Own Incent.: High	13.38	13.63	-2.11	-2.17
	(3.56)	(3.56)	(0.67)	(0.67)
Opp. Incent.: High	-0.06	-0.49	0.21	0.29
	(3.38)	(3.37)	(0.66)	(0.66)
IDS: 3	-1.44	-1.48	-2.23	-2.22
	(3.37)	(3.36)	(0.65)	(0.65)
Intercept	6.13	4.08	14.66	14.9
	(8.98)	(10.56)	(1.66)	(2.01)
Controls	No	Yes	No	Yes
N	834	834	834	834
R^2	0.06	0.07	0.05	0.06

Note: Subj. BR denotes a binary variable taking the value 1 whenever the subject best-responds to her beliefs. Implementation loss denotes the difference between the maximum subjective expected payoff attainable and the subjective expected payoff associated with the action chosen by the subject. Controls refer to socio-demographic data collected, namely age, sex, education and prior exposure to game theory. Heteroskedasticity-robust standard errors in parentheses.

payoff attainable by choosing the best action and the expected payoff associated with the action chosen. As it can be seen in Figure 6, there is a statistically significant distributional shift,[24] with higher own incentives leading to lower implementation losses. The distributions of losses when contrasting opponent's high and low incentives are not significantly different.

We explore the relationship between best response rates and incentives as well as other features of game relying on linear regressions (Table 5). The analysis confirms that indeed higher own incentives increase the best-response rate and decrease the implementation loss, but it also reveals that response times are a relevant factor in explaining both these variables. The coefficients associated with these two variables are positive and statistically significant ($p < .001$) when considering their relation to the subjective best-response rate and negative when considering instead their effect on implementation loss. Moreover, further robustness checks[25] show that these are fairly stable, suggesting that higher own incentives decrease implementation mistakes even controlling for the effort, as proxied by response times, that individuals commit to the decision problem at hand.

4.3. Changes in Beliefs

In the previous sections we establish that incentives affect beliefs, actions and how these are related. In contrast, in this section we examine the ways in which beliefs are affected and point towards mechanisms that could underlie these effects. We show that higher incentives lead to significant changes in beliefs. Beliefs become more responsive, as measured by their distance to the uniform distribution, and their bias also changes. As opposed to what occurs with actions and implementation mistakes, which were orthogonal to treatments other than whether own incentives were high or low, belief patterns depend not only on own incentive level but also the opponent's and the game's strategic complexity. We then propose a mechanism underlying these changes. In particular, we argue that incentives causally affect the effort committed by the individuals and that, in turn, it is effort and not the incentive level directly, that accounts for changes in belief formation.

[24]The first-order stochastic dominance ranking is statistically significant at a 5% level, according to a two-sample Kolmogorov-Smirnov test. See Table 15.
[25]Included in the online appendix: Tables 22 and 23.

We calculate the distance (L2 norm) of belief reports from the uniform distribution and show that this distance increases, in a distributional sense (an FOSD shift). As can be seen in panel 6a of Table 6, this shift is only statistically significant when the opponents have high incentives.[26] This suggests that when opponents face low incentives, subjects believe that their choices are more random, closer to the uniform.

Panel 6b confirms this conjecture: when the opponent's incentives increase, subjects report beliefs farther away from the uniform, although only when the subjects themselves have high incentives. Moreover, this change in beliefs is also associated with a change in gameplay: an increase in the opponent's incentives leads to a decrease by more than 10 pp. in the fraction of subjects whose best-response is the level-1 action (significant at 1%). Game complexity (panel 6c), has no significant effect on belief responsiveness.

Our second main result in this section is that higher own incentives lead to lower beliefs bias (FOSD shift and lower mean), but only when the opponent has high incentives as well. As is possible to observe in panel 7a of Table 7,[27] when the opponent has low incentives, increasing own incentives has no significant effect in the simpler game and has a pernicious effect in the more complex setting, as it results in more biased beliefs and indicates overthinking.

Belief bias is lower when opponents face lower incentives and in more complex games (panels 7b and 7c). When opponents have low incentives subjects believe they are going to choose more randomly, closer to the uniform distribution, which they do (Table 12b). Similarly, greater complexity of the environment would suggest more random gameplay and, again, this is also the case (Table 12c).

If beliefs do change with incentives, this begets the question: why? We suggest a clear mechanism: higher incentives result in more effort committed by subjects which in turn accounts for the changes in beliefs.

[26] Although the uniform distribution is also a mechanical focal point for belief reports, the fact that only when opponents have high incentives are belief reports more distant from the uniform, we then conclude that these patterns are reflective of actual changes in beliefs, beyond possible mechanical influence. We included the figures with the distributions in the online appendix — see Figure 15.

[27] The appendix includes a figure with the distributions — see Figure 16. A test for difference in mean belief bias can be found in the online appendix, Table 26.

Table 6. **Testing for FOSD in Belief Distance from Uniform**

(a) **Own Incentives: High vs. Low**

IDS	Opponent's Incentives	KS Statistic $>_{FOSD}$		KS Statistic $<_{FOSD}$	
2	High	0.27	[<.001]	0.03	[.917]
2	Low	0.11	[.250]	0.02	[.958]
3	High	0.19	[.016]	0.03	[.864]
3	Low	0.11	[.246]	0.02	[.933]

(b) **Opponent's Incentives: High vs. Low**

IDS	Own Incentives	KS Statistic $>_{FOSD}$		KS Statistic $<_{FOSD}$	
2	High	0.22	[.005]	0.02	[.952]
2	Low	0.05	[.779]	0.04	[.781]
3	High	0.16	[.059]	0.07	[.556]
3	Low	0.11	[.274]	0.05	[.700]

(c) **IDS: 2 vs. 3**

Own Incentives	Opponent's Incentives	KS Statistic $>_{FOSD}$		KS Statistic $<_{FOSD}$	
High	High	0.11	[.281]	0.04	[.795]
Low	High	0.04	[.845]	0.14	[.112]
High	Low	0.05	[.700]	0.10	[.326]
Low	Low	0.03	[.905]	0.07	[.584]

Note: The metric is the L2 norm. The statistical significance is given by a two-sample implementation of the Kolmogorov-Smirnov test for first-order stochastic dominance ranking; p-values are given in squared brackets.

Table 7. **Testing for FOSD in Belief Bias**

(a) **Own Incentives: High vs. Low**

IDS	Opponent's Incentives	KS Statistic $>_{FOSD}$		KS Statistic $<_{FOSD}$	
2	High	0.03	[.917]	0.24	[.002]
2	Low	0.01	[.986]	0.12	[.219]
3	High	0.04	[.811]	0.16	[.049]
3	Low	0.18	[.033]	0.05	[.734]

(b) **Opponent's Incentives: High vs. Low**

IDS	Own Incentives	KS Statistic $>_{FOSD}$		KS Statistic $<_{FOSD}$	
2	High	0.39	[<.001]	0.02	[.930]
2	Low	0.51	[<.001]	0.00	[.000]
3	High	0.27	[<.001]	0.06	[.653]
3	Low	0.40	[<.001]	0.01	[.987]

(c) **IDS: 2 vs. 3**

Own Incentives	Opponent's Incentives	KS Statistic $>_{FOSD}$		KS Statistic $<_{FOSD}$	
High	High	0.25	[.001]	0.09	[.420]
Low	High	0.29	[<.001]	0.01	[.989]
High	Low	0.03	[.872]	0.14	[.115]
Low	Low	0.15	[.096]	0.03	[.902]

Note: Belief bias is given by the distance between belief reports and the actual gameplay. The metric is the L2 norm. The statistical significance is given by a two-sample implementation of the Kolmogorov-Smirnov test for first-order stochastic dominance ranking; p-values are given in squared brackets.

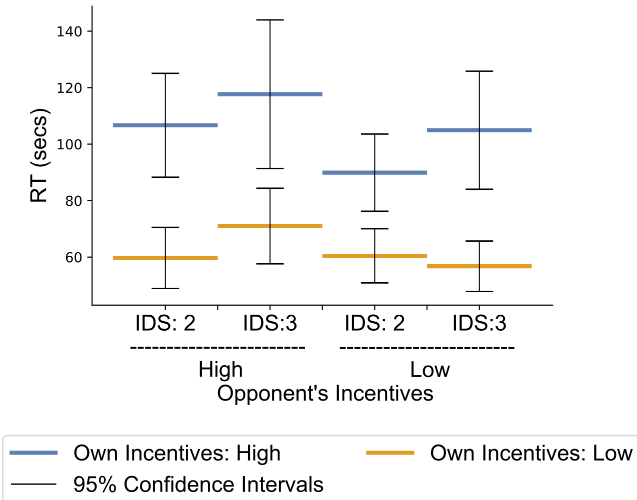

Note: Bootstrapped confidence intervals.

Figure 7. **Average Response Times**

Table 8. **Determinants of Response Times (RT)**

	Dep. Var.: Log RT (1)	(2)
Own Incent.: High	0.45	0.44
	(0.05)	(0.05)
Opp. Incent.: High	0.11	0.11
	(0.05)	(0.05)
IDS: 3	0.03	0.03
	(0.05)	(0.05)
Intercept	3.78	3.32
	(0.05)	(0.11)
Controls	No	Yes
N	834	834
R^2	0.08	0.11

Note: Controls refer to socio-demographic data collected, namely age, sex, education and prior exposure to game theory. Heteroskedasticity-robust standard errors are given in parentheses. Results are robust to using response times in levels.

Figure 7 shows that response times, our proxy for effort, increase significantly with higher own incentives. This is further confirmed by regression analysis (Table 8) showing that higher opponent's incentives result in longer response times ($p < .001$). Furthermore, when we control for response time, the statistical significance of the effects of own incentives on belief responsiveness and belief bias vanishes (Table 9), while response time becomes a leading explanatory factor ($p < .001$). This provides suggestive evidence that the effect of own incentives on beliefs is indirect. Own incentives affect effort and it is effort that relates to changes in beliefs. Furthermore, if people only put in more effort when incentives are high and this is generally understood, then it is reasonable that beliefs do not react if the opponent's incentives are low.

A closer look at column (4) in Table 9 shows that more effort leads to more accurate beliefs only when the opponent has high incentives. When the opponent has low incentives, more effort has no relation with belief bias in the simple game. In the more complex one, more effort actually results in an increase in belief bias ($p < .05$), which highlights that subjects may be overthinking. We therefore conclude that more effort does translate into more responsive beliefs, but not necessarily into more accurate ones.

This stands in contrast with the results in Alos-Ferrer and Buckenmaier (2019) who find that distorting incentives in a dominance-solvable game favoring playing actions consistent with a greater number of iterations of deletion of dominated strategies[28] leads not only subjects to play actions consistent with a higher level of reasoning but also having shorter response times. Again, we observe that a different pattern emerges once one scales up the overall level of incentives: the dominance solution may be played *less often* and, consistent with models of sequential sampling, response times *increase*.

4.4. Implementation Mistakes and Belief Formation

In this section we delve further into how incentives affect actions and beliefs, exploring the extent to which the change in gameplay induced by increasing own incentives is due to implementation mistakes as opposed to belief formation. If it is straightforward to observe that changes in beliefs cannot, by themselves, account for the decrease in dominated gameplay, it is less clear whether all change in gameplay can be explained by implementation mistakes.

On the one hand, because of the properties of the diagonal games,[29] we can classify subjects into sophistication/level-k types, following Nagel (1995). Looking at the choices alone, Figure 8 shows that incentives shift the distribution of inferred cognitive levels towards higher levels of rationality as given by the level-k model. This suggests that indeed individuals shift towards better actions and become more sophisticated in their beliefs.[30] On the other hand, once we take into account belief data, Figure 9 offers a different story. There, we show the frequency with which each action is the best-response to subjects' beliefs, contrasting the case where own incentives are high vs. low. It is striking that while in the 2-IDS game, how often each action is a best-response changes dramatically,[31] in the more complex game, the changes are small and not statistically significant. If higher own incentives lead

[28]This distortion is akin to that in Alaoui et al. (2020) discussed above and is performed on the same variant of the 11-20 game.
[29]See Section 2.2 for details.
[30]Interestingly, the distributions of inferred cognitive levels are not significantly different when contrasting either opponent's incentive treatments or strategic complexity — see Figures 12 and 13 in the appendix.
[31]A Wald test reveals that the changes are statistically significant at 5% level.

Table 9. **Determinants of Changes in Beliefs**

Dep. Var.:	Dist. to Uniform		Belief Bias	
	(1)	(2)	(3)	(4)
Log RT	–	7.69	–	0.74
		(0.64)		(0.62)
Log RT×Own Incent.: High	–	–	–	-0.53
				(0.63)
Log RT×Opp. Incent.: High	–	–	–	-3.01
				(0.60)
Log RT×IDS: 3	–	–	–	1.39
				(0.60)
Own Incent.: High	4.09	0.74	-1.09	-0.85
	(1.12)	(1.09)	(0.51)	(0.53)
Opp. Incent.: High	1.83	0.99	5.41	5.37
	(1.13)	(1.06)	(0.51)	(0.52)
IDS: 3	0.23	0.02	-1.54	-1.45
	(1.13)	(1.05)	(0.51)	(0.51)
Intercept	18.35	24.31	12.96	12.81
	(2.61)	(2.58)	(1.18)	(1.20)
Controls	Yes	Yes	Yes	Yes
N	834	834	834	834
R^2	0.03	0.16	0.15	0.18

Note: Belief bias is given by the distance between belief reports and the actual gameplay. The metric is the L2 norm. Controls refer to socio-demographic data collected, namely age, sex, education and prior exposure to game theory. Heteroskedasticity-robust standard errors are given in parentheses. Results are robust to using response times in levels.

to better responses as we have previously established, in the simpler game the change in beliefs affects which actions are the best-responses.

We propose a parametric decomposition of the relative importance of implementation mistakes and changes in beliefs based on logit quantal response. The logit quantal response model is a convenient one-parameter representation of implementation mistakes (McKelvey and Palfrey 1995). As the parameter multiplies the utility from obtaining the prize, it is not possible to identify these separately and we estimate two noise parameters by maximum likelihood, one for the high own incentives and another for low own incentives. Differently from quantal response equilibrium, we rely not on the actual distribution of gameplay to

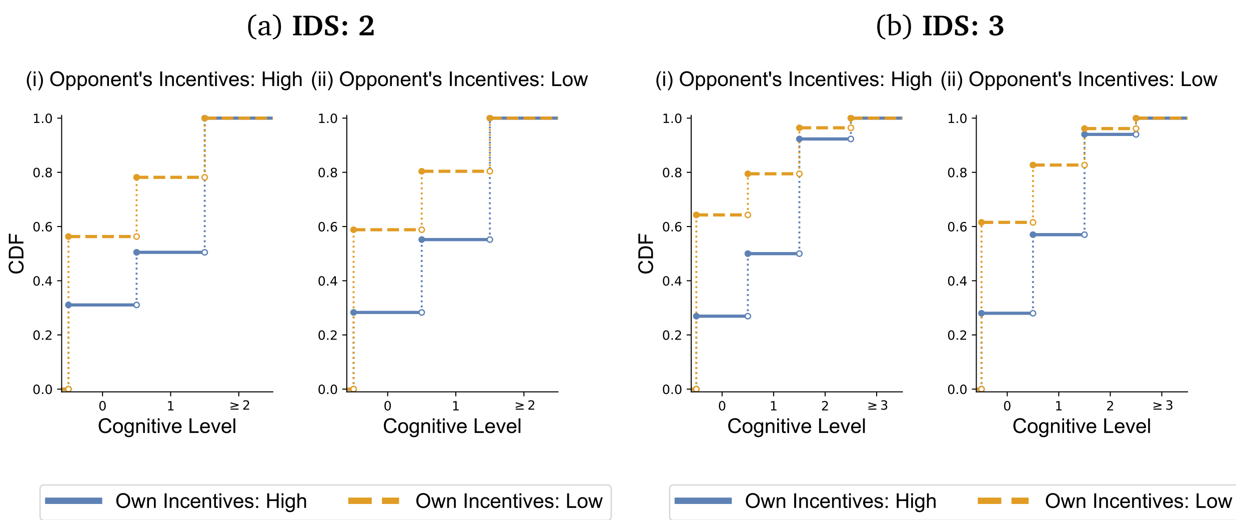

Figure 8. **Distribution of Inferred Cognitive Levels**

Note: The actions are ranked according to the level-k prediction assuming level 0 uniformly randomizes across the available actions (and adjusted accordingly).

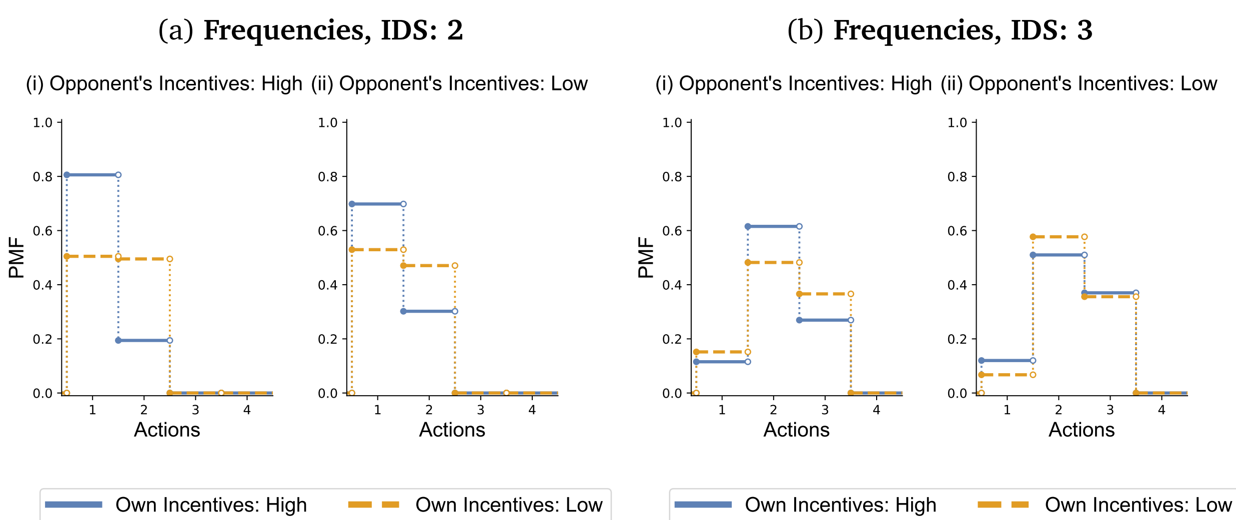

Figure 9. **Distribution of Best-Responses**

Note: This figure shows, for each action, the fraction of subjects for whom that action is a best-response to their beliefs.

compute the expected payoffs associated with each action but use instead the reported beliefs. This is in line with the fact that subjective beliefs and not objective beliefs are a better predictor of behavior (Section 4.2). As expected (see Table 16), the estimate for the noise parameter is statistically significantly higher (at 5%) under high own incentives than under low incentives, reflecting a higher utility associated with the larger prize ($20 vs. $0.50).

Figure 10 illustrates the resulting predictions of average gameplay, which are fairly close to the observed gameplay.[32]

We proceed with a counterfactual exercise to understand the relative importance of the change in the noise parameter, capturing implementation mistakes, as opposed to the change in beliefs in explaining gameplay. Formally, let $\ell(\hat{\sigma}_i, \lambda_j | \overline{\sigma}_k)$ denote the loss — the squared error in prediction of average gameplay $\overline{\sigma}_k$ — obtained from relying on the logit quantal response fit with beliefs $\hat{\sigma}_i$ and noise parameter λ_j, where $i, j, k \in \{H, L\}$. When $k = i$ we say that beliefs match gameplay and that they mismatch when otherwise; the same applies to the noise parameter. For $i, j \in \{H, L\}$, we then define $\Delta\ell_{\hat{\sigma}}(\overline{\sigma}_i, \lambda_j) = \ell(\hat{\sigma}_i, \lambda_j \mid \overline{\sigma}_i) - \ell(\hat{\sigma}_j, \lambda_j \mid \overline{\sigma}_i)$ as the decrease in loss obtained from relying on "correcting" for the use of "wrong" beliefs; $\Delta\ell_{\lambda}(\overline{\sigma}_i, \hat{\sigma}_j)$ is defined analogously. Then, the share of the overall decrease in loss attributable to beliefs when gameplay is $\overline{\sigma}_i$ is given by

$$\rho_{\hat{\sigma}}(\overline{\sigma}_i) := \frac{\Delta\ell_{\hat{\sigma}}(\overline{\sigma}_i, \lambda_H) + \Delta\ell_{\hat{\sigma}}(\overline{\sigma}_i, \lambda_L)}{\Delta\ell_{\hat{\sigma}}(\overline{\sigma}_i, \lambda_H) + \Delta\ell_{\hat{\sigma}}(\overline{\sigma}_i, \lambda_L) + \Delta\ell_{\lambda}(\overline{\sigma}_i, \hat{\sigma}_H) + \Delta\ell_{\lambda}(\overline{\sigma}_i, \hat{\sigma}_L)}$$

and that attributable to the noise parameter is given by $\rho_{\lambda}(\overline{\sigma}_i)$.

Finally, we provide a second counterfactual exercise, now aiming at understanding the relative importance of implementation mistakes and belief formation not in explaining gameplay but instead in explaining the *change* in gameplay induced by higher own incentives. Specifically, we compare the squared prediction error of the change in average gameplay to those we obtain when fixing either the noise parameter or beliefs to a specific treatment.

Table 10 summarizes the results obtained. The decomposition in 10a confirms what Figure 9 already suggested: While implementation mistakes are overwhelmingly responsible for predicting gameplay in the more complex 3-IDS game, both channels — implementation mistakes and belief formation — are important in accounting for gameplay in the simpler 2-IDS. In particular, the improvement in predicting gameplay from using the correct beliefs is fairly stable across the two games when own incentives are low: around 30-35%. However, there is a sharp difference in the contribution of beliefs to explaining gameplay when own incentives are high when comparing the 2-IDS and the 3-IDS games. While in the simpler game, both belief formation and implementation mistakes account for roughly the same

[32] Table 17 in the appendix provides additional details.

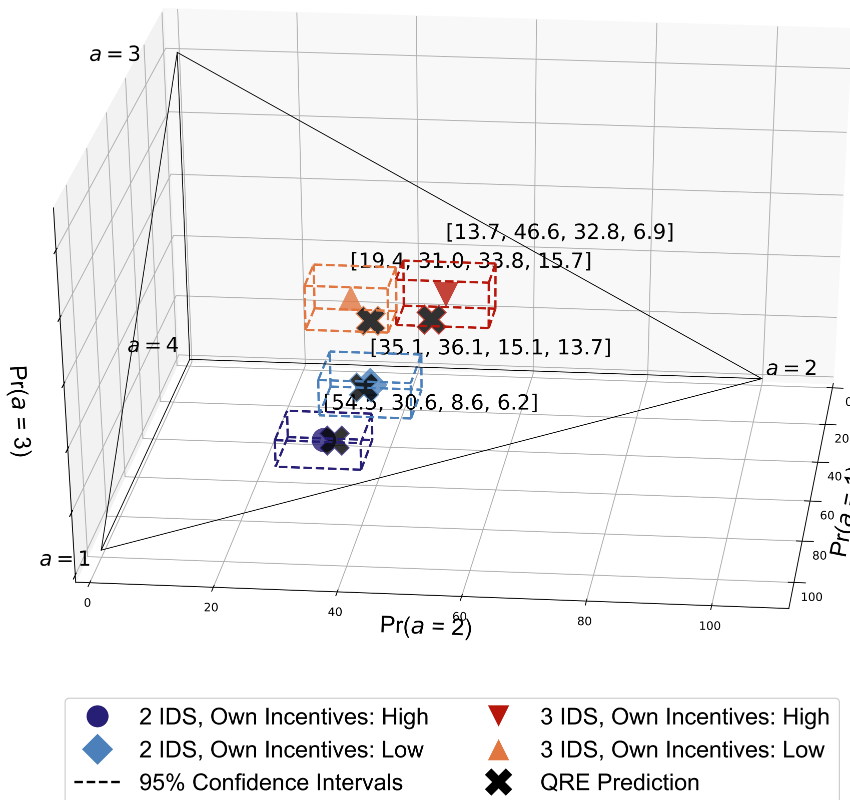

Figure 10. **QRE with Subjective Beliefs**

Table 10. **Relative Importance of Belief Formation vs. Implementation Mistakes**

(a) **Squared Error in Predicting Gameplay**

IDS	Own Incentives	Error with $(\lambda_L, \hat{\sigma}_L)$ $\ell(\lambda_L, \hat{\sigma}_L \vert \bullet)$	Error with $(\lambda_H, \hat{\sigma}_H)$ $\ell(\lambda_H, \hat{\sigma}_H \vert \bullet)$	% Improvement due to ΔBeliefs $\rho_{\hat{\sigma}}$	Noise ρ_λ
2	High	326.53	16.81	50.3	49.7
2	Low	49.65	582.58	35.4	64.6
3	High	288.21	63.92	15.9	84.1
3	Low	54.21	289.58	30.5	69.5

(b) **Squared Error in Predicting Changes in Gameplay**

IDS	Error	Error with Δ Beliefs Fixing Noise		Error with Δ Noise Fixing Beliefs	
		$\Delta\,\hat{\sigma} \mid \lambda_H$	$\Delta\,\hat{\sigma} \mid \lambda_L$	$\Delta\,\lambda \mid \hat{\sigma}_H$	$\Delta\,\lambda \mid \hat{\sigma}_L$
2	4.64	14.01	17.28	10.35	15.82
3	6.67	15.37	16.82	8.82	10.65

improvement in predictive accuracy, in the more complex game the predictive accuracy is mostly due to utilizing the noise parameter matching the treatment condition.

As to what concerns explaining the changes in gameplay induced by varying the own incentive level, we observe a similar phenomenon. In the simpler game, the squared error in predicting the change in gameplay obtained by allowing beliefs to adjust but not the noise parameter, or vice-versa, is always more than 2 times higher when compared to that when we allow both beliefs and the noise parameter to adjust. This suggests that, in simpler environments, both belief formation and implementation mistakes are significantly important channels in explaining how changing incentives affects gameplay. In the more complex environment, the prediction error increases by 30%-60% when allowing the noise parameter to adjust but fixing beliefs. This stands in sharp contrast with the 130-150% increase in prediction that we obtains when we fix the noise parameter and allow beliefs to adjust. Again this points to the fact that when the environment is complex, changes in gameplay seem to be mostly due implementation mistakes.

5. Conclusion

This paper provides evidence that incentives matter in strategic settings.

To do so we rely three main features of the experimental design. First, we propose a modification to the standard use of observers that allows us to change the incentives of a subject while keeping gameplay fixed. Second, we rely on a class of games where it is possible to vary the iterations to the dominance solution while changing the payoffs minimally. Third, we vary incentives by scaling up the prize by a factor of 20 and an effective payoff difference of $19.50.

We find that changing the stakes of a game does lead to significantly different behavior in both action choices and stated beliefs. Moreover, if higher incentives do decrease the frequency of dominated choices, it is not necessarily the case that they increase the frequency with which the dominance solution is played. Whereas in the simple games increasing incentives does lead to an increase in the frequency of play of the dominance solution, in the complex game it decreases. However, this is because, overall, gameplay distribution becomes less random under high incentives.

Finally, we find that effort — through response times — is the driving force behind changes in beliefs. When incentives increase, subjects put in more time to understand the game, and their reported beliefs change and become less close to the uniform. Furthermore, when the opponent's incentives are also high —- and is thus expected to put in more effort — beliefs also become more accurate.

Eliciting beliefs allows us to explore the potential mechanisms underlying the effects of incentives. Our results indicate that two mechanisms are at play when incentives increase: (i) implementation mistakes decrease, leading to better responses to beliefs, and (ii) individuals commit more effort, resulting in changes in beliefs and ultimately in different choices. If more effort leads to are more responsive beliefs, these are not necessarily more accurate, depending crucially on the opponent's incentive level and on environment complexity.

In sum, both implementation mistakes and belief formation explain how incentives affect gameplay. Furthermore, the complexity of the environment emerged as the main factor determining which of these two channels is more important: While in simpler games, both channels are equally important, in more complex games, the decrease in implementation mistakes is the predominant factor.

References

Agranov, M. and P. Ortoleva (2017). Stochastic choice and preferences for randomization. *Journal of Political Economy 125*(1), 40–68.

Alaoui, L., K. A. Janezic, and A. Penta (2020). Reasoning about others' reasoning. *Journal of Economic Theory 189*, 1–51.

Alaoui, L. and A. Penta (2016). Endogenous Depth of Reasoning. *Review of Economic Studies 83*(4), 1297–1333.

Alaoui, L. and A. Penta (2018). Cost-benefit analysis in reasoning. *Working Paper*.

Alos-Ferrer, C. and J. Buckenmaier (2019). Cognitive sophistication and deliberation times. *Working Paper*.

Andersen, S., S. Ertaç, U. Gneezy, M. Hoffman, and J. A. List (2011). Stakes matter in ultimatum games. *American Economic Review 101*(7), 3427–3439.

Arad, A. and A. Rubinstein (2012). The 11-20 money request game: A level-k reasoning study. *American Economic Review 102*(7), 3561–3573.

Brocas, I., J. D. Carrillo, S. W. Wang, and C. Camerer (2014). Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games. *Review of Economic Studies 81*(3), 944–970.

Camerer, C., T.-H. Ho, and J.-K. Chong (2004). A Cognitive Hierarchy Model of Games. *Quarterly Journal of Economics 119*(3), 861–898.

Camerer, C. F. and R. M. Hogarth (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty 19*, 7–42.

Caplin, A. and M. Dean (2013). Behavioral Implications of Rational Inattention with Shannon Entropy. *NBER Working Paper* (August), 1–40.

Caplin, A. and M. Dean (2015). Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review 105*(7), 2183–2203.

Carpenter, J., E. Verhoogen, and S. Burks (2005). The effect of stakes in distribution experiments. *Economics Letters 86*(3), 393–398.

Chen, D. L., M. Schonger, and C. Wickens (2016). oTree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.

Costa-Gomes, M. A. and V. P. Crawford (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review 96*(5), 1737–1768.

Costa-Gomes, M. A. and G. Weizsäcker (2008). Stated Beliefs and Play in Normal Form Games. *Review of Economic Studies 75*(3), 729–762.

Dean, M. and N. Neligh (2019). Experimental tests of rational inattention. *Working Paper*.

Dewan, A. and N. Neligh (2020). Estimating information cost functions in models of rational inattention. *Journal of Economic Theory 187*, 105011.

Friedman, E. and J. Ward (2019). Stochastic choice and noisy beliefs in games: an experiment. *Mimeo*.

Fudenberg, D., R. Iijima, and T. Strzalecki (2015). Stochastic Choice and Revealed Perturbed Utility. *Econometrica 83*(6), 2371–2409.

Fudenberg, D. and A. Liang (2019). Predicting and understanding initial play. *American Economic Review 109*(12), 4112–4141.

Georganas, S., P. J.Healy, and R. A.Weber (2015). On the persistence of strategic sophistication. *Journal of Economic Theory 159*, 369–400.

Goeree, J. K. and C. A. Holt (2001). Ten Little Treasures of Game Theory and Ten Intuitive Contradictions. *American Economic Review 91*(5), 1402–1422.

Goeree, J. K., C. A. Holt, P. Louis, T. R. Palfrey, and B. Rogers (2019). Rank-dependent choice equilibrium: a non-parametric generalization of QRE. In S (Ed.), *Handbook of Research Methods and Applications in Experimental Economics*, pp. 252–268. Edward Elgar Publishing.

Goeree, J. K., C. A. Holt, and T. R. Palfrey (2005). Regular Quantal Response Equilibrium. *Experimental Economics 8*(4), 347–367.

Gonçalves, D. (2020). Diagonal games: A tool for experiments and theory. *Working Paper*.

Grosskopf, B. and R. Nagel (2008). The two-person beauty contest. *Games and Economic Behavior 62*(1), 93–99.

Gächter, S. and E. Renner (2010). The effects of (incentivized) belief elicitation in public goods experiments. *Experimental Economics 13*(3), 364–377.

Heinemann, F., R. Nagel, and P. Ockenfels (2009). Measuring strategic uncertainty in coordination games. *Review of Economic Studies 76*(1), 181–221.

Hoffman, E., K. A. McCabe, and V. L. Smith (1996). On expectations and the monetary stakes in ultimatum games. *International Journal of Game Theory 25*, 289–301.

Holt, C. A. and A. M. Smith (2016). Belief elicitation with a synchronized lottery choice menu that is invariant to risk attitudes. *American Economic Journal: Microeconomics 8*(1), 110–139.

Hossain, T. and R. Okui (2013). The binarized scoring rule. *Review of Economic Studies 80*(3), 984–1001.

Huck, S. and G. Weizsäcker (2002). Do players correctly estimate what others do? Evidence of conservatism in beliefs. *Journal of Economic Behavior and Organization 47*(1), 71–85.

Kneeland, T. (2015). Identifying Higher-Order Rationality. *Econometrica 83*(5), 2065–2079.

Matějka, F. and A. McKay (2015). Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review 105*(1), 272–298.

Mattsson, L.-G. and J. W. Weibull (2002). Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior 41*(1), 61–78.

McKelvey, R. D. and T. R. Palfrey (1995). Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior 10*(1), 6–38.

McKelvey, R. D., T. R. Palfrey, and R. A. Weber (2000, aug). The effects of payoff magnitude and heterogeneity on behavior in 2×2 games with unique mixed strategy equilibria. *Journal of Economic Behavior & Organization 42*(4), 523–548.

Montiel Olea, J. L. and M. Plagborg-Møller (2018). Simultaneous confidence bands: Theory, implementation, and an application to SVARs. *Journal of Applied Econometrics 34*(1), 1–17.

Nagel, R. (1995). Unraveling in Guessing Games: An Experimental Study. *American Economic Review 85*(5), 1313–1326.

Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics 54*(2), 286–295.

Pulford, B. D., A. M. Colman, and G. Loomes (2018). Incentive magnitude effects in experimental games: Bigger is not necessarily better. *Games 9*(1), 1–10.

Rey-Biel, P. (2009). Equilibrium play and best response to (stated) beliefs in normal form games. *Games and Economic Behavior 65*(2), 572–585.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665–690.

Slonim, R. and A. E. Roth (1998). Learning in high stakes ultimatum games: An experiment in the slovak republic. *Econometrica 66*(3), 569–596.

Stahl, D. O. and P. W. Wilson (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization 25*(3), 309–327.

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

Székely, G. J., M. L. Rizzo, and N. K. Bakirov (2007). Measuring and testing dependence by correlation of distances. *The Annals of Statistics 35*(6), 2769–2794.

Appendix.. Supporting Tables and Figures

Table 11. **Choice Frequencies, Average Beliefs and Average Response Times (RT)**

	Incentives		Choice Frequencies				Average Beliefs					
IDS	Own	Opponent's	s_1	s_2	s_3	s_4	s_1	s_2	s_3	s_4	RT (secs)	Obs
2	High	High	57.28	27.18	8.74	6.80	33.01	35.65	18.92	12.42	106.68	103
2	Low	High	35.92	35.92	14.56	13.59	26.77	31.69	23.01	18.53	59.72	103
2	High	Low	51.89	33.96	8.49	5.66	27.94	36.91	20.15	15.01	89.91	106
2	Low	Low	34.31	36.27	15.69	13.73	26.55	31.52	23.83	18.10	60.45	102
3	High	High	14.42	49.04	29.81	6.73	18.40	31.05	36.35	14.19	117.68	104
3	Low	High	19.64	33.04	31.25	16.07	21.94	27.49	32.02	18.55	71.00	112
3	High	Low	13.00	44.00	36.00	7.00	21.28	25.89	33.96	18.87	104.93	100
3	Low	Low	19.23	28.85	36.54	15.38	22.18	27.04	32.73	18.04	56.75	104

Table 12. **Gameplay Distance to the Uniform Distribution**

(a) **Own Incentives: High vs. Low**

	Opponent's	Own Incentives			
IDS	Incentives	High	Low	Difference	
2	High	20.27	10.93	9.34	[.011]
2	Low	19.04	10.34	8.70	[.001]
3	High	16.18	7.28	8.89	[.001]
3	Low	15.41	8.27	7.14	[.001]

(b) **Overall Incentives: High vs. Low**

	Incentives			
IDS	High	Low	Difference	
2	20.27	10.34	9.92	[.010]
3	16.18	8.27	7.90	[.017]

(c) **IDS: 2 vs. 3**

Overall	IDS			
Incentives	2	3	Difference	
High	20.27	16.18	4.09	[.483]
Low	10.34	8.27	2.07	[.216]

Note: The distance to the uniform refers to the L2 norm between the average gameplay and the uniform distribution; p-values are obtained by bootstrapping and given in squared brackets.

Table 13. **Subjective vs. Objective Best-Response (BR) Rates**

	Incentives		BR Rate			
IDS	Own	Opponent's	Subjective	Objective	Difference	
2	High	High	66.02	57.28	8.74	(6.75)
2	Low	High	42.72	35.92	6.80	(6.79)
2	High	Low	64.15	51.89	12.26	(6.73)
2	Low	Low	46.08	34.31	11.76	(6.82)
3	High	High	63.46	49.04	14.42	(6.81)
3	Low	High	46.43	33.04	13.39	(6.48)
3	High	Low	59.00	44.00	15.00	(6.99)
3	Low	Low	45.19	28.85	16.35	(6.60)

Note: Standard errors in parentheses.

Table 14. **Testing for FOSD in Choice Ranking**

Own Incentives: High vs. Low				Opponent's Incentives: High vs. Low				IDS: 2 vs. 3			
	Opponent's	Wald			Own	Wald		Own	Opponent's	Wald	
IDS	Incentives	Statistic		IDS	Incentives	Statistic		Incentives	Incentives	Statistic	
2	High	11.99	[.007]	2	High	1.73	[.631]	High	High	0.17	[.983]
2	Low	11.32	[.010]	2	Low	0.59	[.899]	Low	High	3.52	[.318]
3	High	9.68	[.021]	3	High	0.74	[.864]	High	Low	1.45	[.694]
3	Low	9.49	[.023]	3	Low	0.94	[.815]	Low	Low	2.15	[.542]

Note: p-values in squared brackets.

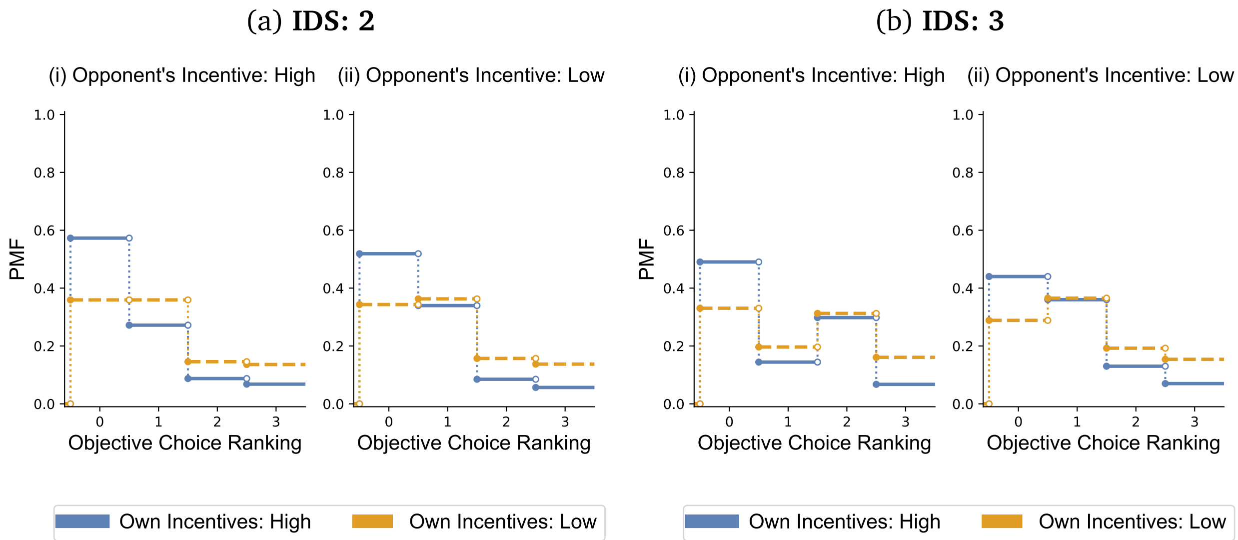

Figure 11. **Distribution of Objective Choice Ranking**

Note: The choices are ranked according to objective expected payoff given the gameplay frequencies, with 0 denoting the event that the subject chose the action with highest objective expected payoff and 3 that with the lowest.

Table 15. **Testing for FOSD in Subjective Payoff Loss**

(a) **Own Incentives: High vs. Low**

IDS	Opponent's Incentives	KS Statistic $>_{FOSD}$		KS Statistic $<_{FOSD}$	
2	High	0.06	[.706]	0.25	[.001]
2	Low	0.01	[.987]	0.18	[.026]
3	High	0.00	[>.999]	0.22	[.004]
3	Low	0.01	[.970]	0.17	[.042]

(b) **Opponent's Incentives: High vs. Low**

IDS	Own Incentives	KS Statistic $>_{FOSD}$		KS Statistic $<_{FOSD}$	
2	High	0.05	[.722]	0.05	[.743]
2	Low	0.10	[.307]	0.04	[.812]
3	High	0.01	[.989]	0.09	[.422]
3	Low	0.04	[.834]	0.06	[.616]

(c) **IDS: 2 vs. 3**

Own Incentives	Opponent's Incentives	KS Statistic $>_{FOSD}$		KS Statistic $<_{FOSD}$	
High	High	0.12	[.222]	0.03	[.915]
Low	High	0.15	[.090]	0.01	[.987]
High	Low	0.08	[.511]	0.05	[.729]
Low	Low	0.14	[.113]	0.07	[.582]

Note: Implementation loss denotes the difference between the maximum subjective expected payoff attainable and the subjective expected payoff associated with the action chosen by the subject. The statistical significance is given by a two-sample implementation of the Kolmogorov-Smirnov test for first-order stochastic dominance ranking; p-values are given in squared brackets.

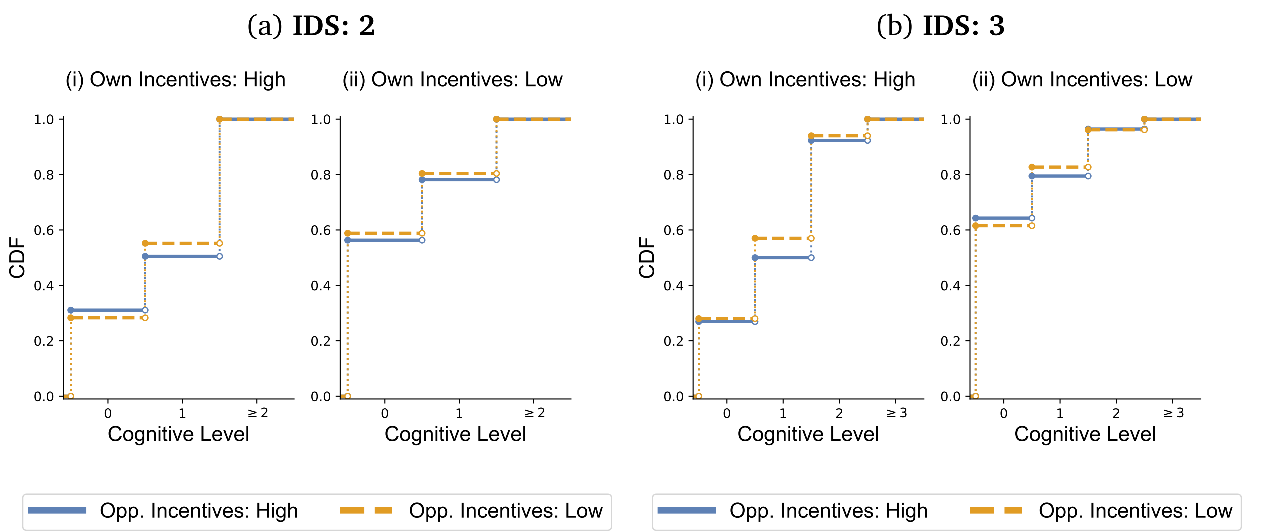

Figure 12. **Distribution of Inferred Cognitive Levels: Opponent's Incentives High vs. Low**

Note: The actions are ranked according to the level-k prediction assuming level 0 uniformly randomizes across the available actions (and adjusted accordingly).

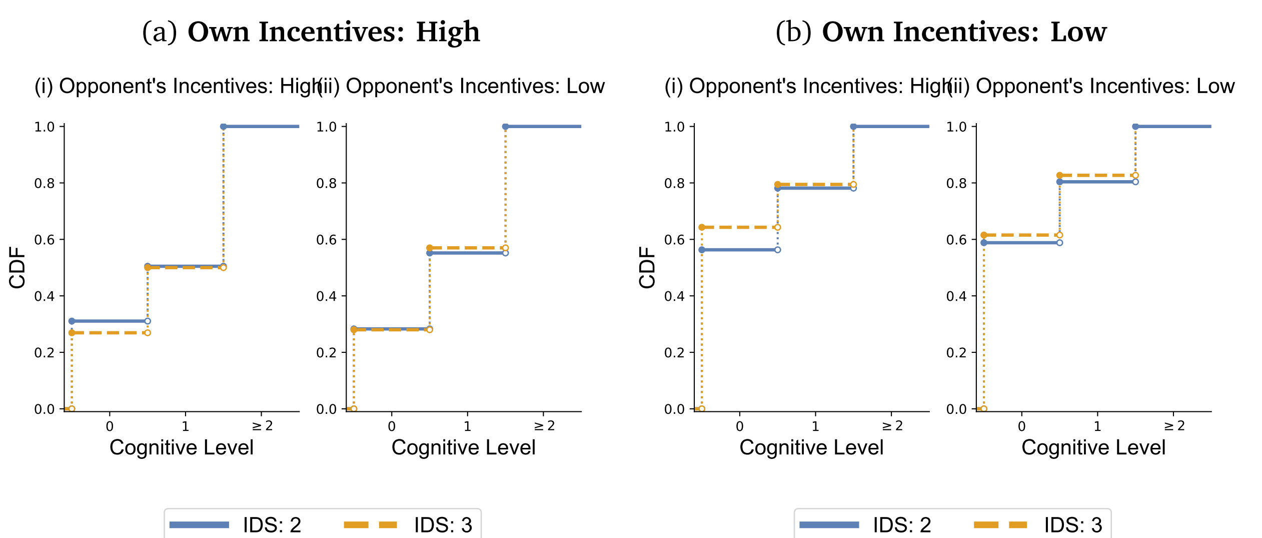

Figure 13. **Distribution of Inferred Cognitive Levels: IDS 2 vs. 3**

Note: The actions are ranked according to the level-k prediction assuming level 0 uniformly randomizes across the available actions (and adjusted accordingly).

Table 16. **Quantal Response Parameter**

Own's Incentives	QRE λ	95% Confidence Interval	
High	9.00	7.14	10.89
Low	5.05	3.49	6.64

Table 17. **Choice Frequencies vs. QRE Predictions**

	IDS	Own Incentives	Action 1	Action 2	Action 3	Action 4
Gameplay	2	High	54.55	30.62	8.61	6.22
QRE	2	High	55.52	32.54	9.17	2.77
Gameplay	2	Low	35.12	36.10	15.12	13.66
QRE	2	Low	39.35	35.55	16.80	8.30
Gameplay	3	High	13.73	46.57	32.84	6.86
QRE	3	High	19.68	44.78	27.88	7.67
Gameplay	3	Low	19.44	31.02	33.80	15.74
QRE	3	Low	22.66	34.74	28.61	13.98

For Online Publication

Online Appendix A.. Additional Tables and Figures

Table 18. **Testing for Treatment Effects**

(a) **Own Incentives: High vs. Low**

IDS	Opponent's Incentives	Choices Wald statistic		Beliefs Wald statistic		Response Times t-statistic	
2	High	10.54	[.032]	19.08	[<.001]	4.34	[<.001]
2	Low	9.66	[.047]	8.98	[.062]	3.48	[<.001]
3	High	9.02	[.061]	12.24	[.016]	3.11	[.002]
3	Low	8.04	[.090]	1.09	[.896]	4.18	[<.001]

(b) **IDS: 2 vs. 3**

Incentives		Choices		Beliefs		Response Times	
Own	Opponent's	Wald statistic		Wald statistic		t-statistic	
High	High	57.03	[<.001]	98.54	[<.001]	-0.68	[.500]
Low	High	13.16	[.011]	21.52	[<.001]	-1.29	[.197]
High	Low	52.37	[<.001]	66.41	[<.001]	-1.19	[.236]
Low	Low	14.95	[.005]	18.35	[.001]	0.56	[.578]

(c) **Opponent's Incentives: High vs. Low**

IDS	Own Incentives	Choices Wald statistic		Beliefs Wald statistic		Response Times t-statistic	
2	High	1.19	[.880]	6.31	[.177]	1.44	[.149]
2	Low	0.08	[.999]	0.41	[.982]	-0.10	[.920]
3	High	0.93	[.920]	14.66	[.005]	0.75	[.455]
3	Low	0.74	[.947]	0.21	[.995]	1.74	[.081]

Note: p-values in squared brackets.

Table 19. **Average Payoffs Obtained**

IDS	Opponent's Incentives	Own Incentives High	Own Incentives Low	Difference	
2	High	42.74	38.35	4.39	(1.47)
2	Low	47.44	43.30	4.14	(1.37)
3	High	40.36	37.99	2.36	(1.02)
3	Low	39.22	36.83	2.39	(0.85)

Note: Standard errors in parentheses.

Table 20. **Distance Correlation**

IDS	Incentives Own	Incentives Opponent's	Test Statistic	
2	High	High	2.22	[<.001]
2	Low	High	0.97	[<.001]
2	High	Low	0.86	[<.001]
2	Low	Low	1.52	[<.001]
3	High	High	2.31	[<.001]
3	Low	High	1.24	[<.001]
3	High	Low	1.37	[<.001]
3	Low	Low	1.06	[<.001]

Note: Distance correlation measures arbitrary dependence between two random vectors (Székely et al. 2007). The tests of independence rely on a permutation test with 10,000 repetitions on the difference; p-values in squared brackets.

Table 21. **Correlation between Change in Gameplay and Beliefs due to different Incentives**

			s_1	s_2	s_3	s_4	Correlation
IDS: 2	Δ Gameplay		22.97	-9.09	-6.95	-6.93	
IDS: 2	Δ Beliefs	Own Incent.: High	5.08	-1.26	-1.23	-2.59	.97
IDS: 2	Δ Beliefs	Own Incent.: Low	0.23	0.16	-0.81	0.42	.25
IDS: 3	Δ Gameplay		-4.81	20.19	-6.73	-8.65	
IDS: 3	Δ Beliefs	Own Incent.: High	-2.88	5.16	2.39	-4.67	.77
IDS: 3	Δ Beliefs	Own Incent.: Low	-0.25	0.45	-0.71	0.51	.44

Note: This table presents the change in gameplay frequencies induced by going from low to high incentives (for all players) and the change in average beliefs from changing the opponent's incentives. The right-most column shows the (Pearson) correlation of the change in beliefs with the associated change in gameplay.

Table 22. **Determinants of Subjective Best-Response Rates**

Dep. Variable:	Subj. BR			
	(1)	(2)	(3)	(4)
Log RT	10.33	10.40	–	12.74
	(2.21)	(2.24)		(2.14)
Own Incent.: High	13.38	13.63	18.16	–
	(3.56)	(3.56)	(3.42)	
Opp. Incent.: High	-0.06	-0.49	0.65	-0.82
	(3.38)	(3.37)	(3.41)	(3.40)
IDS: 3	-1.44	-1.48	-1.19	-1.78
	(3.37)	(3.36)	(3.41)	(3.39)
Intercept	6.13	4.08	38.60	1.78
	(8.98)	(10.56)	(7.52)	(10.61)
Controls	No	Yes	Yes	Yes
N	834	834	834	834
R^2	0.06	0.07	0.04	0.05

Note: Subj. BR denotes a binary variable taking the value 1 whenever the subject best-responds to her beliefs. Controls refer to socio-demographic data collected, namely age, sex, education and prior exposure to game theory. Heteroskedasticity-robust standard errors are given in parentheses.

Table 23. **Determinants of Implementation Loss**

Dep. Variable:	Implementation Loss			
	(1)	(2)	(3)	(4)
Log RT	-1.56	-1.59	–	-1.96
	(0.38)	(0.38)		(0.37)
Own Incent.: High	-2.11	-2.17	-2.87	–
	(0.67)	(0.67)	(0.66)	
Opp. Incent.: High	0.21	0.29	0.12	0.35
	(0.66)	(0.66)	(0.66)	(0.66)
IDS: 3	-2.23	-2.22	-2.26	-2.17
	(0.65)	(0.65)	(0.66)	(0.66)
Intercept	14.66	14.98	9.71	15.35
	(1.66)	(2.01)	(1.51)	(2.02)
Controls	No	Yes	Yes	Yes
N	834	834	834	834
R^2	0.05	0.06	0.04	0.05

Note: Implementation loss denotes the difference between the maximum subjective expected payoff attainable and the subjective expected payoff associated with the action chosen by the subject. Controls refer to socio-demographic data collected, namely age, sex, education and prior exposure to game theory. Heteroskedasticity-robust standard errors are given in parentheses.

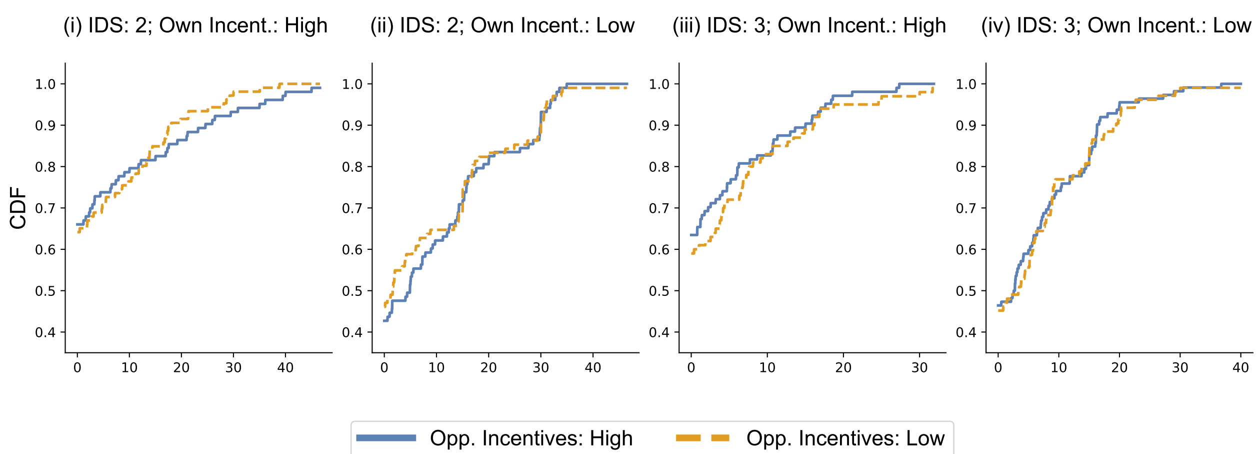

Figure 14. **Distribution of Implementation Loss: Opponent's Incentives**

Note: Implementation loss denotes the difference between the maximum subjective expected payoff attainable and the subjective expected payoff associated with the action chosen by the subject.

Table 24. **Testing for Difference in Average Subjective Payoff Loss**

Own Incentives: High vs. Low				Opponent's Incentives: High vs. Low				IDS: 2 vs. 3			
IDS	Opponent's Incentives	Difference in Means		IDS	Own Incentives	Difference in Means		Own Incentives	Opponent's Incentives	Difference in Means	
2	High	-3.32	(1.59)	2	High	0.78	(1.43)	High	High	2.50	(1.29)
2	Low	-3.50	(1.45)	2	Low	0.60	(1.60)	Low	High	3.31	(1.35)
3	High	-2.51	(0.98)	3	High	-0.89	(0.98)	High	Low	0.84	(1.16)
3	Low	-1.96	(1.11)	3	Low	-0.33	(1.11)	Low	Low	2.38	(1.42)

Note: Standard errors are given in parentheses.

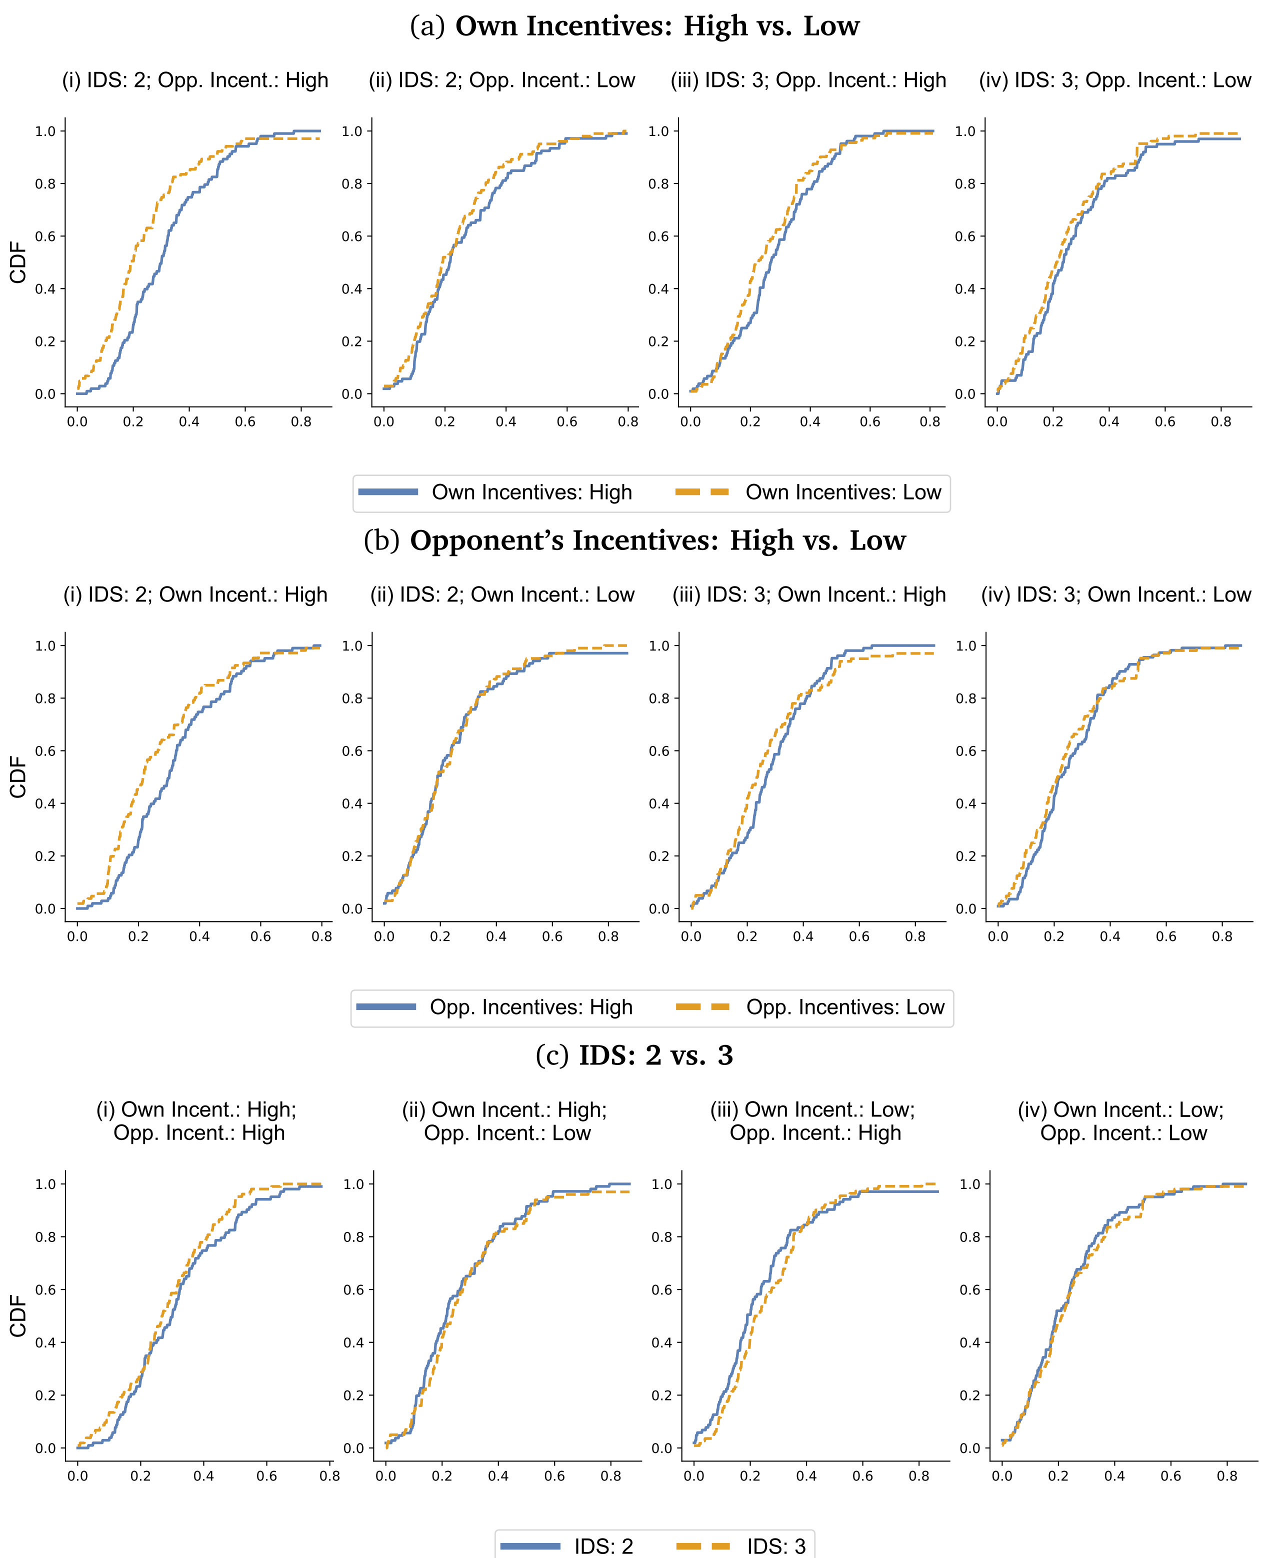

Figure 15. **Distribution of Belief Distance to Uniform**

Note: The metric is the L2 norm.

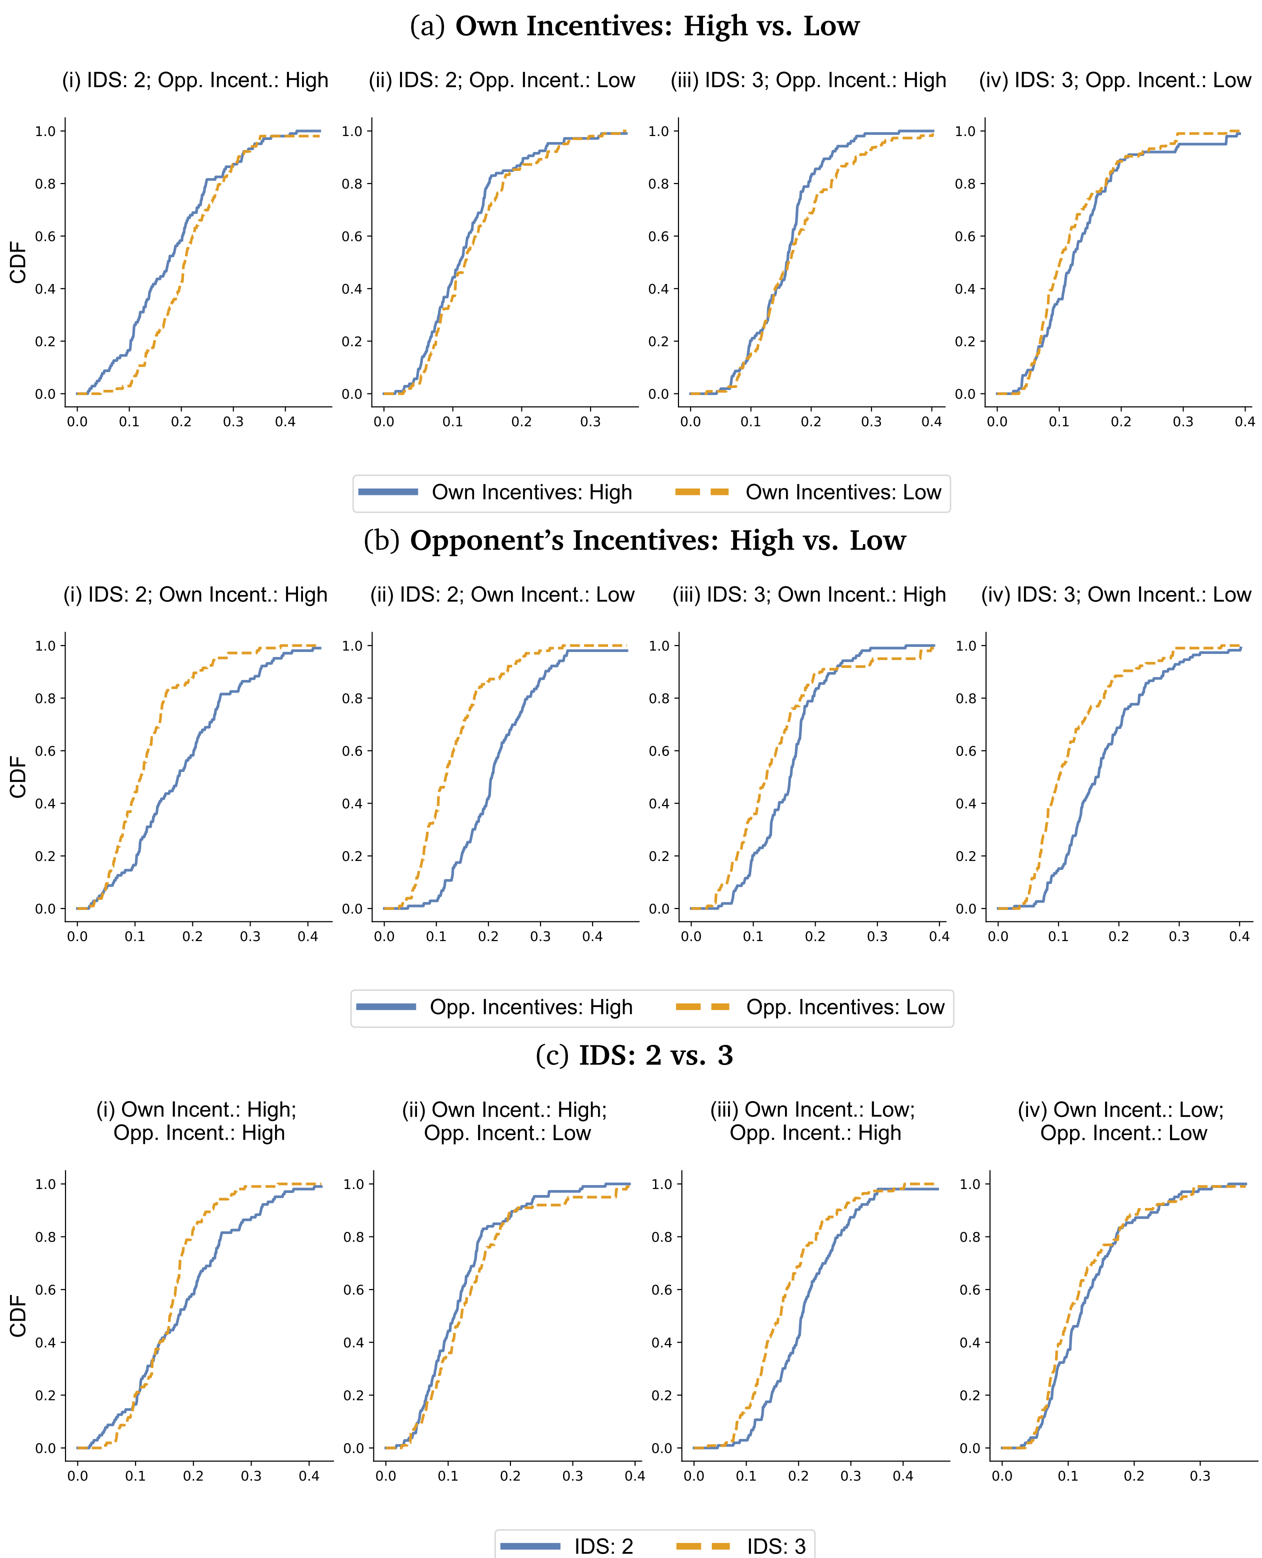

Figure 16. **Distribution of Belief Bias**

Note: Belief bias is given by the distance between belief reports and the actual gameplay. The metric is L2 norm.

Table 25. **Testing for Difference in Average Belief Distance from Uniform**

Own Incentives: High vs. Low				Opponent's Incentives: High vs. Low				IDS: 2 vs. 3			
IDS	Opponent's Incentives	Difference in Means		IDS	Own Incentives	Difference in Means		Own Incentives	Opponent's Incentives	Difference in Means	
2	High	7.55	(2.32)	2	High	5.45	(2.23)	High	High	3.37	(2.09)
2	Low	2.84	(2.23)	2	Low	0.74	(2.32)	Low	High	-1.79	(2.21)
3	High	2.39	(1.97)	3	High	0.60	(2.28)	High	Low	-1.48	(2.42)
3	Low	3.16	(2.40)	3	Low	1.37	(2.11)	Low	Low	-1.15	(2.21)

Note: The metric is the L2 norm. Standard errors are given in parentheses.

Table 26. **Testing for Difference in Average Belief Bias**

Own Incentives: High vs. Low				Opponent's Incentives: High vs. Low				IDS: 2 vs. 3			
IDS	Opponent's Incentives	Difference in Means		IDS	Own Incentives	Difference in Means		Own Incentives	Opponent's Incentives	Difference in Means	
2	High	-3.33	(1.17)	2	High	6.17	(1.10)	High	High	2.51	(1.06)
2	Low	-1.05	(0.90)	2	Low	8.45	(0.99)	Low	High	4.09	(1.03)
3	High	-1.75	(0.89)	3	High	2.11	(0.95)	High	Low	-1.54	(1.00)
3	Low	1.39	(1.02)	3	Low	5.25	(0.96)	Low	Low	0.90	(0.92)

Note: Belief bias is given by the distance between belief reports and the actual gameplay. Standard errors are given in parentheses.

Table 27. **Determinants of Belief Responsiveness**

	Dep. Var.: Belief Distance to Uniform				
	(1)	(2)	(3)	(4)	(5)
Log RT	7.86	9.26	–	7.69	9.07
	(0.61)	(1.28)		(0.64)	(1.35)
Log RT×Own Incent.: High	–	-0.39	–	–	-0.37
		(1.26)			(1.28)
Log RT×Opp. Incent.: High	–	-1.31	–	–	-1.30
		(1.22)			(1.24)
Log RT×IDS: 3	–	-0.97	–	–	-0.95
		(1.23)			(1.24)
Own Incent.: High	–	–	4.09	0.74	0.72
			(1.12)	(1.09)	(1.11)
Opp. Incent.: High	–	–	1.83	0.99	0.98
			(1.13)	(1.06)	(1.07)
IDS: 3	–	–	0.23	0.02	0.01
			(1.13)	(1.05)	(1.06)
Intercept	25.19	25.24	18.35	24.31	24.38
	(2.41)	(2.42)	(2.61)	(2.58)	(2.60)
Controls	Yes	Yes	Yes	Yes	Yes
N	834	834	834	834	834
R^2	0.16	0.16	0.03	0.16	0.16

Note: The metric is the L2 norm. Controls refer to socio-demographic data collected, namely age, sex, education and prior exposure to game theory. Heteroskedasticity-robust standard errors are given in parentheses.

Table 28. **Determinants of Belief Bias**

	Dep. Var.: Belief Bias				
	(1)	(2)	(3)	(4)	(5)
Log RT	-0.27	0.70	–	-0.32	0.74
	(0.33)	(0.64)		(0.32)	(0.62)
Log RT×Own Incent.: High	–	-0.30	–	–	-0.53
		(0.66)			(0.63)
Log RT×Opp. Incent.: High	–	-3.26	–	–	-3.01
		(0.64)			(0.60)
Log RT×IDS: 3	–	1.58	–	–	1.39
		(0.65)			(0.60)
Own Incent.: High	–	–	-1.09	-0.95	-0.85
			(0.51)	(0.52)	(0.53)
Opp. Incent.: High	–	–	5.41	5.45	5.37
			(0.51)	(0.52)	(0.52)
IDS: 3	–	–	-1.54	-1.53	-1.45
			(0.51)	(0.51)	(0.51)
Intercept	14.16	14.29	12.96	12.71	12.81
	(1.27)	(1.26)	(1.18)	(1.21)	(1.20)
Controls	Yes	Yes	Yes	Yes	Yes
N	834	834	834	834	834
R^2	0.02	0.05	0.15	0.15	0.18

Note: Belief bias is given by the distance between belief reports and the actual gameplay. The metric is the L2 norm. Controls refer to socio-demographic data collected, namely age, sex, education and prior exposure to game theory. Heteroskedasticity-robust standard errors are given in parentheses.

Online Appendix B.. Diagonal Games

We here briefly review the main definitions and general properties of diagonal games as they apply to the games used in this paper.

Definition 1 (Diagonal Game)**.** A diagonal game $\Gamma(f_1, f_2, N)$ is a normal-form, finite two-player game $\langle I, S, u\rangle$ where $I = \{1, 2\}$ denotes the set of players with general element i; $S = \{s_1, ..., s_N\}$ denotes the set of actions available to each player and $u = (u_1, u_2)$ with $u_i : S^2 \to \mathbb{R}$ where, for some function $f_i : \mathbb{N} \to \mathbb{R}$, $u_i(s_i, s_j) = f_i(N - i + j)$.

Note that, given the way the payoff function is specific, players' payoffs are constant along the parallels to the main diagonal as shown in the following figure:

		Player j		
		s_1	s_2	s_3
	s_1	$f_i(3)$	$f_i(4)$	$f_i(5)$
Player i	s_2	$f_i(2)$	$f_i(3)$	$f_i(4)$
	s_3	$f_i(1)$	$f_i(2)$	$f_i(3)$

Figure 17. **Player i's Payoffs**

In order to obtain symmetry, as desired, we will focus on the symmetric case where $f_1 = f_2 \equiv f$, writing $\Gamma(f, N)$ for the sake of brevity.

Dominance-solvability and natural manipulation of the number of steps of deletion needed to reach the dominance solution is immediately met by a simple assumption that this class of games imposes on the function f:

Assumption 1. (i) $h := \arg\max_{n\in\mathbb{N}} f(n) - N \geq 1$, (ii) f is single-peaked and (iii) f is strictly increasing up to its maximum.

From here, the following properties are derived:

Proposition 1 (Properties of Dominance-Solvable Diagonal Games (Gonçalves 2020))**.** Suppose Assumption 1 holds. Then the following properties on dominance-solvability hold:

1. If $h \geq N$, then $\Gamma(f, N)$ is solvable in exactly 1 step of simultaneous iterated deletion of strictly dominated strategies, i.e. it is solvable in strictly dominant strategies.

2. If $h \leq N-1$, then $\Gamma(f,N)$ is solvable in exactly $\lceil \frac{N-1}{h} \rceil =: T$ steps of simultaneous iterated deletion of strictly dominated strategies. In each step $t < T$, h actions are deleted.

Finding the dominance solution has a natural dynamics, similar to the beauty contest game (Nagel 1995): in the first round of deleted, the h actions with the highest index are found to be strictly dominated and deleted for both players; in the second round, the surviving h actions with the highest index are then strictly dominated and deleted and so on.

Moreover, this class of games leads to a clear ranking of the players' actions — a linear order on S — which we denote by $\rhd$. This linear order $\rhd$ is such that $s'_i \rhd s_i$ if and only if either s_i is deleted before s'_i or both actions are deleted at the same time, but one action strictly dominates the other for any of the opponent's actions that survived up to that iteration. In particular, actions are ranked according to their indices, that is, $s_n \rhd s_m$ if and only if $n \leq m$. This implies that even actions that are found to be (iteratedly) strictly dominated at a given round of deletion, they are ranked among themselves: every of the deleted actions strictly dominates the following. This contrasts with, for instance, the dominance-solvable version of the 11-20 game proposed by Alaoui et al. (2020), where only one action is deleted at a time and where modifications to allow more than one action to be deleted at each step would preclude being able to rank all actions.

Note that with this assumption alone, quantal response and, more generally, rank-dependent choice equilibrium has a clear testable implication: it predicts that the gameplay frequencies of $s_{N-h}, \ldots, s_N$ are monotonically decreasing. These models postulate that actions associated with higher expected payoffs are played with greater frequency. Then, given that actions are ranked according to strict dominance, for any beliefs the players might hold, the frequencies with which each of these strategies is played is also ranked in the same fashion by these models.

However, unless one is willing to attribute cardinal significance to the payoff values, the level-1 action cannot be generically identified without further assumptions. Adding a further refinement on f is sufficient to be able to identify the chain of both level-k and dominance-k[33] actions. This refinement is given by the following assumption:

Assumption 2. $\forall n > N+h,\ f(n) = \min_{m \in \mathbb{N}} f(m)$.

[33] A dominance-k player does k steps of deletion of strictly dominated strategies and then best-responds to a uniform prior over the remaining strategies. Cf. Costa-Gomes and Crawford (2006).

Relying on these two assumptions we are able to generate games that pin-down how level-k players choose:

Proposition 2 (Level-k Ranking of Strategies (Gonçalves 2020))**.** Suppose Assumptions 1 and 2 hold and let $h < N$. Then $\forall k = 1, \ldots, T-1$, $s_{N-k \cdot h}$ is the strategy corresponding to the level-k player and s_1 that of a level-T, where $T = \lceil \frac{N-1}{h} \rceil$, when a level-0 uniformly randomizes. The indices shift by one for the dominance-k levels, i.e. $s_{N-(k+1) \cdot h}$ corresponds to dominance-k players.

The payoff vector associated with the dominance solution, $(u_1(s_1, s_1), u_2(s_1, s_1)) = (f(N), f(N))$, is neither Pareto dominant nor Pareto dominated. As happens in the specific instances used in this paper, it holds more generally that it corresponds to the values along the main diagonal of the payoff matrix. This precludes this outcome from being used as a coordination or focal point as whenever both players choose the same action they achieve the equilibrium payoff. As we mentioned in Section 2.2, we believe this can help mitigate potential other-regarding concerns that may act as confounding factors when analyzing the effect of incentives and strategic complexity on gameplay.

Finally, note that, within this class of games, the exact payoffs are not pinned down, only their ordinal relations as these properties only depend on f ordinally: positive monotone transformations of f do not affect the properties listed above. Additional properties and proofs are collected in Gonçalves (2020).