Experimental Economics (2026), 1–26
doi:10.1017/eec.2026.10056

SPECIAL ISSUE ARTICLE

Incentives and strategic behaviour: an experiment

Teresa Esteban-Casanelles[1] and Duarte Gonçalves[2]

[1]Department of Political Economy, King's College London, London, UK
[2]Department of Economics, University College London, London, UK
Corresponding author: Teresa Esteban-Casanelles; Email: teresa.estebancasanelles@kcl.ac.uk

(Received 5 February 2025 UTC; revised 26 March 2026 UTC; accepted 25 June 2026 UTC)

Abstract

How do incentive levels affect strategic behaviour? We address this with an experiment that separately identifies own- and opponent-incentive effects in two dominance-solvable games that differ in strategic complexity. Higher own incentives favour more strategically sophisticated actions and increase best responding to stated beliefs. Beliefs shift in a parallel direction and participants expect more sophisticated opponent actions. Furthermore, although higher own incentives increase belief accuracy, opponents with higher incentives are harder to predict. Higher incentives also increase response times, and longer response times are associated with better performance and more sophisticated actions and beliefs. Taken together, the evidence suggests that incentives affect strategic behaviour through two main channels: by reducing payoff-dependent mistakes and by increasing the effort devoted to reasoning, with the returns to that effort shaped by strategic complexity of the environment.

Keywords: behavioural game theory; belief formation; costly cognition; incentives; quantal response; response times

JEL Classifications: C72; C92; D83; D84; D91

1. Introduction

A large experimental literature documents systematic deviations from Nash equilibrium and even from rationalisable behaviour, and shows that behaviour in strategic settings responds to changes in relative incentives.[1] This brought to the fore a variety of behavioural models of play in strategic settings, featuring limited reasoning ability (Camerer et al., 2004; Nagel, 1995; Stahl & Wilson, 1995), payoff-dependent mistakes (McKelvey & Palfrey, 1995), random beliefs (Friedman & Mezzetti, 2005; Friedman, 2022), and costly cognition (Alaoui & Penta, 2016; Gonçalves, 2024). More broadly, these models have been used to explain otherwise-puzzling patterns in settings as diverse as auctions (Bajari & Hortaçsu, 2005), financial bubbles (Moinas & Pouget, 2013), and responses to monetary policy (Farhi & Werning, 2019).

Whilst a large literature shows that behaviour in games responds to changes in relative incentives, much less is known about how the *level* of incentives itself affects strategic behaviour. Strategic settings such as auctions, elections, negotiations, and pricing often differ markedly in the stakes the parties face. Understanding how incentive levels affect strategic behaviour matters for both theory and policy. On the theory side, evidence on how incentive levels affect behaviour can help distinguish whether the key driver is limited reasoning, payoff-dependent mistakes, random beliefs, or costly cognition.

[1]This includes, among many others, Nagel (1995), Ochs (1995), Capra et al. (1999), McKelvey et al. (2000), Costa-Gomes and Crawford (2006), Kneeland (2015), Alaoui et al. (2020), and Friedman and Ward (2022).

© The Author(s), 2026. Published by Cambridge University Press on behalf of the Economic Science Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike licence (http://creativecommons.org/licenses/by-nc-sa/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the same Creative Commons licence is used to distribute the re-used or adapted article and the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

On the policy and design side, many interventions work either by changing the strength of incentives or by changing the information available to decision-makers, and the relative effectiveness of these approaches depends on which channel is operative.

Existing evidence on incentive levels in strategic settings is limited, focuses mostly on whether choices change with incentive levels, and is itself hard to interpret.[2] The evidence is mixed, with choices responding to incentive levels to different extents across strategic environments, which can reflect multiple channels operating at once. Higher incentives make decision mistakes more costly, which can not only reduce mistake-proneness, but may also induce more reasoning and affect the beliefs players form. Moreover, existing designs confound a player's own-incentive effect with indirect effects operating through higher-order beliefs. This is because a player's behaviour may react directly to their own incentive level, through reduced payoff-dependent mistakes or greater effort devoted to reasoning, but it may also react indirectly through beliefs about how opponents respond.

This paper experimentally studies how incentive levels affect strategic behaviour. Our design relies on recent advances in experimental methodology to causally identify how own and opponent incentive levels affect actions, beliefs, and response times in strategic settings. Focusing on dominance-solvable games, we show that higher own incentive levels lead to more strategically sophisticated actions (of higher order of rationality), fewer mistakes, and better responses to stated beliefs. Beliefs are also affected: higher incentives lead participants to expect opponents to play more sophisticated actions. Moreover, whereas higher own incentives improve belief accuracy, opponents with higher incentives are harder to predict precisely. Finally, higher incentive levels increase response time, which, in turn, is positively associated with payoffs and both action and belief sophistication.

Three considerations motivate separating own- and opponent-incentive effects. First, different behavioural models make different predictions about which incentive components should matter.[3] Second, if one varies only overall incentives, then changes in beliefs and higher-order reasoning naturally confound the direct effect of own incentives on behaviour. Finally, since many experimental and applied environments involve asymmetric stakes, understanding separately how behaviour responds to own and opponents' incentive levels is useful both for interpretation and for design.

To study these issues, we provide an experimental design that separately identifies the effects of own and opponents' incentive levels on actions, beliefs, and response times. Participants are randomly assigned to treatments that vary own and the opponents' incentive level and the game played. Our identification strategy relies on a simple opponent-matching protocol adapted from the replacement method (Alaoui & Penta, 2016; Alaoui et al., 2020). This allows us to vary a participant's own incentive level whilst holding fixed the behaviour relevant for their decision problem, and conversely to vary opponents' incentive levels whilst holding fixed the participant's own incentives. The design also fixes higher-order beliefs, in the sense that participants know not only the incentive level faced by their opponent but also the incentive level faced by that opponent's opponent, and so on. We vary incentive levels by randomly assigning participants to a high- or low-bonus treatment, with choices affecting the probability of receiving the corresponding payment bonus. To focus on initial responses and keep per-decision incentives high-powered, participants complete two unincentivised practice rounds without feedback and then make a single incentivised choice.

[2] A significant share of the literature is dedicated to ultimatum bargaining (see, e.g., Andersen et al., 2011; Hoffman et al., 1996; Slonim & Roth, 1998; Straub & Murnighan, 1995), which is prone to other-regarding preferences and cultural norms intervening as a confounding factor (Camerer, 2003). Evidence on other settings includes centipede games (McKelvey & Palfrey, 1992; Rapoport et al., 2003), coordination games (Parravano & Poulsen, 2015), and beauty-contest games (Camerer, 2003, p. 211).

[3] In many standard models, including Nash equilibrium, level-k, cognitive hierarchy, and random beliefs, scaling payoffs leaves predicted behaviour unchanged. By contrast, in models with payoff-dependent mistakes, choices react to own incentives when holding fixed the opponent choices, but beliefs do not. In models with costly reasoning, own incentives matter for both choices and beliefs, which should also react to opponents' incentives.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

We implement this design in two symmetric dominance-solvable normal-form games that differ in strategic complexity. We focus on dominance-solvable games because they provide a clean, theory-grounded ranking of actions and beliefs by strategic sophistication. In the simpler game, the dominance solution is reached after two rounds of maximal iterated elimination of strictly dominated strategies; in the more complex game, it is reached only after three rounds. In both games, level-k actions are identified in a way that is robust to risk attitudes. The games share a number of features with standard environments used to study strategic sophistication — including beauty-contest, undercutting, ring, and related dominance-solvable games (Costa-Gomes & Crawford, 2006; Kneeland, 2015; Nagel, 1995) — which enables us to relate our results to the existing literature. This makes them especially useful for studying whether incentives affect mistakes, higher-order reasoning, and belief formation in a tightly controlled strategic setting. This results in a $2\times2\times2$ between-participant design with random assignment, in which over 1600 participants recruited online play only one game and only once.

The paper features four main sets of results. First, higher own incentives make actions more sophisticated, but the relevant margin depends on game complexity. In both games, higher own incentives reduce dominated play and increase play of the level-2 action. The level-2 action corresponds to the dominance-solution action in the simpler game. In the more complex game, the level-3 action is the dominance-solution action, but its frequency does not change with higher incentives. Opponents' incentive levels have qualitatively similar but substantially weaker effects.

Second, the data support a payoff-sensitivity channel consistent with a subjective version of quantal response. Actions with higher *subjective* expected payoff (as per stated beliefs) are chosen more often, but this payoff-rank monotonicity does not carry over to *objective* expected payoffs (as given by empirical frequency of opponents' actions). Higher own incentives not only substantially increase best-response rates to stated beliefs but also shift choices toward better responses to participants' reported beliefs. A subjective logit quantal-response specification fits the data much better than an analogous specification based on empirical opponent frequencies, underscoring the importance of elicited beliefs for understanding observed choices.

Third, incentives affect beliefs in a way that mirrors actions. Higher own incentives reduce the probability assigned to dominated play and shift beliefs toward more sophisticated opponent behaviour. In the 2-Step game, this shift extends to the dominance-solution action. In the 3-Step game, however, beliefs move primarily toward intermediate level-1 and level-2 actions rather than toward the level-3 dominance solution. The effects of opponent incentive level are qualitatively similar, but sharply muted. Higher own incentives also improve several dimensions of belief accuracy, whereas opponents facing higher incentives become harder to predict precisely.

Fourth, response times provide suggestive evidence on mechanism. Higher own incentives increase response times by over 40% in both games. Additionally, controlling for incentives and individual characteristics as well as longer response times are associated with better performance: higher best-response rates, higher expected payoffs, fewer dominated choices, and more sophisticated beliefs. Finally, we explore how response times relate to both actions and beliefs and find that deliberating for longer is associated with more sophisticated actions being chosen and with more sophisticated beliefs being reported. Specifically, in both games, the frequency of dominated actions monotonically decreases with time, the level-1 action is the modal action at short response times and its frequency decreases with deliberation, whilst the level-2 action frequency increases. Beliefs track these patterns consistently: longer response times are also associated with greater belief mass on both the level-1 and level-2 actions. Interestingly, the frequency of and the belief about the dominance-solution action in the more complex game, the level-3 action, are fairly invariant over time. Noting that response time is an informative, even if imperfect, proxy for cognitive effort (Caplin et al., 2020; Gill & Prowse, 2022), we interpret these patterns as suggestive evidence that higher incentives affect behaviour through both reduced payoff-dependent mistakes and greater effort devoted to reasoning.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

Our robustness analysis shows that the main patterns are stable within our setup. The results replicate across participant pools, remain qualitatively similar under alternative specifications, and are not driven by participants with weaker attention or poorer comprehension. We also discuss the extent to which alternative explanations based on other-regarding preferences may still matter, whilst showing that the treatment effects in the data are difficult to reconcile with such channels.

These findings should nevertheless be interpreted within the scope of a tightly controlled benchmark environment, precisely one in which direct confounds from feedback, learning, and payoff externalities are reduced. We therefore do not claim to speak to all strategic environments. More generally, it is well known that very high or unfamiliar incentive levels can sometimes backfire, so our conclusions should not be extrapolated mechanically to settings with extreme stakes or to environments in which additional motivational channels are likely to be important (Ariely et al., 2009; Camerer & Hogarth, 1999; Enke et al., 2023).

Within those limits, this paper contributes well-identified evidence on how incentive levels affect reasoning, mistakes, beliefs, and response times in a strategic setting. By separately identifying own- and opponent-incentive effects, the paper clarifies through which channels higher incentives change behaviour and how those effects depend on strategic complexity. Taken together, our results suggest that, at familiar incentive levels, higher incentives increase action sophistication both by reducing the propensity to make mistakes and by increasing the incentive to reason further and form better beliefs. At the same time, the fact that action sophistication and belief accuracy increase in the simpler game, but less so or not at all in the more complex game, implies that one should not expect changes in incentives to have uniform effects across strategic environments that differ in complexity. Our findings therefore caution both against inferring universally low levels of sophistication from evidence obtained under low-powered incentives and against relying on predictions that require high levels of sophistication in complex but low-stakes environments.

The remainder of the paper proceeds as follows. Section 2 presents the hypotheses, experimental design, and identification strategy. Sections 3, 4, 5, and 6 report the results on actions, payoff-dependent mistakes, beliefs, and response times. Section 7 discusses alternative explanations, robustness checks, and limitations. Section 8 concludes.

1.1. Related literature

This paper contributes to four related strands of literature.

Strategic sophistication and bounded rationality. A large literature documents persistent deviations from Nash equilibrium and substantial heterogeneity in strategic sophistication. This is true in dominance-solvable environments such as beauty-contest, undercutting, and ring games (Costa-Gomes & Crawford, 2006; Kneeland, 2015; Nagel, 1995), as well as in many other normal-form games (Arad & Rubinstein, 2012; Fudenberg & Liang, 2019). Experimental work also shows that action sophistication is not stable across games and depends importantly on what players believe about their opponents (Agranov et al., 2012; Georganas et al., 2015). These findings motivated several models of bounded rationality in games, the common theme of which is that observed play reflects a combination of shallow reasoning, payoff-dependent mistakes, and costly cognition.[4] This paper studies the effect of incentive *levels* on choices, beliefs, and response times in an environment where actions and beliefs admit a clear ranking by strategic sophistication. The design allows us to distinguish shifts away from dominated play from shifts toward intermediate and dominance-solution actions, and to separate changes in mistake propensity from changes in the depth of reasoning.

[4]Prominent examples include level-k and cognitive hierarchy (Camerer et al., 2004; Nagel, 1995; Stahl, 1993; Stahl & Wilson, 1994, 1995), cursed equilibrium (Eyster & Rabin, 2005), quantal response equilibrium (McKelvey & Palfrey, 1995; Goeree et al., 2005), random- and noisy-belief equilibrium (Friedman & Mezzetti, 2005; Friedman, 2022), endogenous depth of reasoning (Alaoui & Penta, 2016, 2022), and sequential models of costly reasoning (Gonçalves, 2024).

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

Incentives and mistakes in strategic settings: A substantial literature studies how behaviour reacts to changes in *relative* incentives across actions. This includes work on matching-pennies games (Ochs, 1995; McKelvey et al., 2000), traveler's dilemma (Capra et al., 1999), joker games (Melo et al., 2019), and the modified 11–20 game (Alaoui & Penta, 2016; Alaoui et al., 2020). Closest to our paper is Alaoui et al. (2020). They show that changing *relative* incentives in a replacement design shifts behaviour toward lower numbers and emphasise the role of beliefs about opponents' sophistication. The broad lesson from this literature is that players react strongly when one action becomes more attractive than another, and that beliefs about opponents often play a central role in mediating these responses.

Our paper differs in two key respects. First, we study *absolute* incentive levels whilst holding relative incentives between actions fixed. Second, we separately identify own- and opponent-incentive effects on both choices and beliefs. Our findings consistently highlight the central role of beliefs. We also show that, even holding relative incentives fixed, higher own incentives operate through reduced payoff-dependent mistakes and greater effort devoted to reasoning.

The literature on *absolute* incentive levels in strategic settings is much thinner. Existing evidence shows that higher stakes sometimes can move behaviour toward equilibrium benchmarks, for instance in centipede games (McKelvey & Palfrey, 1992; Rapoport et al., 2003). The broader methodological literature has long stressed that stake variation is useful both for assessing robustness and for understanding the trade-off between incentives and decision costs (Camerer & Hogarth, 1999; Smith & Walker, 1993). Our contribution is to provide evidence on absolute incentives in a strategic setting whilst jointly observing choices, elicited beliefs, and response times.

Our results also speak to the literature on stochastic choice in games. Beginning with quantal response equilibrium, a large body of work models strategic behaviour as probabilistic choice — as in discrete choice (Luce, 1959; McFadden, 1974) and costly control models (Fudenberg et al. 2015) — in which better actions are chosen more often, but not always (Goeree et al., 2005; Mattsson & Weibull, 2002; McKelvey & Palfrey, 1995). In this class of models, higher incentives should, under suitable conditions, sharpen the relation between expected payoffs and choice probabilities. We provide evidence in line with this logic, but with an important qualification: monotonicity and best-response improvements are much clearer when expected payoffs are computed using participants' *stated beliefs* than when they are computed using empirical opponent frequencies. This places beliefs at the centre of any explanation based on payoff-dependent mistakes.

Beliefs in strategic settings: Our paper also contributes to the literature that elicits beliefs in games and studies how they vary systematically with the strategic environment. A large empirical literature shows that stated beliefs are informative about behaviour, though far from perfectly aligned with observed choices (Costa-Gomes & Weizsäcker, 2008; Nyarko & Schotter, 2002; Rey-Biel, 2009). Related work shows that eliciting beliefs can itself affect play, including in dominance-solvable environments (Croson, 1999, 2000; Gächter & Renner, 2010), and that beliefs respond systematically to features of the game (Agranov et al., 2012; Friedman & Ward, 2022).

Our paper contributes to this literature in two ways. First, we examine directly how own and opponents' incentive levels affect the beliefs participants form about others' actions. Second, we relate elicited beliefs to observed choices and show that subjective expected payoffs explain behaviour much better than expected payoffs computed from empirical opponent frequencies. In this sense, the paper supports the view that beliefs are informative about strategic behaviour, but also that understanding how incentives affect beliefs is essential for interpreting observed play.

Response times in strategic settings: Finally, our paper relates to the literature using response times to study strategic behaviour. Response times have been used to classify participants, to predict choices and sophistication, and as an informative proxy for cognitive effort (Frydman & Nunnari, 2023; Gill & Prowse, 2022; Proto et al., 2019; Rubinstein, 2007, 2016; Schotter & Trevino, 2021).[5]

[5]See Spiliopoulos and Ortmann (2018) and Clithero (2018) for recent reviews.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

Particularly relevant is Alós-Ferrer and Buckenmaier (2021), who study response times in beauty-contest and 11–20-type games.

We differ by providing a design that separately identifies how own and opponents' incentive levels affect response times, and by relating response times jointly to actions, elicited beliefs, and payoff performance. More broadly, our paper connects the strategic response-time literature to the growing evidence on incentives and effort in non-strategic tasks (Ariely et al., 2009; Caplin et al., 2020; Dean & Neligh, 2023; Enke et al., 2023). Our findings show that higher incentives increase response times, but performance gains are limited by the complexity of the strategic environment.

In short, these connections place the paper at the intersection of the literatures on strategic sophistication, incentives in games, belief formation, and costly reasoning.

2. Hypotheses and experimental design

2.1. Hypotheses

In this section, we present our hypotheses. Throughout, we consider finite normal-form games, $\Gamma = \langle I, A, u\rangle$, where I denotes the set of players, A_i the action set of player i, A_{-i} the action profiles of player i's opponents, $A := \times_i A_i$ the set of action 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 main interest is in how incentive levels affect strategic behaviour. For any two games Γ and $\tilde{\Gamma}$ that are identical in all except that player i's payoffs are scaled up — so that $\tilde{u}_i = \lambda u_i$ for some $\lambda > 1$ — we say that player i has a higher incentive level in $\tilde{\Gamma}$ than in Γ.

Our first hypothesis relates to the effect of incentives on action sophistication. We use rationalisability to formalise this concept. An action is k-rationalisable if it is a best response to some distribution over opponents' $(k-1)$-rationalisable actions, and it is rationalisable if it is k-rationalisable for all k.[6] We say that action a_i is more sophisticated than action a_i' if (i) there exists some k such that a_i is k-rationalisable whereas a_i' is not, or (ii) if these are both k-rationalisable but a_i iteratively strictly dominates a_i'.

Although experiments routinely find positive choice frequencies for non-rationalisable actions (e.g.Costa-Gomes & Weizsäcker, 2008; Rey-Biel, 2009), it remains unclear whether and how incentive levels affect the sophistication of play. Differently from payoff distortions, increasing incentive levels leaves relative incentives unchanged and thus the sets of k-rationalisable action distributions unaffected. Existing models also yield different predictions. Many prominent models predict there to be no changes in behaviour in response to incentive *levels*. This is true of Nash equilibrium (Nash, 1951), level-k (Nagel, 1995; Stahl & Wilson, 1994, 1995), cognitive hierarchy (Camerer et al., 2004), regret minimisation (Bell, 1982; Loomes & Sugden, 1982), uncertainty-aversion equilibrium (Bade, 2011; Klibanoff, 1996), random and noisy-belief equilibria (Friedman & Mezzetti, 2005; Friedman, 2022), and sampling equilibrium (Osborne & Rubinstein, 1998, 2003), among others. In other models, as endogenous depth of reasoning in symmetric games (Alaoui & Penta, 2016), increasing a player's incentive level may affect that player's actions but not others'.[7] 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).

[6] As is well known (Pearce, 1984, Lemma 3), in two-player games we can equivalently define player i's k-rationalisable actions R_i^k as those surviving k rounds of iterated elimination 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 $R_i^0 := A_i$.

[7] For instance, if Γ is a symmetric two-player game and cognitive costs and anchors 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. Beyond symmetry, in two-player games, a player's actions are either affected by their own incentive level (when the player considers the opponent to be at least as sophisticated) or by their opponent's incentive level (when the player considers the opponent to be less sophisticated), but not by both at once.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

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 behaviour accordingly. We formulate this as our first hypothesis:

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

Although testing Hypothesis 1 is in itself a contribution of this paper, we seek to shed light on mechanisms, that is, *why* incentive levels may affect choices.

One possible channel is that higher incentives make mistakes more costly, reducing the frequency of suboptimal choices. This is the logic underlying structural quantal response models, in which the probability of choosing an action depends on the vector of expected payoffs. Both models of stochastic choice based on random utility with additive i.i.d. payoff disturbances (McFadden, 1974; McKelvey & Palfrey, 1995) and those based on additive perturbed utility or costly precision (Fudenberg et al., 2015; Mattsson & Weibull, 2002) predict, under mild conditions, that higher incentives lead better actions being chosen more often (see Appendix C.1). More generally, regular quantal response equilibrium requires actions with higher expected payoffs to be chosen more frequently (Goeree et al., 2005). To assess support for this class of models, we test the following hypothesis:

Hypothesis 2. (Quantal Responding) (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. Previous work shows that beliefs about opponents' sophistication can help explain the choice of non-rationalisable actions. For example, sophistication of play changes depending on whether players are informed about opponents' characteristics associated with sophistication, such as chess ratings (Palacios-Huerta & Volij, 2009) or their educational background (Agranov et al., 2012; Alaoui et al., 2020). In our setting, higher incentives may induce greater effort in understanding the game, leading participants to think more carefully about opponents' payoffs and to assign lower probability to opponents choosing non-rationalisable actions.

The effect of incentive level on beliefs about the opponent may depend on both own and opponent incentive levels. On the one hand, if higher incentives induce more sophisticated play, then participants should expect more sophisticated behaviour from opponents facing higher incentives. On the other hand, if reasoning about the game is costly, then higher own incentives raise the value of forming more accurate beliefs. This motivates our next hypothesis:

Hypothesis 3. (Belief Sophistication) Beliefs about opponents' strategic sophistication increase with (a) own incentive level and (b) opponents' incentive level. Furthermore, (c) belief accuracy increases with own incentive level.

Finally, we consider whether higher incentives induce greater cognitive effort and whether such effort is associated with better decisions. In line with recent models of sequential reasoning in games (Alaoui & Penta, 2016, 2022; Gonçalves, 2024), we conjecture that increasing a player's incentive level induces greater cognitive effort and, in turn, results in better choices. Whilst greater effort need not always translate into longer deliberation, as it may instead take the form of more focused thinking, we follow existing literature in taking response time as an informative proxy for cognitive effort — see Gill and Prowse (2022) and references therein. Our final hypothesis is:

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

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

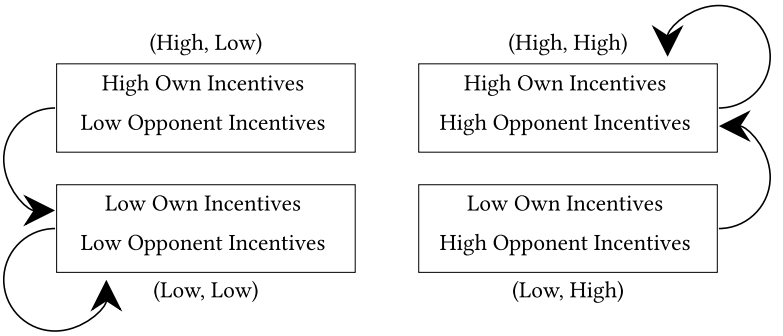


Fig 1. Opponent matching procedure
Notes: This diagram describes the matching procedure implemented. Participants in incentive level groups (High, High) and (Low, Low) have their payoffs depend on an action taken by another random participant in the same group. Participants in incentive level groups (High, Low) and (Low, High) have their payoffs depend on an action taken by a random participant in the groups (Low, Low) and (High, High), respectively.

2.2. Experimental design and identification strategy

Identification strategy: Identifying the effect of incentives on behaviour in strategic settings poses a fundamental challenge. A higher incentive level for player *i* may have a first-order, direct effect on the player's behaviour by motivating them to choose better actions, make fewer mistakes, or forming better beliefs. But it may also generate indirect effects through higher-order reasoning: player *i* may expect opponents to respond to *i*'s higher incentives, and may adjust behaviour accordingly. To study the effect of incentive levels on strategic behaviour, it is then important to separate first-order, direct effects from higher-order, indirect effects, operating through beliefs about others' responses.

Our identification strategy relies on a simple opponent-matching protocol, summarised in Figure 1. Participants are assigned both a role in the game and an incentive group, where the group determines their own incentive level and the incentive level of the opponent whose behaviour is relevant for their decision problem. First, we have two incentive groups — (High, High) and (Low, Low) in Figure 1 — in which all players face the same incentive level and know their opponents do as well. Second, we have another two incentive groups — (High, Low) and (Low, High) in Figure 1 — in which participants are matched with opponents with a different incentive level. The opponents themselves are randomly drawn from the other two groups, so that participants in (High, Low) are matched with opponents from (Low, Low), and participants in (Low, High) are matched with opponents from (High, High). This design allows us (i) to hold fixed the opponent's behaviour whilst varying own incentives, thereby identifying the effect of own incentives on choices and beliefs; and (ii) to hold fixed own incentives whilst varying opponents' incentives, thereby identifying how beliefs and behaviour respond to the incentive level faced by opponents. Importantly, this procedure also fixes higher-order incentives: it determines not only the incentive level faced by a participant's opponent, but also the incentive level faced by that opponent's opponent, and so on.[8]

This approach corresponds to a version of the replacement method featured in Alaoui and Penta (2016) and Alaoui et al. (2020), and is also related to observer-based designs such as Huck and Weizsäcker (2002). A key feature of the design is that, for participants in the mixed-incentive groups (High, Low) and (Low, High), the action they choose does not enter the payoff determination of the opponent whose behaviour is payoff-relevant for them. To avoid making this feature unique to those groups, our implemented matching protocol is so that the same property holds throughout the

[8]Note that simply informing player *i* about the opponent's incentive level would not suffice, because player *i* may also care about the incentive level faced by that opponent's opponent when forming beliefs about behaviour. By contrast, our matching procedure fixes the entire higher-order incentive structure relevant for strategic reasoning.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

		Player 2			
	Actions	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)

		Player 2			
	Actions	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)

Fig 2. Games (a) 2-Step game (b) 3-Step game

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) elimination of strictly dominated strategies to obtain the strategy prescribed by the dominance solution, whereas the game in panel (b) takes 3 steps.

design: the participant whose action is payoff-relevant for player i is never someone whose own payoff is determined by i's action.[9]

Given this matching procedure and random assignment across incentive groups, identification of causal effects of incentive levels is obtained by comparing outcomes — choices, beliefs, and response times — across participants who differ in own incentives or in opponents' incentives.

Games: Participants faced one of the two normal-form games shown in Figure 2. Both are dominance-solvable and have similar payoff structures. The game in panel (a) requires two rounds of (maximal and simultaneous) iterated elimination of strictly dominated strategies to reach the dominance-solution action, whereas the game in panel (b) requires three rounds.

The games have the useful property that actions can be ranked naturally by strategic sophistication. First, if a_{n+1} is k-rationalisable, then so is a_n. Second, if both a_n and a_{n+1} are k- but not $(k+1)$-rationalisable, then a_n (iteratedly) strictly dominates a_{n+1}.[10] Accordingly, we will throughout treat a_n as more strategically sophisticated than a_{n+1}. We note that, assuming a level-0 player uniformly randomises, level-k actions are also uniquely pinned down in a way that does not depend on participants' risk attitudes. In the 2-Step game, a_2 corresponds to the level-1 action regardless of risk preferences, and a_1 to the level-2 action; both a_3 and a_4 are strictly dominated. In the 3-Step game, the level-1, -2, and -3 actions are, respectively, a_3, a_2, and a_1.

The preceding discussion shows why these games are useful for our purposes. Because they admit a clean, theory-grounded ranking of actions and beliefs by strategic sophistication, they allow us to study whether incentives affect mistakes, higher-order reasoning, and belief formation in a tightly controlled environment. In addition, we specifically chose four-action games because this is the smallest action space that allows us to vary the number of elimination rounds with only minor payoff changes, and we imposed symmetry to improve statistical power and reduce data requirements. The games are related to 11–20 (Arad & Rubinstein, 2012) and other undercutting games (Costa-Gomes & Crawford, 2006; Georganas et al., 2015; Nagel, 1995) in that they induce ordered steps of strategic reasoning, but differ in that our games are dominance-solvable and admit a clean ranking of actions by strategic sophistication.

Other design details and logistics: The experiment implemented a $2 \times 2 \times 2$ design, corresponding to own incentive level, opponents' incentive level, and the game played (2-Step or 3-Step game). Participants were assigned uniformly at random to one of the eight treatments and played only once.

The experiment used binary-lottery incentives and game payoffs corresponded to the probability of receiving a prize of x rather than \$2.00. We presented exactly the same game representation across

[9]Related designs with this property include ring games (Kneeland, 2015).

[10]For instance, in the 2-Step game (Figure 2(a)), (i) a_4 is strictly dominated by a_3, which, in turn, is strictly dominated by a_2, and (ii) after eliminating a_3 and a_4, a_2 is iteratedly strictly dominated by a_1. Similarly, in 3-Step game (Figure 2(b)), (i) a_4 is strictly dominated by a_3; (ii) once a_4 is deleted, a_3 becomes iteratedly strictly dominated by a_2; and (iii) after deleting a_3, again a_2 is iteratedly strictly dominated by a_1.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

incentive conditions. This allows us to isolate the effect of incentives whilst holding the strategic structure fixed and avoids potential confounds arising from presenting different payoff numbers due to perceptual or salience effects. In the high-incentive treatment, x=\$22.00; in the low-incentive treatment, x=\$2.50. We chose the high incentive level to maximise treatment variation; the low incentive level nevertheless exceeds the expected payment per game in several related experiments.[11]

The experimental design induces incentive-compatible elicitation of both actions and beliefs under weak assumptions. To this end, we (i) restricted the experiment to a single paid round, thereby also ruling out non-feedback learning (Weber, 2003), and (ii) paid either the action task or the belief task, but not both, to avoid hedging concerns (Blanco et al., 2010). Actions selected determine game payoffs, which correspond to a probability of getting the prize. The order of rows and columns was randomised and participants were informed of this. Beliefs $b_i = (b_{i,1}, b_{i,2}, b_{i,3}, b_{i,4}) \in \Delta(A_i)$, interpreted as reported probabilities that the opponent chooses each action, were incentivised using a binarised scoring rule (Hossain & Okui, 2013). Choices and beliefs were elicited simultaneously to reduce the mismatch between choices and beliefs as suggested by existing evidence (see Costa-Gomes & Weizsäcker, 2008, Figure 4), and one of them was randomly selected for payment. In the instructions, we also explicitly explained that the payment rule was designed to reward choosing the action that maximises expected points given one's beliefs, as well as to report beliefs truthfully.

We took several steps to reduce the scope for focal-point effects and other-regarding motives to mechanically account for the observed treatment differences. First, we chose games in which the dominance-solution outcome is neither Pareto dominant nor Pareto dominated, as is true for most outcomes in the games. Second, for most outcomes, the same payoff pair is associated with multiple action profiles. Third, because higher incentives can increase coordination on focal points (Parravano & Poulsen, 2015), we randomly shuffled the order of rows and columns, reducing the salience of visually symmetric payoff configurations and action labels. And fourth, the matching protocol was random and asymmetric, and this was explicitly explained to participants. We return to these concerns in Section 7.1, where we discuss the extent to which the treatment effects are consistent with alternative explanations based on other-regarding preferences.

The experiment proceeded as follows: (i) participants received instructions and completed comprehension and attention checks; (ii) participants played two unincentivised practice rounds without feedback;[12] (iii) own and opponent incentive levels were revealed;[13] (iv) actions and beliefs were elicited simultaneously; and (v) participants completed a short socio-demographic survey. Instructions and interface screenshots are provided in Online Appendix A.

We recruited 834 participants on Amazon Mechanical Turk (MTurk) in 8-10 and 15 January 2020 and, for replication and robustness purposes, 880 participants on Prolific in 22-23 January 2026, ensuring 100 participants in each treatment. This additional sample also allows us to assess whether the main treatment patterns are robust to a less salient emphasis on incentives.[14] In both samples, participants were adults residents in the United States of America with approval rates of at least 95%. We find only minor differences across platforms and therefore report results for the pooled sample; we expand on this in Section 7.2. Participants faced no time limit. Average duration and hourly earnings were approximately 22 minutes and \$24.97 on MTurk, and 17 minutes and \$23.43 on Prolific.

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

[12] Both practice rounds were four-action two-player games, one with a strictly dominant action and one with no pure-strategy Nash equilibrium.

[13] Only 36 participants — fewer than 1.9% of the sample — dropped out after learning their incentive condition, and attrition was nearly balanced across own-incentive treatments: 20 in low incentives and 16 in high incentives.

[14] Following the editor's suggestion, we revised the instructions and interface to reduce the salience of incentive levels and simplified the explanation of the matching procedure.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

Table 1. Action frequency, average beliefs distributions, and response time

Game	Incentives		Action Frequency				Average Belief Distribution				Time
Steps	Own	Opponent	1	2	3	4	1	2	3	4	(secs)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
2	High	High	56.59	28.29	8.78	6.34	34.02	36.11	17.51	12.35	97.1
2	High	Low	49.76	35.07	9.48	5.69	30.46	34.74	19.91	14.90	82.1
2	Low	High	37.33	36.41	14.75	11.52	28.43	32.89	21.63	17.06	62.2
2	Low	Low	31.86	39.71	17.16	11.27	27.12	31.11	23.97	17.80	65.5
3	High	High	15.32	48.20	30.18	6.31	20.36	29.78	35.17	14.70	113.5
3	High	Low	13.70	42.92	37.44	5.94	21.45	26.80	33.37	18.37	93.8
3	Low	High	16.38	33.62	37.93	12.07	22.70	27.99	30.55	18.77	69.9
3	Low	Low	14.22	32.35	42.65	10.78	22.79	26.77	31.39	19.05	54.6

Notes: This table reports treatment-level summary statistics. Each participant is assigned to one treatment, defined by the game (2-Step or 3-Step game), own incentive level (high or low), and opponents' incentive level (high or low). Columns (4)–(7) report action frequencies, columns (8)–(11) report mean belief distributions, and column (12) reports mean response time in seconds. Participants play one game and make one incentivised decision.

Table 2. Incentive level, dominance, and dominated play (Hypothesis 1)

	Dominance Play		Dominated Play	
	2 Steps	3 Steps	2 Steps	3 Steps
	(1)	(2)	(3)	(4)
Own Incent.	19.48***	-1.01	-12.86***	-5.27***
	(3.43)	(2.47)	(2.84)	(1.99)
Opp. Incent.	5.95*	1.75	-0.71	0.37
	(3.42)	(2.45)	(2.82)	(1.94)
Controls	Yes	Yes	Yes	Yes
R^2	0.06	0.02	0.05	0.02
N	837	877	837	877

Notes: This table reports estimates from the regression specified in equation (1), where the dependent variable is an indicator for whether the participant chose the dominance-solution action (columns (1) and (2)) or a strictly dominated action (columns (3) and (4)). High Own/Opponent Incentives are indicators for whether the participant and their opponent face a high incentive level. 2 Steps and 3 Steps denote the two games in the experiment; see Figure 2. Controls refer to participants' socio-demographic characteristics. Heterocedasticity-robust standard errors are reported in parentheses; *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

We also collected socio-demographic information, including age and education; see Appendix A.6 for details.

3. Incentives and action sophistication

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

Our first observation is that participants do not always play the dominance-solution action, a_1. Table 1 shows substantial heterogeneity in action frequencies across the eight treatments. Dominance play ranges from 13.7% to 56.6%, whilst dominated play — corresponding to actions a_3 and a_4 in the 2-Step game and action a_4 in the 3-Step game — ranges from 5.9% to 28.4%. The raw frequencies already suggest that own incentives matter more in the 2-Step game, whereas in the 3-Step game dominance play remains low across all treatments.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

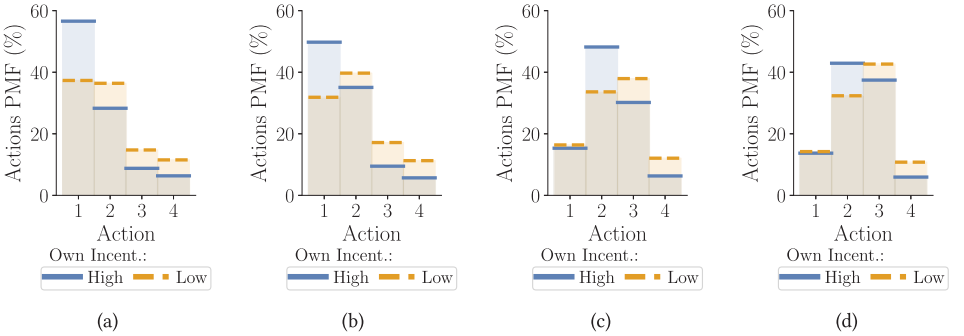


Fig 3. Incentive level and action frequency: own incentives (a) 2-Step game; high opp. incentive level (b) 2-Step game; low opp. incentive level (c) 3-Step game; high opp. incentive level (d) 3-Step game; low opp. incentive level
Notes: The panels exhibit the action frequency (probability mass function) for high and low own incentive levels, for different games (2-Step game, (a) and (b), and 3-Step game, (c) and (d)) and holding fixed opponent incentive level (High, (a) and (c), or Low, (b) and (d)).

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

$$y_i = \beta_0 + \beta_1 H_i + \beta_2 H_{-i} + \text{Controls}_i + \epsilon_i, \tag{1}$$

where H_i is an indicator variable equal to 1 if participant i faces a high own incentive level, and H_{-i} is an analogous indicator equal to 1 if participant i's opponent faces a high incentive level. Controls refer to participants' socio-demographic characteristics.[15] When testing for the effect of incentive levels on dominance play, the dependent variable y_i is an indicator equal to 1 if participant i chose the dominance-solution action and 0 otherwise. When testing the effect of incentives on dominated play, we use the analogous indicator for whether the participant chose a strictly dominated action. Table 2 summarises the results.

Higher own incentives increase action sophistication, but the margin along which this occurs depends on the complexity of the game. In the 2-Step game, higher own incentives increase dominance play by 19.67 percentage points (pp) and reduce dominated play by 12.98 pp. In the 3-Step game, higher own incentives do not increase dominance play, but still reduce dominated play by 5.37 pp. Opponents' incentive levels have much weaker effects: they are associated with a small increase in dominance play in the 2-Step game, but otherwise do not have statistically significant effects. One interpretation is that higher incentives help participants avoid clearly inferior actions even in the more complex game, but are not sufficient to increase dominance play in the 3-Step game.

We next consider the full distribution of actions, shown in Figures 3 and 4.[16] For each game and opponent-incentive condition, higher own incentives shift the distribution of actions toward more strategically sophisticated actions (Mann-Whitney U rank tests, $p < .05$). In the 2-Step game, this shift satisfies first-order stochastic dominance, whereas in the 3-Step game the shift is still toward greater sophistication but not in an FOSD sense. When considering opponents' incentive levels, the data suggest a shift away from the level-1 action (a_2 in the 2-Step game and a_3 in

[15]Throughout, we use the same set of controls. Results are robust to omitting controls and to including the interaction between own and opponents' incentive levels. We report the full regression tables with the interaction term in Appendix A.7. Where relevant, we also discuss the full treatment-level conditioning shown in the figures.

[16]We report the empirical CDF in Appendix A.2 — see Figures 11 and 12.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

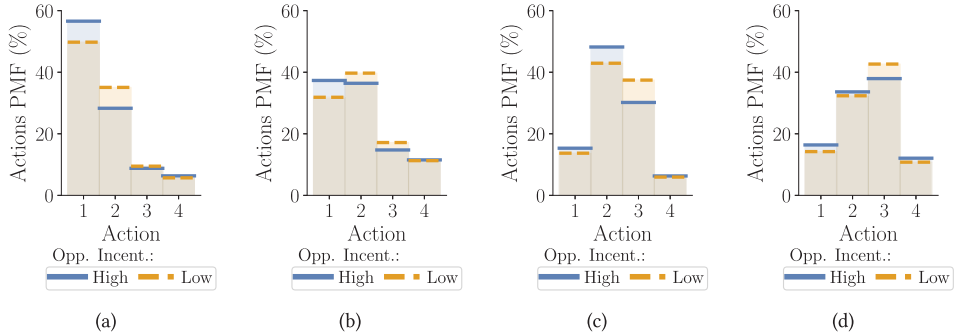


Fig 4. Incentive level and action frequency: opponent incentives (a) 2-Step game; high own incentive level (b) 2-Step game; low own incentive level (c) 3-Step game; high own incentive level (d) 3-Step game; low own incentive level
Notes: The panels exhibit the action frequency (probability mass function) for high and low opponent incentive levels, for different games (2-Step game, (a) and (b), and 3-Step game, (c) and (d)) and holding fixed own incentive level (High, (a) and (c), or Low, (b) and (d)).

Table 3. Incentive level and best-response rate (Hypothesis 2a)

	Subjective BR		Objective BR	
	2 Steps (1)	3 Steps (2)	2 Steps (3)	3 Steps (4)
Own Incent.	20.55*** (3.45)	15.71*** (3.41)	19.48*** (3.43)	12.71*** (3.33)
Opp. Incent.	5.81* (3.44)	3.45 (3.38)	5.95* (3.42)	3.17 (3.31)
Controls	Yes	Yes	Yes	Yes
R^2	0.06	0.04	0.06	0.04
N	837	877	837	877

Notes: This table reports the results of the regression specified in equation (1), where the dependent variable is an indicator for whether the participant best responds to their reported beliefs — that is, chooses the action that maximises subjective expected payoff according to those beliefs (columns (1) and (2)) — or best responds to the empirical frequency of opponents' actions — that is, an objective best response (columns (3) and (4)). High Own/Opponent Incentives correspond to indicators for whether the participant 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 participants' socio-demographic characteristics. Heterocedasticity-robust standard errors are given in parentheses; *, **, and *** denote p-values < 0.1, < 0.05, and < 0.01, respectively.

the 3-Step game) toward more sophisticated actions, but these shifts are not statistically significant overall.

Overall, the evidence provides clear support for Hypothesis 1a. Higher own incentives make actions more sophisticated: they substantially increase dominance play in the simpler 2-Step game and reduce dominated play in both games. This suggests that observed strategic sophistication — as captured, for example, by types in level-*k* and cognitive hierarchy models — depends on the incentive *level* players face. By contrast, the evidence for Hypothesis 1b is weaker. The treatment distributions are suggestive of greater sophistication when opponents face higher incentives, but the corresponding regression estimates are generally small and statistically insignificant.

4. Incentives and quantal responding

One possible reason for the behavioural changes documented above is that, when own incentives are higher, mistakes become more costly and participants therefore choose better responses more often. This is the logic underlying Hypothesis 2.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

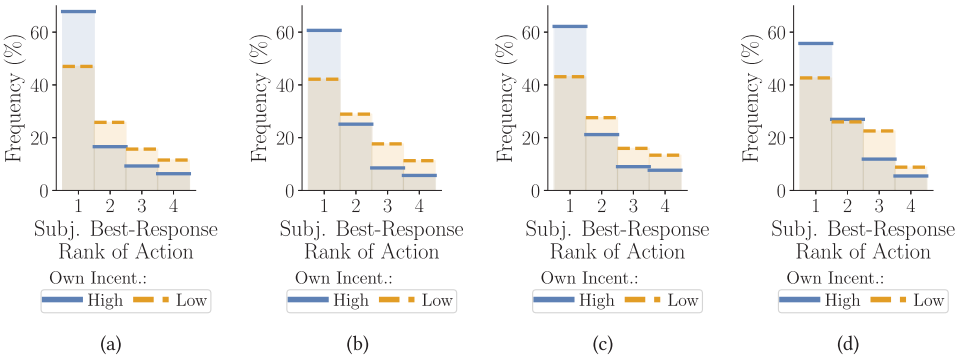


Fig 5. Incentive Level and Subjective Best-Response Rank (Hypothesis 2b) (a) 2-Step game; high opp. incentive level (b) 2-Step game; low opp. incentive level (c) 3-Step game; high opp. incentive level (d) 3-Step game; low opp. incentive level
Notes: The panels show the distribution of the subjective rank of the chosen action for high and low own incentive levels, separately by game (2-Step game, (a) and (b), and 3-Step game, (c) and (d)) and holding fixed opponent 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 maximise subjective expected payoffs.

We first test whether incentive levels affect the likelihood that participants (i) best respond to their stated beliefs (subjective best responses), and (ii) best respond to the observed frequency of opponents' play (objective best responses). In Table 3, we report estimates from a linear probability model using a specification analogous to equation (1). The results show that higher own incentives substantially increase subjective best-response rates in both games: by 20.83 pp in the 2-Step game and by 15.77 pp in the 3-Step game.[17] The effect on objective best-response rates is similar, though somewhat smaller, at 19.67 and 12.81 pp, respectively. Opponents' incentive levels do not significantly affect subjective best-response rates; they have only a modest positive effect on objective best-response rates in the simpler 2-Step game.

A central property of quantal response is monotonicity: actions with higher expected payoffs should be chosen more often. We then examine next how incentives affect the distribution of choices by subjective expected-payoff rank. We define the subjective best-response rank of the chosen action to be *n* if the chosen action has the *n*-th highest subjective expected payoff, according to the participant's reported beliefs, among the four available actions. Figure 5 shows the choice frequencies by subjective rank.

Participants do choose actions with higher *subjective* expected payoff more often, but the same is not true when choices are ranked instead by *objective* expected payoff. In every treatment, actions with higher subjective expected payoff are chosen more frequently than actions with lower subjective expected payoff. Moreover, higher own incentives shift choices toward better responses to stated beliefs (Mann-Whitney U rank tests, $p < .01$); in addition, the data do not reject a monotone likelihood-ratio shift, consistent with broad classes of quantal-response models (see Appendix C.1). By contrast, when one defines an analogous rank using empirical opponent frequencies and objective expected payoffs, monotonicity is visibly violated; see Figure 10 in Appendix A.1. Overall, the evidence provides strong support for Hypothesis 2, especially when best responding is evaluated relative to participants' stated beliefs rather than to empirical opponent frequencies.

To assess how well quantal responding can account for the data, we fit logit quantal-response models and compare their fitted choice frequencies with the observed data. In Table A.5, we report the

[17] Best-response rates to stated beliefs are around 40-50% under low own incentives and 60-70% under high own incentives; this corresponds to the frequency of choosing the action with subjective rank 1 in Figure 5. These figures are comparable to those observed in earlier experiments with two-player three-action games, such as Costa-Gomes and Weizsäcker (2008) and Rey-Biel (2009).

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

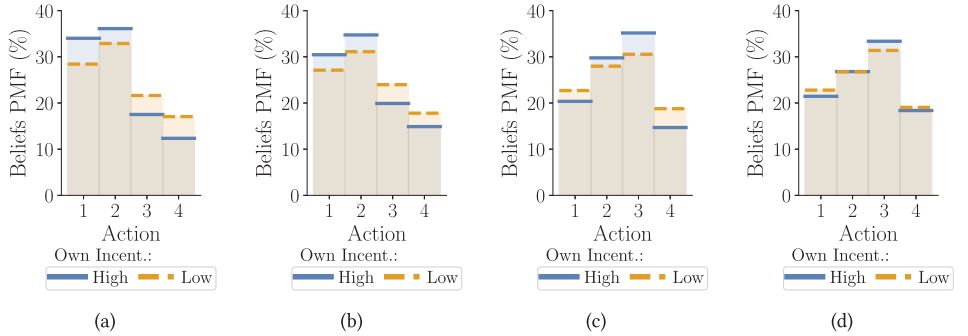


Fig 6. Incentive level and beliefs: own incentives (a) 2-Step game; high opp. incentive level (b) 2-Step game; low opp. incentive level (c) 3-Step game; high opp. incentive level (d) 3-Step game; low opp. incentive level
Notes: The different panels exhibit the mean of the reported belief distributions for high and low own incentive levels, for different games (2-Step game, (a) and (b), and 3-Step game, (c) and (d)) and holding fixed opponent incentive level (High, (a) and (c), or Low, (b) and (d)).

predicted action frequencies obtained by maximum likelihood estimation.[18] We compare the data with two fitted specifications: an objective logit quantal-response model, in which expected payoffs are computed using empirical opponent choice frequencies, and a subjective logit model, in which expected payoffs are computed using stated beliefs. In both specifications, the estimated payoff-sensitivity parameter is higher when own incentives are higher, as one would expect. However, the subjective model fits the data much better: its mean squared error is more than 40% lower than that of the objective model, underscoring the importance of beliefs in accounting for observed choices. In this sense, our findings are consistent with Alaoui et al. (2020) in pointing to a central role for beliefs — which we also directly elicit — but differ in showing that, even holding relative incentives fixed, higher own incentives also operate through a reduction in payoff-dependent mistakes. These findings also point to beliefs as a key part of the mechanism, which motivates the next section's direct analysis of how incentives affect the beliefs participants form.

5. Incentives and belief sophistication

The previous section showed that, whilst payoff-dependent mistakes are an important channel through which incentives affect choices, beliefs remain central to understanding observed behaviour. In this section, we examine if incentive levels affect the beliefs participants form (Hypothesis 3).

Mean reported belief distributions by treatment are shown in Figures 6 and 7.[19] In both games, similarly to what occurred with action frequencies, higher own incentives shift mean beliefs toward level-1 and level-2, and reduce the perceived likelihood of dominated play. In the simpler 2-Step game, where the dominance solution is the level-2 action, higher own incentives shift mean beliefs toward greater opponent sophistication in a first-order stochastic dominance sense. In the more complicated 3-Step game, however, higher own incentives reduce the perceived likelihood of dominated play without increasing the perceived likelihood of the dominance-solution action, which is there the level-3 action. Instead, beliefs shift toward assigning greater probability to the intermediate level-1 and -2 actions. More generally, higher incentives also make beliefs less uniform.[20] The effects of opponents' incentive levels on reported beliefs are qualitatively similar, but less pronounced. Importantly, beliefs are less responsive to opponents' incentive levels when own incentives are low.

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

[19]We report the corresponding mean cumulative belief distributions in Appendix A — see Figures 13 and 14.

[20]The L1 distance of individual belief reports from the uniform distribution decreases in incentive levels; see Table 8 in Appendix A.3.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

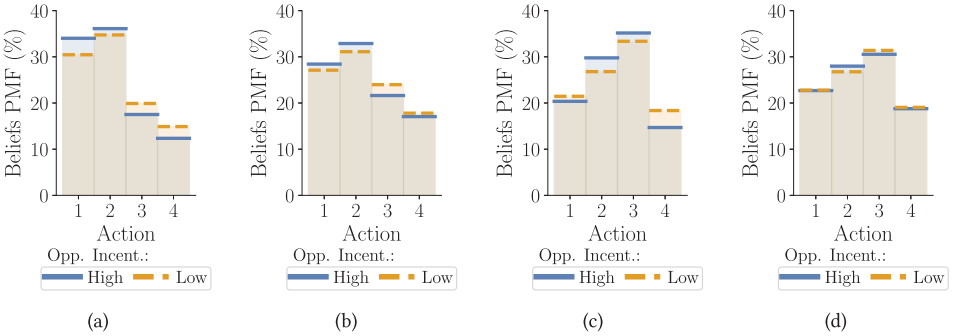


Fig 7. Incentive level and beliefs: opponent incentives (a) 2-Step game; high own incentive level (b) 2-Step game; low own incentive level (c) 3-Step game; high own incentive level d) 3-Step game; low own incentive level
Notes: The different panels exhibit the mean of the reported belief distributions for high and low opponent incentive levels, for different games (2-Step game, (a) and (b), and 3-Step game, (c) and (d)) and holding fixed own incentive level (High, (a) and (c), or Low, (b) and (d)).

Table 4. Incentive level and belief in opponent sophistication (Hypotheses 3a-b)

	Belief in Dominance Play		Belief in Dominated Play	
	2 Steps (1)	3 Steps (2)	2 Steps (3)	3 Steps (4)
Own Incent.	4.68***	-1.85*	-7.64***	-2.36***
	(1.16)	(0.97)	(1.27)	(0.83)
Opp. Incent.	1.99*	-0.67	-3.73***	-1.89**
	(1.15)	(0.96)	(1.24)	(0.81)
Controls	Yes	Yes	Yes	Yes
R^2	0.05	0.02	0.09	0.04
N	837	877	837	877

Notes: This table reports the results of the regression specified in equation (1), where the dependent variable is the reported belief that the opponent chooses the dominance-solution action, $b_{i,1}$, in columns (1) and (2), or the reported belief that the opponent chooses a strictly dominated action, $b_{i,3}+b_{i,4}$ in column (3) and $b_{i,4}$ in column (4). High Own/Opponent Incentives correspond to indicators for whether the participant 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 participants' socio-demographic characteristics. Heterocedasticity-robust standard errors are given in parentheses; *, **, and *** denote p-values < 0.1, < 0.05, and < 0.01, respectively.

To assess the effect of incentives on expected sophistication more formally, we estimate the effect of own and opponents' incentive levels on the reported probability that the opponent chooses the dominance-solution action and on the reported probability that the opponent chooses a strictly dominated action. Table 4 shows that higher own incentives lead participants to attribute greater sophistication to their opponents, but in a nuanced way. In the 2-Step game, higher own incentives increase the reported probability that the opponent chooses the dominance-solution action and decrease the reported probability of dominated play. In the 3-Step game, by contrast, higher own incentives reduce the perceived probability of dominated play, but do not increase the perceived probability of dominance play. Opponents' incentive levels have weaker effects. Participants tend to believe that opponents facing higher incentives are less likely to choose dominated actions, but the effects on beliefs about dominance play are small and statistically insignificant. Overall, the evidence provides partial support for Hypotheses 3a-b.

We next test Hypothesis 3c, namely whether higher incentives lead participants to form more accurate beliefs. Table 5 reports regressions for three measures of belief bias. Columns (1) and (2) consider the absolute difference between the participant's reported belief that the opponent chooses the dominance-solution action and the corresponding empirical frequency. Columns (3) and (4) use the

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

Table 5. Incentive level and belief accuracy (Hypothesis 3c)

	|Belief - Opponent Action Frequency|				|Subj. - Obj. Payoff|	
	Dominance Play		Dominated Play			
	2 Steps (1)	3 Steps (2)	2 Steps (3)	3 Steps (4)	2 Steps (5)	3 Steps (6)
Own Incent.	-1.80**	-0.63	-5.30***	-1.67**	-4.87***	-1.72**
	(0.85)	(0.73)	(0.88)	(0.65)	(0.91)	(0.83)
Opp. Incent.	16.39***	-1.03	5.75***	1.43**	8.78***	5.99***
	(0.85)	(0.74)	(0.87)	(0.64)	(0.90)	(0.82)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.33	0.03	0.12	0.03	0.16	0.08
N	837	877	837	877	837	877

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 realised 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 distance between subjective expected payoffs (computed using reported beliefs) and objective expected payoffs (computed using observed action frequencies). High Own/Opponent Incentives correspond to indicators for whether the participant 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 participants' socio-demographic characteristics. Heterocedasticity-robust standard errors are given in parentheses; *, **, and *** denote p-values < 0.1, < 0.05, and < 0.01, respectively.

analogous measure for strictly dominated play. Columns (5) and (6) consider the absolute difference between subjective expected payoffs, computed using reported beliefs, and objective expected payoffs, computed using observed opponent frequencies: $\sum_{n=1}^{4} |u_i(a_n, b_i) - u_i(a_n, \overline{\sigma}_{-i})|$, where $u_i(a_n, b_i)$ denotes participant i's expected payoff from action a_n under their reported beliefs b_i and $u_i(a_n, \overline{\sigma}_{-i})$ denotes the corresponding expected payoff under the empirical frequency of opponents' actions $\bar{\sigma}_{-i}$. Higher own incentives generally improve belief accuracy. They significantly reduce errors in predicting dominated play in both games and also reduce the overall discrepancy between subjective and objective expected payoffs in both games. For beliefs about dominance play, the improvement is significant only in the simpler 2-Step game. Interestingly, opponents facing higher incentives seem harder to predict: bias in beliefs tends to increase. Taken together with Table 4, this suggests that participants broadly understand the direction in which opponents' behaviour changes with incentives, but misjudge the magnitude of those changes.

Overall, higher own incentives tend to make beliefs both more sophisticated and more accurate, but these effects differ across games and across belief measures. Higher own incentives tend to shift beliefs toward attributing greater sophistication to opponents, especially in the simpler 2-Step game, and they improve several dimensions of belief accuracy. At the same time, beliefs about opponents facing higher incentives become harder to predict precisely. Thus, beliefs are an important part of the mechanism through which incentives affect choices. This motivates our next analysis of whether higher incentives also affect cognitive effort, as proxied by response times.

6. Incentives and response time

We proceed to examining whether incentive levels affect response time. Higher own incentives raise the marginal value of reducing mistakes and refining beliefs, and may therefore induce participants to devote more effort to the decision problem. If that additional effort takes the form of longer deliberation, one would expect response times to increase, as conjectured in Hypothesis 4a. We examine the mechanism underlying this hypothesis by regressing log response time on the incentive treatments, using the same baseline specification as before. Since greater effort need not always manifest itself in longer deliberation, we therefore interpret the analysis in this section as providing suggestive evidence on mechanism rather than a direct test of cognitive effort.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

Table 6. Incentive level and response time (Hypothesis 4a)

	Log(Time)	
	2 Steps (1)	3 Steps (2)
Own Incent.	0.37***	0.37***
	(0.05)	(0.05)
Opp. Incent.	0.03	0.19***
	(0.05)	(0.05)
Controls	Yes	Yes
R^2	0.09	0.12
N	837	877

Notes: This table reports the results of the regression specified in equation (1), where the dependent variable is the participant's log response time (in seconds). High Own/Opponent Incentives correspond to indicators for whether the participant 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 participants' socio-demographic characteristics. Heterocedasticity-robust standard errors are given in parentheses; *, **, and *** denote p-values < 0.1, < 0.05, and < 0.01, respectively.

As Table 6 shows, higher own incentives increase response times substantially in both games: the log-point estimates are 0.38 in the 2-Step game and 0.37 in the 3-Step game, corresponding to increases of roughly 40-45%. Opponents' incentives have a mixed effect. In the 2-Step game they have essentially no effect on response time, whereas in the 3-Step game they significantly increase response time by almost 20%. These patterns are consistent with the view that higher incentives increase the value of deliberation, although they should not be interpreted as showing that response time is itself a direct measure of cognitive effort. Relative to the existing literature, our findings are in line with the data in Alós-Ferrer and Buckenmaier (2021), where scaling up payoffs tends to increase response time in most variants of the 11-20 games discussed,[21] as well as with non-strategic evidence in Caplin et al. (2020) and Enke et al. (2023), where higher incentives increase response times, even though performance gains depend on task complexity.

We then examine whether longer response times are associated with better performance, that is, whether longer response times are associated with (i) higher objective best-response rates and (ii) higher expected payoffs. As reported in Table 7, longer response times are indeed associated with a higher rate of best responses to opponent action frequencies (columns (1) and (2)) as well as an increase in expected payoffs (columns (3) and (4)). The association with expected payoffs is statistically strong but economically modest. The findings support Hypothesis 4b, although it is important to emphasise that the estimates here are associative rather than causal. Moreover, as in the existing literature, response time is an informative but imperfect proxy for cognitive effort, so the evidence in this section should be read as suggestive of mechanism rather than as a direct causal test of cognitive effort.

We conclude with a descriptive analysis of how actions and beliefs vary with response time. For actions, we estimate a multinomial logit model relating choices to response time, incentive treatments, and individual characteristics, and then recover fitted choice frequencies as a function of

[21] Alós-Ferrer and Buckenmaier (2021) use several different versions of the 11-20 game and, for each of these, they scale up payoffs and subtract a constant. We used the published dataset to explore the issue of the effect of incentive levels on response time via linear regression of the latter on the former. The overall average effect is also positive and significant. Separating by type variant, we find that a higher bonus (1) decreases response time only in the low-cost version of their "baseline" game and the high-cost version of the "extreme" variant, (2) has no effect in the low-cost version of the "flat" variant there is no effect, and (3) has a positive and significant effect on the other five versions in their data — high-cost "baseline", low-cost "extreme", high-cost "flat", and low- and high-cost "SOCP". These conclusions are unaffected by controlling or not for participant socio-demographic characteristics and using response time in levels or logs.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

Table 7. Response time, best responses, and payoffs (Hypothesis 4b)

	Objective BR		Log(Obj. Exp. Payoff)	
	2 Steps	3 Steps	2 Steps	3 Steps
	(1)	(2)	(3)	(4)
Log(Time)	15.88***	11.42***	0.10***	0.06***
	(2.16)	(2.06)	(0.01)	(0.01)
Own Incent.	13.59***	8.46**	0.08***	0.03**
	(3.50)	(3.40)	(0.02)	(0.01)
Opp. Incent.	5.50*	1.03	-0.14***	-0.03*
	(3.31)	(3.29)	(0.02)	(0.01)
Controls	Yes	Yes	Yes	Yes
R^2	0.12	0.07	0.16	0.08
N	837	877	837	877

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 participant 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 participants' socio-demographic characteristics. Heterocedasticity-robust standard errors are given in parentheses; *, **, and *** denote p-values < 0.1, < 0.05, and < 0.01, respectively.

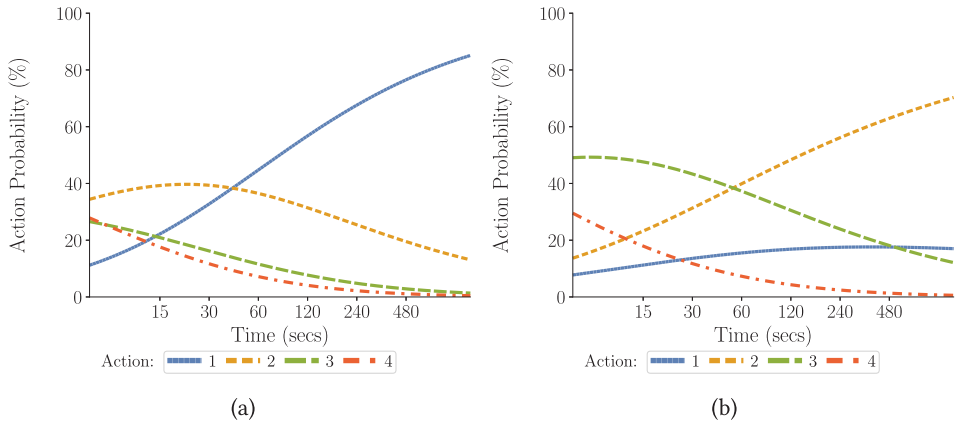


Fig 8. Action frequency and response time (a) 2-Step game (b) 3-Step game

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 participants' choices, on the one hand, and response time, incentive treatments, and participants' socio-demographic characteristics, estimated separately for each game. The action frequency conditional on response times is given by marginalising over the other regressors. 2-Step and 3-Step games denote the different games in the experiment (see Figure 2).

response time by marginalising over the remaining covariates. For beliefs, we use nonparametric kernel regression to estimate the mean belief vector, $b_i = (b_{i,1}, b_{i,2}, b_{i,3}, b_{i,4})$, conditional on response time, incentive treatments, and individual characteristics. We then use the estimated conditional distribution to recover predicted mean belief distribution as a function of response time, holding the remaining covariates fixed at their sample means.

Figures 8 and 9 show how action probabilities and belief distributions change with response time, revealing two main patterns.[22] In the simpler 2-Step game, longer response times are associated with

[22]In Appendix A.4, 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 9 and 10.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

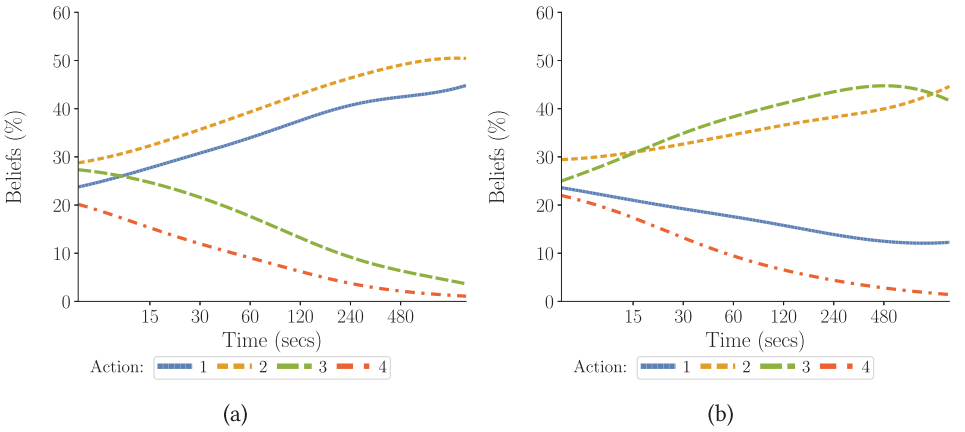


Fig 9. Beliefs and response time (a) 2-Step game (b) 3-Step game

Notes: This figure shows the predicted relation between reported belief distributions and response time (in seconds). The predictions are given by kernel estimation of the relation between participants' reported beliefs, on the one hand, and response time, incentive treatments, and participants' socio-demographic characteristics, estimated separately for each game. The expected beliefs conditional on response times are obtained by using the estimated joint conditional probability distribution over a grid on Σ_{-i} and marginalising over covariates other than response times. 2-Step and 3-Step games denote the different games in the experiment (see Figure 2).

greater sophistication in both actions and beliefs. Panel (a) of Figure 8 shows that the predicted frequency of dominated actions, a_3 and a_4, declines quickly with response time, whereas a_2 — the level-1 action — first mildly increases and then declines. Throughout, the frequency of choice of the level-2 action, a_1, increases monotonically in response time. Panel (a) of Figure 9 similarly shows that longer decisions are associated with lower beliefs in opponents choosing dominated actions and with higher beliefs in both the level-1 action (a_2) and the level-2 action (a_1). In the 3-Step game (panels (b) in Figures 8 and 9), we observe analogous patterns between actions and beliefs that correspond to the same level of sophistication. In actions there is a shift from the dominated action (a_4) and the level-1 action (a_3) toward the level-2 action (a_2), while play of the level-3 dominance-solution action (a_1) remains low and nearly flat. 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-Step game belief in dominance play — which corresponds to the level-3 action, a_1 — is lower for longer response times.

In short, the evidence suggests that faster decisions are associated with the intuitive level-1 actions and beliefs closer to uniform, whereas longer decisions are associated with fewer dominated choices, more weight on level-1 and level-2 opponent play, and a shift toward level-2 actions.[23] What does not increase with response time in the 3-Step game is the frequency of the level-3 dominance solution.

The estimation exercise and the positive relation between own incentives and response time point to a coherent mechanism. When own incentives are higher, the expected gain from reducing mistakes and refining beliefs is larger, so participants have a stronger incentive to invest in further reasoning. Consistent with this interpretation, higher own incentives are associated with longer response times, and longer response times are associated with more strategically sophisticated beliefs and actions. This is in line with models of sequential reasoning (e.g. Alaoui & Penta, 2016; Gonçalves, 2024), in which players decide whether to exert additional cognitive effort by weighing its expected benefits against its cognitive costs.

[23]This is consistent with the data in Alós-Ferrer and Buckenmaier (2021), where longer response times are associated with more sophisticated choices. This is clear for the beauty contest game from their Figure 2. For the 11-20 games, the paper does not report the effect of response time on choices; using the published data, we regress choices on response time and find that longer response times are associated with more strategically sophisticated choices, in the level-*k* sense.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

7. Alternative explanations, robustness, and limitations of the findings

In this section, we discuss alternative explanations, robustness checks, and the natural limitations in interpreting the findings.

7.1. Alternative explanations

Although we cannot rule out all forms of other-regarding motives, the evidence suggests that they are unlikely to be the main driver of our findings. We sought to minimise confounding factors due to other-regarding preferences. As discussed in Section 2.2, we designed the experiment to reduce the scope for other-regarding preferences to mechanically account for the treatment effects, including through the choice of games, the way they are displayed, and the matching protocol. Moreover, many models of other-regarding preferences cannot generate behavioural changes from symmetrically scaling *everyone's* payoffs (Charness & Rabin, 2002; Fehr & Schmidt, 1999; Levine, 1998),[24] though asymmetric treatments could still matter if participants react to relative payoff differences.

A related concern is that participants assigned to low own incentives may become disappointed when their opponent has high incentives, and that this disappointment may lead to more random or less strategic behaviour. The data do not support this interpretation. If anything, choices are weakly *more* sophisticated, not less, when the opponent has higher incentives than when both face low incentives; see Figures 12b and 12d. In addition, even under low own incentives, participants best respond to their stated beliefs more than 40% of the time (Figure 5), well above the 25% rate implied by random choice, and this rate is stable relative to opponents' incentives (Table 3).[25]

7.2. Robustness checks

Participant pools: MTurk was the benchmark platform for online economics experiments in 2020, when the first experiment was run, whereas Prolific was the most common platform at the time of the replication. Online participant pools are particularly suitable for our purposes for two reasons. First, participants face a salient opportunity cost of time, as alternative paid tasks are readily available. This makes variation in incentives and response times especially meaningful. Second, online implementation avoids frictions that a laboratory setting would create for response-time analysis, since participants can complete and leave the task without having to wait for or disrupting others.[26]

We conducted a robust replication in 2026 on Prolific.[27] Although there are some level differences across the two participant pools,[28] the treatment effects are consistent across MTurk and Prolific; see Appendix B.2.[29] We conclude the main findings hold across participant pools and are robust to a less salient presentation of the incentive manipulation.

Attention and understanding: Another possible concern is that our findings are driven by inattention and misunderstanding of the tasks. We took several steps to promote understanding: instructions were designed to be as clear and transparent as possible, participants could familiarise

[24] Exceptions do exist, for example Bolton and Ockenfels (2000).

[25] More generally, the contrast between the 2-Step and 3-Step games is difficult to reconcile with a simple disappointment-based account. If the main effect of asymmetric incentives were merely to induce frustration or randomisation, one would not expect the systematic shift away from dominated actions and toward intermediate non-dominated actions that we observe, especially in the more complex game.

[26] Instead, lab participants would (i) need to stay until the end of the experiment, significantly affecting their opportunity cost of time, or (ii) would be able to leave early, disrupting others, or (iii) would need to come one by one, implying very significant time demands to recruit hundreds of participants.

[27] Following the editor's suggestion, the design was kept the same except that instructions were simplified and incentive levels were de-emphasised both in the instructions and in the interface to mitigate treatment-demand concerns.

[28] The most salient difference is that Prolific participants tend to play dominated actions less often and the dominance solution slightly more often, but otherwise react in the same way to incentive levels.

[29] Another recent paper, Gonçalves et al. (2025), reports an experiment run on MTurk and a replication on Prolific in similar time periods, 2020 and 2025 respectively, and likewise finds consistent treatment effects across platforms.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

themselves with the interface during the instructions, and incentive levels were revealed only after the instruction stage, so attentiveness to the instructions should not differ systematically across treatments. The data also suggest that participants understood the task. Best-response rates to stated beliefs range from roughly 40-50% under low own incentives and 60-70% under high own incentives. These rates are comparable to those observed in laboratory studies with two- and three-action games (Costa-Gomes & Weizsäcker, 2008; Friedman & Ward, 2022; Nyarko & Schotter, 2002; Rey-Biel, 2009). Moreover, the median time spent on the paid round was about one minute and the median time spent on the instruction page was 9 minutes, indicating that participants engaging seriously with both the instructions and the paid task.

At the end of the instructions, participants completed a comprehension quiz designed both to consolidate understanding and to emphasise the key aspects of the experiment; the median number of mistakes was two. Because all participants had to complete the quiz correctly before proceeding, the quiz ensured a minimum level of understanding. We nevertheless use initial quiz performance to assess whether the main results are driven by participants who at first misunderstood the task. In Appendix B.1, we reproduce the main regressions interacting the treatment indicators with the number of mistakes in the quiz. Initial quiz performance makes little difference to the qualitative conclusions and, if anything, the treatment effects are somewhat stronger among participants who made fewer mistakes.

7.3. Extreme incentives

Although higher incentives in our design tend to increase strategic sophistication, caution is warranted when extrapolating beyond the range of incentives we study. There may be an optimal level of incentives, deviations from which are detrimental to performance, as posited by the Yerkes and Dodson (1908) 'law'. As discussed by Camerer and Hogarth (1999), for simple or well-learned tasks stronger incentives may lead to better performance, but extreme incentive levels can generate stress and impair cognition, limiting or even worsening performance. Ariely et al. (2009) and Enke et al. (2023) provide evidence that performance in individual decision-making tasks involving memory and reasoning can deteriorate under extreme incentives. In strategic settings, evidence on this phenomenon is more limited; if anything, extreme incentives have been found to shift behaviour toward benchmark equilibrium predictions in ultimatum bargaining (Andersen et al., 2011), a setting arguably strongly affected by cultural norms and other-regarding preferences.

In our design, the low incentive treatment was broadly in line with payment levels commonly used in related experiments, whilst the high incentive treatment substantially scaled those payments up but still remained within a realistic range. Our findings should therefore be interpreted as evidence about moderate incentive variation rather than about arbitrarily small or arbitrarily large stakes. The effect of scaling incentives, is likely to depend not only on the strategic environment, but also on the absolute level of incentives. At extremely low stakes, even large proportional changes may remain behaviourally negligible, and at very high stakes further increases may also have little additional effect or could even become counterproductive. That said, it is plausible that the main qualitative patterns we document — greater action and belief sophistication, better responses to stated beliefs, and longer response times — would continue to hold within a reasonable range.

8. Conclusion

This paper provides well-identified evidence that incentive levels affect behaviour in strategic settings. Higher own incentives reduce dominated play in both games. In the simpler 2-Step game, they also increase play of the dominance-solution action, corresponding to the level-2 action. In the more complex 3-Step game, however, they mainly shift behaviour away from dominated and

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

level-1 actions toward the level-2 action, without increasing play of the dominance-solution action, which there corresponds to the level-3 action. Reported beliefs shift in a similar direction and also become more accurate along several dimensions. Choices and beliefs also react to opponent incentive levels in a consistent manner, although more weakly. At the same time, beliefs about opponents with higher incentives become less accurate, indicating that although higher opponent incentives shift behaviour in predictable directions, the magnitude of the effect is harder to predict than noisier play.

Response times provide suggestive evidence on mechanism. Higher own incentives are associated with substantially longer response times in both games, and longer response times are in turn associated with more sophisticated choices and beliefs. Faster decisions are concentrated on the intuitive level-1 action and on beliefs closer to uniform, whereas longer decisions are associated with fewer dominated choices, greater weight on level-1 and level-2 opponent play, and a shift toward level-2 actions. What longer response times do not deliver in the 3-Step game is an increase in play of the level-3 dominance solution.

In summary, the evidence points to two main channels through which own incentives affect strategic behaviour: a reduction in payoff-dependent mistakes and an increase in effort devoted to reasoning, for which response time provides an informative, though imperfect, proxy.

Two directions for future research seem particularly important. First, it would be useful to better understand the limits of incentives in fostering strategic sophistication, and how those limits depend on the complexity of the strategic environment. If, on the one hand, there is naturally a limit to what one can achieve in thinking about others' behaviour through introspection alone, on the other hand such a limit surely depends on (or is informative about) how complex the environment is. This can shed light on what constitutes a simple strategic environment, which is a crucial consideration in designing successful mechanisms.

Second, response times deserve further attention in the study of strategic behaviour. They are often straightforward to collect and, as our results suggest, can be informative about both beliefs and choices. Further work developing game-theoretic models that incorporate response times, and establishing more systematically how response times relate to behaviour across environments, could significantly deepen our understanding of observed strategic behaviour.

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/eec.2026.10056.

Replication package. The data and code underlying this research is available on Zenodo at https://dx.doi.org/10.5281/zenodo.20609182.

Acknowledgements. 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, Essex, SWEET, the SAE Symposium, the EEA-ESEM Congress, the Barcelona Summer Forum, the ESA Conference, and the EAYE Annual Meeting for valuable comments. We are particularly grateful to Yeon-Koo Che, Terri Kneeland, and Mark Dean for their many insightful conversations. We are also indebted to the editor, Marta Serra-Garcia, and two anonymous referees whose comments and suggestions contributed significantly to improving this article. This material is based upon work supported by the National Science Foundation under Grant Number 1949395.

References

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

Alaoui, L., Janezic, K. A., & Penta, A. (2020). Reasoning about others' reasoning. *Journal of Economic Theory*, *189*, 1–51. https://doi.org/10.1016/j.jet.2020.105091.

Alaoui, L., & Penta, A. (2016). Endogenous depth of reasoning. *Review of Economic Studies*, *83*(4), 1297–1333. https://doi.org/10.1093/restud/rdv052.

Alaoui, L., & Penta, A. (2022). Cost-benefit analysis in reasoning. *Journal of Political Economy*, *130*(4), 881–925. https://doi.org/10.1086/718378.

Alós-Ferrer, C., & Buckenmaier, J. (2021). Cognitive sophistication and deliberation times. *Experimental Economics*, *24*, 558–592. https://doi.org/10.1007/s10683-020-09672-w.

Andersen, S., Ertaç, S., Gneezy, U., Hoffman, M., & List, J. A. (2011). Stakes matter in ultimatum games. *American Economic Review*, *101*(7), 3427–3439. https://doi.org/10.1257/aer.101.7.3427.

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

Ariely, D., Gneezy, U., Loewenstein, G., & Mazar, N. (2009). Large stakes and big mistakes. *Review of Economic Studies*, *76*(2), 451–469. https://doi.org/10.1111/j.1467-937X.2009.00534.x.

Bade, S. (2011). Ambiguous act equilibria. *Games and Economic Behavior*, *71*(2), 246–260. https://doi.org/10.1016/j.geb.2010.04.006.

Bajari, P., & Hortaçsu, A. (2005). Are structural estimates of auction models reasonable? Evidence from experimental data. *Journal of Political Economy*, *113*(4), 675–741. https://doi.org/10.1086/432138.

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, *30*(5), 803–1022. https://doi.org/10.1287/opre.30.5.961.

Blanco, M., Engelmann, D., Koch, A. K., & Normann, H. -T. (2010). Belief elicitation in experiments: Is there a hedging problem? *Experimental Economics*, *13*, 412–438. https://doi.org/10.1007/s10683-010-9249-1.

Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, *90*(1), 166–193. https://doi.org/10.1257/aer.90.1.166.

Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

Camerer, C., Ho, T. -H., & Chong, J. -K. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, *119*(3), 861–898. https://doi.org/10.1162/0033553041502225.

Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty*, *19*, 7–42. https://doi.org/10.1023/A:1007850605129.

Caplin, A., Csaba, D., Leahy, J., & Nov, O. (2020). Rational inattention, competitive supply, and psychometrics. *Quarterly Journal of Economics*, *135*(3), 1681–1724. https://doi.org/10.1093/qje/qjaa011.

Capra, C., Goeree, J. K., Gomez, R., & Holt, C. A. (1999). Anomalous behavior in a traveler's dilemma? *American Economic Review*, *89*(3), 678–690. https://doi.org/10.1257/aer.89.3.678.

Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, *117*(3), 817–869. https://doi.org/10.1162/003355302760193904.

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

Costa-Gomes, M. A., & Crawford, V. P. (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, *96*(5), 1737–1768. https://doi.org/10.1257/aer.96.5.1737.

Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal form games. *Review of Economic Studies*, *75*(3), 729–762. https://doi.org/10.1111/j.1467-937X.2008.00498.x.

Croson, R. T. A. (1999). The disjunction effect and reason-based choice in games. *Organizational Behavior and Human Decision Processes*, *80*(2), 118–133. https://doi.org/10.1006/obhd.1999.2846.

Croson, R. T. A. (2000). Thinking like a game theorist: Factors affecting the frequency of equilibrium play. *Journal of Economic Behavior & Organization*, *41*(3), 299–314. https://doi.org/10.1016/S0167-2681(99)00078-5.

Dean, M., & Neligh, N. (2023). Experimental tests of rational inattention. *Journal of Political Economy*, *131*(12), 3415–3461. https://doi.org/10.1086/725174.

Enke, B., Gneezy, U., Hall, B., Martin, D. C., Nelidov, V., Offerman, T., & van de Ven, J. (2023). Cognitive biases: Mistakes or missing stakes? *The Review of Economics and Statistics*, *105*(4), 818–832. https://doi.org/10.1162/rest_a_01093.

Eyster, E., & Rabin, M. (2005). Cursed equilibrium. *Econometrica*, *73*(5), 1623–1672. https://doi.org/10.1111/j.1468-0262.2005.00631.x.

Farhi, E., & Werning, I. (2019). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review*, *109*(11), 3887–3928. https://doi.org/10.1257/aer.20171400.

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, *114*(3), 817–868. https://doi.org/10.1162/003355399556151.

Friedman, E. (2022). Stochastic equilibria: Noise in actions or beliefs? *American Economic Journal: Microeconomics*, *14*(1), 94–142. https://doi.org/10.1257/mic.20190013.

Friedman, E., & Ward, J. (2022) Stochastic choice and noisy beliefs in games: An experiment. *Working Paper*. 1–77. https://doi.org/10.2139/ssrn.4190338.

Friedman, J., & Mezzetti, C. (2005). Random belief equilibrium in normal form games. *Games and Economic Behavior*, *51*(2), 296–323. https://doi.org/10.1016/j.geb.2003.08.004.

Frydman, C., & Nunnari, S. (2023) Coordination with cognitive noise. *Working Paper.* https://doi.org/10.2139/ssrn.3939522.

Fudenberg, D., Iijima, R., & Strzalecki, T. (2015). Stochastic choice and revealed perturbed utility. *Econometrica*, *83*(6), 2371–2409. https://doi.org/10.3982/ECTA12660.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

Fudenberg, D., & Liang, A. (2019). Predicting and understanding initial play. *American Economic Review*, *109*(12), 4112–4141. https://doi.org/10.1257/aer.20180654.
Georganas, S., Healy, P. J., & Weber, R. A. (2015). On the persistence of strategic sophistication. *Journal of Economic Theory*, *159*, 369–400. https://doi.org/10.1016/j.jet.2015.07.012.
Gill, D., & Prowse, V. (2022). Strategic complexity and the value of thinking. *Economic Journal*, *133*, 761–786. https://doi.org/10.1093/ej/ueac070.
Goeree, J., Holt, C., & Palfrey, T. (2005). Regular quantal response equilibrium. *Experimental Economics*, *8*(4), 347–367. https://doi.org/10.1007/s10683-005-5374-7.
Gonçalves, D. (2024) Sequential sampling equilibrium. *Working Paper*. 1–53. https://doi.org/10.48550/arXiv.2212.07725.
Gonçalves, D., Libgober, J., & Willis, J. (2025). Retractions: Updating from complex information. *Review of Economic Studies*, *93*(1), 476–516. https://doi.org/10.1093/restud/rdaf032.
Gächter, S., & Renner, E. (2010). The effects of (incentivized) belief elicitation in public goods experiments. *Experimental Economics*, *13*(3), 364–377. https://doi.org/10.1007/s10683-010-9246-4.
Hoffman, E., McCabe, K. A., & Smith, V. L. (1996). On expectations and the monetary stakes in ultimatum games. *International Journal of Game Theory*, *25*, 289–301. https://doi.org/10.1007/BF02425259.
Hossain, T., & Okui, R. (2013). The binarized scoring rule. *Review of Economic Studies*, *80*(3), 984–1001. https://doi.org/10.1093/restud/rdt006.
Huck, S., & Weizsäcker, G. (2002). Do players correctly estimate what others do? Evidence of conservatism in beliefs. *Journal of Economic Behavior and Organization*, *47*(1), 71–85. https://doi.org/10.1016/S0167-2681(01)00170-6.
Klibanoff, P. (1996). Uncertainty, decision and normal form games. *Working Paper*. 1–39.
Kneeland, T. (2015). Identifying higher-order rationality. *Econometrica*, *83*(5), 2065–2079. https://doi.org/10.3982/ECTA11983.
Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, *1*(3), 593–622. https://doi.org/10.1006/redy.1998.0023.
Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, *92*(368), 805–824. https://doi.org/10.2307/2232669.
Luce, D. (1959). *Individual choice behavior*. John Wiley.
Mattsson, L. -G., & Weibull, J. (2002). Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior*, *41*(1), 61–78. https://doi.org/10.1016/S0899-8256(02)00014-3.
McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Karembka (Ed.), *Frontiers in Econometrics* (pp. 105–142). Academic Press.
McKelvey, R. D., & Palfrey, T. R. (1992). An experimental study of the centipede game. *Econometrica*, *60*(4), 803–36. https://doi.org/10.2307/2951567.
McKelvey, R. D., & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, *10*(1), 6–38. https://doi.org/10.1006/game.1995.1023.
McKelvey, R. D., Palfrey, T. R., & Weber, R. (2000). The effects of payoff magnitude and heterogeneity on behavior in 2 x 2 games with unique mixed strategy equilibria. *Journal of Economic Behavior & Organization*, *42*(4), 523–548. https://doi.org/10.1016/s0167-2681(00)00102-5.
Melo, E., Pogorelskiy, K., & Shum, M. (2019). Testing the quantal response hypothesis. *International Economic Review*, *60*, 53–74. https://doi.org/10.1111/iere.12344.
Moinas, S., & Pouget, S. (2013). The bubble game: An experimental study of speculation. *Econometrica*, *81*(4), 1507–1539. https://doi.org/10.3982/ECTA9433.
Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, *85*(5), 1313–1326.
Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics*, *54*(2), 286–295. https://doi.org/10.2307/1969529.
Nyarko, Y., & Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica*, *70*, 971–1005. https://doi.org/10.1111/1468-0262.00316.
Ochs, J. (1995). Games with unique, mixed strategy equilibria: An experimental study. *Games and Economic Behavior*, *10*(1), 202–17. https://doi.org/10.1006/game.1995.1030.
Osborne, M., & Rubinstein, A. (1998). Games with procedurally rational players. *American Economic Review*, *88*(4), 834–847.
Osborne, M., & Rubinstein, A. (2003). Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, *45*(2), 434–444. https://doi.org/10.1016/S0899-8256(03)00147-7.
Palacios-Huerta, I, & Volij, O. (2009). Field centipedes. *American Economic Review*, *99*(4), 1619–1635. https://doi.org/10.1257/aer.99.4.1619.
Parravano, M., & Poulsen, O. (2015). Stake size and the power of focal points in coordination games: Experimental evidence. *Games and Economic Behavior*, *94*, 191–199. https://doi.org/10.1016/j.geb.2015.05.001.
Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, *52*(4), 1029–1050. https://doi.org/10.2307/1911197.
Proto, E., Rustichini, A., & Sofianos, A. (2019). Intelligence, personality, and gains from cooperation in repeated interactions. *Journal of Political Economy*, *127*(3), 1351–90. https://doi.org/10.1086/701355.

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press

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

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

Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *The Economic Journal*, *117*(523), 1243–1259. https://doi.org/10.1111/j.1468-0297.2007.02081.x.

Rubinstein, A. (2016). A typology of players: Between instinctive and contemplative. *Quarterly Journal of Economics*, *131*(2), 859–890. https://doi.org/10.1093/qje/qjw008.

Schotter, A., & Trevino, I. (2021). Is response time predictive of choice? An experimental study of threshold strategies. *Experimental Economics*, *24*(1), 87–117. https://doi.org/10.1007/s10683-020-09651-1.

Slonim, R., & Roth, A. E. (1998). Learning in high stakes ultimatum games: An experiment in the slovak republic. *Econometrica*, *66*(3), 569–596. https://doi.org/10.2307/2998575.

Smith, V. L., & Walker, J. M. (1993). Monetary rewards and decision cost in experimental economics. *Economic Inquiry*, *31*, 245–261. https://doi.org/10.1111/j.1465-7295.1993.tb00881.x.

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

Stahl, D. O. (1993). Evolution of smart$_n$ players. *Games and Economic Behavior*, *5*(4), 604–617. https://doi.org/10.1006/game.1993.1033.

Stahl, D. O., & Wilson, P. W. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization*, *25*(3), 309–327. https://doi.org/10.1016/0167-2681(94)90103-1.

Stahl, D. O., & Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, *10*(1), 218–254. https://doi.org/10.1006/game.1995.1031.

Straub, P. G., & Murnighan, J. K. (1995). An experimental investigation of ultimatum games: information, fairness, expectations, and lowest acceptable offers. *Journal of Economic Behavior & Organization*, *27*(3), 345–364. https://doi.org/10.1016/0167-2681(94)00072-M.

Weber, R. A. (2003). 'Learning' with no feedback in a competitive guessing game. *Games and Economic Behavior*, *44*(1), 134–144. https://doi.org/10.1016/S0899-8256(03)00002-2.

Yerkes, R. M., & Dodson, J. D. (1908). The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative Neurology and Psychology*, *18*(5), 459–482. https://doi.org/10.1002/cne.920180503.

Cite this article: Esteban-Casanelles, T., & Gonçalves, D. (2026). Incentives and strategic behaviour: an experiment. *Experimental Economics*, 1–26. https://doi.org/10.1017/eec.2026.10056

https://doi.org/10.1017/eec.2026.10056 Published online by Cambridge University Press