Complexity and Higher Order Rationality: An Experimental Study

Jiaying Li*

Department of Economics, University of Edinburgh,UK

August 9, 2025

Abstract

We experimentally study how game complexity influences individuals' strategic sophistication. In a series of refined Ring Games, subjects interact with computer-controlled opponents programmed to play the equilibrium strategy, allowing us to isolate the effect of complexity on ability-bound players. We vary complexity along three dimensions: informational complexity, structural complexity, and the salience of strategic considerations. All three dimensions significantly affect observed Level-k behaviour, albeit in distinct ways. Prior exposure to game theory further enhances strategic sophistication. These findings indicate that strategic sophistication is shaped by the structure of the decision environment and by prior experience, rather than being a fixed cognitive trait.

Keywords: Behavioural Game Theory, Higher Order Rationality, Level-K, Bounded Rationality
JEL Codes: C70, C91

*University of Edinburgh, Edinburgh, Scotland, UK. E-mail: Jiaying.Li@ed.ac.uk. Telephone:+4407529781396

1 Introduction

An important question in behavioural and experimental economics is whether individuals' strategic sophistication is an inherent trait or a context-dependent outcome shaped by the decision environment. While previous studies have documented substantial variation in individuals' revealed levels of higher-order rationality across different strategic settings (e.g., Georganas, 2015), the mechanisms driving this variation remain unclear. Strategic decision-making often requires individuals to engage in iterative reasoning, but the extent to which they can do so may be constrained by the complexity of the game they face. Understanding this relationship is crucial, as it informs both theoretical models of bounded rationality and practical applications in mechanism design, policy-making, and market behaviour. While the role of complexity in decision-making has been widely acknowledged (Oprea, 2024), its specific influence on strategic sophistication in game-theoretic environments has yet to be systematically examined.

In this paper, we design a novel N-Player N-Choice Ring Game, building on the framework of Kneeland (2015), to examine how game characteristics shape players' strategic sophistication under the Level-k theory. Our design ensures that players at different levels of reasoning make distinct choices, providing a straightforward and transparent identification strategy. To control for subjects' beliefs, we have human participants interact with computer-controlled opponents programmed to play the Nash equilibrium strategy, thereby fixing subjects' expectations about their opponents and eliciting their best responses at their highest level of sophistication. This experimental design allows us to systematically vary game complexity and identify its causal effect on observed strategic sophistication. We also contribute methodologically, as the use of this refined game and the identification strategy improves the efficiency of identifying higher-order rationality compared to the original framework. While the original framework requires eight rounds of play, our game only needs to be played twice.

This paper provides novel evidence that individuals' observed strategic sophistication is systematically influenced by the complexity of the game they are playing. Across a range of games, we find that subjects exhibit, on average, a shift of more than one level of reasoning, a substantial effect given that the highest identifiable level in most of our games is Level-4. The proportion of Level-4 or above players ranges from 24.2% to 67.7%. These findings challenge the conventional view that strategic sophistication is an inherent trait of individuals or endogenously decided by the game stake or belief in other players (Alaoui and Penta, 2015). Instead, we argue that the revealed level of strategic sophistication is context-dependent and shaped by the characteristics of the game itself.

A key feature of our experimental design is the elimination of belief-based confounds. Prior research suggests that observed cognitive levels depend not only on individuals' inherent reasoning abilities but also on their beliefs about the sophistication of others (Agranov et al., 2012; Alaoui et al., 2020). To isolate the causal effect of complexity on ability to perform higher-order iteration, our experiment replaces strategic interaction between human subjects with interactions against computerized opponents. This design ensures that observed differences in reasoning levels are not driven by expectations about others' behaviours but rather reflect the intrinsic cognitive demands imposed by the game itself.

We further disentangle game complexity into three distinct dimensions: informational complexity, structural complexity, and the presence of salient strategies. Our results reveal that increased informational complexity reduces subjects' observed strategic sophistication, whereas structural complexity enhances it[1]. Notably, the positive effect of structural complexity is amplified when informational complexity is high, suggesting an interaction between these two dimensions. Additionally, we find that the existence of a dominant strategy significantly increases subjects' observed levels of reasoning. These findings highlight the multifaceted nature of game complexity and its pivotal role in shaping strategic sophistication.

Finally, we provide evidence that individuals with greater exposure to formal game theory tend to exhibit higher levels of strategic reasoning and are more likely to select equilibrium strategies. This challenges the traditional assumption that equilibrium behaviour emerges naturally and instead suggests that it may be acquired through learning and experience. Given the relative paucity of research on this topic, our findings open an important avenue for future work on the role of education and training in the development of strategic reasoning abilities.

[1]Some players' true reasoning levels may not be revealed in games where the Nash equilibrium requires only a lower level of reasoning. Increasing the level required to reach the Nash equilibrium naturally raises the overall average level. However, beyond this mechanical effect, we also observe an increase in reasoning levels among players who were not restricted in games with lower structural complexity.

This paper contributes to several strands of literature, foremost among them being the research on bounded rationality. The limitations of Nash equilibrium in explaining experimental results from games of initial play have long been recognized (Camerer, 2003). A widely adopted alternative is to assume heterogeneous levels of strategic sophistication among subjects, a framework formalized in Level-k and Cognitive Hierarchy models (Nagel, 1995; Costa-Gomes et al., 2001; Camerer et al., 2004).

However, an increasing body of literature has documented that observed strategic sophistication is not consistent across different types of games. Georganas et al. (2015)find that while reasoning levels are stable within a given family of games, neither the absolute levels nor the ranking of individuals' sophistication exhibit cross-game consistency. Similarly, Hyndman et al. (2022) studies a set of normal-form games and reports that 77% of subjects do not maintain stable reasoning levels across different game environments. These findings suggest that an individual's strategic sophistication is not fixed but context-dependent, though the precise ways in which it depends on context remain unexplored

To address this inconsistency, our study builds on the framework of Kneeland (2015), who is the first to introduce the network ring game as a method for identifying subjects with different levels of thinking. While Kneeland's framework provides a rigorous approach to distinguishing strategic sophistication, our paper advances the identification strategy by simplifying the original design while preserving its key strengths, such as the independence of Level-0 specification (See Section 3.1 for details). This refinement enables a more practical and scalable investigation of how different dimensions of game complexity shape strategic reasoning within a unified experimental setting.

Beyond an individual's intrinsic ability to engage in higher-order reasoning, another crucial determinant of observed strategic sophistication is the belief about the sophistication of others. Agranov et al. (2012) provide direct evidence that subjects' choices depend on their beliefs about their opponents' reasoning levels by experimentally controlling for beliefs in the 2/3 guessing game. In particular, their treatment condition includes manipulating participants' believes by playing against computer-generated opponents. Alaoui and Penta (2015) extend the analysis to an 11-20 game, showing that the depth of reasoning is endogenously determined by a combination of the player's cognitive ability, their beliefs about opponents' sophistication, and the payoffs, while keeping the game's sophistication exogenous. This highlights the importance of belief formation in shaping strategic behaviour.

Similar observations have been made in the context of the Ring Game. Researchers have distinguished between belief-bound and ability-bound subjects, recognizing that observed reasoning levels may not solely reflect cognitive limitations but also individuals' expectations about others' reasoning capacity. Jin (2021)introduces additional steps in the reasoning process compared to the one used in Kneeland's to separate belief-bound players from ability-bound ones and finds that both cognitive ability and beliefs jointly determine observed strategic sophistication.

In contrast to Jin's approach, which attempts to separate belief-bound and ability-bound players ex-post, our study follows Agranov et al. (2012) and Alaoui and Penta (2015)by directly manipulating players' beliefs through experimental treatments. By systematically controlling for beliefs, we directly focus on ability-bound players and also provide indirect evidence on the role of belief formation in strategic sophistication.

Meanwhile, a growing body of literature highlights the importance of complexity in decision-making, yet much of this research has focused on individual decision problems rather than strategic interactions. Our paper contributes to this gap by examining how different dimensions of game complexity influence higher-order rationality in strategic content.

The complexity of a game can manifest in multiple ways. A well-established strand of research in individual decision-making, particularly in financial contexts, has shown that complexity systematically shifts behaviours away from equilibrium predictions. Breunig et al. (2021) conducted an investment experiment on a representative group of German households and found that more complicated investment options cause systematically weaker responses to incentives. Similarly, Abeler and Jäger (2015) experimentally studied how complexity in tax systems affects workers' behavioral responses. They found that workers underreact to tax changes in more complex environments, suggesting that cognitive constraints distort decision-making under complexity. Oprea (2020) further demonstrates that implementing rules is cognitively costly and that individuals exhibit an aversion to complexity, which affects adherence to optimal decision-making.

Beyond its effect on observed strategic sophistication, our study also examines perceived complexity by analyzing deliberation time. Response time has frequently been used as a measure of cognitive effort in decision-making tasks (e.g., Alós-Ferrer and Buckenmaier (2020) in the 11-20 game) and has been viewed as a subjective measure of complexity. Grabiszewski and Horenstein (2022) use response time to gauge tree complexity in decision trees, illustrating how deliberation duration can capture cognitive constraints. Similarly, Huck and Weizsäcker (1999) find that individuals prefer simpler decision tasks, as evidenced by their behaviour in lottery choices.

Another body of literature relevant to this paper examines the impact of experience on strategic sophistication. Specifically, whether the Level-k behaviour is a natural result of profit-maximizing or it can be taught or triggered by experience in game theory remains unclear. Marchiori et al. (2021), using eye-tracking to study a normal-form game, found that comprehensive feedback from previous games significantly enhances strategic complexity. Xu (2022) studied whether iterative reasoning can be taught in lab and finds out that after a tutorial on how to perform equilibrium, 45% player learned to choose the equilibrium option. Game theory experience can be viewed as analogous, as players with more experience in game theory tend to undergo systematic learning in the discipline.

The rest of the paper is structured as follows. Section 2 introduces the conceptual background. Section 3 outlines the experimental design. Section 4 presents the results, and Section 5 discusses their implications and concludes.

2 Conceptual Background

2.1 Ring Games

The Ring Game is a type of network game characterized by sequential dependencies, designed to explore decision-making and strategic interaction in circular or interdependent environments. It was first introduced by Kneeland (2015) as an empirical framework for studying strategic reasoning. Compared to the traditional Level-k literature, the Ring Game significantly reduces the reliance on structural assumptions. For instance, distinguishing Level-k reasoning in the classic beauty contest game requires strong assumptions about the distribution of Level-k types to infer higher-order rationality from choice data. Kneeland's framework eliminates the need for these assumptions and avoids dependence on the specification of level-0 behavior.

This methodological advance provides a robust foundation for examining the relationship between higher-order rationality and game complexity. However, Kneeland's framework, which involves only human players, captures higher-order rationality as influenced by both players' intrinsic abilities (e.g., cognitive capacities and best-response capabilities) and their beliefs about opponents. These beliefs include expectations about opponents' strategies and behaviors.

2.2 Belief Bound and Ability Bound

For a player identified as Level-k (Lk), there are two possible explanations. First, their reasoning ability is constrained to allow precisely k iterations of strategic thinking. Second, their behavior reflects the best response to the belief that their opponent is an L(k-1) player. We define the former as ability-bound players, whose cognitive capacities have reached their upper limit, preventing further iterations of reasoning. These players represent the true focus of our study on Lk-level reasoning. In contrast, we define the latter as belief-bound players, whose capabilities may allow for additional iterations of reasoning. However, due to their belief that their opponents are of L(k-1), their best response behaviorally aligns with Lk, without necessitating further reasoning steps. For these players, the upper limit of their cognitive ability remains unobservable, as it is masked by their belief structure. This distinction is critical for our analysis, as our framework aims to isolate and study ability-bound Lk players, whose strategic reasoning is shaped by inherent cognitive constraints rather than beliefs about their opponents.

When focusing on the effects of complexity on higher-order rationality, it becomes essential to disentangle the impact of ability limitations from belief-related factors. This distinction is particularly crucial because discussing beliefs about others becomes irrelevant when players themselves struggle to reason effectively in complex environments. Building on Kneeland's framework, Jin (2019) addressed this limitation by introducing a sequential game that adds an additional reasoning step. This allows the differentiation between ability-bounded and belief-bounded players, as ability-bounded players fail to perform at level-k reasoning under the added complexity.

While Jin's method successfully distinguishes belief-bounded from ability-bounded players, it inherits several limitations of the Kneeland framework. For example, players are required to play the same game in different roles, implicitly assuming that their level of reasoning (or difficulty of iterated reasoning) remains consistent across positions. Additionally, this framework assumes no learning effects, a strong assumption given the repeated use of the same payoff table. Over successive rounds, players may learn strategic reasoning patterns, leading to potential misclassification of reasoning levels during identification.

Another limitation arises from the framework's reliance on multiple plays to identify higher-order reasoning. To identify level-k reasoning, at least k distinct plays are required. This means that increasing the complexity of the task would demand even more plays. Since our primary focus is on within-subject performance across varying levels of complexity, higher required rounds become impractical, both in terms of experimental feasibility and the potential for learning effects.

To address these challenges, we developed a novel identification strategy based on Kneeland's framework, aiming to eliminate learning effects at the individual level and maintain practical experimental design. We also employed a robot player approach to distinguish belief-bounded from ability-bounded players, similar to the method used by Xu (2022) to identify ability bound. Instead of having participants repeatedly interact with each other, as in Kneeland's setup, we introduce computer-controlled players programmed to always select the equilibrium choice.

Our game retains the structure of Kneeland's dominant-solvable design, where a unique Nash equilibrium exists. The robot players, by playing equilibrium strategies, fully exploit the iterative reasoning limits of their assigned positions. For example, in a 4-player game Player 4, having a dominant strategy, only needs to exhibit L1 reasoning to choose the equilibrium option. Players 3, 2, and 1 must demonstrate L2, L3, and L4 reasoning, respectively, to play their optimal choices at equilibrium.

By programming robot players to consistently perform at their reasoning capacity in all positions, we eliminate the influence of participants' beliefs about their opponents. During instructions and comprehension tests, participants are informed of the robot players' decision rules. Rather than explicitly describing the game as dominant-solvable, we use neutral language, explaining that robot players analyze all payoff tables, assume all players aim to maximize their payoff and select the highest achievable outcome. This explanation conveys the robots' decision-making processes without leading participants to anticipate the dominant-solvable nature of the game.

This design fixes participants' beliefs about their opponents, ensuring that choices reflect only their reasoning capacities. For players with Lk reasoning, where $k \geq N$ or the ability to do iterative learning is greater than the highest identifiable level (N) of the game, they can identify the only one dominant strategy of the game and find out the optimal choice using backward induction, which is the equilibrium strategy. For $k \leq N$, the best response reflects the maximal level of iterative reasoning achievable. By removing the influence of belief-bounded players, our results capture only ability-bounded participants, isolating the effect of complexity on reasoning ability.

2.3 Our design

This study builds on the framework of the Ring Game introduced by Kneeland (2015) to investigate subjects' higher-order rationality, introducing several key innovations. In the original framework, the game involves four players arranged in a closed-ring structure, each selecting from three possible options. A player's payoff depends on their own choice and the choice of their direct opponent in the ring—Player 1 is paired with Player 2, Player 2 with Player 3, and so forth, with Player 4 looping back to Player 1. The structure allows all payoffs to be captured using four 3×3 payoff matrices. Subjects need to play each role, from Player 1 to Player 4, twice before they can be identified at a specific level, requiring a total of eight rounds for identification.

Our experiment modifies this framework by introducing an N-player, n-choice Ring Game, where N corresponds to the highest level of reasoning (Lk) that the game can identify. There are n choices, and we aim to identify individuals of types L1 to LN, exact identification is feasible when N=n, provided individuals of each Level selects a distinct choice from the set of n options. For example, in a two-player (S_1, S_2), two choice game (U, D), the payoff scheme can be designed so that S_1 will only choose U and S_2 only choose D. By looking at the choice, we then know which player is of type S_1 and S_2. This structure ensures that individuals with varying reasoning levels will select different options, allowing their reasoning level (Lk) to be inferred directly from their choice. To reduce the possibility of misidentification, we follow Kneeland's approach by requiring the game to be played twice with relabeling. Unlike the original framework, which uses eight rounds to specify a subject's level, our design allows the identification process to be completed within two rounds, significantly improving the efficiency of analyzing higher-order rationality.

The innovation retains the conceptual elegance of the original framework while improving its efficiency and allowing for precise manipulation of game complexity.

Additionally, the experiment employs a within-subject design, providing better control over individual effects and isolating the impact of game complexity on rationality. These enhancements reduce practical constraints, such as excessive rounds, while maintaining rigorous analysis of strategic reasoning.

3 Experiment Design

Our experiment implements a variation of the Ring Game described in Section 2, employing a within-subject design in which human subjects interact with pre-programmed robot players.

3.1 Ring Games

In this experiment, we identify subjects' Lk levels using a set of 18 N-player, n-choice dominant-solvable Ring Games (N=n), where N or n varies by treatment (see Section 3.3). These games are designed to capture key factors influencing complexity in decision-making, as established in the behavioural literature. Each game is presented using payoff matrices that summarize all possible choices and their corresponding payoffs based on different choice combinations. Below is a screenshot of one of the games used in our experiment.

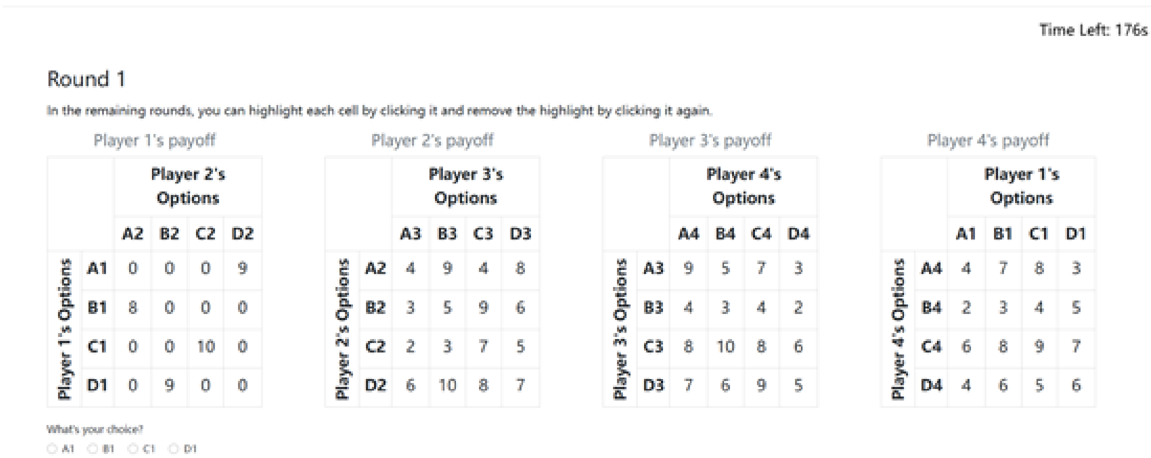

Player 1's payoff

Player 1's Options	Player 2's Options: A2	B2	C2	D2
A1	0	0	0	9
B1	8	0	0	0
C1	0	0	10	0
D1	0	9	0	0

Player 2's payoff

Player 2's Options	Player 3's Options: A3	B3	C3	D3
A2	4	9	4	8
B2	3	5	9	6
C2	2	3	7	5
D2	6	10	8	7

Player 3's payoff

Player 3's Options	Player 4's Options: A4	B4	C4	D4
A3	9	5	7	3
B3	4	3	4	2
C3	8	10	8	6
D3	7	6	9	5

Player 4's payoff

Player 4's Options	Player 1's Options: A1	B1	C1	D1
A4	4	7	8	3
B4	2	3	4	5
C4	6	8	9	7
D4	4	6	5	6

Figure 1: Screenshot of one of the Ring Games played in the main section

In these matrices, subjects always assume the role of the row player, while the columns represent the potential choices of their direct opponent. The table displays all possible choices for each individual and the corresponding payoffs. For example, if Player 1 selects B1 and Player 2 selects A2, Player 1 receives a payoff of 8. If Player 4 selects A4, their payoff is 7. Throughout the experiment, human subjects always play as Player 1, while their opponents are robot players programmed to always select the equilibrium strategy. Subjects are informed of this rule before the experiment begins and can refer to printed instructions at any time.

In each round, human and robot players make their decisions simultaneously. However, subjects do not receive feedback on the robot players' choices; they only learn their own payoffs at the end of the experiment. Each game lasts between 30 seconds and 3 minutes. The "Submit" button becomes available only after 30 seconds, and a timer is displayed on the top right throughout the round. When only 30 seconds remain, the timer turns red. If a subject fails to make a selection before time expires, they are automatically advanced to the next game. Subjects can highlight any cell by clicking on it, with choice cells highlighted in green and payoff cells highlighted in yellow.

To control for individual differences, the experiment employs a within-subject design. Each subject participates in 9 groups of games, with each group containing 2 Ring Games, totaling 18 Ring Games. This structure follows the Natural Exclusivity condition used in Kneeland (2015), where the only difference between the two games within a group is the location of specific options (see Section 3.2 on the identification strategy). The 18 games consist of: 2 games with 3 players and 3 choices, 2 games with 5 players and 5 choices, and 14 games with 4 players and 4 choices.

3.2 Identification strategy

As discussed in Section 2, this paper introduces an innovative method to identify subjects' level of reasoning which increase the efficiency of identification. Rather than imposing assumptions on subjects, we embed the necessary identification mechanism directly within the game's payoff matrices. Let $\mathcal{S} = s_1, s_2, ..., s_n$ be the set of n available choices, and let $\mathcal{L} = L_1, L_2, ..., L_N$ represent the set of individuals we aim to identify, where each L_i follows a strategy σ_i mapping to a choice in $\mathcal{S}$. Exact identification is feasible if and only if $N = n$ and the mapping $\sigma : \mathcal{N} \to \mathcal{S}$ is injective, meaning each individual selects a distinct choice from S. Since the mapping is injective, each choice each choice uniquely corresponds to a single individual type. Therefore, knowing an individual's choice $s \in \mathcal{S}$ directly reveals their type $L_i \in L$.

Player 1's Payoff

Player-1's Choice \ Player-2's Choice	A2	B2	C2	D2
A1	4	2	6	9
B1	8	7	3	5
C1	3	8	10	7
D1	6	9	5	4

Player 2's Payoff

Player-2's Choice \ Player-3's Choice	A3	B3	C3	D3
A2	4	9	4	8
B2	3	5	9	6
C2	2	3	7	5
D2	6	10	8	7

Player 3's Payoff

Player-3's Choice \ Player-4's Choice	A4	B4	C4	D4
A3	9	5	7	3
B3	4	3	4	2
C3	8	10	8	6
D3	7	6	9	5

Player 4's Payoff

Player-4's Choice \ Player-1's Choice	A1	B1	C1	D1
A4	4	7	8	3
B4	2	3	4	5
C4	6	8	9	7
D4	4	6	5	6

Figure 2: Example of one Ring Game

Using the Ring Game illustrated in Figure 2 as an example, the payoff matrices represent one of the baseline games used in this study. All players have access to all payoff matrices. Human subjects always act as Player 1 (P1), while the remaining players are computer-controlled (robot) players programmed to act as equilibrium players. Among the four players, Player 4 (P4) has a dominant strategy, and the game's unique equilibrium can be derived via backward induction. For this game, there exists only one Nash equilibrium (B1, A2, D3, C4). At the Position of P4, an L1 player has sufficient intelligence to recognize the dominant strategy C4. While at the position of P3, and P2, it requires at least L2 and L3 respectively to recognize the dominant strategy of P4 is C4 and iteratively figure out the best response D3, A2. At the position of P1, which is the only position played by our human subjects, it requires at least L4 to identify the only dominant strategy among all the players and best respond to this by choosing B1. So, if a subject chose B1 at position P1, we know that he/she is at least an L4 player.

Since our human subject only plays the role of P1, we still need the identification strategy for L0 to L3 player at P1. The identification strategy begins with the assumption that L0 players are non-strategic, selecting options randomly. L1 players, by contrast, best respond to L0 behaviour or employ heuristic decision rules such as risk dominance or aiming for the highest possible payoff. The payoff matrix is designed to ensure that these rules lead to the same choice. For instance, in the payoff matrix above, option C1 satisfies multiple criteria: it has the highest expected payoff when assuming Player 2 (P2) chooses randomly, it aligns with the risk-dominant strategy, and it contains the maximum payoff value of 10 in the first matrix.

L2 players best respond to an opponent being one level lower than theirs, that is L2 players at P1 must consider their own payoff matrix as well as P2's and believe their opponent P2 is at L1. A rational L1 player at P2, will repeat the decision process of L1 as discussed previously and choose D2. Therefore, the best response for L2 at P1 will be to choose A1. At this stage, this L2 player is only an ability-bounded L2 player and cannot be a belief-bounded L3 or above player, meaning their behaviour only reflects their reasoning capacity rather than they are best responding to some incorrect beliefs. Use an example to illustrate this, when a belief-bounded L3 player at P1 chose A1, this can be a best response of believing P2 is L2 and choosing D2, based on the level-k literature, this also implies that P3 believes P2 do this because his opponent P3 is of L1 and is choosing B3. However, that's where the conflict arises, if P3 is of L1, he should have the ability to reason out that B3 is a never best response at P3 and shouldn't choose B3 as a rational player. This means that an L2 player at P2 should also know this and update his behavioural accordingly, and an L3 player at P1 should also update this and never choose A1. Therefore, if a player chooses A1 at P1, he cannot be a belief-bounded L3 player, but only an ability-bounded L2 player. Because only L2 player at P1 will believe his opponent is of L1 and only look at the payoff matrices of P1 and P2 but not P3. As long as the player checked the payoff matrices of P3, he should figure out that B3 should never be chosen.

The identification of level-3 players follows a similar pattern. L3 players can perform up to three iterations of reasoning. Such players, assuming P3 selects C3, would infer that P2 will choose B2, leading to D1 as P1's response. Importantly, this behaviour remains "ability-bound," as one additional iteration would reveal that B4 is not a best response. Given that P4 has a dominant strategy, the only optimal response for P1 is B1. Consequently, subjects selecting B1 can be identified as level-4 or higher.

The iteration path of different Lk player is highlighted in different colours, equilibrium choice or the choice path for players of L4 or above is highlighted in yellow, L3 players' iteration path is in green, L2's path in purple and L1's in blue. The only dominant strategy of all the payoff matrices is circled in yellow. In summary, for human players who always play at P1, if they are of level-4 or above, they should go through all the payoff matrices and iteratively figure out that only P4 has a dominant strategy, and have the ability to do backward induction to best response at P1 position. The iteration path for L4 or above player is C4-D3-A2-B1, once figure out the

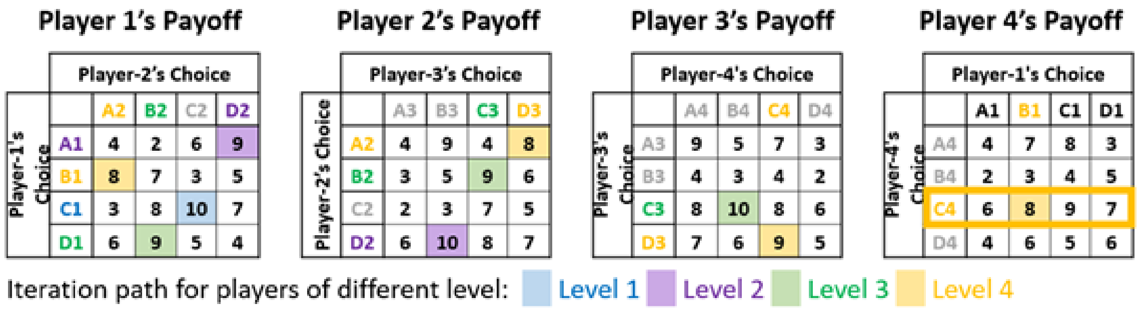

Figure 3: Iteration Path for players of different Levelwhere C4 for Player 4 is the only dominant strategy among the 4 matrices

dominant strategy of C4, L4 player has sufficient cognitive ability to do best respond at the position of P3, P2, which results in belief P3 will choose D3 and P2 will chose A2, and finally, as P1, will choose B1. Similarly, L3 player can only do three times of iteration and go through the path C3-B2-D1, and L2 player's path is D2-A1, while L1 player will assume their opponent P2 playing non-strategically and best respond by choosing the risk dominant strategy C1.

One concern in this framework is the limited number of choices, which may lead to misidentification. To address this, the study incorporates Kneeland's (2015) natural exclusivity restriction that subjects do not respond to payoff changes exceeding their capacity for reasoning, by relabelling options and repeating the game, expanding the choice portfolio from 4 to 16. Since the player of Lk will not be affected by the changing in Lk+1's choice, or in our game, an Lk player at P1 will not take into account the payoff table of Pk+1's. As shown in the screenshots below, from Round 1 to Round 2, we changed the order of P4's dominant strategy, in Round 1, P4's dominant strategy is C4 while in Round 2, P4's dominant strategy is A4.

To retain the attractive features of our Ring Games that we can lead people of different levels to different choices, payoff matrices for P1 and P3 are also relabelled. Otherwise, some risk-dominant choices will correspond to the equilibrium choice and we will not be able to separate belief-bound players and ability-bound players. For example, if we only change the order of P4's dominant strategy as following P4's payoff table from Round 2, and keep using the other three payoff matrices from Round 1, we will have the equilibrium choice D2 of P2 the same as the risk dominant strategy, which will violate our identification strategy. And since we don't change the numbers inside the payoff matrices but purely the order of the payoff matrices, this will not

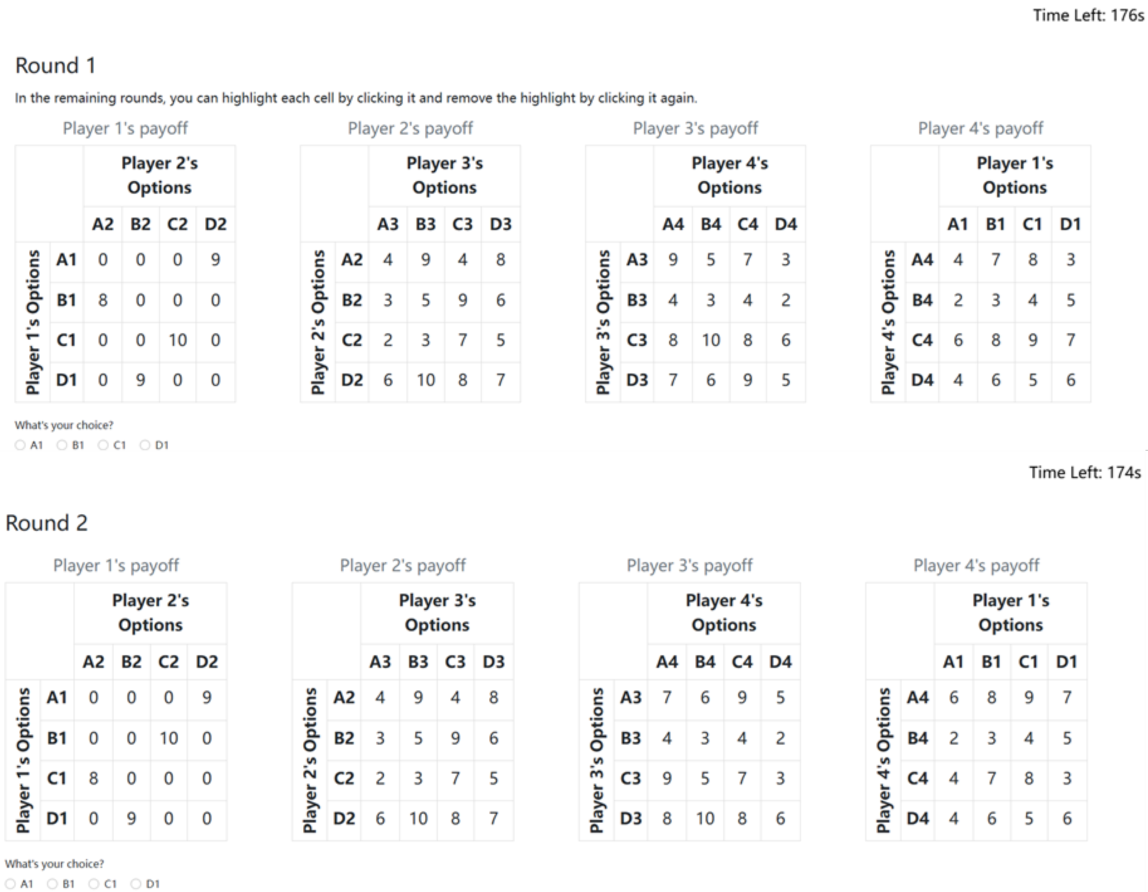

Time Left: 176s

Round 1

In the remaining rounds, you can highlight each cell by clicking it and remove the highlight by clicking it again.

Player 1's payoff

Player 1's Options \ Player 2's Options	A2	B2	C2	D2
A1	0	0	0	9
B1	8	0	0	0
C1	0	0	10	0
D1	0	9	0	0

Player 2's payoff

Player 2's Options \ Player 3's Options	A3	B3	C3	D3
A2	4	9	4	8
B2	3	5	9	6
C2	2	3	7	5
D2	6	10	8	7

Player 3's payoff

Player 3's Options \ Player 4's Options	A4	B4	C4	D4
A3	9	5	7	3
B3	4	3	4	2
C3	8	10	8	6
D3	7	6	9	5

Player 4's payoff

Player 4's Options \ Player 1's Options	A1	B1	C1	D1
A4	4	7	8	3
B4	2	3	4	5
C4	6	8	9	7
D4	4	6	5	6

What's your choice?

○ A1 ○ B1 ○ C1 ○ D1

Time Left: 174s

Round 2

Player 1's payoff

Player 1's Options \ Player 2's Options	A2	B2	C2	D2
A1	0	0	0	9
B1	0	0	10	0
C1	8	0	0	0
D1	0	9	0	0

Player 2's payoff

Player 2's Options \ Player 3's Options	A3	B3	C3	D3
A2	4	9	4	8
B2	3	5	9	6
C2	2	3	7	5
D2	6	10	8	7

Player 3's payoff

Player 3's Options \ Player 4's Options	A4	B4	C4	D4
A3	7	6	9	5
B3	4	3	4	2
C3	9	5	7	3
D3	8	10	8	6

Player 4's payoff

Player 4's Options \ Player 1's Options	A1	B1	C1	D1
A4	6	8	9	7
B4	2	3	4	5
C4	4	7	8	3
D4	4	6	5	6

What's your choice?

○ A1 ○ B1 ○ C1 ○ D1

Figure 4: An Example of Applying Natural Exclusive Restriction

affect the complexity we want to measure and also prevent subjects from learning from the same payoff matrices.

Moreover, since we assume L0 players are choosing randomly, only subjects who exhaust their time without selecting an option are classified as L0. For conservative analysis, a subject's final level for each game is assigned as the lower level when ambiguous behaviour is observed during playing the two rounds of game in one group, thereby minimizing the risk of misidentification. Following table summarize the choice portfolio for Lk players when playing the above two rounds of games. We are using a conservative identification strategy, there will be a lower probability of being identified at a higher level.

L0	L1	L2	L3	L4
Time out in any round	(C1,B1)	(A1,A1)	(D1,C1)	(B1,D1)
	(C1,A1)	(B1,A1)	(D1,D1)	
	(C1,C1)	(D1,A1)	(B1,C1)	
	(C1,D1)	(A1,B1)		
	(A1,C1)	(A1,D1)		
	(B1,C1)			
	(D1,C1)			

Table 1: Summary of LK Player's Choice Profile for the game in Figure 4 (with each pair of choices indicating decisions in Round 1 and Round 2, respectively)

3.3 Treatment

Figure 5 presents the detailed payoff matrices used in our experiment.

The first and most important hypothesis we test is:

Hypothesis 1: Subjects' higher-order rationality is affected by the complexity of a game.

Building on our identification strategy, the central idea of this study is to systematically vary the complexity of the game and analyze how individuals adjust their behaviour in response. Since we employ a within-subject design, we can largely eliminate individual heterogeneity and isolate changes in iterative reasoning behaviour. To further investigate subjects' reasoning process, we include a Revealing Game treatment at the end of the experiment. In this task, all numbers in the payoff matrices are initially hidden, and subjects must click on each cell to reveal its value[2]. The payoff matrices in this condition share the same structure as those in Baseline 2, with the only difference being the additional step of clicking to uncover the payoffs. This design allows us to track the decision-making process in real time, capturing how subjects explore the strategy space and engage with the available information.

This study further investigates three key facets of complexity and tests the corresponding hypotheses, which we label as Structural Complexity, Informational Com-

[2]Each cell needs to be clicked only once; once revealed, the number remains visible for the rest of the task.

	Player 1's Payoff					Player 2's Payoff					Player 3's Payoff					Player 4's Payoff				
		A2	B2	C2	D2		A3	B3	C3	D3		A4	B4	C4	D4		A1	B1	C1	D1
	A1	0	0	0	9	A2	0	0	0	8	A3	9	0	0	0	A4	0	0	0	0
Baseline 1	B1	8	0	0	0	B2	0	0	9	0	B3	0	0	0	0	B4	0	0	0	0
(B1)	C1	0	0	10	0	C2	0	0	0	0	C3	0	10	0	0	C4	9	8	9	9
	D1	0	9	0	0	D2	0	10	0	0	D3	0	0	9	0	D4	0	0	0	0
	A1	4	2	6	9	A2	4	9	4	8	A3	9	5	7	3	A4	4	7	8	3
Baseline 2	B1	8	7	3	5	B2	3	5	9	6	B3	4	3	4	2	B4	2	3	4	5
(B2)	C1	3	8	10	7	C2	2	3	7	5	C3	8	10	8	6	C4	6	8	9	7
	D1	6	9	5	4	D2	6	10	8	7	D3	7	6	9	5	D4	4	6	5	6
	A1	0	0	0	9	A2	0	9	0	8	A3	9	0	7	0	A4	0	7	8	0
Half Zero 1	B1	8	7	0	0	B2	0	0	9	0	B3	0	0	0	0	B4	0	0	0	5
(HZ1)	C1	0	8	10	7	C2	0	0	7	0	C3	8	10	8	6	C4	6	8	9	7
	D1	6	9	0	0	D2	6	10	8	7	D3	7	0	9	0	D4	0	6	0	6
	A1	4	2	6	9	A2	0	0	0	8	A3	9	0	0	0	A4	4	7	8	3
Half Zero 2	B1	8	7	3	5	B2	0	0	9	0	B3	0	0	0	0	B4	2	3	4	5
(HZ2)	C1	3	8	10	7	C2	0	0	0	0	C3	0	10	0	0	C4	6	8	9	7
	D1	6	9	5	4	D2	0	10	0	0	D3	0	0	9	0	D4	4	6	5	6
	A1	0	0	0	9	A2	4	9	4	8	A3	9	5	7	3	A4	4	7	8	3
Salience of	B1	8	0	0	0	B2	3	5	9	6	B3	4	3	4	2	B4	2	3	4	5
Iteration **(SI)**	C1	0	0	10	0	C2	2	3	7	5	C3	8	10	8	6	C4	6	8	9	7
	D1	0	9	0	0	D2	6	10	8	7	D3	7	6	9	5	D4	4	6	5	6
Salience of	A1	4	2	6	9	A2	4	9	4	8	A3	9	5	7	3	A4	0	0	0	0
Dominant	B1	8	7	3	5	B2	3	5	9	6	B3	4	3	4	2	B4	0	0	0	0
Strategy	C1	3	8	10	7	C2	2	3	7	5	C3	8	10	8	6	C4	6	8	9	7
(SD)	D1	6	9	5	4	D2	6	10	8	7	D3	7	6	9	5	D4	0	0	0	0

	A1	9	0	0	A2	0	10	0	A3	0	0	0
3×3	B1	0	10	0	B2	9	0	0	B3	0	0	0
	C1	0	0	8	C2	0	0	9	C3	9	9	8

	Player 1's Payoff						Player 2's Payoff						Player 3's Payoff						Player 4's Payoff						Player 5's Payoff					
		A2	B2	C2	D2	E2		A3	B3	C3	D3	E3		A4	B4	C4	D4	E4		A5	B5	C5	D5	E5		A1	B1	C1	D1	E1
	A1	3	7	4	9	2	A2	4	6	6	8	5	A3	9	8	3	5	2	A4	9	7	5	3	6	A5	3	2	8	2	7
	B1	9	2	5	5	3	B2	3	7	9	3	7	B3	3	2	2	3	4	B4	2	3	4	2	3	B5	4	3	8	3	6
5×5	C1	4	6	9	7	2	C2	4	8	5	7	9	C3	5	10	4	6	5	C4	7	10	6	6	5	C5	6	5	9	7	8
	D1	6	9	8	5	4	D2	6	10	4	5	6	D3	7	4	9	7	2	D4	5	8	9	3	7	D5	5	3	2	6	5
	E1	7	5	6	6	10	E2	2	5	3	2	3	E3	6	3	4	8	3	E4	6	5	7	4	8	E5	2	4	7	4	4

Figure 5: Summary of Treatment

Each row summarizes the payoff tables for one group of games, with each group containing two games. The games within a group differ only in the labeling of the dominant strategy of the last player.

plexity, and Salience of Strategy.

Hypothesis 2: Informational complexity decreases the revealed level of subjects, and only the general level of informational complexity matters. Different methods of changing informational complexity do not affect the subject's higher-order rationality.

Informational complexity refers to the level of uncertainty and dispersion within the strategic environment. While structural complexity alters the strategy space by expanding the number of possible combinations, informational complexity influences how players perceive and navigate a fixed strategy space. We manipulate informational complexity by introducing zero-payoff cells into the payoff matrices. Since all players aim to maximize their payoffs, they are expected to avoid options containing zero payoffs and assume that others will do the same. This effectively eliminates certain choice combinations, thereby reducing the uncertainty of the game.

To control for informational complexity, we construct three levels by varying the proportion of zero payoffs. As shown in Figure 5, Baseline 1 contains the highest proportion of zero payoffs, Half Zero 1 and Half Zero 2 represent intermediate levels, and Baseline 2 contains none. The two Half Zero treatments also serve as a robustness check for different distributions of zero payoffs[3]. These variations allow us to assess how the overall level and structure of informational complexity influence strategic reasoning.

Hypothesis 3: Structural complexity decreases the revealed level of subjects.

Structural complexity in strategic interactions refers to the extent to which a game's architecture, defined by the number of players and available choices, affects decision-making. In our framework, the number of players is always equal to the number of choices to ensure proper identification. As a result, structural complexity is captured by N, where each game consists of N players and N choices. Instead of fixing $N = 4$, we introduce variation by incorporating games with $N = 3$ and $N = 5$, as shown in the 3×3 and 5×5 treatments in Figure 5.

Structural complexity is important because it exponentially increases the strategy space. In a 3×3 game, there are 27 possible choice combinations; in a 4×4 game, this increases to 256; and in a 5×5 game, the number of possible combinations escalates to 3,125. As structural complexity increases, subjects are required to engage in deeper levels of iterative reasoning to reach equilibrium. This may exceed their cognitive capacity and lead them to rely more heavily on heuristics. In contrast, in less complex environments, players may find it easier to anticipate others' strategies and exhibit behaviour closer to equilibrium.

We distinguish between a mechanical effect, where increasing N raises the average observed level because players with K higher than N are indistinguishable from those with K=N when N is low, and a broader effect beyond this mechanism. We argue that the latter is stronger, leading to an overall decrease in observed levels despite

[3]To ensure the identification of reasoning levels, the proportion of zero payoffs in Half Zero 1 and Half Zero 2 is not exactly the same.

the mechanical effect.

Hypothesis 4: The two treatments, emphasizing either the Dominant Strategy or Iteration, increase subjects' revealed level of reasoning.

We further test whether highlighting certain strategic considerations can promote higher-order thinking. We refer to these treatments as Salience of Dominant Strategy and Salience of Iteration, reflecting the intended emphasis at the design stage. These labels are used as descriptive shorthand, rather than as definitive claims about the underlying behavioural mechanisms.

In both treatments, subjects face one payoff matrix from Baseline 1 (the simpler game) and three from Baseline 2 (the more complex game). In the Salience of Iteration condition, the player's own payoff matrix (e.g., Player 1's) is taken from the simpler baseline. This makes their own optimal response relatively transparent and shifts the primary difficulty to anticipating the opponent's behavior, thereby encouraging deeper iterative reasoning.

In contrast, the Salience of Dominant Strategy treatment simplifies the payoff matrix of the opponent (e.g., Player 4), making their dominant strategy more obvious. This reduces uncertainty about the opponent's likely move, potentially facilitating reasoning about one's own best response. If difficulty in predicting the opponent's strategy is a major barrier to higher-order reasoning, this treatment is expected to mitigate that complexity.

3.4 Complexity Measurements

In this study we employ both ex-ante and ex-post measurements of complexity to capture the subjective and objective dimensions of game complexity. The three ex-ante measurements closely correspond to the three treatments in the experiment. We also include ex-post measurements including response time and click history to measure the perceived complexity of subjects.

Ex-ante measurements

We measure Structural Complexity using the highest identifiable level- k(Lk) in each game, denoted as N, which captures both the number of available choices per player and the number of players involved. Additionally, we introduce a dummy variable to indicate whether a particular game design makes certain strategies salient.

The key distinction in our ex-ante complexity measures lies in our use of entropy as a measure of informational complexity, replacing the zero-payoff rate. The zero-payoff rate is inherently restrictive, as it can only be applied when zero-payoff outcomes exist within the payoff matrix. In contrast, entropy provides a more generalizable approach, allowing for broader application across different strategic environments. Originally introduced by Shannon (1948) in information theory, entropy quantifies the uncertainty or unpredictability of a probability distribution. In the context of the Ring Game, entropy serves as a natural measure of informational complexity, capturing the extent of uncertainty players face when making strategic decisions. Formally, the formula for the entropy H(x) of a game is

$$H(x) = -\sum_{i} P(x_i) \log_2 P(x_i) \tag{1}$$

Where x_i is the number in all the payoff matrices in one game, $P(x_i)$ represents the frequency of each value in all matrices.

Higher entropy reflects greater informational complexity, as payoffs are more dispersed and less predictable, increasing the difficulty of identifying optimal strategies. Conversely, lower entropy indicates a more structured and predictable environment, facilitating decision-making. By employing entropy to analyze informational complexity, we can systematically compare different strategic environments and assess how uncertainty influences decision-making processes.

Ex-post Measurement

Ex-post measurements primarily capture subjects' perceived complexity of the game. While these perceptions do not causally affect revealed higher-order rationality, they reflect the subjective difficulty players experience. To examine this, we track the time subjects spend on each game and record their click history for each cell.

Table2 summarizes the different dimensions of complexity for each game in the experiment. Notably, the zero-payoff rate is monotonic, entropy is not. As a result, we observe that although the zero-payoff rate in the SI and SD treatments is not the lowest among all games, entropy suggests that these games exhibit the highest level of informational complexity. This discrepancy arises from the fundamental nature of entropy, which captures the dispersion of payoffs within the matrices. Since we do not differentiate between numerical values in different tables, increasing the number of zero-payoff entries—as in the transition from B2 to SI and SD—results in a more uniform distribution of numbers in payoff matrices. As a consequence, the frequency of each payoff value becomes more evenly spread, leading to greater entropy and, therefore, higher informational complexity.

Game	Entropy	Zero-Payoff Rate	N	Salience of Iteration	Salience of Dominant Strategy	Average Time (s)	Average Click (per table)
B1	1.07	78.13%	4	0	0	54.96	2.17
B2	3.10	0.00%	4	0	0	72.81	2.26
3×3	1.40	66.67%	3	0	0	45.69	1.88
5×5	3.08	0.00%	5	0	0	86.65	2.22
HZ1	2.23	48.44%	4	0	0	73.57	2.39
HZ2	2.82	40.63%	4	0	0	84.72	2.40
SI	3.22	18.75%	4	1	0	86.89	2.61
SD	3.22	18.75%	4	0	1	63.46	2.13
Revealing Game	3.10	0.00%	4	0	0	82.01	12.89

Table 2: Summary of Complexity of Game

3.5 Implementation

The experiment was conducted in the Behavioural Laboratory at the University of Edinburgh (BLUE) in April 2024. A total of 127 students, including both undergraduates and postgraduates from various academic disciplines, participated in the study. The experiment was implemented using oTree (Chen et al., 2016).

Before beginning the experiment, subjects received detailed instructions about the games and payment structure, including a comprehensive explanation of the decision rules governing the robot players. They were provided with printed instructions throughout the session and were required to pass a comprehension test on robot behaviour and understanding of the game before proceeding. Out of the 127 participants, 124 successfully passed the test and completed all tasks.

During the experiment, subjects could highlight a cell with a single click and remove highlights by clicking again. Each participant played all 18 Ring Games, with the order of the main games randomly predetermined as follows: SI – HZ1 – B1 – 3×3 – 5×5 – SD – B2 – HZ2. To account for potential order effects, 65 participants played the games in this sequence, while 59 played them in reverse order. After completing the main games, subjects proceeded to the Revealing Game.

Following the experimental tasks, participants completed a Cognitive Reflection

Test (CRT) consisting of four questions, as well as a demographic questionnaire collecting information on age, gender, and academic background. Additionally, subjects self-reported their previous experience with game theory using a five-point scale: from one to five, the degree of understanding is 'Never Heard of It', 'Heard of It', "Know It but Don't Know How to Use It", 'Know It and Can Use It', 'Good at Using It'.

Subject's monetary payment consists of three parts, first is a £4 show-up fee, second is a performance-based payment tied to their choices in the experiment, and the third is a CRT-based payment, where subjects earned £0.50 for each correctly answered question. Specifically, four out of the 18 Ring Games were randomly selected for payment, and the actual values in the payoff tables were converted into monetary rewards at a rate of £0.25 per point. On average, participants earned £11.49, and the experiment lasted approximately 45 minutes per session.

4 Results

4.1 Overall Result

We observe considerable variations in subjects' higher-order rationality with the change of a game's complexity. Table 3 summarizes the identified levels in each game, and the average level across all games in our study is 2.49. The variation of the identified level appears both in the distribution of level and the average level.

The lowest average level, 2.12, is observed in Half Zero Game 1 (HZ1), where half of the payoff values in the matrices are zero. In contrast, the highest average level, 3.18, is found in Baseline Game 1 (B1), where most payoffs are zero. Given that the highest identifiable level in all games is Level 5, and most games have a maximum of Level 4, the difference of more than one level between the lowest and highest average suggests a significant impact of complexity on subjects' higher-order rationality. In general, average reasoning levels are higher in Baseline Game 1 and the 5×5 Game, whereas Baseline Game 2, Half Zero 1, and Salience of Iteration 1 exhibit lower levels.

We also observe two distinct patterns of the distribution of level. The first group of games features the equilibrium player as the largest subgroup, including Baseline Game 1, the 3×3 game, the 5×5 game, SD, and the Revealing Game. Among these, Baseline Game 1 has the highest proportion of level-4 players, with 67.7% of subjects

Game	L0	L1	L2	L3	L4	L5	Average Level
B1	1.6%	18.5%	8.1%	4.0%	67.7%	N/A	3.18
B2	4.8%	45.2%	11.3%	8.9%	29.8%	N/A	2.14
3×3	0.0%	24.2%	4.8%	71.0%	N/A	N/A	2.47
5×5	6.5%	29.0%	4.8%	4.8%	20.2%	34.7%	3.07
HZ1	2.4%	51.6%	7.3%	8.9%	29.8%	N/A	2.12
HZ2	11.3%	35.5%	8.9%	10.5%	33.9%	N/A	2.20
SI	9.7%	34.7%	11.3%	20.2%	24.2%	N/A	2.15
SD	3.2%	38.7%	8.9%	4.8%	44.4%	N/A	2.48
Revealing Game	4.0%	28.2%	12.1%	14.5%	41.1%	N/A	2.60

Table 3: Frequency of levels and the average level in each game

displaying Level 4 reasoning. While the 5×5 game has the smallest proportion of level-4 players (about 20%), it still contains the second-largest proportion of players at level-4 or above among all the games (54.9%). The 3×3 game, on the other hand, has the largest proportion of Level 3 players, displaying a distribution pattern similar to that of Baseline Game 1.

The second group comprises games in which the level-1 player is the largest subgroup, including Baseline Game 2, Half-Zero 1 & 2, and SI. These games exhibit similar distributions in the proportion of players at levels other than L1, where L4 players constitute the second-largest group. Across these games, reasoning levels reveal a clear bimodal pattern, with most players clustering at either level-1 or level-4. This highlights two distinct approaches among participants: some rely solely on their own payoff matrices without considering others' strategies, while others fully recognize the necessity of higher-order reasoning and demonstrate the capability for level-4 reasoning or above.

These differences in pattern of distribution also suggest an interesting finding, that is subjects seldom get stuck at the middle level. Compared with other forms of games, which have been used to study higher-order rationality, especially the beauty contest, where people seldom get to play the equilibrium. This might suggest that the structure of the Ring game naturally nudges people to do more iterations. This finding also supports the findings in the literature that subject levels are not constant across different types of game.

The impact of game complexity is also evident at the individual level. Only 5 participants maintained the same reasoning level across all 18 rounds, each consistently demonstrating level-1, and 7 Nash equilibrium players. In contrast, the remaining 112 participants displayed variability in their reasoning levels, raising the question of whether these variations resulted from accidental errors. One concern is that a player might consistently hold a constant reasoning level but occasionally select an incorrect answer by mistake, leading to a lower revealed level.

To examine this possibility, we assume that a participant's true reasoning level corresponds to the mode of their behavior across all games. As the mode is the most frequently observed level, this assumption minimizes the estimated mistake rate. Mistake rates were then calculated based on how often a participant's choices deviated from their most-played level. Among the 124 participants, 5 showed no deviations, and another 5 deviated in only one game[4]. However, 76.6% of participants deviated in at least three games, and the average mistake rate was 34.6% (considering only the seven 4-player games). A mistake rate of over one-third is far too high to be purely accidental.

Further evidence against the error hypothesis comes from logistic regression analysis conducted at the individual level. At the 95% confidence level, we can reject the null hypothesis that participants' choices come from the same distribution, meaning that their choices do not consistently align with their mode. Together, these findings indicate that the observed variability in reasoning levels cannot be attributed solely to errors; rather, it reflects genuine shifts in participants' reasoning in response to the complexities of different games.

These findings strongly support our central argument that players' reasoning levels are inherently not constant and that game complexity primarily influences those with intermediate levels of strategic sophistication. The observed variation highlights the importance of considering individual differences when analyzing strategic behavior, as the response to increasing complexity is not uniform across participants.

Result 1: Subject levels vary with game complexity. Participants' reasoning levels are influenced by the complexity of the games, with notable variations observed across different conditions.

[4]No deviation: 5 players (all L1 players); One deviation: 5 players (all players' mode level are L1); Two deviation: 19 players (7 of them are Nash equilibrium player)

4.2 Informational Complexity

We begin by examining the primary factor influencing reasoning levels—informational complexity. This effect is most evident when comparing revealed levels in Baseline Game 1 (B1) and Baseline Game 2 (B2). The sole difference between these two games is their informational complexity. While B2 consistently exhibits higher complexity than B1, these two games represent two extreme conditions—one where the zero-payoff rate is minimized and another where it is maximized.

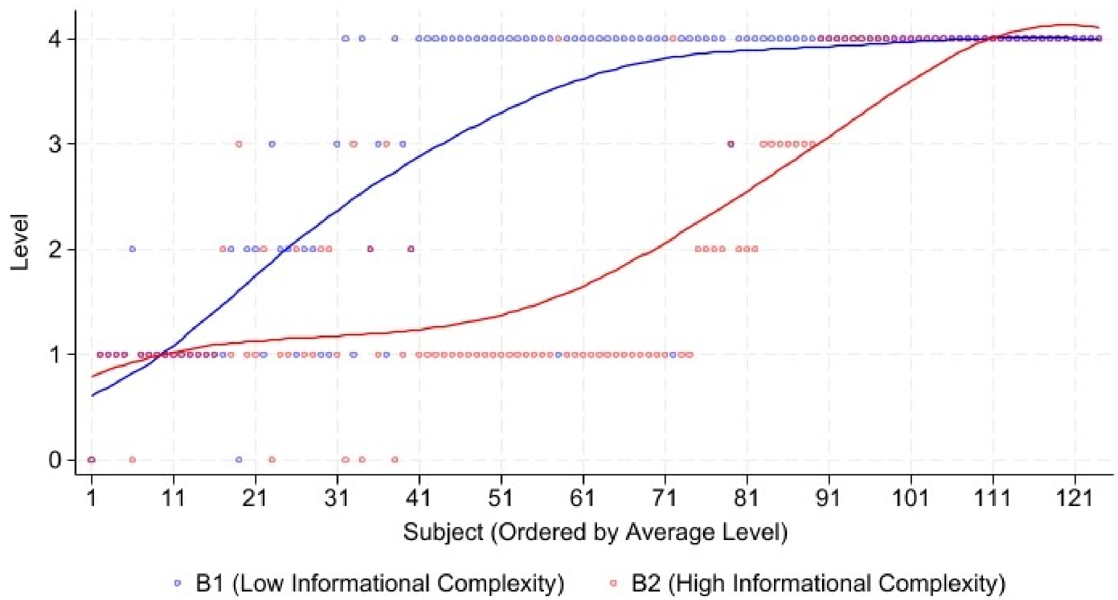

Figure 6: Level by Subject in Baseline Game 1 (B1) and Baseline Game 2 (B2) and the Nonparametric Trend of Level as a Function of Subject (LOWESS Smoothing)

In Figure.6, each dot represents an individual's level in one game, blue for Baseline Game 1, and red for Baseline Game 2. On the x-axis, individuals are ordered based on their average level across both games. The solid lines show the nonparametric trends. The gap between the two lines illustrates the treatment effect across the distribution. When the informational complexity decrease, there is a huge jump of subjects's level. 49.2% of players have displayed an increase in revealed level. And 25.8% of all players jumped from L1 to L4, which is a very huge increase. Suggesting that in the easier game, these players are capable of playing the equilibrium but can only do one rounds of iteration in the more complex game. This significant drop further reinforces Result 1, demonstrating that individuals' revealed strategic sophistication is not stable across different games. Only 42.7% of players' level remained the same, which is the complexity invariant player, and most of them are of level 1 or level 4. We seldom witness subjects remaining at the middle level.

One possible explanation is that informational complexity influences strategic sophistication by altering the minimum level of reasoning required to solve a game. Lower informational complexity reduces the number of decision paths a player must consider before making a choice, making the elimination of strictly dominated strategies more apparent. For a simple decision-making heuristic that evaluates all potential choice paths and selects the one yielding the highest payoff, B1 presents only four viable paths, as each option has only one nonzero payoff to evaluate. In contrast, in B2, where payoffs are more evenly distributed, the number of possible paths expands exponentially to 4^4, significantly increasing cognitive demands.

Consistent with this mechanism, the distribution of inferred reasoning levels in our experiment is heavily concentrated at the extremes—level-1 and level-4—with relatively few subjects classified as intermediate types. This pattern suggests that iterative reasoning may operate more like a threshold skill: once subjects understand how to perform strategic iteration, they are able to carry out multiple steps naturally, often reaching the highest level needed for the task. In other words, subjects do not appear to get "stuck" at intermediate levels; rather, once they acquire the right reasoning approach, they tend to execute it fully.

Moreover, some participants may begin with limited understanding of strategic reasoning and default to simpler heuristics (e.g., level-1 behavior). However, once they figure out the logic or recognize the value of iterated elimination, they are capable of applying it repeatedly without much difficulty. This may indicate that the skill of strategic iteration is not gradually developed but rather binary—either one has it or one does not. Alternatively, intermediate reasoning levels may represent unstable or transitional cognitive states that are passed through quickly. Finally, the task structure may not sufficiently reward intermediate levels, further encouraging a bimodal pattern in reasoning behavior.

Figure 7: Distribution of Subjects' Level across Baseline Game 1, Baseline Game 2, Half Zero Game 1, and Half Zero Game 2 (where the horizontal and vertical axes represent subjects' level, and the numbers indicate the number of players.)

While the B1–B2 comparison captures the impact of informational complexity under extreme conditions, additional insights can be drawn from the Half Zero Treatment, which introduces intermediate levels of informational complexity between B1 and B2. Figure 7 illustrates how individual reasoning levels change in games with different proportions of zero payoffs. Each dot represents one pattern of level change, and the number beside each dot represents the number of subjects displaying this pattern.

In the upper two graphs, 47.6% of observations lie below the 45-degree line, indicating that participants' reasoning levels tended to decline when moving from the game with the highest proportion of zero payoffs (B1) to the Half-Zero conditions (HZ1 and HZ2). On average, 55.4% of level-4 players in B1 exhibited lower reasoning levels in HZ1 and HZ2. This decline is comparable to the 58.3% observed when moving from B1 to B2. The lower two graphs show that as zero payoffs are removed, the distribution of reasoning levels becomes more dispersed. The decline in level-4 players is less pronounced from B1 to the Half-Zero games than from B1 to B2. For example, when comparing HZ1 and B2, 27 participants reduced their reasoning level, while 29 increased it, suggesting a nearly balanced shift. A similar trend appears between HZ2 and B2, with slightly more participants showing a decrease (33) than an increase (29).

These patterns suggest diminishing marginal effects of informational complexity. The drop in reasoning level is more substantial when complexity rises from a low baseline, while additional increases in complexity have weaker effects once the initial complexity is already high.

Another key finding is that the distribution of zero payoffs does not affect strategic sophistication. In other words, the method used to reduce informational complexity has no significant impact on reasoning levels. This is reflected in the figure, where the distributions in the left and right panels appear nearly identical.

4.3 Structural Complexity

The effect of structural complexity can be checked through the comparison of revealed reasoning levels in Baseline Game 2 with the 5×5 game and Baseline Game 1 with the 3×3 game. The following two figures summarize player behaviour across the four games, both indicating that higher structural complexity leads to increased revealed reasoning levels. This finding is particularly interesting because it contradicts our initial hypothesis that greater structural complexity would reduce participants' reasoning levels.

Two key mechanisms underlying this result can be identified from the figures. The first is a structural limitation inherent in the Ring Game, where players whose reasoning level exceeds the equilibrium level are indistinguishable through their behaviour from those at the equilibrium, as their choices coincide.

For instance, in the 3×3 game, a level-3 player can already play the equilibrium strategy. Therefore, any participant with a higher reasoning level will make the same choice, making it impossible to distinguish them. As a result, in games with lower structural complexity, the highest observed reasoning level may conflate equilibrium-level players with those possessing even greater strategic depth.

This issue is particularly pronounced in Figure 8(a), where results from Baseline Game 1 and the 3×3 game are compared. Both games feature low informational complexity, but differ in structural complexity. We find that 80 participants classified as level-4 in B2 could only be identified as level-3 in the 3×3 game. A similar pattern

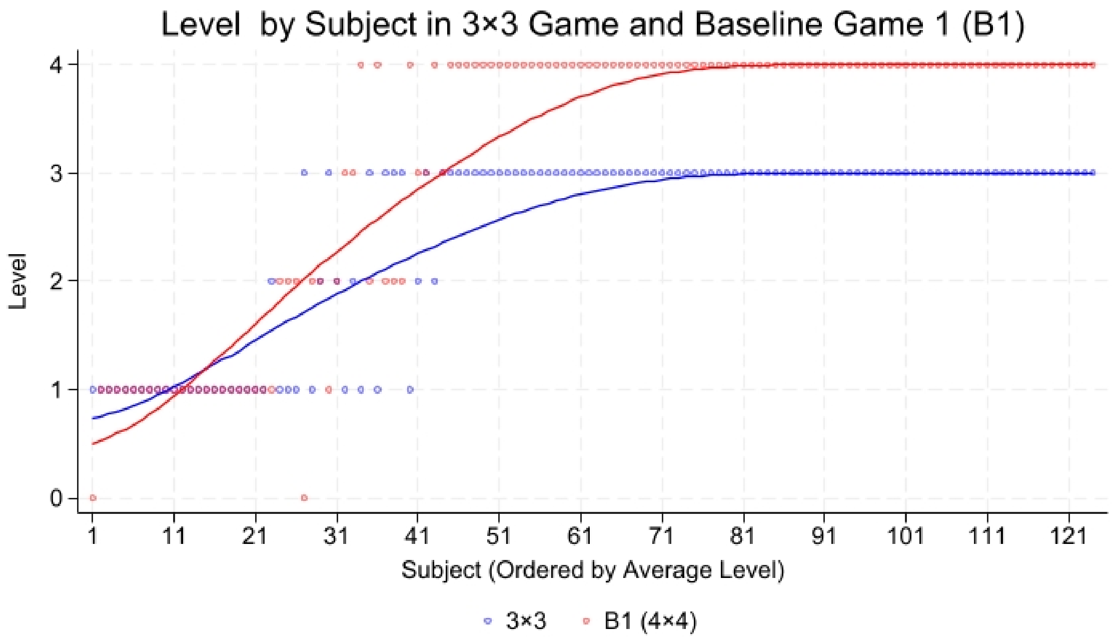

(a) Low Informational Complexity

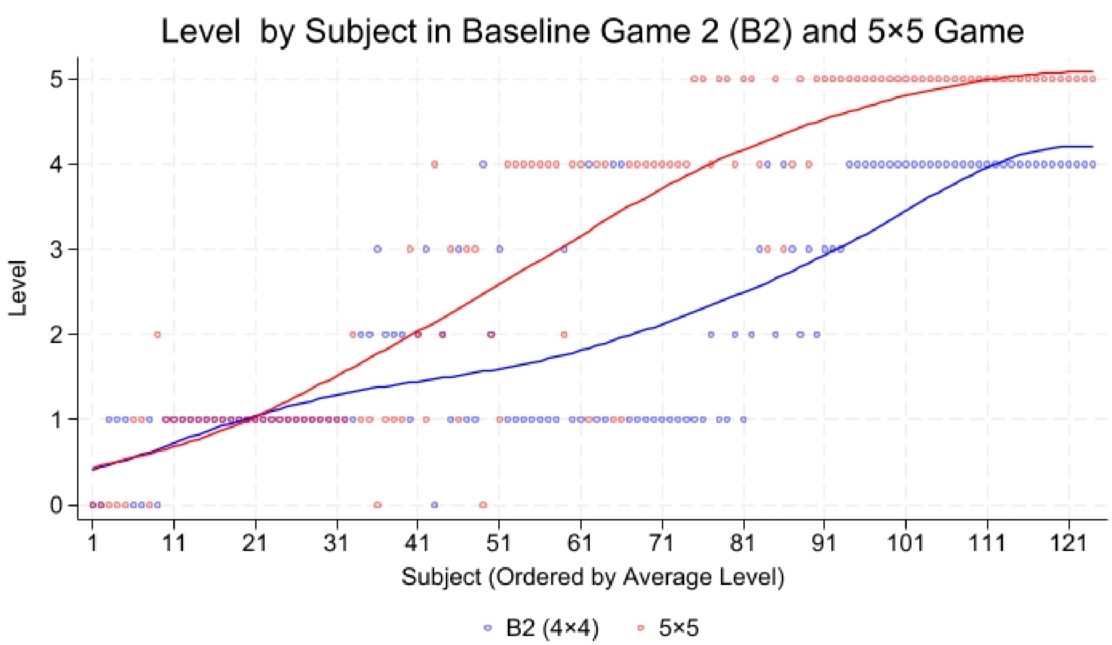

(b) High Informational Complexity

Figure 8: Strategic sophistication under varying structural complexity and the Non-parametric Trend of Level as a Function of Subject (LOWESS Smoothing).

emerges when comparing Baseline Game 2 with the 5×5 game, where computational complexity is high.

However, this reflects a limitation of the measurement framework rather than an actual decline in reasoning ability. Focusing on unrestricted players, we identify a second mechanism driving the increase in observed reasoning levels when structural complexity rises. Among the 87 participants below the equilibrium level in Baseline Game 2 (i.e., those at level 0 to 3), over half (51.7%) exhibited an increase in their revealed reasoning level. This suggests that, beyond the confounding effect of the Ring structure, the impact of increasing structural complexity is both genuine and substantial.

Two possible mechanisms may explain this effect. First, greater structural complexity may render additional information more salient, encouraging participants to engage in higher-order reasoning they might neglect in simpler environments. Second, players may rely on heuristics to simplify decision-making. When structural complexity increases, these heuristics may become ineffective, forcing participants to engage in more deliberate and iterative reasoning processes.

4.4 Statistical Analysis

To analyze the determinants of higher-order rationality, we estimate a panel data ordered probit regression where the dependent variable, observed Level (K), represents a player's level of strategic sophistication. A higher K corresponds to a greater depth of iterative reasoning. The key independent variables include Entropy, which captures the informational complexity of the game, and Structural Complexity, defined as N in an N-player, N-choice game, representing the maximum identifiable reasoning level within each game. Additionally, we include two dummies for the two treatment emphasizing certain strategy, which are labeled as Salience of Dominant Strategy (SD) and Salience of Iteration (SI). CRT measures the number of correctly answered questions in the Cognitive Reflection Test, serving as a proxy for cognitive ability. Finally, G_2 to G_4+ are dummy variables capturing prior experience in game theory, with the baseline category being participants with the least exposure to game theory.

Result 2. Higher-order reasoning declines with increases in informational complexity, but the effect is non-linear and plateaus at higher levels.

	(1)	(2)	(3)
	Level	Level	Level
Entropy	-2.638***	-2.637***	-2.637***
	(-7.57)	(-7.57)	(-7.57)
$Entropy^2$	0.534***	0.534***	0.534***
	(6.57)	(6.57)	(6.57)
N	0.714***	0.713***	0.713***
	(6.83)	(6.83)	(6.82)
SD	-0.0183	-0.0182	-0.0184
	(-0.17)	(-0.17)	(-0.17)
SI	-0.401***	-0.402***	-0.403***
	(-2.83)	(-2.84)	(-2.84)
CRT	0.160**	0.134**	0.129**
	(2.47)	(2.13)	(2.07)
G_2		0.157	0.167
		(0.84)	(0.89)
G_3		0.418**	0.415**
		(2.28)	(2.26)
G_4+		0.871***	0.842***
		(3.76)	(3.57)
age			0.0138
			(0.71)
gender			0.139
			(0.82)
reverse			-0.0555
			(-0.38)
N	1116	1116	1116

t statistics in parentheses

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

Table 4: Observed Level in all games.

Different methods of reducing complexity do not significantly influence strategic sophistication.

The regression results reveal a strong non-linear relationship between informational complexity and higher-order reasoning, which aligns with the behavioural patterns observed in the graphical analysis. The coefficient for Entropy is significantly negative (-2.638, $p < 0.001$), indicating that as the informational complexity of a game increases, subjects become significantly less likely to engage in higher levels of reasoning. The positive and significant coefficient on Entropy squared (0.534) indicates diminishing marginal effects, with the decline in reasoning levels becoming less pronounced at higher levels of complexity.

Additional regression results also confirm the graphical observation that the allocation of zero payoffs does not affect subjects' revealed reasoning levels. The difference between the two Half-Zero treatments is not statistically significant ($p = 0.872$), supporting the conclusion that how zero payoffs are distributed does not influence strategic sophistication.

Result 3: Increasing structural complexity promotes strategic sophistication.

The coefficient on Structural Complexity is positive and highly significant (0.714, $p < 0.001$), indicating that for a game of higher structural complexity, participants are more likely to engage in deeper strategic thinking. This pattern aligns with the graphical analysis and reflects the average effect of the two underlying mechanisms discussed in Section 4.3. First, the structure of the Ring Game creates an upper bound on revealed reasoning levels: equilibrium-level players and those with even higher reasoning ability make the same choices, making them indistinguishable. Second, for players whose reasoning levels are below equilibrium (i.e., unrestricted players), an actual increase in revealed level is also observed, suggesting that structural complexity indeed induces more rounds of iterations, possibly by drawing attention to additional information. This result underscores that higher-order rationality is not only an inherent trait but also shaped by the decision environment.

Result 4: The strategy-oriented interventions yielded diverging effects. The two treatments, which were designed to emphasize different strategic approaches, show contrasting effects and contradict our initial hypothesis that both interventions would promote higher-order thinking. The Dominant Strategy condition had no measurable impact, and the Iteration condition significantly reduced the level of reasoning.

In the Dominant Strategy condition, the payoff matrix containing the dominant strategy was made more prominent. We find no significant effect on participants' reasoning levels, suggesting that the change in average reasoning level can be fully explained by other features of the game, particularly informational complexity. Our results suggest that participants' failure to engage in higher-order reasoning cannot be attributed to difficulties in identifying a dominant strategy.

In the Iteration condition, the intervention was implemented by simplifying the first payoff matrix. It was intended to encourage deeper iterative thinking by making it clear that no dominant strategy exists in participants' own payoff matrix, thereby prompting them to engage in more steps of reasoning. However, the results suggest that this intended mechanism did not function as expected. Rather than facilitating more iterations, the simplified matrix may have conveyed the impression that less iterative reasoning was necessary.

Result 5: Strategic sophistication increases with game theory experience.

The regression results show a clear upward trend in reasoning levels as subjects' experience with game theory increases. Compared to the baseline group (G1), the coefficient for G2 is positive (0.157) but not statistically significant, suggesting that limited exposure does not lead to measurable gains in reasoning ability. For G3, the effect becomes significant (0.418, $p < 0.05$), indicating that moderate experience is associated with a modest improvement in strategic thinking. The strongest effect is observed for G4+, with a large and highly significant coefficient (0.871, $p < 0.001$), showing that extensive experience leads to substantial gains in higher-order reasoning. These findings support the view that strategic sophistication can be cultivated through formal training, and is not solely determined by innate ability. The monotonic increase across experience levels further highlights the role of learning in shaping complex reasoning.

We also observe that score in cognitive reflection test is positively and significantly associated with higher-order reasoning, with an estimated coefficient of 0.160 ($p = 0.014$). This result supports the bounded rationality framework, where cognitive ability acts as a constraint on reasoning depth. While informational complexity alters the upper bound of the cognitive demands required to engage in strategic reasoning, performance on the CRT serves as an indicator of an individual's inherent cognitive capacity. Given that the CRT is widely regarded as a measure of cognitive reflection and deliberative thinking, it is unsurprising that individuals who score higher on this test exhibit greater strategic sophistication.

4.5 Response Time

The revealed Lk level reflects the actual effect of game complexity on strategic reasoning. However, it is also essential to examine whether a similar effect applies to perceived complexity as experienced by subjects. As discussed in Section 3, response time and click history serve as proxies for perceived complexity.

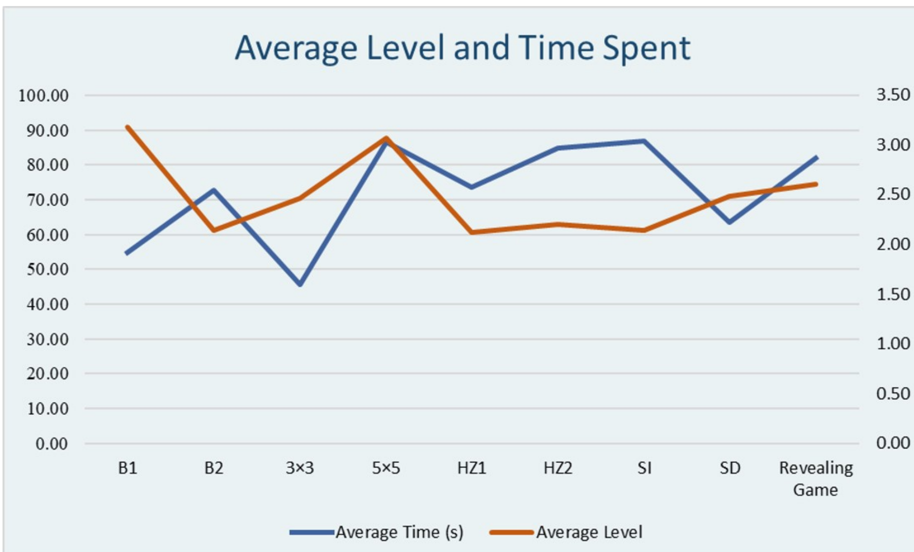

Figure 9: Average Level and Time Spent (seconds) in each game

Figure 9 presents the average observed level alongside the average time spent per game. A negative correlation emerges between time spent and Lk level, suggesting that the impact of game complexity on perceived complexity mirrors its effect on actual complexity. In simpler games, subjects spend less time while attaining higher reasoning levels, whereas in more complex games, they spend more time but exhibit lower strategic sophistication.

The linear regression results provide further insights into the determinants of decision time in strategic reasoning tasks. As reported in Table.5, the coefficient on entropy is 44.1, and on entropy squared is -7.6, indicating that there's a non-linear relationship between informational complexity and time spent. This result is consistent with prior research showing that more complex strategic environments impose greater cognitive demands, leading to longer deliberation times. Similarly, N, which represents the highest distinguishable level of reasoning within a game, has a positive and highly significant effect on decision time, suggesting that games requiring deeper strategic reasoning necessitate greater cognitive effort, thereby extending response times.

By contrast, the treatment condition SD, which enhances the salience of the dominant strategy, significantly reduces decision time. This suggests that when a dominant strategy is made more salient, participants identify and apply it more efficiently, reducing cognitive load and expediting decision-making. In contrast, SI, which emphasizes the salience of iterative reasoning, also has a significant effect (9.38), implying that making iterative reasoning more prominent requires more time to think. Additionally, CRT, which measures cognitive reflection ability, is not significantly associated with decision time (1.231), indicating that individuals with higher cognitive reflection scores do not necessarily spend more or less time on the task.

These findings present an important contrast to the results from the ordered probit regression on strategic sophistication. While the ordered probit estimates indicated that entropy negatively affected level-k reasoning, the linear regression results reveal that higher entropy significantly increases decision time. This suggests that more complex games not only hinder the ability to engage in higher-order reasoning but also prolong decision-making due to elevated cognitive demands. Meanwhile, the significant effect of SD in reducing response time reinforces the idea that emphasizing a dominant strategy improves decision-making efficiency, even if it does not necessarily induce higher levels of reasoning.

When we look at the effect of Game Theory experience on time spent, the estimated coefficients for G_2, G_3, and G_4+ are all negative, suggesting that individuals with more exposure to game theory tend to spend less time making their decisions compared to the baseline group where subject has no game theory experience(G_1). However, the effect is only statistically significant for G_3, where the coefficient of -9.09 indicates a significant reduction in decision time for participants with moderate game theory experience. In contrast, the estimates for G_2 (-2.87) and G_4p (-4.76) are not statistically significant, suggesting that for individuals with either limited or extensive experience, decision time does not differ systematically from the baseline. The absence of a significant reduction in decision time for the most experi-

	(1) Time	(2) Time	(3) Time
Entropy	44.10***	44.10***	44.10***
	(4.28)	(4.27)	(4.27)
$Entropy^2$	-7.579***	-7.579***	-7.579***
	(-3.01)	(-3.01)	(-3.00)
N	11.92***	11.92***	11.92***
	(5.45)	(5.44)	(5.44)
SD	-14.06***	-14.06***	-14.06***
	(-4.73)	(-4.73)	(-4.72)
SI	9.377*	9.377*	9.377*
	(1.86)	(1.85)	(1.85)
CRT	1.231	1.545	1.161
	(0.70)	(0.90)	(0.74)
G_2		-4.505	-2.872
		(-0.86)	(-0.55)
G_3		-9.441**	-9.093**
		(-2.16)	(-2.09)
G_4+		-7.563	-4.758
		(-1.20)	(-0.84)
age			0.745
			(1.26)
gender			-6.108
			(-1.43)
reverse			-9.519**
			(-2.51)
cons	-37.27***	-34.18**	-39.40*
	(-2.67)	(-2.46)	(-1.89)
N	1116	1116	1116

t statistics in parentheses

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

Table 5: Time Spent (seconds) in all games.

enced group suggests that greater sophistication does not necessarily equate to faster decision-making. These findings contrast with the earlier results on Lk, where greater game theory experience was associated with higher strategic sophistication. While the G_4+ group (pooling the most experienced participants) exhibited the highest reasoning levels in the probit model, their decision times in the linear regression are not significantly different from the baseline. This suggests that while highly experienced participants reach higher reasoning levels, they do not necessarily do so more quickly. Instead, the most pronounced reduction in decision time occurs for those in the G_3 category, implying that individuals with moderate experience may have reached an optimal balance between efficiency and strategic depth, allowing them to process complex games faster without sacrificing reasoning quality.

These results underscore a key distinction between strategic sophistication and cognitive efficiency. While greater experience in game theory facilitates higher levels of reasoning, its effect on decision speed follows a non-linear pattern. Moderate experience enhances efficiency, but extensive experience does not necessarily reduce deliberation time, potentially due to increased sensitivity to game complexity or a more deliberate approach to decision-making. More broadly, the findings highlight the need to differentiate between cognitive effort (as captured by decision time) and actual strategic reasoning (as reflected in level-k sophistication) when analysing behaviours in strategic environments.

4.6 Revealing Game

Although the revealing game was first designed to unravel the thinking process of subjects, the results suggest some interesting findings. Baseline Game 2 and the Revealing Game have the same features or have the same complexity, but the action of revealing the payoff still seems to affect the performed strategic reasoning. Figure 10 shows the change in levels in these two games. Although the majority (51.6%) of subjects stayed at the same level, 37% of subjects' revealed level has increased. This can be explaind by the action of revealing or the learning effect, as this treatment has been played at after playing all the other games.

Figure 11 presents the heatmap of net click history in the Revealing Game, calculated as the number of times each cell was highlighted minus the number of times it was un-highlighted. This measure allows us to track the payoff tables each

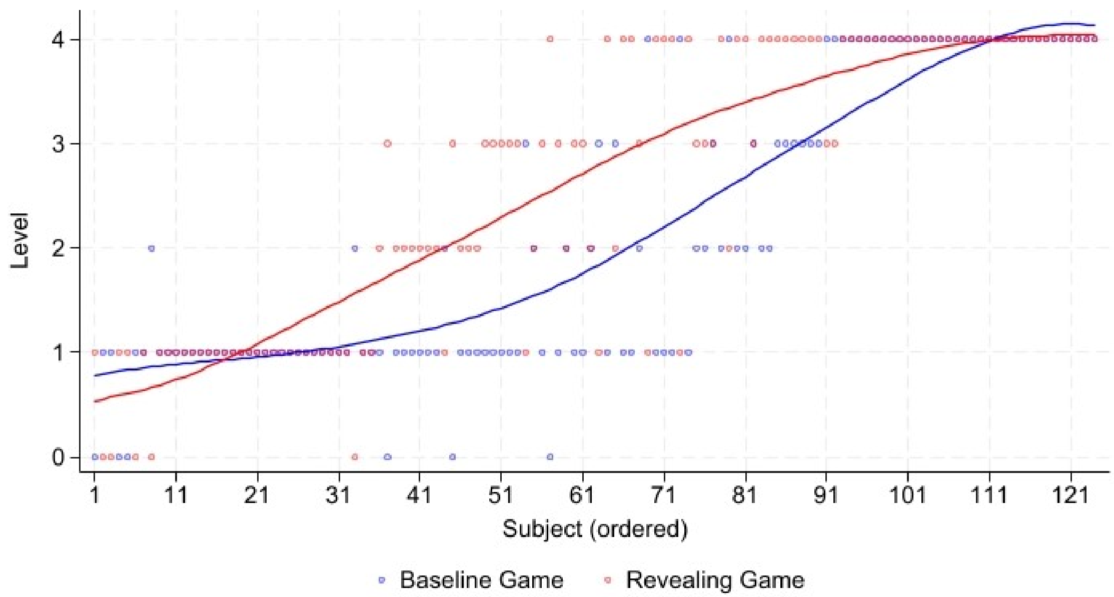

Figure 10: Level by Subject in Baseline Game 2 (B2) and the Revealing Game, with Nonparametric Trend of Level as a Function of Subject (LOWESS Smoothing)

player referenced when making decisions, providing insight into the information used in the decision-making process.

The heatmap reveals that Table 1 and Table 4 are the most frequently clicked payoff matrices, corresponding to the decision environments of Player 1 and Player 4, respectively. Nearly all cells in these two tables exhibit high click frequencies, with Table 4 registering the highest overall number of interactions. In contrast, in Table 2 and Table 3, equilibrium choices receive substantially more clicks than non-equilibrium cells, which are rarely selected. This pattern is consistent with the observed distribution of reasoning types in the Revealing Game, where L4 and L1 players are predominant, while L2 and L3 players are relatively scarce. Since L4 players engage in backward induction, they also tend to highlight equilibrium choices in Table 2 and Table 3.

A particularly notable finding is that nearly every participant interacted with all cells in Table 4, suggesting that even L1 players systematically considered this table when making decisions. This result challenges standard theoretical predictions, which posit that ability-bounded players should focus only on their own payoff table and disregard those of higher-order players. For instance, L1 players are theoretically

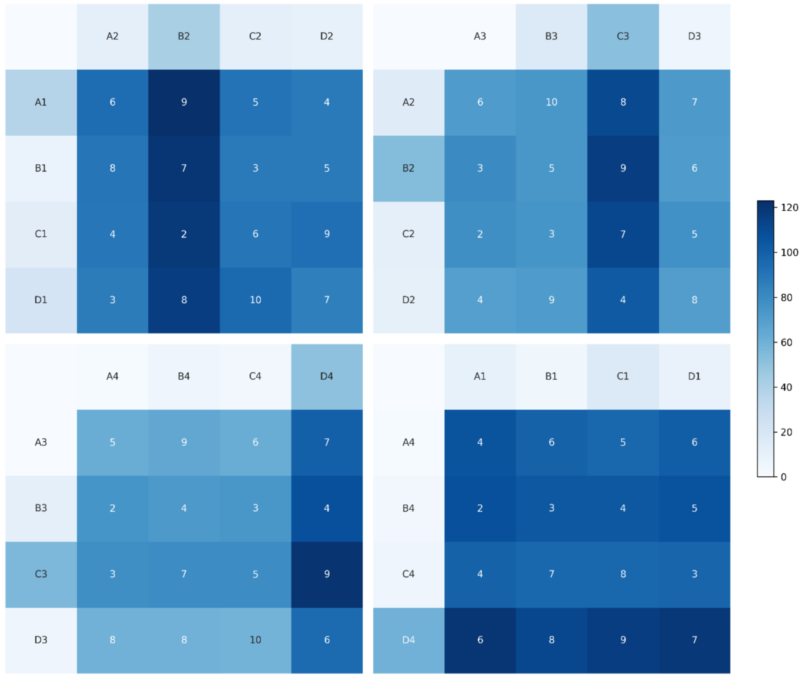

Figure 11: Heatmap of Net Click history

The four tables correspond to the payoff matrices for Player 1 to Player 4 in the Revealing game, where the numbers are hidden at first and subjects need to click on the cell to view the number, from left to right, the top two graphs correspond to the payoff matrices for Player 1 and 2, and the bottom two are for Player 3 and 4.

expected to attend only Table 1 without incorporating information from Table 2. However, the highlighted data contradicts this assumption, indicating that nearly all players, regardless of their reasoning level, engaged with Table 4.

One potential explanation for this pattern is that in prior rounds, Player 4 consistently had a dominant strategy, prompting participants to examine Table 4 in its entirety. However, engaging in higher-order iterative reasoning may still exceed the cognitive limits of L1 players. Even after highlighting Table 4, they may be unable to successfully perform backward induction to infer the equilibrium strategy. This suggests that while lower-level players recognize the importance of Player 4's decisions, their ability to integrate this information into a fully strategic response remains limited.

5 Conclusion

This paper examines how strategic sophistication in human subjects is influenced by game complexity, a dimension not fully captured by traditional bounded rationality theories. By focusing on ability-bounded players and employing computer-controlled opponents to manipulate beliefs, we establish that three key dimensions of game complexity, entropy, the upper bound of iterative reasoning, and the salience of a strategy, play a crucial role in shaping strategic behavior. Our findings provide new empirical evidence on the determinants of observed strategic sophistication, offering a potential explanation for cross-game variation in revealed reasoning levels and opening new avenues for unifying results across different classes of games in the literature.

Beyond revealed strategic sophistication, our analysis of perceived complexity provides novel insights into the cognitive processes underlying strategic decision-making. We show that decision time and revealed reasoning levels are not linearly correlated: higher-level players tend to reach decisions more quickly, suggesting that they rely on different heuristics or mental models rather than merely possessing superior cognitive ability. This challenges the conventional view that higher-order reasoning is solely a function of cognitive capacity and highlights the role of strategic efficiency in decision-making. Specifically, we show that game complexity serves as a constraint on individuals' ability to engage in higher-order reasoning: if a game's required level of iterative reasoning exceeds a subject's cognitive capacity, their revealed sophistication is bounded. Conversely, when the required level is below their capacity, their full reasoning ability is more likely to be expressed.

From a methodological perspective, our framework for identifying higher-order rationality proves to be both robust and parsimonious, simplifying the inference process while maintaining identification power. Moreover, our belief-controlled experimental design yields systematically higher reasoning levels compared to standard human-interaction settings, reinforcing the notion that belief formation is a critical determinant of strategic sophistication. This aligns with existing literature while providing direct experimental evidence on the role of belief manipulation in shaping behavior.

Our findings also raise important questions for future research. In our experimental setting, players' observed reasoning levels evolve dynamically with game complexity, suggesting that belief formation is a more intricate and fragile process than previously assumed. If individuals struggle to form accurate beliefs, their ability to best respond optimally may be systematically constrained. Future studies should further investigate the interplay between complexity, belief formation, and strategic adaptation, particularly in environments where beliefs must be updated over repeated interactions. Understanding these mechanisms is essential for refining models of strategic behavior and enhancing the predictive power of behavioral game theory.

Acknowledgments. I would like to thank Ed Hopkins, Paweł Gola, Tatiana Kornienko, Miguel Costa-Gomes, Yaoyao Xu, Darija Barak-Halatova, Jinge Liu and audiences at SGPE conference 2024, European Meeting of the Economic Science Association in Helsinki, 1st Summer School in Experimental and Behavioral Economics, Workshop 'Computational Complexity of Decision Making' in Melbourne, and 2025 Behavioral, Experimental and Theoretical Economics Conference for helpful discussions and comments. This project is funded by the Scottish Economics Society and Edinburgh Future Institution-The University of Edinburgh. The reported experiment is a part of a project that was pre-registered on the AEA RCT registry, https://doi.org/10.1257/rct.13275-1.0

References

Abeler, J. and Jäger, S. (2015). Complex tax incentives. *American Economic Journal: Economic Policy*, 7(3):1–28.

Agranov, M., Potamites, E., Schotter, A., and Tergiman, C. (2012). Beliefs and endogenous cognitive levels: An experimental study. *Games and Economic Behavior*, 75(2):449–463.

Alaoui, L., Janezic, K. A., and Penta, A. (2020). Reasoning about others' reasoning. *Journal of Economic Theory*, 189:105091.

Alaoui, L. and Penta, A. (2015). Endogenous depth of reasoning. *The Review of Economic Studies*, 83(4):1297–1333.

Alós-Ferrer, C. and Buckenmaier, J. (2020). Cognitive sophistication and deliberation times. *Experimental Economics*, 24.

Breunig, C., Huck, S., Schmidt, T., and Weizsäcker, G. (2021). The standard portfolio choice problem in germany. *The Economic Journal*, 131(638):2413–2446.

Camerer, C. (2003). *Behavioral Game Theory Experiment in Strategic Interaction*, volume 32.

Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A cognitive hierarchy model of games*. *The Quarterly Journal of Economics*, 119(3):861–898.

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

Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235.

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

Grabiszewski, K. and Horenstein, A. (2022). Measuring tree complexity with response times. *Journal of Behavioral and Experimental Economics*, 98:101876.

Huck, S. and Weizsäcker, G. (1999). Risk, complexity, and deviations from expected-value maximization: Results of a lottery choice experiment. *Journal of Economic Psychology*, 20(6):699–715.

Hyndman, K., Terracol, A., and Vaksmann, J. (2022). Beliefs and (in)stability in normal-form games. *Experimental Economics*, 25(4):1146 – 1172.

Jin, Y. (2021). Does level-k behavior imply level-k thinking? *Experimental Economics*.

Kneeland, T. (2015). Identifying higher-order rationality. *Econometrica*, 83(5):2065–2079.

Marchiori, D., Di Guida, S., and Polonio, L. (2021). Plasticity of strategic sophistication in interactive decision-making. *Journal of Economic Theory*, 196:105291.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5):1313–1326.

Oprea, R. (2020). What makes a rule complex? *American Economic Review*, 110(12):3913–51.

Xu, Y. (2022). Can iterative reasoning be taught in the lab? *Working Paper*.