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

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

Abstract

How and why do incentive levels affect strategic behavior? This paper examines an experiment designed to identify the causal effect of scaling up incentives on choices and beliefs in strategic settings by holding fixed opponents' actions. In dominance-solvable games, higher incentives increase action sophistication and best-response rates and decrease mistake propensity. Beliefs tend to become more accurate with higher own incentives in simple games. However, opponents with higher incentive levels are harder to predict: while beliefs track opponents' behavior when they have higher incentive levels, beliefs about opponents also become more biased. We provide evidence that incentives affect cognitive effort and that greater effort increases performance and predicts choice and belief sophistication. Overall, the data lends support to combining both payoff-dependent mistakes and costly reasoning.

Keywords: Game Theory; Level-k; Quantal Response; Incentives; Belief Formation; Response Times.

JEL Classifications: C72, C92, D83, D84, D91.

[*] Department of Political Economy, King's College London; teresa.estebancasanelles@kcl.ac.uk.

[†] Department of Economics, University College London; duarte.goncalves@ucl.ac.uk.

We thank Larbi Alaoui, Marina Agranov, Guy Aridor, Alessandra Casella, Kfir Elias, Evan Friedman, Navin Kartik, Rosemarie Nagel, Ryan Oprea, Antonio Penta, Jacopo Perego, Silvio Ravaioli, Evan Sadler, Karl Schlag, Leeat Yariv, Michael Woodford, and the participants at Columbia, SWEET-Princeton 2020, SAE Symposium 2020, EEA-ES Congress 2021, and Barcelona Summer Forum 2022 for valuable comments. We are particularly grateful to Yeon-Koo Che, Terri Kneeland, Mark Dean for their many insightful conversations. This material is based upon work supported by the National Science Foundation under Grant Number 1949395.

First posted draft: 27 February 2020. *This draft*: 2 April 2023.

1. Introduction

Deviations from Nash equilibrium and rationalizable behavior have been abundantly documented, both within and across games. For the former, within-game evidence suggests the existence of systematic differences to Nash equilibrium predictions, including nontrivial proportions of non-rationalizable and even dominated choices. For the latter, existing evidence mostly relates to how increasing *relative* incentives for a player to choose a given action increases its observed frequency.[1] The evidence has then brought to the fore a variety of new models of behavior in strategic settings in which players have a limited reasoning ability as in level-k (Nagel, 1995; Stahl and Wilson, 1995) and cognitive hierarchy (Camerer et al., 2004), make mistakes as in quantal response equilibrium (McKelvey and Palfrey, 1995), or face cognitive costs as in endogenous depth of reasoning models (Alaoui and Penta, 2016) and sequential sampling equilibrium (Gonçalves, 2022).

If the underlying rationale for observed behavior is that decision-making in strategic settings is cognitively challenging — people need to form an understanding of the environment, how others may act, and what to do — incentive levels can decisively impact strategic behavior. In non-strategic decision-making, for instance, higher incentive levels or stakes increase exerted effort and tend to lead to better choices in simple problems (Dean and Neligh, 2022; Caplin et al., 2020).[2] In strategic settings, what makes a choice better or worse depends on others' behavior, which renders the effect of incentives on behavior ambiguous. Existing models suggest two mechanisms through which incentive levels may operate on behavior. On the one hand, when one's own incentive level is higher, mistakes such as choosing dominated actions are more costly and it is reasonable to expect them to be less frequent, in line with quantal response equilibrium. On the other hand, higher own incentive levels increase the marginal benefit to exert cognitive effort, which may entail forming different beliefs about others' behavior and thus choose differently — the main mechanism operating in endogenous depth of reasoning and sequential sampling equilibrium. Moreover, it is not only one's own incentive level that matters: when opponents' incentive level is higher, one may expect opponents' to make fewer mistakes and be more strategically sophisticated, and then react accordingly.

[1] See, among many others, McKelvey and Palfrey (1992), Goeree and Holt (2001) and Costa-Gomes and Weizsäcker (2008). Camerer (2003) presents on early overview of experimental evidence.

[2] The evidence is nuanced though. Recent evidence by Enke et al. (2021) shows that higher incentive levels does increase response time — a proxy for exerted effort — but performance increases only in simple enough tasks.

Existing literature, however, lacks clearly identified evidence of both of the effects of incentive levels on strategic behavior and the underlying mechanisms. Note how, even if scaling up a player's incentive level would, all else constant, lead to better choices, more effort, and more accurate beliefs, opponents may naturally react to incentive levels — either because opponents' incentive level is also higher, or because they anticipate the change in the player's behavior. Then, even when observing that increasing overall incentive levels (scaling up payoffs) results in greater action sophistication (e.g. McKelvey and Palfrey, 1992; Rapoport et al., 2003; Camerer, 2003, ch. 1.3, 5.2), it is still unclear what changes in behavior are due to a player's own incentive level, their opponents' incentive level, and which mechanisms underlie the changes in strategic behavior.

This paper experimentally examines the causal effect of incentive levels on both choices and beliefs in strategic settings. Focusing on dominance solvable games, we show that higher own incentive levels lead to more sophisticated behavior (lower rate of dominated choices) and a higher best-response rate to reported beliefs. Furthermore, with higher own incentive levels, reported beliefs also denote greater sophistication: opponents are expected to play more sophisticated actions and beliefs are less biased. Finally, higher incentive levels significantly increase response time, which is in turn we find to be positively associated with expected payoffs and both action and belief sophistication.

Our design separately identifies the effect of own and opponents' incentive levels on strategic behavior. For this, players are assigned to treatments specifying whether the incentive level they and their opponents' face are both high, low, or one is high and the other is low. We use as baseline opponent choices from groups in which both own and opponent incentive treatment are the same. Subjects with a given incentive level are then asked to play against a randomly selected opponent from the baseline treatment group with either the high or low incentive level, similar to an observer method (Huck and Weizsäcker, 2002; Alaoui et al., 2020). This allows us to freely vary a subjects' own incentives level while holding fixed their opponents' choices and incentive level, as well as resolving all uncertainty about the incentives their opponents face.[3] We vary the incentive level by randomly assigning subjects to have either a high or a low bonus, which they can increase the probability of earning with the choices the make in the experiment.

We focus on two two-player dominance-solvable games in which actions are linearly ordered

[3]Specifically, this simple setup not only resolves subjects' beliefs about their opponents' incentive level, it also informs them of the incentive level of their opponents' opponents and so on.

by iterated dominance. One game is arguably 'simpler' than the other in that it takes fewer steps of iterated elimination of strictly dominated actions to reach to the dominance solution. In both games, level-k actions are identified in a manner that is robust to uncertainty aversion. The games share a number of features with standard dominance-solvable games (such as 11-20, undercutting games, and ring games), which enables us to relate our results to the existing literature. In focusing on initial responses and to keep incentives per game high-powered, we provide unincentivized two practice rounds without feedback and a single incentivized round. We then obtain a 2×2×2 between-subject design with random assignment, in which over 800 subjects recruited online play only one game and only once.

Our first set of results show that higher own incentive level increases the sophistication of play and decreases mistake propensity. Specifically, the frequency of dominated actions is halved with higher own incentive levels and that of the level 2 action increases. However, if in the simpler game, higher own incentives do increase the frequency of the dominance solution (the level 2 action), this slightly decreases in the more complicated game, for which the dominance solution corresponds to the level 3 action. The data suggests this is due to lower mistake propensity and choices becoming less random: in line with intuition from stochastic choice models and quantal response, higher own incentive levels induce subjects to best respond to their beliefs more often and their choices are more responsive to differences in expected payoffs as calculated using stated beliefs.

We then consider if and how subjects' beliefs react to incentive levels. We find that in the simpler game — but not in the more complicated one — higher *own* incentives entail assigning lower probability to opponents with high incentives choosing dominated actions, and to players forming a more accurate estimate of the expected payoffs to the different actions. Moreover, while beliefs on the choices of low-incentive opponents are closer to the uniform and thus easier to predict, it is harder with high-incentive opponents. Consequently, reported beliefs on the frequency of dominance play and of dominated choices are more biased, that is, farther away from realized frequencies. Similarly, expected payoffs implied by reported beliefs are also more biased relative to empirical expected payoffs when opponents have higher incentive levels. This suggests that the effect of own and opponent incentive levels operate in opposite directions on belief accuracy even in a dominance-solvable game.

Our third set of results pertains to the effect of incentive levels on cognitive effort and how it relates to strategic behavior. By increasing the marginal benefit to reasoning, own incentive levels may affect beliefs (and therefore choices), as predicted by models in which players engage in a cost-benefit analysis of reasoning such as in endogenous depth of reasoning or in sequential sampling equilibrium. Similarly to existing literature (Caplin et al., 2020; Alós-Ferrer and Buckenmaier, 2021; Frydman and Nunnari, 2023, e.g.), we proxy for cognitive effort by examining response time data and find that higher own incentive levels increase effort by more than 40%. Additionally, controlling for incentives and individual characteristics, response times are associated with performance — for instance, a 40% increase in response time increases the likelihood of best responding to the empirical frequency of opponent behavior by 5.2 percentage points. Finally, we explore how response times relate to both actions and beliefs and find that, controlling for incentives and individual characteristics, longer response times are associated with more sophisticated actions being chosen and with more sophisticated beliefs being reported. In particular, longer response times are associated with believing in a higher frequency of play for the level 2 action, which corresponds to the modal action in both games regardless of the incentive treatment. In contrast, the belief in the frequency of play of the dominance solution in the more complicated game — in which it does not correspond to the level 2 action — is fairly invariant with respect to response time.

Taken together, our results suggest incentive levels increase action sophistication both by increasing the incentive to reason further and form better beliefs and by decreasing the propensity to make mistakes. The fact that action sophistication and belief accuracy increases in the simpler game but less so or not at all in the more complicated game suggests that one should not expect changes in incentives to have the same effect in strategic environments with different degrees of complexity. Furthermore, our results caution against positing universal low levels of sophistication based on experimental evidence with low-powered incentives, as well as against relying on predictions requiring high levels of sophistication in complex but low-stakes environments.

1.1. Related Literature

Our paper contributes to the study of strategic sophistication in games, of how incentives relate to performance and mistakes in games, and of how response time relates to choices.

Strategic Sophistication. There is longstanding evidence that different individuals seem to exhibit different levels of strategic sophistication, which seems to help explain deviations from Nash equilibrium predictions. This is so in dominance-solvable games — such as beauty contests (Nagel, 1995), undercutting games (Costa-Gomes and Crawford, 2006), and ring games (Kneeland,

2015) — but also in games that are not dominance-solvable (see e.g. Arad and Rubinstein, 2012; Fudenberg and Liang, 2019). but also especially — but not only — in dominance-solvable games. Motivated by the observation that the level of strategic sophistication for a given individual is not stable across games (Georganas et al., 2015) and depends on beliefs about the opponents (Agranov et al., 2012), the discussion has since moved toward better understanding the effects of incentives on strategic sophistication, with recent models endogenizing this relation (e.g. Alaoui and Penta, 2016; Gonçalves, 2022).

Closest to this paper is the recent work by Alaoui et al. (2020), which shows that changing relative incentives so as to make the structure of iterated dominance more salient leads to a first-order stochastic dominance shift of choices toward higher levels of strategic sophistication, as given by the order of rationalizability of a given action. In contrast, this paper keeps relative incentives between actions fixed and examines the effect of incentive level on both choices and beliefs. While there is evidence for greater strategic sophistication due to greater cognitive effort exerted in forming beliefs about the opponents' behavior, we also find that the effect of incentive levels on reducing mistakes is also a major channel that needs to be taken into account.

Incentives and Mistakes. Choices in strategic settings tend to have a dimension of payoff-dependent stochasticity — an idea originating in discrete choice models (Luce, 1959; McFadden, 1974) — and action frequencies are closely related to their associated expected payoffs: those with higher expected payoffs are not always chosen, but they are chosen with higher probability. This monotonicity of choice frequencies with respect to payoffs is premise is the hallmark of quantal response equilibrium (McKelvey and Palfrey, 1995; Goeree et al., 2005) and other models of costly optimization in games (Mattsson and Weibull, 2002), which originated a wealth of applications — see (Holt et al., 2016) for a survey. In line with the main premise of this class of models, we find evidence for the posited monotonicity in dominance solvable games, with a crucial difference: monotonicity holds considering expected payoffs not according to the objective empirical frequency of actions, but with respect to each subjects' reported beliefs. Moreover, although subjects do not perfectly best-respond to stated beliefs,[4] we find that a higher own incentive levels sharpens the association between (subjective) expected payoffs and choice frequencies, and subjects best-respond more often to their beliefs.

This paper is also related to work relating incentive levels and performance. Recent evidence from

[4]Something which has been noted in prior work (Costa-Gomes and Weizsäcker, 2008; Rey-Biel, 2009, e.g.)

individual decision-making problems shows that, in line with costly cognition models (Matějka and McKay, 2015; Caplin and Dean, 2015), with higher incentives, people perform better (Dean and Neligh, 2022). However, this effect seems to depend on the complexity of the task; for instance, Enke et al. (2021) find that choice performance improves with significantly higher incentives in a simple task, but not in more complicated ones. In strategic settings, higher stakes lead rejection rates close to the subgame perfection prediction of zero in ultimatum games (Andersen et al., 2011), and also seem to entail greater sophistication in centipede games, with fewer 'passes' (McKelvey and Palfrey, 1992; Rapoport et al., 2003), but do not necessarily lead behavior closer to Nash equilibrium predictions in games with a unique mixed strategy equilibrium (McKelvey et al., 2000; Camerer, 2003, ch. 3). As argued earlier, our contribution relative to these papers is to provide causally identified evidence of the effect of incentive levels on both choices, beliefs and response times within a strategic setting.

Response Times in Strategic Settings. Finally, our work is related to a burgeoning literature which studies response time in strategic settings, either using it as a way to classify subjects (Rubinstein, 2016), predict choice (Schotter and Trevino, 2021) or strategic sophistication (Alós-Ferrer and Buckenmaier, 2021; D.Gill and Prowse, 2022), or as a proxy for cognitive effort (e.g. Rubinstein, 2007; Proto et al., 2019; Frydman and Nunnari, 2023) (see Spiliopoulos and Ortmann (2018); Clithero (2018) for a review). The most related paper within this is that of Alós-Ferrer and Buckenmaier (2021), who examine response times in beauty contest and different variants of the 11-20 game with a unique mixed strategy Nash equilibrium. The authors find that high levels of sophistication are associated with longer response time and that *distorting* incentives in favor of 'undercutting' leads smaller numbers — arguably associated with higher sophistication, despite the fact the variants chosen not being dominance-solvable — and to shorter response times. Differently from these studies, we provide a clearly identified results on how own and opponents' incentive levels affect to response time. Additionally, we provide suggestive evidence of the association between response time and strategic sophistication, relying both on choices and reported beliefs, supporting a mechanism that relatives incentives to cognitive effort and this, in turn, affecting belief formation and choices.

2. Hypotheses and Experimental Design

2.1. Hypotheses

In this section, we lay out our hypotheses. Throughout we focus on finite normal-form games, $\Gamma = \langle I, A, \rangle$, where I denotes the set of players, A_i the actions available for player i, A_{-i} action profiles for player i's opponents, $A := \times A_i$ the set of all strategy profiles, and $u_i : A \to \mathbb{R}$ player i's payoff function, with $u := \{u_i\}_{i \in I}$. As is conventional, we extend payoffs to the space of probabilities over actions, and write $-i$ to denote player i's opponents.

Our first hypothesis pertains the effect of incentives on action sophistication. We rely on rationalizability to describe action sophistication. Recall that an action is **k-rationalizable** if there is a distribution over $(k-1)$-rationalizable actions of opponents against which it is a best response. Relying on Pearce (1984, Lemma 3),[5] we can equivalently define player i's k-rationalizable actions R_i^k as those surviving k rounds of iterated deletion of strictly dominated strategies: $R_i^k := \{a_i \in A_i \mid \nexists \sigma_i \in \Delta(A_i) : u_i(\sigma_i, a_{-i}) > u_i(a_i, a_{-i}), \forall a_{-i} \in \times_{j \neq i} R_j^{k-1}\}$, with $R_i^0 := A_i$. Actions are **rationalizable** if they are k-rationalizable for all k. We say that an action a_i is more sophisticated than another a_i' if there is a k for which the former is k-rationalizable but the latter is not. Further, the above discussion can be extended to distributions of actions.

While nonrationalizable actions are seen to be chosen with positive probability in experiments (e.g. Costa-Gomes and Weizsäcker, 2008; Rey-Biel, 2009), it remains unclear if incentive levels affect the sophistication of play. For any two games $\Gamma, \tilde{\Gamma}$ that are identical in all except that payoffs for player i are scaled up, $\tilde{u}_i = \lambda u_i$ with $\lambda > 1$, we say that player i's incentive level is higher. Differently from payoff distortions, increasing incentive levels leaves relative incentives unchanged and thus the sets of k-rationalizable action distributions unaffected. Theoretical models also disagree on predictions. In many models — such as Nash equilibrium, level-k (Stahl and Wilson, 1994, 1995; Nagel, 1995), and cognitive hierarchy (Camerer et al., 2004) — increasing a player's incentive level has no effect on the observed distribution of actions. In other models, as endogenous depth of reasoning in symmetric games (Alaoui and Penta, 2016), increasing a player's incentive level may affect that player's actions but not other players.[6] And other models still make no clear prediction on how action sophistication is affected by incentive levels, as is the case of regular quantal response equilibrium (Goeree et al., 2005).

[5] While this relies on allowing for correlated distributions over opponents' actions, in two-player games — as in our experimental setting — this will not play a role.

[6] For instance, if Γ is a symmetric two-player game and cognitive costs are symmetric, scaling up a player's incentive level increases would lead the player to take more steps of reasoning but would not affect predictions about their opponent's actions.

Yet, intuition would suggest that by increasing a player's incentive level, not only that player's action sophistication increases, but also that opponents' can to some extent anticipate this and adjust their behavior. We formulate this as our first hypothesis:

Hypothesis 1 (Action Sophistication)**.** Action sophistication increases in (a) own incentive level and (b) the opponents' incentive level.

While testing Hypothesis 1 is in itself a contribution of this paper, we seek to speak to mechanisms, that is, *why* the incentive level affects choices.

One plausible channel for the effect of incentives on choices is that, by scaling up incentives, mistakes become more costly, and thus players are more likely to choose suboptimal actions less often. This is the logic underlying quantal response which posits the probability a given action is chosen depends on the player's vector of expected payoffs. Both models of stochastic choice based on random utility with additive iid payoff disturbances (McFadden, 1974; McKelvey and Palfrey, 1995) and on additive perturbed utility or costly precision (Mattsson and Weibull, 2002; Fudenberg et al., 2015) would predict that, all else equal, a higher incentive level for a player entails them choosing the best alternative more often. More generally, quantal response equilibrium would also predict that, for any game, actions with higher expected payoffs are chosen more often (Goeree et al., 2005). In assessing support for this class of models, we will test the following hypothesis:

Hypothesis 2 (Mistakes)**.** (a) Best-response rates increase in own incentive level. (b) Actions with higher expected payoff are chosen more often.

Another possible channel through which incentives may affect choices is by affecting players' beliefs. It has been shown that players' beliefs about the sophistication of their opponents may partially explain play of nonrationalizable actions. For example, sophistication of play changes depending on whether players are informed about opponents' characteristics that are associated with sophistication, such as chess ratings (Palacios-Huerta and Volij, 2009) or their educational background (Agranov et al., 2012; Alaoui et al., 2020).

The effect of incentive level on beliefs about the opponent may depend on whose incentive level changes. On the one hand, if incentives entail higher sophistication of play, then when opponents face a higher incentive level, one should believe their play to be more sophisticated. This suggests a pattern akin to what is observed when changing *relative* incentives: if increasing the relative incentive for a player to choose an action increases the frequency with which it is chosen (e.g. Ochs, 1995; McKelvey et al., 2000; Goeree and Holt, 2001), opponents then not only adjust behavior accordingly, they also report beliefs predicting this empirical pattern (Friedman and Ward, 2022). On the other hand, since reasoning about the game is costly, if one's own incentive level is higher, there is then a higher value to thinking and forming better beliefs. As it has been previously found that people underestimate their opponents' sophistication, higher own incentives can contribute to correct this misperception. These insights constitute the core of our next hypothesis:

Hypothesis 3 (Belief Sophistication)**.** The belief in opponents' strategic sophistication increases in both (a) higher own incentive levels and (b) higher opponents' incentive levels. Furthermore, (c) beliefs are more accurate with higher own incentives.

Finally, we consider extent to which incentive levels affect exerted cognitive effort, and how it, in turn, affects behavior. In line with recent models of sequential reasoning in games (Alaoui and Penta, 2016; Gonçalves, 2022), we conjecture that increasing a player's incentive level induces them to exert greater cognitive effort and ultimately resulting in better choices. To operationalize testing of this mechanism, we follow existing literature and proxy cognitive effort through response times (see e.g. Alós-Ferrer and Buckenmaier, 2021; Frydman and Nunnari, 2023). Our last hypothesis is then:

Hypothesis 4 (Response Time)**.** (a) Response time increases in one's own incentive level and (b) expected payoff increases with response time.

2.2. Experimental Design and Identification Strategy

Identification Strategy. Testing the effects of incentive levels on a player's choices and beliefs in strategic settings poses a fundamental identification problem, since it requires to hold fixed their opponents' behavior. If a higher incentive level for player i affects their behavior and this is understood by their opponents, then it may also affect the opponents' behavior. Hence, changes in player i's behavior conflate the direct effect of the change in the incentive level and an indirect effect due to their perception of their opponents' reaction to it.

Our identification strategy relies on a simple opponent matching protocol, summarized in Figure 1. Subjects are assigned a role in a game as well as an incentive group that determines their own incentives and their opponents' incentives. First, we have two incentive groups in which all players in all roles have either high or low incentives and know their opponents also play

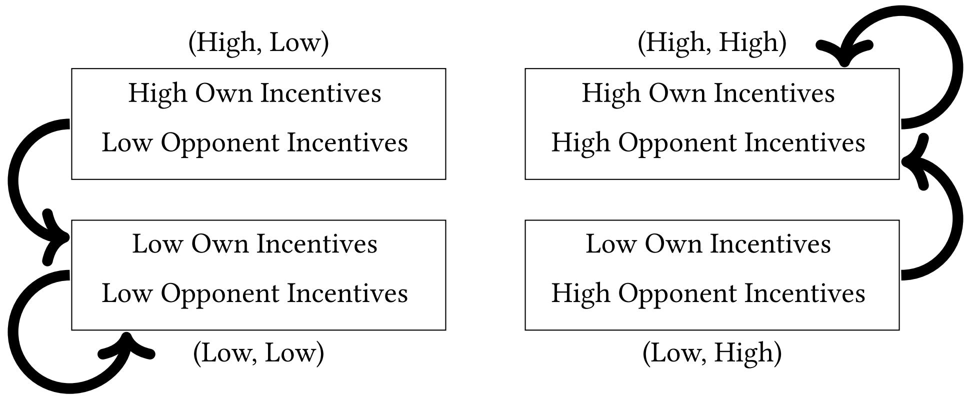

Figure 1: Opponent Matching Procedure

Notes: This diagram describes the matching procedure implemented. Subjects in incentive level groups (High, High) and (Low, Low) have their payoffs depend on an action taken by another random subject in the same group. Subjects in incentive level groups (High, Low) and (Low, High) have their payoffs depend on an action taken by a random subject in the groups (Low, Low) and (High, High), respectively.

the same game with the same incentives. These correspond to groups (High, High) and (Low, Low) in Figure 1. Second, we have another two incentive groups, with high or low incentives, that are matched with opponents with the opposite incentive level, low or high, corresponding to groups (High, Low) and (Low, High) in Figure 1. This ensures that we are able (i) to hold fixed opponents' behavior and compare how own incentive levels affect choices and beliefs, and (ii) to hold fixed one's own incentive level and infer how opponents' incentive level affects their beliefs and choices. Furthermore, this procedure not only fixes the incentive level a player's opponents have, but also their opponents' opponents, and so on.[7] Such methodology is akin to having having subjects as observers as in Huck and Weizsäcker (2002) and, more recently, in Alaoui et al. (2020). We also designed the opponent matching procedure so that the strategic incentives across treatments was symmetric in that subject's actions did not affect their opponents' payoffs — a necessity for subjects in incentive groups (High, Low) and (Low, High). With the described matching procedure and random assignment to incentive level groups, we need but to compare variables of interest — choices, beliefs, and response times — across subjects with different own incentive levels or opponent incentive levels to identify the causal effect of incentive levels.

[7] Note that simply informing player i of their opponents' incentive levels won't do, since this would not provide information about the incentive level of their opponents' opponents, which could be relevant for player i to form beliefs about their opponents' behavior. In contrast, this matching procedure ensures that if player i's opponents have a given incentive level, then their opponents have the same incentive levels, fixing the player's higher order beliefs.

	Actions	Player 2 a_1	a_2	a_3	a_4
Player 1	a_1	40, 40	70, 30	80, 20	10, 10
	a_2	30, 70	40, 40	70, 30	80, 20
	a_3	20, 80	30, 70	40, 40	70, 30
	a_4	10, 10	20, 80	30, 70	40, 40

(a) 2 Steps

	Actions	Player 2 a_1	a_2	a_3	a_4
Player 1	a_1	40, 40	70, 30	10, 20	10, 10
	a_2	30, 70	40, 40	70, 30	10, 20
	a_3	20, 10	30, 70	40, 40	70, 30
	a_4	10, 10	20, 10	30, 70	40, 40

(b) 3 Steps

Figure 2: Games

Notes: This figure exhibits the games used in the experiment. Both are symmetric, two-player dominance-solvable games, with the game in panel (a) taking 2 steps of iterated (maximal and simultaneous) deletion of strictly dominated strategies to obtain the strategy prescribed by the dominance solution, whereas the game in panel (b) takes 3 steps.

Games. Subjects faced one of the two normal-form games exhibited in Figure 2. The two games are dominance solvable and exhibit similar payoffs. The game in panel (a) takes two iterations of (maximal and simultaneous) deletion of strictly dominated strategies to reach the dominance solution strategy, and that in panel (b) takes three iterations. We opted for four-action games since this is the smallest number of actions that allows us to contrast how the number of iterations of deletion of dominated actions interacts with the incentive level with minor modifications in payoffs alone. Symmetry was imposed in order to improve statistical power and minimize data collection. Both exhibit similar payoffs and similar in spirit to 11-20 (Arad and Rubinstein, 2012) and undercutting games more generally (Nagel, 1995; Costa-Gomes and Crawford, 2006).

All actions in both games can be ranked by iterated dominance. Specifically, actions deleted within the same round of iterated deletion of strictly dominated actions are also ordered in terms of dominance relation. For instance, in ‘2 Steps’ (Figure 2(a)), (i) a_4 is strictly dominated by a_3, which in turn is strictly dominated by a_2, and (ii) upon deletion of a_3 and a_4, a_2 is iteratedly strictly dominated by a_1. Similarly, in ‘3 Steps’ (Figure 2(b)), (i) a_4 is strictly dominated by a_3; (ii) deleting a_4 renders a_3 iteratedly strictly dominated by a_2, and (iii) deleting a_3, again a_2 is iteratedly strictly dominated by a_1.

The above implies that actions are ranked in terms of strategic sophistication in a simple manner. First, if a_{n+1} is k-rationalizable, then so is a_n. Second, if both a_n,a_{n+1} are k- but not $(k+1)$-rationalizable, then a_n (iteratedly) strictly dominates a_{n+1}. As such, we will throughout take action a_n as more strategically sophisticated than a_{n+1}. We note that level-k actions are also uniquely pinned-down assuming level 0 uniformly randomizes, regardless of subjects’ risk attitudes. Specifically, in the 2 Steps game, given level 0 uniformly randomizes, a_2 corresponds to the level 1 action regardless of risk preferences, and a_1 to level 2 action; both a_3 and a_4 are strictly dominated. In the 3 Steps game, the level 1 action is now a_3, level 2 a_2, and a_1 the level 3 action.

Lastly, we note the games were also chosen so as to mitigate concerns that payoffs would act as focal coordination points. First, we note that all payoff vectors are associated to multiple action profiles, and the dominance payoff (40,40) is too. Although the iterative structure of the game is apparent Figure 2, we randomly shuffled rows and columns. Second, the dominance payoff is neither Pareto dominated nor Pareto dominant, as are the majority of payoff vectors. Third, the above-described random opponent matching protocol additionally contributes to alleviate concerns about focal coordination and other-regarding preferences.

Other Design Details and Logistics. The experiment implemented a 2×2×2 design, corresponding to own and opponent's assignment (or, alternatively, to the 4 incentive level groups) and to the specific games subjects played — 2 or 3 Steps as in Figure 2. Subjects were sorted into one of the 8 treatments uniformly at random, and play only once.

The experiment was incentivized via a binary lottery so that payoffs corresponded to the probability of getting paid a prize of \$$x$ versus \$2.00. In the high own incentive level treatment, \$$x$ corresponded to \$22.00 and in the low incentive level to \$2.50. We chose the high incentive level to maximize treatment effect, while the low incentive level is higher than the expected payment per game in other experiments.[8]

We elicited both actions and beliefs about their opponents' action distribution. Actions selected in the game, $a_i \in A_i := \{a_1, a_2, a_3, a_4\}$, and payoffs obtained with opponents randomly selected as per the matching procedure correspond to a probability of getting the prize. The order of rows and columns was randomized and subjects were informed of this. Beliefs b_i correspond to the probability opponents are to choose any given action — i.e. $b_i = (b_{i,1}, b_{i,2}, b_{i,3}, b_{i,4}) \in \Delta(A_i)$ — and were incentivized via a binarized scoring rule (Hossain and Okui, 2013).[9] Choices and beliefs were elicited simultaneously and one of them was randomly selected for payment.

The experiment proceeded as follows: (i) instructions were provided together with comprehension questions and attention controls, (ii) subjects played two unincentivized practice rounds

[8]For instance, the average payments per game in Alaoui et al. (2020) and Fudenberg and Liang (2019) were €0.88 and \$0.93, respectively.

[9]Specifically, given their belief report b_i, the associated probability of getting the prize was $1 - \sum_{n=1}^{4} (\overline{\sigma}_{i,n} - b_{i,n})^2/2$, where $\overline{\sigma}_{i,n}$ corresponds to the empirical frequency of action a_n among the subject's opponents.

without feedback,[10] (iii) own and opponent's incentive levels were revealed, (iv) actions and beliefs were elicited simultaneously, (v) subjects answered a brief questionnaire on sociodemographics and received payment information. Screenshots of the interface and instructions are provided in Online Appendix B.

We targeted 100 subjects per treatment, having recruited 834 subjects, with the treatment with fewer subjects having exactly 100 subjects. Sessions were conducted on 9-10 and 15-17 January 2020 on Amazon Mechanical Turk, constraining potential workers to be adults based in the United States. The subjects faced no time constraint; the average duration was about 22 minutes and average earnings $24.97 per hour. Sociodemographic data collected consisted of age, sex, education, and prior exposure to game theory (see Online Appendix A for details); samples across treatments were balanced across all sociodemographic variables.

3. Incentives and Strategic Sophistication

We first examine if and how incentive levels affect observed action sophistication (Hypothesis 1).

Our first observation is that subjects do not always play the dominance solution action, a_1. From Table 1 — which exhibits action frequencies and average beliefs for each of the 8 treatments — one can observe the different treatments entail a wide variation in dominance play, ranging from 13.0% to 57.3%, more than four times over. Similarly, there is quite significant variation in dominated play (corresponding to actions a_3 and a_4 in the 2 Steps game, and action a_4 in 3 Steps), between 6.7% and 16.1%.

We test whether incentive levels affected the frequency of dominance and dominated play. We consider the following specification:

$$y_i = \beta_0 + \beta_1 H_i + \beta_2 H_{-i} + \text{Controls}_i + \epsilon_i \quad (1)$$

where H_i is an indicator variables equal to 1 if subject i's treatment had a high own incentive level, and H_{-i} an analogous indicator for whether subject i's opponent had a high incentive level. Controls refer to the subjects' age, sex, education, and prior exposure to game theory. In testing for the effects of incentive levels on dominance play, the dependent variable, y_i, corresponds to an indicator equaling 1 if subject i chose the dominance solution action and zero if otherwise. For testing the effect of incentives on dominated play, we take an analogous indicator for whether

[10] Both were four-action two-player games, one with a strictly dominant action, another with no pure-strategy Nash equilibrium.

	Incentive Level		Action Frequency				Mean Belief Reports				Obs.
Game	Own	Opponent	a_1	a_2	a_3	a_4	b_1	b_2	b_3	b_4	N
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
2 Steps	High	High	0.573	0.272	0.087	0.068	0.330	0.357	0.189	0.124	103
2 Steps	High	Low	0.519	0.340	0.085	0.057	0.279	0.369	0.201	0.150	106
2 Steps	Low	High	0.359	0.359	0.146	0.136	0.268	0.317	0.230	0.185	103
2 Steps	Low	Low	0.343	0.363	0.157	0.137	0.265	0.315	0.238	0.181	102
3 Steps	High	High	0.144	0.490	0.298	0.067	0.184	0.311	0.363	0.142	104
3 Steps	High	Low	0.130	0.440	0.360	0.070	0.213	0.259	0.340	0.189	100
3 Steps	Low	High	0.196	0.330	0.312	0.161	0.219	0.275	0.320	0.186	112
3 Steps	Low	Low	0.192	0.288	0.365	0.154	0.222	0.270	0.327	0.180	104

Table 1: Action Frequency and Mean Beliefs

Notes: This table shows the action frequency and the mean beliefs by each of the treatments. Each subject is allocated to a given treatment — a game (2 Steps or 3 Steps), and an incentive level group specifying the subject's incentive level (high or low) and their opponent's incentive level (high or low). Subjects play only one game and only once.

the subject chose a strictly dominated action. Table 2 summarizes the results.

We find that while higher own incentives significantly decreases the frequency of dominated play, it only increases the frequency of the dominance solution action in the 2 Steps game. Specifically, dominated play decreases in both games — by 14.1 percentage points (pp) in the 2 Steps game and 9.5 pp in the 3 Steps game (columns (3) and (4)). However, if dominance play increases by over 20 percentage points (pp) in the 2 Steps game with higher own incentives, it decreases by 6.6 pp in the game taking 3 Steps of iterated deletion of dominated actions. If one understands the number of steps needed to reach the dominance solution as a measure of how complicated the game is from a strategic reasoning standpoint, these results suggest that more complex games require a higher incentive level in order to achieve the same rate of dominance play. A second observation is that if own incentive levels matter for the frequency of dominance and dominated play, the opponents' do not.

We then consider the overall distribution of actions — shown in Figure 3. For each game and opponent's incentive level, the distribution of actions is different across own incentive levels (Wald tests with p-value<.05). For the 2 Steps game, the action distribution is shifted toward more sophisticated actions in a first-order stochastic dominance manner (Mann-Whitney U rank test p-values <.01) when own incentives are higher. For the 3 Steps game, this is no longer the case: the decrease in the frequency of action a_1 observable in panels (c) and (d) of Figure 3 suggests the

	Dominance Play		Dominated Play	
	(1)	(2)	(3)	(4)
	2 Steps	3 Steps	2 Steps	3 Steps
High Own Incent.	0.204***	-0.066*	-0.141***	-0.095***
	(0.048)	(0.038)	(0.041)	(0.032)
High Opp. Incent.	0.013	0.008	0.015	0.004
	(0.049)	(0.037)	(0.040)	(0.032)
Controls	Yes	Yes	Yes	Yes
R-Squared	0.14	0.05	0.10	0.09
Observations	414	420	414	420

Heteroskedasticity-robust standard errors in parentheses.

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

Table 2: Incentive Level, Dominance and Dominated Play (Hypothesis 1)

Notes: This table shows the results for the regression specified in equation 1, considering as dependent variable an indicator for whether the subject chose the dominance solution (columns (1) and (2)) or a strictly dominated action (columns (3) and (4)). High Own/Opponent Incentives correspond to indicators for whether the subject and their opponent face a high incentive level. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Controls refer to the subjects' age, sex, education, and prior exposure to game theory.

distributions are not ranked in stochastic dominance (Mann-Whitney U rank test p-values >.25). Still, the distributions are noticeably different,[11] less flat and assigning higher probability to a_2, the level-2 action.

In short, we find that higher own incentive levels do unambiguously entail a lower frequency of strictly dominated play, but only increase action sophistication in a stochastic dominance in the simpler 2 Steps game. This lends partial support for Hypothesis 1a. One would therefore need that strategic sophistication — described by types in level-k and cognitive hierarchy models — to depend on the incentive *level* players face. Against Hypothesis 1b, we find no evidence that opponent's incentive level affects action sophistication.

4. Incentives and Mistakes

One reason that could underlie the change in behavior observed is that with high own incentive levels mistakes are more costly in terms of payoffs and thus subjects make fewer mistakes. This

[11] A Fisher exact test confirms that the distributions are statistically significantly different: p-values for panels (a), (b), (c), and (d) are, respectively, .023, .017, .055, and .040.

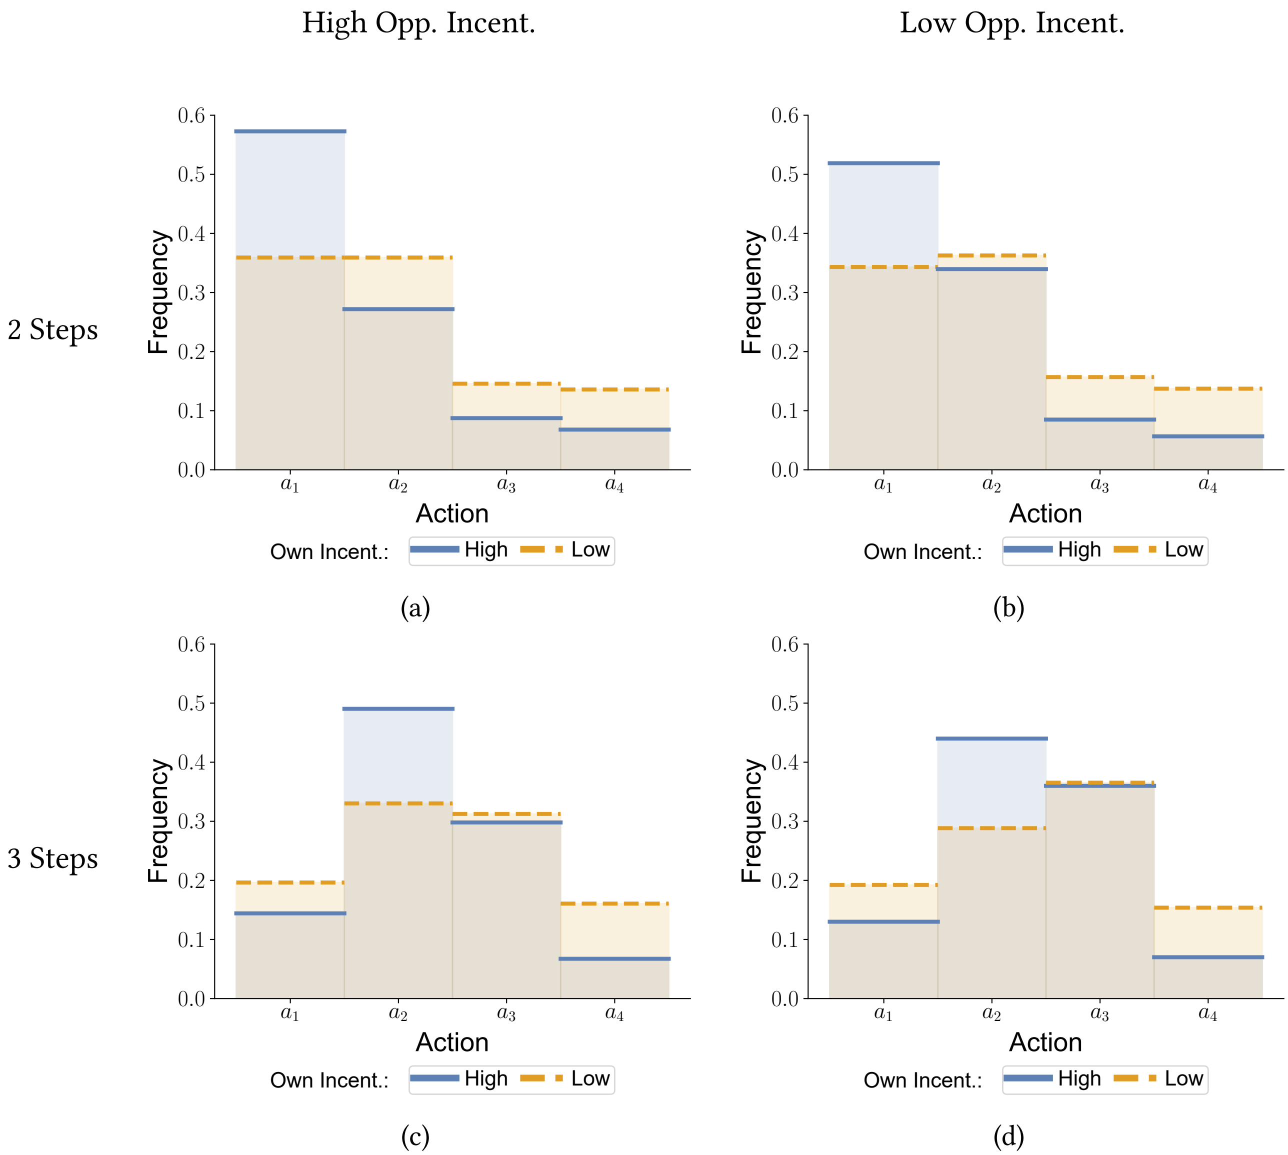

Figure 3: Incentive Level and Action Frequency

Notes: The panels exhibit the action frequency for high and low own incentive levels, for different games (2 Steps, (a) and (b), and 3 Steps, (c) and (d)) and holding fixed the opponent's incentive level (High, (a) and (c), or Low, (b) and (d)). 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2).

is crux of Hypothesis 2.

We first test whether subjects are more or less likely to (i) best respond to their stated beliefs (subjective best responses), and (ii) best respond to the observed frequency of play (objective best responses). In Table 3, we report results from estimating a linear probability model using an analogous specification to that in equation 1. The results show that in fact higher own incentive

	Subjective BR		Objective BR	
	(1)	(2)	(3)	(4)
	2 Steps	3 Steps	2 Steps	3 Steps
High Own Incent.	0.219***	0.151***	0.204***	0.042
	(0.049)	(0.051)	(0.048)	(0.044)
High Opp. Incent.	-0.021	0.027	0.013	-0.195***
	(0.050)	(0.050)	(0.049)	(0.044)
Controls	Yes	Yes	Yes	Yes
R-Squared	0.11	0.06	0.14	0.10
Observations	414	420	414	420

Heteroskedasticity-robust standard errors in parentheses.

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

Table 3: Incentive Level and Best-Response Rate (Hypothesis 2a)

Notes: This table shows the results for the regression specified in equation 1, considering as dependent variable an indicator for whether the subjects best respond to their beliefs — i.e. chose the action that maximizes expected payoffs according to their reported beliefs (columns (1) and (2)) — and for if subjects best respond to the empirical frequency of actions of their opponents — i.e. objective best responses (columns (3) and (4)). High Own/Opponent Incentives correspond to indicators for whether the subject and their opponent face a high incentive level. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Controls refer to the subjects' age, sex, education, and prior exposure to game theory.

levels increase the likelihood of subjects best responding to their reported beliefs by 21.9 pp in the 2 Steps game and 15.1 pp in the 3 Steps game.[12] However, own incentive levels only increase the likelihood of best responding to the empirical frequency of play in the simpler 2 Steps game, and not in the 3 Steps game.

A fundamental precept of regular quantal response equilibrium is that actions with higher expected payoff are chosen more often, a property called 'monotonicity.' We then assess how incentives affect mistake propensity by looking at the distribution of choices based on the subjective expected payoff. For this, we define the subjective rank of the chosen action to be n if the action chosen has the n-th highest subjective expected payoff (according to reported beliefs) of all four actions. In the panels of Figure 4 we display the action frequency by subjective rank.

Subjects indeed choose actions with higher *subjective* expected payoff more often, but the same is not true when considering instead *objective* expected payoff. We note that the actions with

[12] Best-response rates to stated beliefs in the data are between 40-50% in the low own incentive level treatment and 60-70% in the high own incentive level treatment — this corresponds to the frequency of the choosing the action with subjective rank 1 as exhibited in Figure 4. These figures are comparable to what has been observed in two-player three-action games, e.g. Costa-Gomes and Weizsäcker (2008) and Rey-Biel (2009).

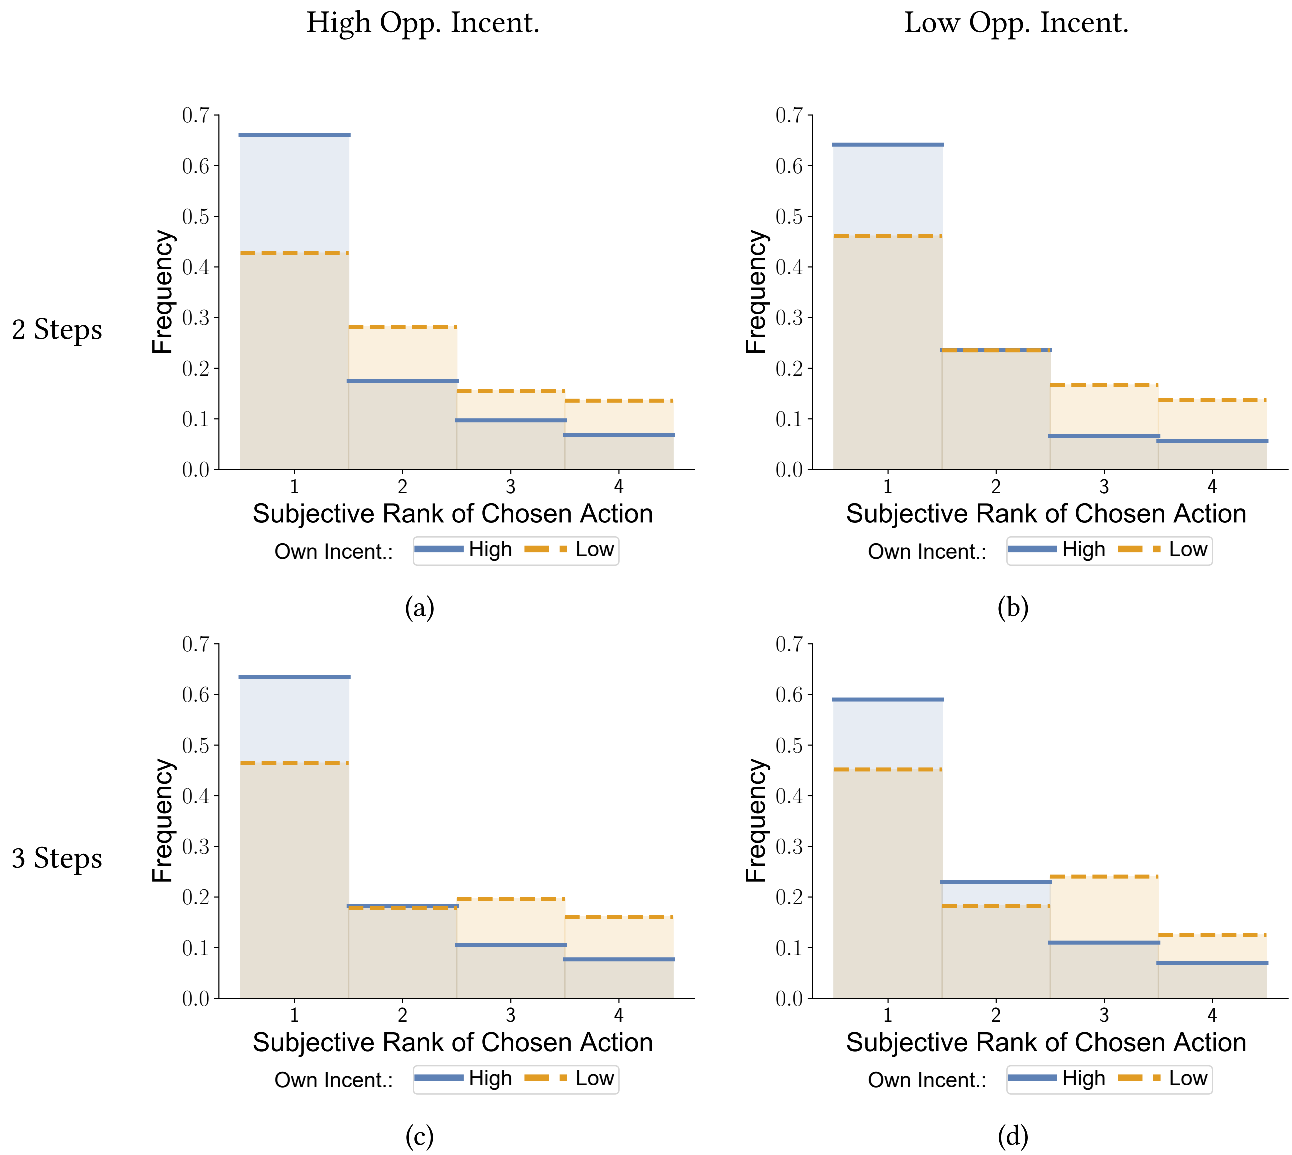

Figure 4: Incentive Level and Subjective Mistakes (Hypothesis 2b)

Notes: The different panels exhibit the frequency of the subjective rank of the action chosen by the subject for high and low own incentive levels, for different games (2 Steps, (a) and (b), and 3 Steps, (c) and (d)) and holding fixed the opponent's incentive level (High, (a) and (c), or Low, (b) and (d)). The subjective rank of an action is n if the action entails the n-th highest subjective expected payoffs according to the reported beliefs; e.g. actions with subjective rank 1 are those that maximize subjective expected payoffs. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2).

higher subjective expected payoff are taken more often than those with lower expected payoff in all treatments. In contrast, using empirical frequencies and objective expected payoffs to define an analogous objective rank of chosen actions, there are flagrant violations of monotonicity —

see Figure 8 in Appendix A.

We further note that not only do higher own incentive levels decrease mistakes — subjective best response rates increase — but also decrease the severity of mistakes. We note that distribution of the subjective rank of the action chosen shifts in a first-order stochastic dominance sense with higher own incentive levels regardless of the game and the incentive level of opponents.[13] In sum, we find decisive support in favor of Hypothesis 2.

In assessing whether mistakes alone are able to rationalize the data, we fit a logit quantal response equilibrium to the data. In Appendix A we provide the predicted action frequencies obtained via maximum likelihood estimation.[14] We contrast the estimates with the data as well as with estimates from a subjective logit quantal response model, i.e. a model of logit choice using subjective expected payoffs. Both in logit quantal response equilibrium and in our subjective logit quantal response model, the scaling factor affecting payoffs is higher with higher own incentive levels, as expected. While both do well in the 2 Steps game, regardless of whether subjects face, logit quantal response equilibrium overpredicts action sophistication in the 3 Steps game. Incorporating subjects' reported beliefs significantly improves the fit to data: in subjective logit quantal response, the loglikelihood increases 28%.

5. Incentives, Beliefs, and Strategic Reasoning

In the previous section we observed that, while mistakes are an important channel to explain how the incentive level changes behavior, subjects' beliefs are a crucial element in rationalizing what a mistake is. In this section, we consider Hypothesis 3, that is, if incentive levels directly affect beliefs.

Average belief reports per treatment are depicted in Figure 5. Similarly to what occurred in action frequencies, while higher own incentives lead to mean beliefs shifting toward assigning greater sophistication to opponents (in a significant stochastic dominance sense) in the simpler 2 Steps game, in the more complicated 3 Steps game, beliefs become less uniform when opponents have high incentives but are indistinguishable when they have low incentives.

Akin to the analysis for action frequencies, we estimate the effect of incentive levels on subjects'

[13] A Mann-Whitney U rank test yields p-values for panels (a), (b), (c), and (d) of, respectively, .002, .001, .013, and .004.

[14] Since treatments in which subjects and their opponents have different incentive levels correspond to off-equilibrium play, we restrict attention to treatments in which incentive levels are the same.

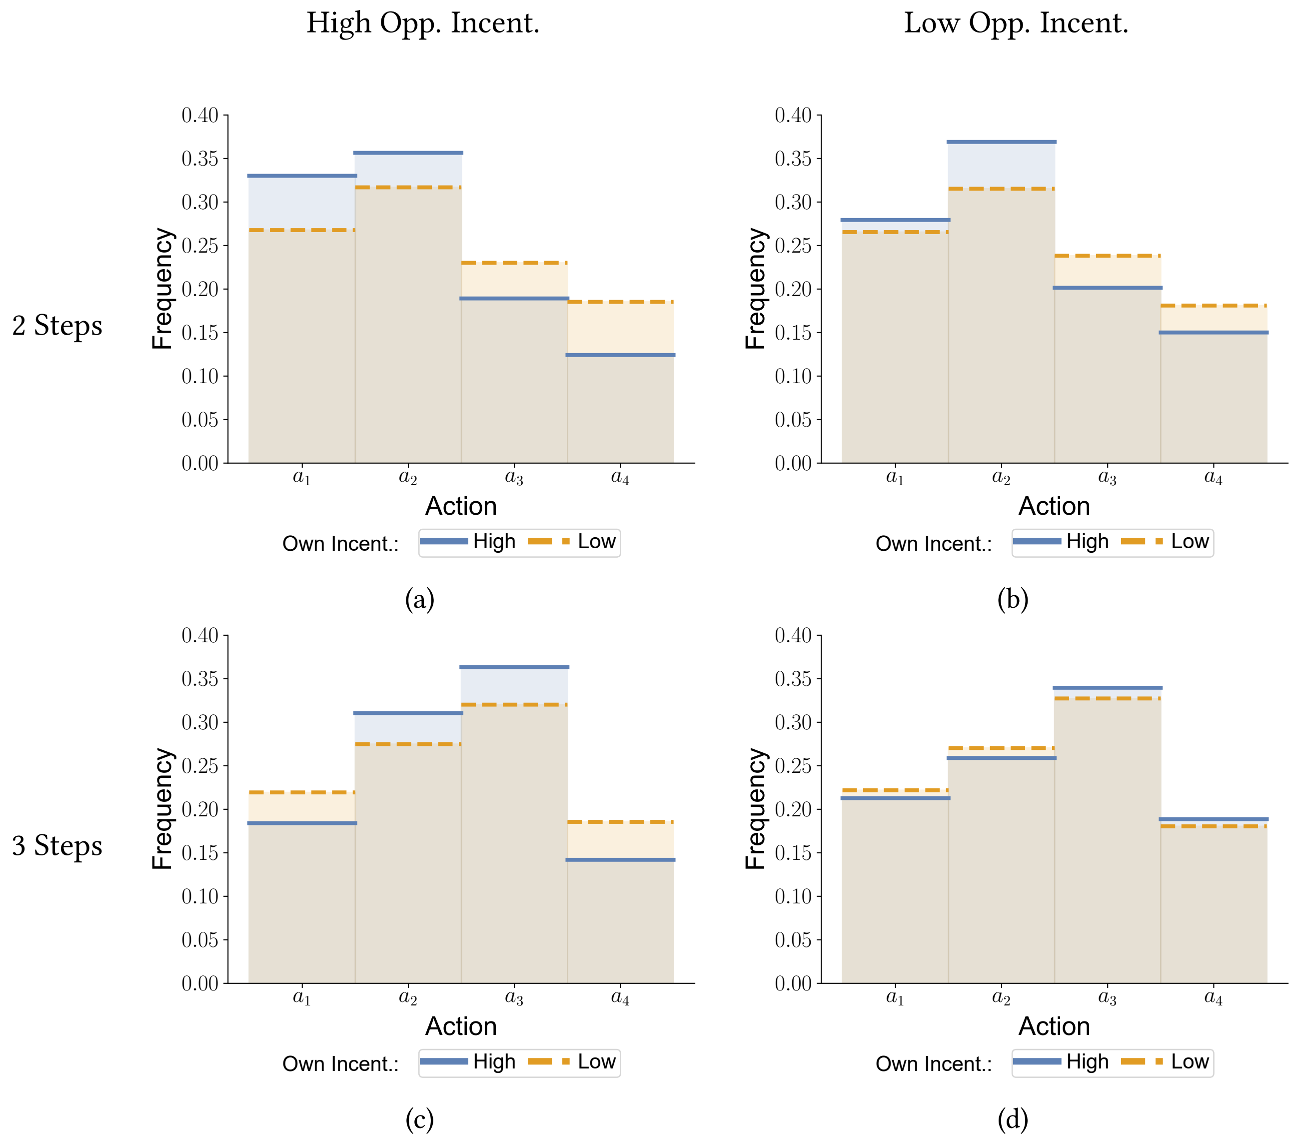

Figure 5: Incentive Level and Mean Beliefs

Notes: The different panels exhibit the mean reported beliefs for high and low own incentive levels, for different games (2 Steps, (a) and (b), and 3 Steps, (c) and (d)) and holding fixed the opponent's incentive level (High, (a) and (c), or Low, (b) and (d)). 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2).

implied belief of opponents playing the dominance solution action and of playing a strictly dominated action. Table 4 confirms that in fact a higher *own* incentive level entails believing in greater *opponent* sophistication. Therefore, we conclude in favor of Hypothesis 3a.

Although the sign of the effect of opponent's incentive level goes in the same direction, the effects are not significant due to lack of statistical power. It is however obvious from Figure 5 that

	Belief Dominance Play		Belief Dominated Play	
	(1)	(2)	(3)	(4)
	2 Steps	3 Steps	2 Steps	3 Steps
High Own Incent.	0.035**	-0.021	-0.074***	-0.019
	(0.016)	(0.014)	(0.018)	(0.013)
High Opp. Incent.	0.020	-0.020	-0.020	-0.016
	(0.017)	(0.014)	(0.018)	(0.013)
Controls	Yes	Yes	Yes	Yes
R-Squared	0.10	0.06	0.15	0.06
Observations	414	420	414	420

Heteroskedasticity-robust standard errors in parentheses.

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

Table 4: Incentive Level and Belief in Opponent Sophistication (Hypothesis 3a,b)

Notes: This table shows the results for the regression specified in equation 1, considering as dependent variable the reported belief in dominance play, $b_{i,1}$, (columns (1) and (2)) or the reported belief in strictly dominated play, $b_{i,3} + b_{i,4}$ in column (3) and $b_{i,4}$ in column (4). High Own/Opponent Incentives correspond to indicators for whether the subject and their opponent face a high incentive level. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Controls refer to the subjects' age, sex, education, and prior exposure to game theory.

the changes in beliefs entailed by a higher own incentive level are more substantial when the opponent's incentive level is high. Indeed in the 3 Steps game, when the opponent's incentive level is low, there is no statistical difference in mean beliefs when varying own incentive levels. This highlights that higher own incentive levels lead to a higher belief in opponent sophistication only when the opponent has high incentives as well, lending support to Hypothesis 3b.

We then turn to the question of whether beliefs become more accurate (Hypothesis 3c). In Table 5, we show the results for regressions considering three different types of biases in reported beliefs with an otherwise identical to before. In columns (1) and (2) we consider the absolute difference between the subject's belief that their opponent plays the dominance action a_1 and the corresponding observed empirical frequency. Columns (3) and (4) use an analogous measure for bias but considering the bias in beliefs about strictly dominated play by opponents. The last two columns, (5) and (6), take for dependent variable the bias in subjective expected payoffs as given by the L1-norm, i.e. $\sum_{n=1}^{4} |u_i(a_n, b_i) - u_i(a_n, \overline{\sigma}_{-i})|$, where $u_i(a_n, b_i)$ denotes subject i's expected payoff to action a_n as per their reported beliefs b_i, and $u_i(a_n, \overline{\sigma}_{-i})$ the subject's expected payoff to the same action but considering the empirical frequency of opponents' actions $\overline{\sigma}_{-i}$. Throughout we find that when the opponents have higher incentive levels, subjects' beliefs are more biased,

	|Belief - Opponent Action Frequency|				Bias in Subjective Expected Payoffs	
	Dominance Play		Dominated Play			
	(1)	(2)	(3)	(4)	(5)	(6)
	2 Steps	3 Steps	2 Steps	3 Steps	2 Steps	3 Steps
High Own Incent.	0.005	0.011	-0.040***	-0.015	-4.345***	0.558
	(0.011)	(0.010)	(0.011)	(0.010)	(1.282)	(1.403)
High Opp. Incent.	0.037***	0.062***	0.033***	0.021**	4.222***	9.520***
	(0.011)	(0.009)	(0.012)	(0.010)	(1.277)	(1.416)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
R-Squared	0.14	0.14	0.14	0.06	0.17	0.15
Observations	414	420	414	420	414	420

Heteroskedasticity-robust standard errors in parentheses.

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

Table 5: Incentive Level and Belief Accuracy (Hypothesis 3c)

Notes: This table shows the results for the regression specified in equation 1, considering as dependent variable the absolute difference between reported belief in dominance play by opponents and its realized frequency, $|b_{i,1} - \overline{\sigma}_{-i,1}|$, (columns (1) and (2)), or the analogous absolute difference for strictly dominated play, $|b_{i,3}+b_{i,4}-\overline{\sigma}_{-i,3}-\overline{\sigma}_{-i,4}|$ in column (3) and $|b_{i,4}-\overline{\sigma}_{-i,4}|$ in column (4). In columns (5) and (6) the dependent variable is the L1-norm of the difference between subjective expected payoffs (according to reported beliefs) and objective expected payoffs (according to observed action frequencies). High Own/Opponent Incentives correspond to indicators for whether the subject and their opponent face a high incentive level. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Controls refer to the subjects' age, sex, education, and prior exposure to game theory.

indicating that actions by high incentive opponents are harder to predict. In contrast, the effect of subjects' own incentive level on belief accuracy is mixed. Own incentives have no detectable effect in the more complicated 3 Steps game, but they do seem to significantly decrease bias in beliefs both in terms of frequency of dominated play — by 4 percentage points (column (3)) and, more importantly, in terms of expected payoffs — by 4.35 payoff points (column (5)).

6. Incentives and Response Time

Lastly, we turn to the effect of incentive levels on cognitive effort as proxied by subjects' response time. If subjects face higher own incentive levels, then the marginal value to exerting cognitive effort increases and one would expect a higher response time, as conjectured in Hypothesis 4a. We test this hypothesis by regressing log response time (in seconds) on incentive level treatments as before. As the results shown in Table 6 indicate, when own incentive levels are higher, response

	Log(Response Time)	
	(1)	(2)
	2 Steps	3 Steps
High Own Incent.	0.478***	0.362***
	(0.074)	(0.082)
High Opp. Incent.	0.077	0.132*
	(0.076)	(0.078)
Controls	Yes	Yes
R-Squared	0.17	0.18
Observations	414	420

Heteroskedasticity-robust standard errors in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Incentive Level and Response Time (Hypothesis 4a)

Notes: This table shows the results for the regression specified in equation 1, considering as dependent variable the subjects' log response time (in seconds). High Own/Opponent Incentives correspond to indicators for whether the subject and their opponent face a high incentive level. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Controls refer to the subjects' age, sex, education, and prior exposure to game theory.

times increase by over 40%.

Whether or not this greater cognitive effort is valuable is not immediate. We examine whether response times are associated with (1) the rate of (objective) best responses and (2) the (objective) expected payoffs as given by the empirical frequency of opponents' play. Table 7 shows that longer response times are associated to a higher rate of best responses (columns (1) and (2)) as well as an increase in expected payoffs (columns (3) and (4)), albeit the elasticity of expected payoffs to response time is small. In both cases, we control not only for incentive level treatments but also for individual characteristics such as age, education level, and field of education, and prior exposure to game theory, all of which would be susceptible to affect response times. We consider the findings support Hypothesis 4b, although it is important to note that here the effects are measures of association and not causally identified.

We conclude by conducting a descriptive analysis of how actions taken and beliefs reported varied with response time. For actions, we estimated a multinomial logistic model, regressing actions on response time, incentive treatments, and individual characteristics, and obtained the fitted conditional distribution with respect to response times by marginalizing over the other regressors. For beliefs, we relied on nonparametric kernel estimation of the mean beliefs, $b_i =$

	Objective BR		Log(Expected Payoff)	
	(1)	(2)	(3)	(4)
	2 Steps	3 Steps	2 Steps	3 Steps
Log(Response Time)	0.131***	0.075***	0.084***	0.065***
	(0.033)	(0.028)	(0.020)	(0.013)
High Own Incent.	0.142***	0.015	0.080**	0.040*
	(0.052)	(0.044)	(0.032)	(0.022)
High Opp. Incent.	0.003	-0.205***	-0.079***	-0.045**
	(0.048)	(0.043)	(0.028)	(0.021)
Controls	Yes	Yes	Yes	Yes
R-Squared	0.17	0.12	0.16	0.12
Observations	414	420	414	420

Heteroskedasticity-robust standard errors in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Response Time, Best Responses, and Payoffs (Hypothesis 4b)

Notes: This table examines the relation between, on the one hand, log response times (in seconds), and on the other objective best responses to the empirical frequency of opponents' actions (columns (1) and (2)) and log expected payoffs, where expectations are taken also with respect to the empirical frequency of opponents' actions (columns (3) and (4)). High Own/Opponent Incentives correspond to indicators for whether the subject and their opponent face a high incentive level. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Controls refer to the subjects' age, sex, education, and prior exposure to game theory.

$(b_{i,1}, b_{i,2}, b_{i,3}, b_{i,4})$, *conditional* on response time. Specifically, we first estimated the distribution of beliefs b_i conditional on response time, incentive treatments, and individual characteristics. We then used the estimated conditional distribution to obtain estimates for the expected beliefs vector conditional on response time, using the mean values for the other covariates. The results are shown in Figures 6 and 7.[15]

Our findings uncover two patterns in the data. First, that in the simpler 2 Steps game longer response times are associated with both greater action and belief sophistication. In panel (a) of Figure 6 we can see that the predicted frequency of dominated actions, a_3 and a_4, declines quickly with response time, while a_2 (corresponding to the level-1 action) first increases and then declines. Throughout, the frequency of choice of the dominance solution action a_1 increases monotonically in response time. The relation between (mean) reported beliefs and response time,

[15]In Appendix A, we also provide linear regressions for each action $a_{i,n}$ and belief about a given action $b_{i,n}$ individually on the same set of covariates confirming the trends depicted — see Tables 8 and 9.

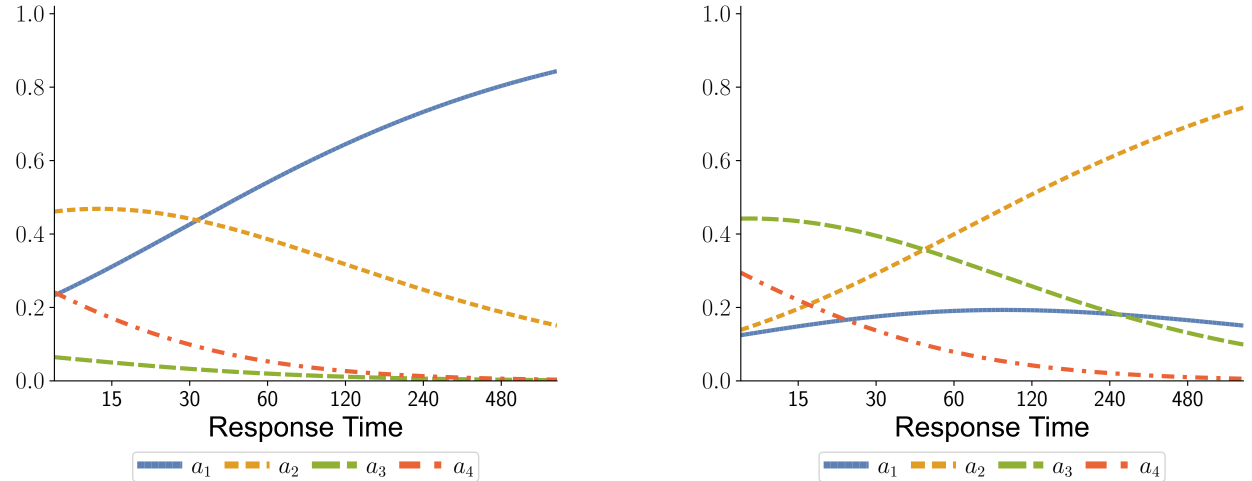

(a) 2 Steps

(b) 3 Steps

Figure 6: Action Frequency and Response Time

Notes: This figure shows the predicted relation between action frequency and response time (in seconds). The predictions are given by multinomial logit estimation of the relation between subjects' choices, on the one hand, and response time, incentive treatments, and individual characteristics, estimated separately for each game. The action frequency conditional on response times is given by marginalizing over the other regressors. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Individual characteristics refer to the subjects' age, sex, education, and prior exposure to game theory.

shown in panel (a) of Figure 7 also indicates that longer decisions are associated with a lower belief in opponents choosing dominated actions. However, belief in both the level-1 action (a_2) and the dominance solution (level-2 action, a_1) for longer response times.

Second, in the 3 Steps game, we observe analogous patterns between actions and beliefs that correspond to the same level of sophistication. As it is possible to observe in panel (b) of Figure 6, the frequency of dominated actions decreases for longer response times, similarly to what occurs in panel (a). Moreover, the relation between choice frequency of a_2 and a_3 and response time in the 3 Steps game are strikingly similar to those of actions a_1 and a_2, corresponding to the level-2 and level 1 actions in each game, respectively. Similarly to the patterns in the 2 Steps game, the frequency of the level 1 action (a_3 in 3 Steps, a_2 in 2 Steps) declines slowly, while the level-2 action (a_2 in 3 Steps, a_1 in 2 Steps) is predicted to be chosen with overwhelming frequency when choices take longer. The choice frequency of the dominance solution is, however, rather invariant with respect to response time, remaining low. Beliefs follow similar patterns (Figure 7b): belief in opponent dominated play is lower and belief in both level 1 and level 2 play is higher at later response times. Again a notable difference is that in the 3 Steps game belief in dominance play

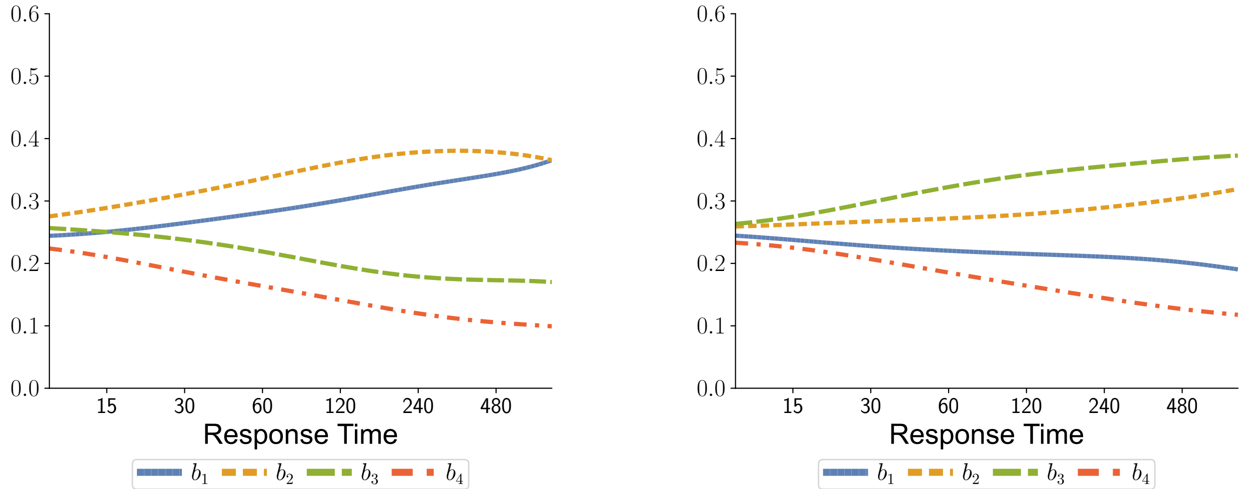

(a) 2 Steps

(b) 3 Steps

Figure 7: Beliefs and Response Time

Notes: This figure shows the predicted relation between mean beliefs and response time (in seconds). The predictions are given by kernel estimation of the relation between subjects' choices, on the one hand, and response time, incentive treatments, and individual characteristics, estimated separately for each game. The expected beliefs conditional on response times is obtained by using the estimated joint conditional probability distribution over a grid on the simplex $\Delta(A_i)$ and marginalizing over covariates other than response times. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Individual characteristics refer to the subjects' age, sex, education, and prior exposure to game theory.

— which corresponds to the level 3 action a_1 — is lower for longer response times.

In short, the level 1 action are the more frequent in fast decisions, being the intuitive choice (maximizing expected payoff against a uniform distribution), and their frequency declines in favor of level 2 action, while dominated choices also decline with response time. Relatedly, fast choices are associated to beliefs being closer to uniform — supporting level 1 actions as a best response — and over time belief in dominated play by opponents declines while belief in both level 1 and 2 play increases.

The estimation exercise together with the relation between incentive levels and response time points toward a consistent mechanism: higher own incentive levels entail higher cognitive effort, proxied by longer response times, which results in both more strategically sophisticated beliefs and actions. This lends support to the intuition that the changes in beliefs and the increase in action sophistication are, at least in part, resulting from greater effort being exerted by the subjects, as suggested by models of sequential reasoning (e.g. Alaoui and Penta, 2016; Gonçalves, 2022), in which players decide whether or not to reason further based on a cost-benefit analysis.

7. Conclusion

This paper provides well-identified evidence that incentive levels affect behavior in strategic settings. Own incentives are crucial in determining mistake propensity and inducing greater cognitive effort, which in turn translates into more accurate beliefs. Beliefs change toward assigning greater sophistication and better predictions of opponent behavior in simple environments, but for more complex environments the differences are minor. While this could suggest that one possibility is that subjects have effective bounds on their ability to reason further, it may also plausibly indicate that to promote similar changes in strategic sophistication sharper increases in incentives are needed in more complex environments than in simpler ones.

Although to a lesser degree, behavior also reacts to opponents' incentive levels. For instance, higher incentives only affect beliefs about opponents' behavior if they also have high incentives. Additionally, reported beliefs both track changes in opponents' choices when opponents' incentive level increases *and* become more biased, denoting that high incentives entail the loss of predictability that noisier (more uniform) behavior confers.

Overall, the greater the incentives, the more the effort exerted, as proxied by response times, which is seen to be related with both greater strategic sophistication of choices and of reported beliefs. Fast response times are associated with beliefs closer to being uniform and to higher frequencies of level 1 actions, while slow response times to beliefs assigning lower probability to dominated play and greater probability to levels 1 and 2 actions as well as being associated with to higher frequency of level 2 choices.

In summary, both payoff-dependent mistakes and exerted cognitive effort (expressed through response times) seem to be important channels to explain how strategic behavior is affected by incentive levels. Furthermore, the effect of increasing incentive levels from low to higher on strategic behavior seems to be more pronounced for simpler environments. Naturally, were our low incentive treatments already enough high-powered, we expect that increasing them further would instead result in more significant effects on less simple environments.

We hope that these insights open exciting new avenues for research. We find two possibilities particularly interesting. First, we believe it is important to explore the limits of incentives in fostering greater strategic sophistication and how this relates to environment complexity. If, on the one hand, there is naturally a limit to what one can achieve in thinking about others' behavior through introspection alone, on the other hand such a limit is surely dependent (or informative about) the environment's degree of complexity. This would also be informative of what constitutes a simple environment, which is of interest for the purpose of designing successful mechanisms. Second, since response times are crucially related to other dimensions of behavior in strategic settings, we consider their study a primary issue. Not only are response times often cheap to collect, they are often reliable indicators of strategic behavior. Further developing game theoretic models and applications that incorporate response times could significantly expand our understanding of observed strategic behavior. Similarly, producing more exploratory experiments to establish basic patterns on how response times relate to behavior and features of the environment is certainly also crucial in guiding theory.

8. References

Agranov, M., E. Potamites, A. Schotter, and C. Tergiman. 2012. "Beliefs and endogenous cognitive levels: An experimental study." *Games and Economic Behavior* 75 (2): 449–63. 10.1016/j.geb.2012.02.002. 5, 8

Alaoui, L., K. A. Janezic, and A. Penta. 2020. "Reasoning about Others' Reasoning." *Journal of Economic Theory* 189 1–51. 10.1016/j.jet.2020.105091. 2, 5, 8, 10, 12

Alaoui, L., and A. Penta. 2016. "Endogenous Depth of Reasoning." *Review of Economic Studies* 83 (4): 1297–1333. 10.1093/restud/rdv052. 1, 5, 7, 9, 26

Alós-Ferrer, C., and J. Buckenmaier. 2021. "Cognitive Sophistication and Deliberation Times." *Experimental Economics* 24 558–592. 10.1007/s10683-020-09672-w. 4, 6, 9

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. 10.1257/aer.101.7.3427. 6

Arad, A., and A. Rubinstein. 2012. "The 11-20 money request game: A level-k reasoning study." *American Economic Review* 102 (7): 3561–3573. 10.1257/aer.102.7.3561. 5, 11

Camerer, C. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton University Press. 1, 2, 6

Camerer, C., T.-H. Ho, and J.-K. Chong. 2004. "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics* 119 (3): 861–898. 10.1162/0033553041502225. 1, 7

Caplin, A., D. Csaba, J. Leahy, and O. Nov. 2020. "Rational Inattention, Competitive Supply, and Psychometrics." *Quarterly Journal of Economics* 135 (3): 1681–724. 10.1093/qje/qjaa011. 1, 4

Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review* 105 (7): 2183–2203. 10.1257/aer.20140117. 6

Clithero, J. A. 2018. "Response times in economics: Looking through the lens of sequential sampling models." *Journal of Economic Psychology* 69 61–86. 10.1016/j.joep.2018.09.008. 6

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. 10.1257/aer.96.5.1737. 4, 11

Costa-Gomes, M. A., and G. Weizsäcker. 2008. "Stated Beliefs and Play in Normal Form Games." *Review of Economic Studies* 75 (3): 729–762. 10.1111/j.1467-937X.2008.00498.x. 1, 5, 7, 17

Dean, M., and N. Neligh. 2022. "Experimental Tests of Rational Inattention." *Journal of Political Economy* Forthcoming. 1, 6

D.Gill, and V. Prowse. 2022. "Strategic Complexity and the Value of Thinking." *Economic Journal* 133 761–86. 10.1093/ej/ueac070. 6

Enke, B., U. Gneezy, B. Hall, D. C. Martin, V. Nelidov, T. Offerman, and J. van de Ven. 2021. "Cognitive Biases: Mistakes or Missing Stakes?" *NBER Working Paper* 28650 1–77. 10.3386/w28650. 1, 6

Friedman, E., and J. Ward. 2022. "Stochastic Choice and Noisy Beliefs in Games: an Experiment." *Working Paper* 1–77. 9

Frydman, C., and S. Nunnari. 2023. "Coordination with Cognitive Noise." *Working Paper.* 10.2139/ssrn.3939522. 4, 6, 9

Fudenberg, D., R. Iijima, and T. Strzalecki. 2015. "Stochastic Choice and Revealed Perturbed Utility." *Econometrica* 83 (6): 2371–2409. 10.3982/ECTA12660. 8

Fudenberg, D., and A. Liang. 2019. "Predicting and Understanding Initial Play." *American Economic Review* 109 (12): 4112–4141. 10.1257/aer.20180654. 5, 12

Georganas, S., P. J. Healy, and R. A. Weber. 2015. "On the Persistence of Strategic Sophistication." *Journal of Economic Theory* 159 369–400. 10.1016/j.jet.2015.07.012. 5

Goeree, J., C. Holt, and T. Palfrey. 2005. "Regular Quantal Response Equilibrium." *Experimental Economics* 8 (4): 347–367. 10.1007/s10683-005-5374-7. 5, 7, 8

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. 10.1257/aer.91.5.1402. 1, 9

Gonçalves, D. 2022. "Sequential Sampling Equilibrium." *Working Paper* 1–53. 10.48550/arXiv.2212.07725. 1, 5, 9, 26

Holt, C. A., J. K. Goeree, and T. Palfrey. 2016. *Quantal Response Equilibrium: A Stochastic Theory of Games.* Princeton University Press. 5

Hossain, T., and R. Okui. 2013. “The Binarized Scoring Rule.” *Review of Economic Studies* 80 (3): 984–1001. 10.1093/restud/rdt006. 12

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. 10.1016/S0167-2681(01)00170-6. 2, 10

Kneeland, T. 2015. “Identifying Higher-Order Rationality.” *Econometrica* 83 (5): 2065–2079. 10.3982/ECTA11983. 4

Luce, D. 1959. *Individual choice behavior.* John Wiley. 5

Matějka, Filip, and Alisdair McKay. 2015. “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model.” *American Economic Review* 105 (1): 272–298. 10.1257/aer.20130047. 6

Mattsson, L.-G., and J. Weibull. 2002. “Probabilistic choice and procedurally bounded rationality.” *Games and Economic Behavior* 41 (1): 61–78. 10.1016/S0899-8256(02)00014-3. 5, 8

McFadden, D. 1974. “Conditional logit analysis of qualitative choice behavior.” In *Frontiers in Econometrics*, edited by Karembka, P. 105–142, Academic Press. 5, 8

McKelvey, R. D., and T. R. Palfrey. 1992. “An Experimental Study of the Centipede Game.” *Econometrica* 60 (4): 803–36. 10.2307/2951567. 1, 2, 6

McKelvey, R. D., and T. R. Palfrey. 1995. “Quantal Response Equilibria for Normal Form Games.” *Games and Economic Behavior* 10 (1): 6–38. 10.1006/game.1995.1023. 1, 5, 8

McKelvey, R. D., T. R. Palfrey, and R. Weber. 2000. “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. 10.1016/s0167-2681(00)00102-5. 6, 9

Nagel, R. 1995. “Unraveling in Guessing Games: An Experimental Study.” *American Economic Review* 85 (5): 1313–1326. 1, 4, 7, 11

Ochs, J. 1995. “Games with Unique, Mixed Strategy Equilibria: An Experimental Study.” *Games and Economic Behavior* 10 (1): 202–17. 10.1006/game.1995.1030. 9

Palacios-Huerta, I., and O. Volij. 2009. “Field Centipedes.” *American Economic Review* 99 (4): 1619–35. 10.1257/aer.99.4.1619. 8

Pearce, D. 1984. “Rationalizable Strategic Behavior and the Problem of Perfection.” *Econometrica* 52 (4): 1029–1050. 10.2307/1911197. 7

Proto, E., A. Rustichini, and A. Sofianos. 2019. “Intelligence, Personality, and Gains from Cooperation in Repeated Interactions.” *Journal of Political Economy* 127 (3): 1351–90. 10.1086/701355. 6

Rapoport, A., W. E. Stein, J. E. Parco, and T. E. Nicholas. 2003. "Equilibrium play and adaptive learning in a three-person centipede game." *Games and Economic Behavior* 42 (3): 239–65. 10.1016/S0899-8256(03)00009-5. 2, 6

Rey-Biel, P. 2009. "Equilibrium play and best response to (stated) beliefs in normal form games." *Games and Economic Behavior* 65 (2): 572–585. 10.1016/j.geb.2008.03.003. 5, 7, 17

Rubinstein, A. 2007. "Instinctive and Cognitive Reasoning: A Study of Response Times." *The Economic Journal* 117 (523): 1243–1259. 10.1111/j.1468-0297.2007.02081.x. 6

Rubinstein, A. 2016. "A Typology of Players: Between Instinctive and Contemplative." *Quarterly Journal of Economics* 131 (2): 859–90. 10.1093/qje/qjw008. 6

Schotter, A., and I. Trevino. 2021. "Is Response Time Predictive of Choice? An Experimental Study of Threshold Strategies." *Experimental Economics* 24 (1): 87–117. 10.1007/s10683-020-09651-1. 6

Spiliopoulos, L., and A. Ortmann. 2018. "The BCD of response time analysis in experimental economics." *Experimental Economics* 21 (2): 383–433. 10.1007/s10683-017-9528-1. 6

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. 10.1016/0167-2681(94)90103-1. 7

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. 10.1006/game.1995.1031. 1, 7

Appendix A. Supporting Tables and Figures

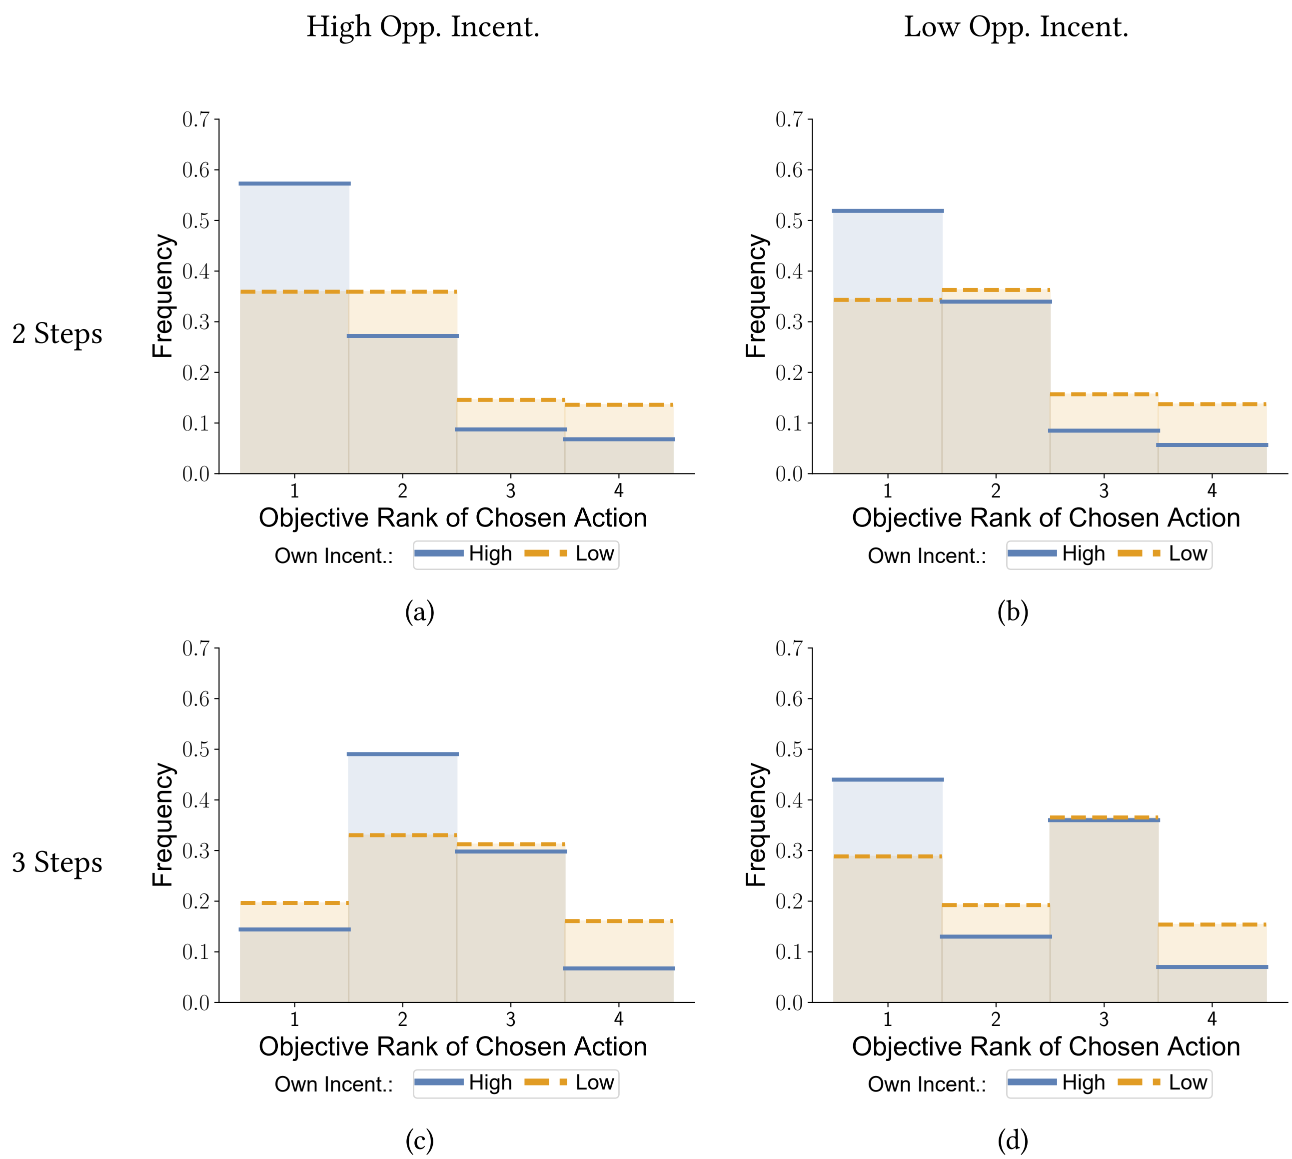

Figure 8: Incentive Level and Objective Mistakes (Hypothesis 2b)

Notes: The different panels exhibit the frequency of the objective rank of the action chosen by the subject for high and low own incentive levels, for different games (2 Steps, (a) and (b), and 3 Steps, (c) and (d)) and holding fixed the opponent's incentive level (High, (a) and (c), or Low, (b) and (d)). The objective rank of an action is n if the action entails the n-th highest objective expected payoffs according to the observed action frequency; e.g. actions with objective rank 1 are those that maximize (objective) expected payoffs. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2).

	a_1		a_2		a_3		a_4	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	2 Steps	3 Steps	2 Steps	3 Steps	2 Steps	3 Steps	2 Steps	3 Steps
Log(Response Time)	0.131***	0.024	-0.021	0.131***	-0.053**	-0.083***	-0.057***	-0.072***
	(0.033)	(0.025)	(0.032)	(0.029)	(0.022)	(0.028)	(0.021)	(0.020)
High Own Incent.	0.140***	-0.072*	-0.050	0.110**	-0.041	0.030	-0.048	-0.068**
	(0.052)	(0.039)	(0.050)	(0.050)	(0.033)	(0.047)	(0.032)	(0.032)
High Opp. Incent.	0.004	0.005	-0.028	0.029	0.004	-0.047	0.020	0.013
	(0.048)	(0.038)	(0.048)	(0.048)	(0.032)	(0.046)	(0.028)	(0.032)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-Squared	0.17	0.05	0.07	0.13	0.08	0.10	0.08	0.12
Observations	414	420	414	420	414	420	414	420

Heteroskedasticity-robust standard errors in parentheses.

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

Table 8: Action Frequency and Response Time

Notes: This table shows the results for the association between choice of a particular action a_n and response time (in seconds) and incentive level treatments, controlling for individual characteristics. Different columns refer to linear probability models relating the choice of an action in a game (2 Steps in columns (1), (3), (5), and (7); 3 Steps in columns (2), (4), (6), and (8)). Errors across regressions are correlated by construction and the table is to be taken as describing measures of association supporting the patterns described in Figure 6. High Own/Opponent Incentives correspond to indicators for whether the subject and their opponent face a high incentive level. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Controls refer to the subjects' age, sex, education, and prior exposure to game theory.

	b_1		b_2		b_3		b_4	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	2 Steps	3 Steps	2 Steps	3 Steps	2 Steps	3 Steps	2 Steps	3 Steps
Log(Response Time)	0.049***	-0.022**	0.054***	0.019***	-0.045***	0.056***	-0.057***	-0.053***
	(0.010)	(0.009)	(0.010)	(0.007)	(0.008)	(0.011)	(0.006)	(0.008)
High Own Incent.	0.013	-0.013	0.014	0.007	-0.010	0.005	-0.017	0.001
	(0.016)	(0.014)	(0.016)	(0.011)	(0.012)	(0.018)	(0.011)	(0.012)
High Opp. Incent.	0.015	-0.017	-0.004	0.030***	-0.004	-0.004	-0.007	-0.010
	(0.016)	(0.014)	(0.016)	(0.011)	(0.012)	(0.018)	(0.010)	(0.012)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-Squared	0.14	0.07	0.16	0.08	0.19	0.11	0.25	0.16
Observations	414	420	414	420	414	420	414	420

Heteroskedasticity-robust standard errors in parentheses.

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

Table 9: Beliefs and Response Time

Notes: This table shows the results for the association between belief b_n in opponents choosing a particular action a_n and response time (in seconds) and incentive level treatments, controlling for individual characteristics. Different columns refer to linear regressions relating the reported beliefs referring to a particular action in a given game (2 Steps in columns (1), (3), (5), and (7); 3 Steps in columns (2), (4), (6), and (8)). Errors across regressions are correlated by construction and the table is to be taken as describing measures of association supporting the patterns described in Figure 7. High Own/Opponent Incentives correspond to indicators for whether the subject and their opponent face a high incentive level. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Controls refer to the subjects' age, sex, education, and prior exposure to game theory.

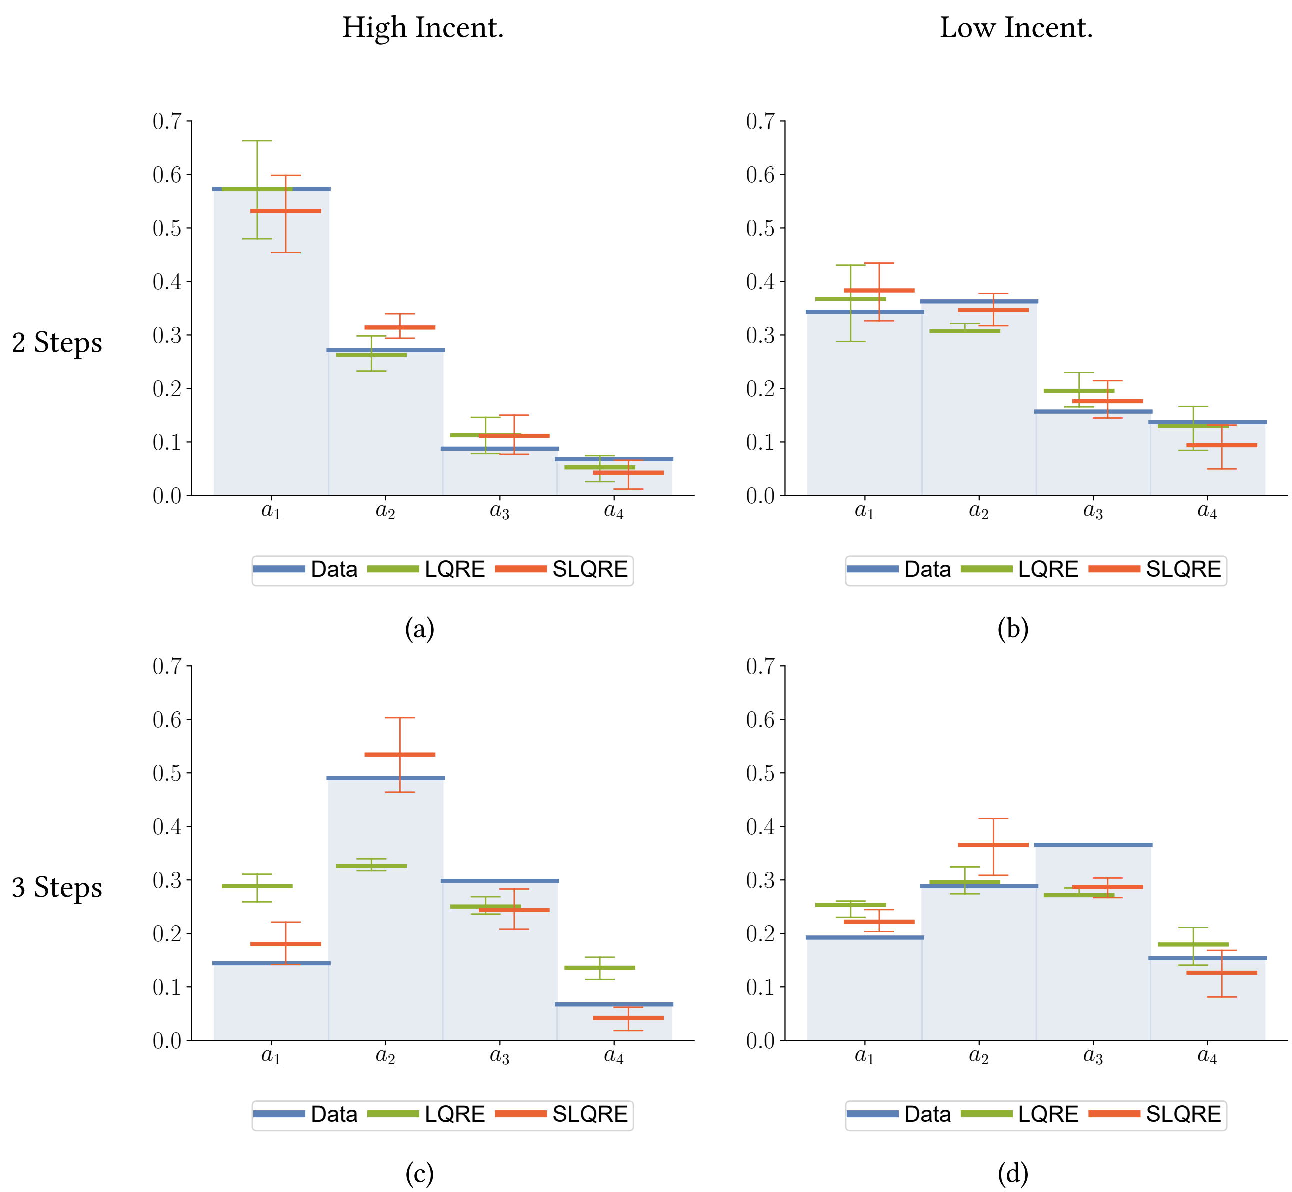

Figure 9: Logit Quantal Response Equilibrium: Maximum Likelihood Estimation

Notes: The different figure displays the empirical frequency of choices in the data (blue), the predicted choice frequencies obtain via maximum likelihood estimation of logit quantal response equilibrium (green), and the predicted choice frequencies obtain via maximum likelihood estimation of subjective logit quantal response (red). Subjective logit quantal response corresponds to logistic multinomial regression using subjective expected payoffs according to subjects' reported beliefs. Different panels correspond to different incentive levels (High, (a) and (c), and Low, (b) and (d)) and to different games (2 Steps, (a) and (b), and 3 Steps, (c) and (d)). We restrict observations to cases in which own and opponents' incentive levels are the same, since cases for which they do not match would correspond to off-equilibrium play. 2 Steps and 3 Steps denote the different games in the experiment (see Figure 2). Asymptotically consistent 95% confidence intervals obtained through bootstrapping with 10,000 replications are given for logit quantal response equilibrium choice frequencies (green whiskers) and for subjective logit quantal response (red whiskers).

Online Appendix A. Additional Tables and Figures

A.1. Sample Characteristics

The questionnaire asked basic demographic information: age, sex, education level, education field, and prior exposure to game theory. 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; 44.36% of the subjects identified themselves as women and 55.64% as men. Reported education ISCED levels are presented in Table 10, and education fields in Table 11. 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. We verified that 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).

	Education Level						
ISCED Level	≤2	3	4	5	6	7	8
Count	8	107	43	228	352	93	3

Table 10: Subjects per Education ISCED Level

Notes: 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.

Education Field	Count
Agriculture, Forestry, Fisheries and Veterinary	6
Arts and Humanities	128
Business, Administration and Law	146
Computer Science, Information and Communication Technologies	130
Economics	29
Education	42
Engineering, Manufacturing and Construction	34
Generic	107
Health and Welfare	54
Mathematics and Statistics	21
Natural Sciences	44
Services (Transport, Hygiene and Health, Security and Other)	39
Social Sciences and Journalism	54

Table 11: Subjects per Education Field

Notes: This classification corresponds to ISCED field education classification, to which economics was added as a separate category.

Online Appendix B. Instructions and Interface

Below we reproduce screenshots with the instructions, practice rounds, the only main (incentivized) round, and the final questionnaire.

Instructions

WELCOME!

After you start the experiment, please focus and avoid multitasking or taking breaks.
Beyond the instructions, the task is extremely short.
This is extremely important for our research.

Please settle in and click the Start button to continue with the instructions.

Start

Introduction

Outline

You are about to participate in an experiment on the economics of decision-making.

You can take 5min or 45min completing this experiment: the time you take depends exclusively on how much effort you want to put in to make a good decision to get the bonus.
Just by completing the experiment, you will secure a **minimum of $2.00**.
You can earn up to a **maximum of $22.00** from this short experiment, depending on your performance.

You will begin, on the next screen, with the instructions. Please read them carefully.
At the end of the instructions there will be questions to check that you understand how the experiment works.
Upon answering these questions correctly, you will proceed to the main task.

The main task has one single round. Your payment will depend on your performance in the main task.
The goal of the experiment is to study how people reason and act in contexts of strategic interaction.

Before the main task begins there will be two practice rounds for you to familiarize yourself with the interface.

After the main task, there is a brief questionnaire.

'Bot'-Detection

This task is designed for humans and cannot be fulfilled using automated answers.

You will be asked to prove you are complying with this requirement by transcribing words at random points in this task. The text will be as legible as the text in these instructions. Any human able to read this text will be able to read the words for transcription, but a 'bot' will not.

You will be allowed 3 attempts and 1 minute per attempt. If you fail to transcribe a word three times, the task will be immediately terminated and you will not be able to complete it nor receive payment.

Quitting the Task

You can quit the task at any moment, but if you do not complete it you will not receive any payment.

Additional Information

In the experiment you will answer questions which ask you to choose between different options. Your responses to this experiment will be used to study how people choose when the outcome depends on others' choices as well.

No identifying data about you will be made available and all data we store will be anonymized. All data and published work resulting from this experiment will maintain your individual privacy.

Please understand that your participation in this experiment is voluntary. If you consent to participate, you have the right to withdraw consent or discontinue participation at any time. For additional details, see this information sheet.

Next

Instructions

Click on the section to show/hide.

Payoff Structure

In this experiment, you are going to be asked to make choices and predictions. Your bonus will depend on the choices and predictions you make and on how other participants choose.
How much you get from your choices depends on the payoff structure that defines your action payoffs.

A payoff structure describes how many "Action Points" each agent obtains depending on the actions chosen by the two agents.

		Orange			
		Action A	Action B	Action C	Action D
Blue	Action 1	10 60	20 30	40 20	70 10
	Action 2	90 10	80 70	20 70	40 20
	Action 3	20 20	50 80	10 40	20 90
	Action 4	30 30	40 10	50 60	80 80

For example, the above figure is a payoff structure.
There are two agents, **Blue** and **Orange**.
- Each row corresponds to an action that the **Blue** agent can take and each column corresponds to an action the **Orange** agent can take.
- The "Action Points" that agents obtain depend on what both choose.
Each cell corresponds a unique combination of a choice of an action (row) by **Blue** and the choice of an action (column) by **Orange**, determining the "Action Points" they get in that case.
In each row (column) you can see all the "Action Points" for the **Blue** (**Orange**) agent associated with choosing that row (column), which also depend on what **Orange** (**Blue**) agent chooses.

Blue's "Action Points" are blue and in the lower left corner of the cell.
Orange's "Action Points" are orange and in the upper right corner of the cell.

Example 1: If **Blue** chooses action 1 and **Orange** chooses action C, then **Blue** gets 20 "Action Points" and **Orange** gets 40 "Action Points".
Example 2: If **Blue** chooses action 2 and **Orange** chooses action C, then **Blue** gets 70 "Action Points" and **Orange** gets 20 "Action Points".
Both agents will have 4 actions they can choose from.

"Action Points" and Choices: Examples

Suppose you are **Blue**.
If you think **Orange** agents choose action A, then action 1 gives you the most "Action Points".
If instead you think **Orange** agents choose action B, then action 3 gives you the most "Action Points".
If you think that on average 50% of the **Orange** agents choose action B and 50% choose action C, then action 2 will give you the most "Action Points" on average.
What **Orange** agents do on average is often crucial to choose the action that gives you the most action points on average.
You can also think about **Orange**'s choices: **Orange** agents would never want to choose action C because no matter what they think **Blue** agents will do, another action delivers more "Action Points".
Also, **Orange** agents would only choose action A if they thought **Blue** agents are very likely to choose action 2.

You will now have a practice round to understand "hands-on" how your and the other agent's choices affect how many points you get.

Next

PRACTICE ROUND

		Orange			
		Action A	Action B	Action C	Action D
Blue	Action 1	30 70	70 30	40 40	20 80
	Action 2	70 30	90 10	80 20	40 40
	Action 3	40 40	80 20	70 30	30 70
	Action 4	20 80	40 40	30 70	10 90

Click to select your action.

You can only proceed after you provide a guess and select an action.

Instructions

Click on the section to show/hide.

Payoff Structure

Opponents

In the task, you will face a payoff structure.
Participants other than you will face the same payoff structure.

The **Other participants** are sorted into one of two groups:
- **HIGH STAKES** group; and
- **LOW STAKES** group.
Other participants in the **HIGH STAKES** group have a bonus of **$20.00**.
Other participants in the **LOW STAKES** group have a bonus of **$0.50**.
Some are **Blue** agents, others **Orange** agents.
Within each group, the **Other participants** are randomly matched with an opponent of a different color.
This means that **Other participants** always face opponents with **the same STAKES**.

The **Other participants**:
- Are asked (1) to *guess* what opponents (participants of different color) with **the same STAKES** choose on average and (2) to choose an *action*.
- Are compensated the same way as you are (details below), getting "Action Points" and "Guess Points".
- Do not know who they are matched with nor others' choices.

You will be a **Blue** agent and you will be randomly matched with an **Orange** opponent.
Your **Orange** opponent will be from either the **HIGH STAKES** group ($20.00 bonus) or the **LOW STAKES** group ($0.50 bonus) with equal probability.
You will have **HIGH STAKES** ($20.00 bonus) or **LOW STAKES** group ($0.50 bonus) with equal probability.
You may or may not have the same STAKES as your **Orange** opponent.

There are then four possible situations:

You have **HIGH STAKES** ($20.00 bonus) **Orange** opponent is from the **HIGH STAKES** group ($20.00 bonus)	**You** have **HIGH STAKES** ($20.00 bonus) **Orange** opponent is from the **LOW STAKES** group ($0.50 bonus)
You have **LOW STAKES** ($0.50 bonus) **Orange** opponent is from the **HIGH STAKES** group ($20.00 bonus)	**You** have **LOW STAKES** ($0.50 bonus) **Orange** opponent is from the **LOW STAKES** group ($0.50 bonus)

You will be asked:
1. To *guess* what your **Orange** opponents in the same STAKES group as your opponent choose on average; and
2. To choose an *action*.

Depending on how correct your guess is (details below), you will get "Guess Points".
Depending on **Your** action and your **Orange** opponent's action, you will get "Action Points" according to the payoff structure.
Your **Orange** opponent is matched to some other **Blue** participant, and therefore:
- **Your** action will not affect your **Orange** opponent's "Action Points";
- But your **Orange** opponent's actions does affect your "Action Points".

Before you make guesses and choose your action, you will be informed of:
- Your **Orange** opponent STAKE group;
- **Your** own STAKES.

"Action Points" and "Guesses Points"

You can get Points either from the actions you choose ("Action Points") or from your guesses as to what **Orange** opponents may do ("Guess Points").
The Points used to implement the **bonus** are either the "Action Points" or the "Guess Points", chosen randomly with equal probability.
You will only get Points from your actions or your guesses, not from both.

The "Action Ponts" you obtain depend **Your** action and your **Orange** opponent's action.
You will get "Action Points" according to the payoff structure.

The "Guess Points" you obtain depend on how accurate your guess is.
You are asked to guess the probability that **Orange** agents such as your opponent are to choose any given action.
If you guess correctly, you get 100 "Guess Points".
If you guess completely incorrectly, you get 0 "Guess Points".
The further away your guess is from the average of others' choices, the fewer "Guess Points" you get.

The rule that determines your "Guess Points" is set so that it is always in your best interest to state what you believe.

The exact rule that determines your points from your guess is as follows:
If y_i is the fraction of subjects who chose action $i = A, B, C, D$, and g_i is your corresponding guess divided by 100 -- so that both y_i and g_i are between 0 and 1 --,

$$\text{Guess Points} = 100 \times \left(1 - \frac{1}{2}\left((y_A - g_A)^2 + (y_B - g_B)^2 + (y_C - g_C)^2 + (y_D - g_D)^2\right)\right)$$

Points, Stakes and Bonus

The Points you get determine the probability with which you will get a **bonus**.
Example: if you get 75 Points, you will get the **bonus** with 75% probability and nothing with 25% probability.

How big the **bonus** is depends on the **STAKES** level you are assigned to.

The bonus for participants assigned to the **HIGH STAKES** level is of **$20.00**.
In this case: Points = Probability of getting a $20.00 bonus.

The bonus for participants assigned to the **LOW STAKES** level is of **$0.50**.
In this case: Points = Probability of getting a $0.50 bonus.

Every participant will be assigned to either the **HIGH STAKES** level or the **LOW STAKES** level randomly and with equal probability (50%-50%).

You should pay attention to whether you are in the **HIGH STAKES** level or the **LOW STAKES** level.
Mistakes and hasty decisions when you are assigned to the **HIGH STAKES** level may severely decrease the probability of getting the $20.00 bonus.
Mistakes and hasty decisions when you are assigned to the **LOW STAKES** level may severely decrease the probability of getting the $0.50 bonus.

It is important to know the following:

- The participants are randomly matched and no one knows others' choices or guesses.
No one knows who they are going to be matched with.

- Other participants have the same compensation scheme as you and observe the same payoff structure.
The order of the payoff structures and labels of the actions are randomized.

- **Your** choices do not affect your **Orange** opponent's "Action Points".
But your **Orange** opponent's action does affect **Your** "Action Points".

- Your **Orange** opponent, their opponent, their opponent's opponent, and so on, all have been assigned to the **same STAKES** group.
But **You** may or may not have been assigned to the same **STAKES** level as your **Orange** opponent.

- Participants are recruited under the same conditions as you.

- It is crucial to pay attention to **STAKES** level of your **Orange** opponent as well as the **STAKES** level **You** were assigned to.

The **STAKES** level determines the bonus.

Interface

This is how the interface looks like:

You have
HIGH STAKES (**$20.00 bonus**).

Your **Orange** opponent is from the
LOW STAKES group (**$0.50 bonus**).

You can only proceed after you provide a guess and select an action.
The payoff structure will appear momentarily.

		Orange			
		Action A	**Action B**	**Action C**	**Action D**
Blue	Action 1	10 60	20 30	40 20	70 10
	Action 2	90 10	80 70	20 70	40 20
	Action 3	20 20	50 80	10 40	20 90
	Action 4	30 30	40 10	50 60	80 80
	Probability Guess	%	%	%	%

What is the probability that **Orange** agents
in the LOW STAKES group ($0.50 bonus)
are to choose any given action?

The accuracy of your guess will determine your "Guess Points".

Use the sliders above to give your answer.

One **Orange** agent will be randomly selected as your opponent
from the LOW STAKES group ($0.50 bonus).

Your **Orange** opponent's action and **Your** action will determine your "Action Points".

Click to select your action.

Instructions

Duration and payments

The main task has ONLY ONE round plus 2 practice rounds.
The experiment concludes with a quick questionnaire.

We expect it takes you
< 10 min Instructions
20sec - 40 min Main task (depending on you entirely)

< 1 min Questionnaire
Total: 10-55 min.

You will get $2.00 just by completing the experiment.
You can get up to $22.00 depending on the **STAKES** level you are assigned to and how you perform on the main task.
Remember: hasty choices may substantially harm your chances of getting a high bonus.

Details on how the bonus is computed

At the start of the experiment, you will be randomly sorted into HIGH STAKES or LOW STAKES, with equal probability (50-50%).
The group of your opponent is then selected: HIGH STAKES group or LOW STAKES group, with equal probability (50-50%).
Your opponent is randomly chosen among all other participants in that group, with equal probability.

Your and your opponent's action choices determine your "Action Points" according to the payoff structure you both observe.
Your guess and the actions of participants (other than you) in the group your opponent is drawn from determine your "Guess Points" according to the formula above.

Randomly and with equal probability, either your "Action Points" or your "Guess Points" are chosen to determine your bonus.
A number y between 0 and 100, is generated uniformly at random (all numbers have equal probability).
If that number y is smaller than your Points, that is, with probability = Points (%) and you have **HIGH STAKES**, you get $20.00.
If that number y is smaller than your Points, that is, with probability = Points (%) and you have **LOW STAKES**, you get $0.50.
If that number y is larger than your Points, you get no bonus.

Next

Instructions

Click on the section to show/hide.

Payoff Structure

Opponents

"Action Points" and "Guesses Points"

Points, Stakes and Bonus

It is important to know the following:

- The participants are randomly matched and no one knows others' choices or guesses.
No one knows who they are going to be matched with.

- Other participants have the same compensation scheme as you and observe the same payoff structure.
The order of the payoff structures and labels of the actions are randomized.

- **Your** choices do not affect your **Orange** opponent's "Action Points".
But your **Orange** opponent's actions does affect **Your** "Action Points".

- Your **Orange** opponent, their opponent, their opponent's opponent, and so on, all have been assigned to the **same STAKES** group.
But **You** may or may not have been assigned to the same **STAKES** level as your **Orange** opponent.

- Participants are recruited under the same conditions as you.

- It is crucial to pay attention to **STAKES** level of your **Orange** opponent as well as the **STAKES** level **You** were assigned to.
The **STAKES** level determines the bonus.

Duration and payments

Details on how the bonus is computed

Questions

You must answer the following questions correctly before you can proceed.

Refer to the following hypothetical interface screenshot:

You have
HIGH STAKES (**$20.00 bonus**).

Your **Orange** opponent is from the
LOW STAKES group (**$0.50 bonus**).

You can only proceed after you provide a guess and select an action.
The payoff structure will appear momentarily.

		Orange			
		Action A	**Action B**	**Action C**	**Action D**
Blue	**Action 1**	10 60	20 30	40 20	70 10
	Action 2	90 10	80 70	20 70	40 20
	Action 3	20 20	50 80	10 40	20 90
	Action 4	30 30	40 10	50 60	80 80
	Probability Guess	%	%	%	%

What is the probability that **Orange** agents
in the LOW STAKES group ($0.50 bonus)
are to choose any given action?

The accuracy of your guess will determine your "Guess Points".

Use the sliders above to give your answer.

One **Orange** agent will be randomly selected as your opponent
from the LOW STAKES group ($0.50 bonus).

Your **Orange** opponent's action and **Your** action will determine your "Action Points".

Click to select your action.

Instructions

1.
What are your stakes?

◯ High stakes.

◯ Low stakes.

2.
What are your opponent's stakes?

◯ High stakes.

◯ Low stakes.

3.
How much is the high stakes bonus?

- $0.20.
- $0.50.
- $10.00.
- $20.00.

4.
How much is the low stakes bonus?
- $0.20.
- $0.50.
- $10.00.
- $20.00.

5.
If all the Orange agents choose action B, which action gives you the highest number of points?
- Action 1.
- Action 2.
- Action 3.
- Action 4.

6.
If 80% of the Orange agents chooses action A and 20% chooses action D, which action gives you the highest number of points on average?
- Action 1.
- Action 2.
- Action 3.
- Action 4.

7.
If 50% of the Blue agents choose action 1 and 50% choose action 2, which action gives the Orange agents the highest number of points on average?
- Action A.
- Action B.
- Action C.
- Action D.

8.
Your opponent's action matters for your Action Points but your action does not affect your opponent's Action Points.
- True.
- False.

9.
You are going to be randomly matched with a Orange agent. That Orange agent is going to be randomly matched with some other Blue agent.
- True.
- False.

10.
Your opponent and your opponent's opponent (and so on) always have the same stakes level.
- True.
- False.

Check Answers

Captchas

Bot Detection - Attempt 1

Type the following word or phrase into the box below, then press 'Next'. Answers are not case-sensitive.
You have three attempts. If you fail all three attempts, the task will end and you will not be paid.
You have one minute per attempt.

Payoff Structure

Next

59

Practice Rounds

PRACTICE

You will now play two rounds as a practice, to familiarize yourself with the interface.

Your choices and guesses in these rounds do not affect your pay.

Next

PRACTICE ROUND

You have
HIGH STAKES (**$20.00 bonus**).

Your **Orange** opponent is from the
HIGH STAKES group (**$20.00 bonus**).

You can only proceed after you provide a guess and select an action.
The payoff structure will appear momentarily.

PRACTICE ROUND

You have
HIGH STAKES (**$20.00 bonus**).

Your **Orange** opponent is from the
HIGH STAKES group (**$20.00 bonus**).

You can only proceed after you provide a guess and select an action.
The payoff structure will appear momentarily.

		Orange			
		Action A	Action B	Action C	Action D
Blue	Action 1	90 / 10	40 / 40	80 / 20	70 / 30
	Action 2	70 / 30	20 / 80	40 / 40	30 / 70
	Action 3	80 / 20	30 / 70	70 / 30	40 / 40
	Action 4	40 / 40	10 / 90	30 / 70	20 / 80
	Probability Guess	14.3%	65.5%	18.0%	2.2%

What is the probability that **Orange** agents
in the HIGH STAKES group ($20.00 bonus)
are to choose any given action?

The accuracy of your guess will determine your "Guess Points".

Use the sliders above to give your answer.

One **Orange** agent will be randomly selected as your opponent
from the HIGH STAKES group ($20.00 bonus).

Your **Orange** opponent's action and **Your** action will determine your "Action Points".

Click to select your action.

Click next to submit your choice and guess.
You will not be able to go back.

Next

Instructions

Main Task

MAIN TASK

The main experiment will now start.

By completing the experiment, you will secure a $2.00 completion fee,
(in addition to the bonus you may get) regardless of how well you do.

In addition to the completion fee, you can get a bonus depending on how well you do.
Remember that Points = Probability of getting the bonus.

The size of the bonus you can get depends on whether you are assigned to
- the **HIGH STAKES**, where you can get a **$20.00 bonus**; or
- the **LOW STAKES**, where you can get a **$0.50 bonus**.

You will be assigned to one of these with equal probability (50%-50%).
Click Next to find out which you've been assigned to.

Next

You are going to enter the MAIN TASK!

You have
HIGH STAKES (**$20.00 bonus**).
Your choices and guesses count towards a **$20.00 bonus.**

This is the only round.
Remember that Points = Probability of getting the bonus.

Your **Orange** opponent is from the
LOW STAKES group (**$0.50 bonus**).

By completing the experiment, you will secure a $2.00 completion fee,
(in addition to the bonus you may get) regardless of how well you do.
So be sure to complete the experiment to get the completion fee.

You have
HIGH STAKES (**$20.00 bonus**).

Your **Orange** opponent is from the
LOW STAKES group (**$0.50 bonus**).

You can only proceed after you provide a guess and select an action.
The payoff structure will appear momentarily.

You have
HIGH STAKES (**$20.00 bonus**).

Your **Orange** opponent is from the
LOW STAKES group (**$0.50 bonus**).

You can only proceed after you provide a guess and select an action.
The payoff structure will appear momentarily.

		Orange			
		Action A	Action B	Action C	Action D
Blue	Action 1	10 / 10	30 / 70	40 / 40	20 / 10
	Action 2	40 / 40	10 / 20	10 / 10	70 / 30
	Action 3	30 / 70	70 / 30	10 / 20	40 / 40
	Action 4	20 / 10	40 / 40	70 / 30	30 / 70
	Probability Guess	%	%	%	%

What is the probability that **Orange** agents
in the LOW STAKES group ($0.50 bonus)
are to choose any given action?

The accuracy of your guess will determine your "Guess Points".

Use the sliders above to give your answer.

One **Orange** agent will be randomly selected as your opponent
from the LOW STAKES group ($0.50 bonus).

Your **Orange** opponent's action and **Your** action will determine your "Action Points".

Click to select your action.

Instructions

Questionnaire

Questionnaire: Socio-Demographics

Please enter your age:

Please state your sex:

- Male
- Female

What is the HIGHEST LEVEL OF EDUCATION that you COMPLETED in school?

- None or Primary Education: Primary School (grades 1-6)
- Lower Secondary Education: Middle School or some High School incomplete
- Upper Secondary Education: High School
- Business, technical, or vocational school AFTER High School
- Some college or university qualification, but not a Bachelor
- Bachelor or equivalent
- Master or Post-graduate training or professional schooling after college (e.g. law or medical school')
- Ph.D or equivalent

Choose the field that best describes your PRIMARY FIELD OF EDUCATION.

- Generic
- Arts and Humanities
- Social Sciences and Journalism
- Economics
- Education
- Business, Administration and Law
- Computer Science, Information and Communication Technologies
- Natural Sciences
- Mathematics and Statistics
- Engineering, Manufacturing and Construction
- Agriculture, Forestry, Fisheries and Veterinary
- Health and Welfare
- Services (Transport, Hygiene and Health, Security and Other)

Did you ever learn Game Theory?

- Yes, as an graduate course.
- Yes, as an undergraduate course.
- Yes, as an online course, Summer school course or similar.
- Yes, for professional reasons.
- Yes, searching the web / out of personal interest.
- No.

Next

You must answer each question before you can continue.

Payment

Payment

As explained in the instructions, you will be paid a completion fee and a bonus.

The completion fee is made of $2.00 to be paid upon conclusion, regardless of how you did.
$0.50 out of the $2.00 completion fee will be paid immediately, as soon as your HIT is approved.
The remainder, $1.50, will be paid together with the bonus.
We expect to be quick (less than 48h) in approving and paying the bonus.

For the main task, you will be matched with another subject and the actions that you both chose in that game will determine how many points you will get from your choices.
How far off your guess is from the average choices of the other participants will determine how many points you will get from your guesses.
One the two will be chosen with equal probability as your final points.
Those points will then be used as the probability of winning a bonus corresponding to the stakes level you were assigned to.

As you were assigned to the HIGH STAKES level, your points count towards the probability of getting a $20.00 bonus.
You will earn the bonus with Probability (%) = Points.

For instance, if you make 90 points, you will earn the bonus with 90% probability.

We will now ask you to complete a comments section.

Click 'Next' to continue to the comments section.

Next