Complexity and Higher Order Rationality:an Experimental Study*

Jiaying Li†

February 28, 2025

Abstract

We experimentally show that game complexity systematically influences individuals' strategic sophistication within a refined Ring Game framework. Our novel identification strategy for higher-order rationality enables us to examine three dimensions of complexity—Informational Complexity, Structural Complexity, and the Salience of Strategy—within a unified experimental setting. We are able to identify the ability bounds of players by controlling for subjects' beliefs about others, requiring them to interact only with a computerized opponent. Our findings reveal that all three complexity dimensions significantly impact observed Level-k behaviour, albeit in distinct ways. Participants tend to exhibit higher levels of strategic sophistication when payoff matrix entropy is lower, when the Nash equilibrium requires more iterative reasoning, and when a salient dominant strategy is present. Moreover, prior experience with strategic games plays a crucial role in individuals' ability to engage in iterative reasoning. These results highlight that strategic sophistication is neither an inherent nor a fixed trait but is instead shaped by the complexity of the decision environment and prior experience.

Keywords: Behavioural Game Theory, Higher Order Rationality, Level-K, Bounded Rationality

JEL Codes: C70, C91

*This project is funded by the Scottish Economics Society and Edinburgh Future Institution. 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

†University of Edinburgh

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.

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.

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

[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.

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. 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 Lk-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 Lk-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. This strategy introduces a robot player approach to distinguish belief-bounded from ability-bounded players. 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 classic framework of the Ring Game introduced by Kneeland (2015) to investigate subjects' higher-order rationality, introducing several key innovations to address limitations in the original design. In Kneeland's 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.

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. For example, in a four-player game (N=4), each player is given four choices (N=4), enabling complete separation of player decisions through the design of payoff matrices. 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. Unlike traditional frameworks, this revised design allows the identification of reasoning levels within a single round, significantly reducing the number of rounds required for analysis.

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, where 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.

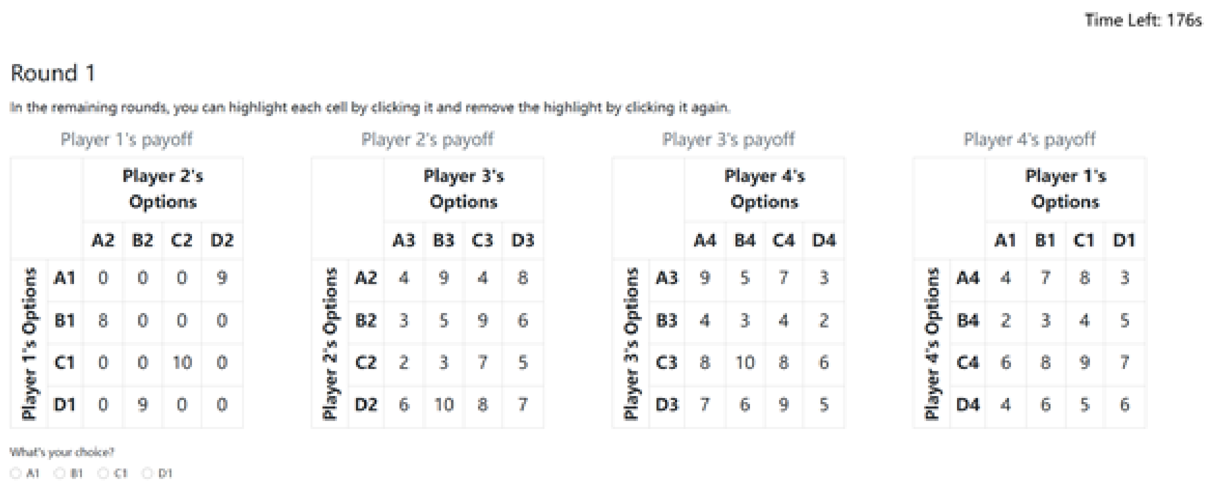

Figure 1: Screenshot of one of the Ring Games played in the main section

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. Rather than imposing assumptions on subjects, we embed the necessary identification mechanism directly within the game's payoff matrices. The specially designed N-player, N-choice Ring Games offer several key advantages. First, the number of choices available to each player, N, is equal to the total number of players, ensuring that the highest identifiable level of reasoning is LN. In theory, this allows us to separate L1 through LN subjects by designing games where their choices systematically diverge based on their iterative reasoning ability. Second, the structure of the Ring Game ensures that boundedly rational Lk players do not engage in an additional round of reasoning beyond their cognitive limit. This enables us to construct payoff matrices that guide subjects of different Lk levels toward distinct choices.

Finally, in our dominance-solvable Ring Game, only one player has a dominant strategy, and the unique Nash equilibrium is derived through backward induction. Critically, the reasoning steps required to reach equilibrium vary across players, depending on their strategic role in the game. Consequently, observing a player's choice allows us to infer their level-k reasoning ability with precision.

Using the Ring Game illustrated in Figure 1 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

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

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.

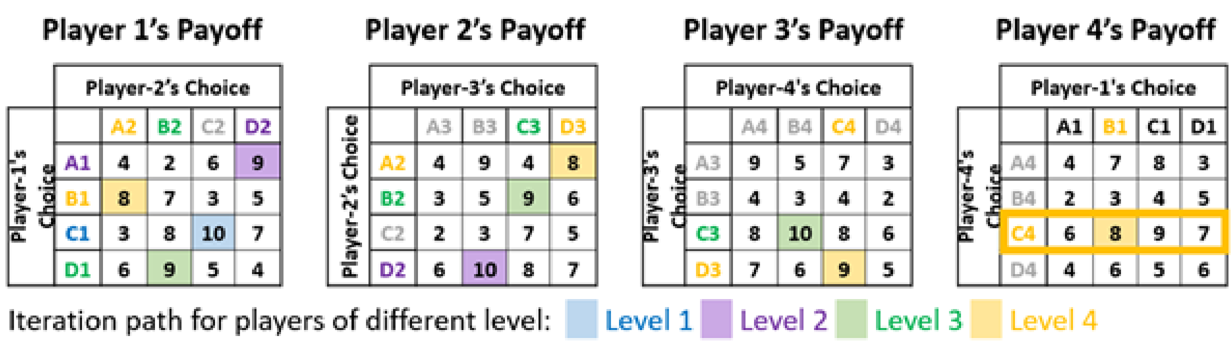

Figure 3: Iteration Path for players of different Levelwhere C4 for Player 4 is the only dominant strategy among the 4 matrices

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 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.

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

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

Moreover, since we don't make any assumptions about the behaviour of L0 players and assume they are choosing randomly, only subjects who exhaust their time without selecting an option are classified as L0 players. 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,A1)	(A1,A1)	(D1,B1)	(B1,B1)
	(C1,B1)	(B1,A1)	(B1,D1)	
	(A1,C1)	(D1,A1)	(D1,D1)	
	(B1,C1)	(A1,B1)		
	(C1,C1)	(A1,D1)		
	(D1,C1)			
	(C1,D1)			

Table 1: Summary of LK Player's Choice Profile

3.3 Treatment

Building upon our identification strategy, this study examines how the complexity of games influences subjects revealed higher-order rationality. The key idea 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.

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

This study further investigate three key faces of complexity, namely the Structural Complexity, Informational Complexity and the Salience of Strategy.

Structural Complexity

Structural complexity in strategic interactions refers to the extent to which a game's structure—defined by the number of players and available choices—affects decision-making. In our setting, the number of players always equals the number of available choices for identification purposes, meaning that structural complexity is captured by N, where each game consists of N players and N choices. Instead of fixing N=4, we introduce additional variation by incorporating games with N=3 and N=5.

Structural complexity is important because it exponentially expands the strategy space. In a 3×3 game, there are only 27 possible choice combinations, whereas 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 must engage in deeper levels of iterative reasoning. This may push some players beyond their cognitive constraints, leading to bounded rationality and reliance on heuristic-based decision-making. Conversely, in less complex strategic environments, players may find it easier to infer others' intentions and converge more closely to equilibrium behaviour. Understanding the role of structural complexity is therefore essential for explaining when and why individuals exhibit strategic sophistication or systematic biases in interactive decision-making.

Hypothesis 2: Structural complexity decreases the revealed level of subjects.

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 the mechanical effect.

Informational Complexity

Informational complexity refers to the level of uncertainty and dispersion in the strategic environment. While structural complexity alters the strategy space by expanding the number of possible combinations, informational complexity affects 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 seek to maximize their payoffs, they will naturally avoid selecting options that contain a zero payoff and will assume that others do the same. As a result, certain choice combinations become effectively eliminated from the strategy space, reducing the uncertainty of the game.

To control for informational complexity, we first define two Baseline Games (B1 and B2)—both 4-player, 4-choice games. Baseline Game 2 (B2) represents the highest informational complexity within the 4-player, 4-choice games, as its payoff matrix contains no zero-payoff entries. Baseline Game 1 (B1), by contrast, replaces as many payoffs as possible with zero, dramatically reducing the 256 possible choice combinations to 36.

We further manipulate informational complexity by increasing the proportion of zero-payoff cells to approximately 50%. To ensure the identification strategy remains valid, in some cases, the exact zero-payoff rate deviates slightly from this threshold. Additionally, we introduce two variations in the allocation of zero-payoff cells:

Half-Zero Game 1 (HZ1): Zero-payoff cells are distributed evenly across all four payoff matrices.

Half-Zero Game 2 (HZ2): Zero-payoff cells are concentrated in the middle two matrices, combining elements from both B1 and B2.

These variations allow us to assess how different distributions of informational complexity influence strategic reasoning.

Hypothesis 3: Computational complexity decreases the revealed level of subjects, and only the general computational complexity matters, different methods of allocating the zero-payoff do not affect the subject's higher-order rationality.

Salience of Dominant Strategy and Iteration This treatment examines how the salience of dominant strategies and the salience of iterative reasoning affect decision-making. In these games, subjects play against opponents using one payoff matrix from B1 and three from B2.

The salience of Iteration Game (SI): Player 1's payoff matrix is taken from B1, making iterative reasoning the optimal response. Since Player 1's best response is relatively clear, the key challenge is to infer the opponent's strategy, encouraging deeper iterative reasoning.

The salience of Dominant Strategy Game (SD): Player 4's payoff matrix is taken from B1, making Player 4's dominant strategy more salient. This, in turn, reduces uncertainty regarding the opponent's behaviour, as subjects can more easily predict Player 4's actions.

By contrasting these two conditions, we can evaluate how players adjust their reasoning when the need for iteration is explicit versus when a dominant strategy is readily identifiable.

Hypothesis 4: Salience of Dominant Strategy or Iteration increase the revealed level of subjects.

Revealing Game

Finally, we introduce a Revealing Game to further examine how subjects process and utilize available information. In this treatment, payoffs are initially hidden, requiring subjects to click on each cell to reveal the values of the payoff. This design allows us to track the decision-making process in real-time, capturing how subjects explore the strategy space and weigh different options.

Figure 6 presents the detailed payoff matrices used in our experiment.

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.

Baseline 1 (B1)

	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
B1	8	0	0	0	B2	0	0	9	0	B3	0	0	0	0	B4	0	0	0	0
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

Baseline 2 (B2)

A1	4	2	6	9	A2	4	9	4	8	A3	9	5	7	3	A4	4	7	8	3
B1	8	7	3	5	B2	3	5	9	6	B3	4	3	4	2	B4	2	3	4	5
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

Half Zero 1 (HZ1)

A1	0	0	0	9	A2	0	9	0	8	A3	9	0	7	0	A4	0	7	8	0
B1	8	7	0	0	B2	0	0	9	0	B3	0	0	0	0	B4	0	0	0	5
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

Half Zero 2 (HZ2)

A1	4	2	6	9	A2	0	0	0	8	A3	9	0	0	0	A4	4	7	8	3
B1	8	7	3	5	B2	0	0	9	0	B3	0	0	0	0	B4	2	3	4	5
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

Salience of Iteration (SI)

A1	0	0	0	9	A2	4	9	4	8	A3	9	5	7	3	A4	4	7	8	3
B1	8	0	0	0	B2	3	5	9	6	B3	4	3	4	2	B4	2	3	4	5
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 Dominant Strategy (SD)

A1	4	2	6	9	A2	4	9	4	8	A3	9	5	7	3	A4	0	0	0	0
B1	8	7	3	5	B2	3	5	9	6	B3	4	3	4	2	B4	0	0	0	0
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	0	0	0	0

3×3

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

5×5

	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
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

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

$$H(x) = -\sum_i P(x_i) \log_2 P(x_i) \quad (1)$$

Where x_i is the numbers 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.

Table 1 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.

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

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

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

The average identified level across all games in our study is 2.49, closely aligning with the 2.31 average level reported in Kneeland's experiment. Figure 7 presents the distribution of individual-level average reasoning levels and their standard deviations, where each dot represents an individual participant. The distribution of average reasoning levels is fairly uniform across the range from 1 to 4, indicating substantial heterogeneity in strategic sophistication. Additionally, most individuals exhibit a standard deviation above 1, suggesting that the observed sophistication is highly sensitive to game complexity.

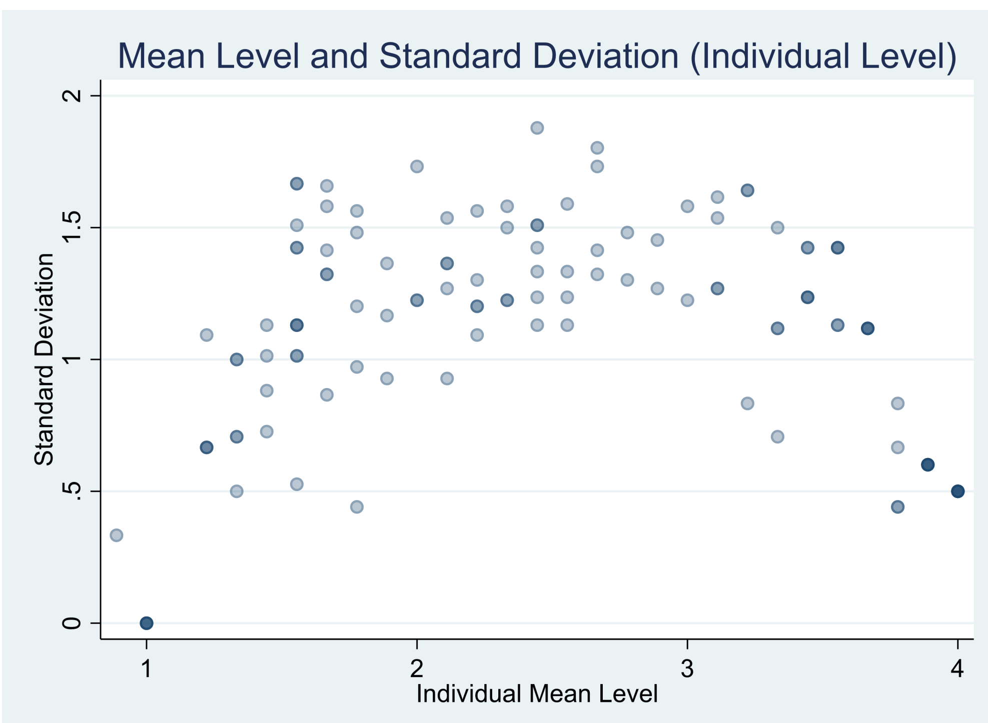

Figure 6: Individual Level and Standard Deviation

These findings strongly support our central argument that players' reasoning levels are inherently unstable 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.

We show this result further from the game-level data. First, let's look at the game-level statistics. Table 2 summarizes the identified levels in each game. We observe considerable variation in the average reasoning levels across different games. 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.

The varying complexities of the games affect not only the overall reasoning levels achieved by players but also the distribution of levels across games. Table 2 illustrates the distribution of subjects' reasoning levels in different games.

We observe two distinct patterns of strategic sophistication across different games. The first group consists of games in which equilibrium players are 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 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, in-

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

cluding 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.

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[2]. However, 76.6% of participants deviated in at least three games, and the average mistake rate was 34.6% (considering only the

[2]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)

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.

4.2 Informational Complexity

Result 2: Higher-order rationality decreases as informational complexity increases, while different methods of reducing informational complexity unchanged the effect on strategic sophistication.

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 key difference between these two games lies solely in their informational complexity: regardless of whether we measure it using the zero-payoff rate or the entropy of the payoff matrix, B2 consistently exhibits higher complexity than B1. In contrast, other aspects of complexity remain identical across the two games. B1 and B2 thus represent two extreme conditions—one where the zero-payoff rate is minimized and another where it is maximized.

A striking pattern emerges from this comparison: nearly half of the participants (49.2%) fall below the 45-degree line, indicating that the majority exhibited lower reasoning levels in the more complex B2 compared to the simpler B1. This significant drop further reinforces Result 1, demonstrating that individuals' revealed strategic sophistication is not stable across different games. Turning to specific changes in reasoning levels, the most common pattern was maintaining Level 4 reasoning in both games, observed in 35 out of 124 subjects. However, the second-largest group (32 subjects) experienced a drastic decline, dropping from Level 4 in B1 to Level 1 in B2.

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

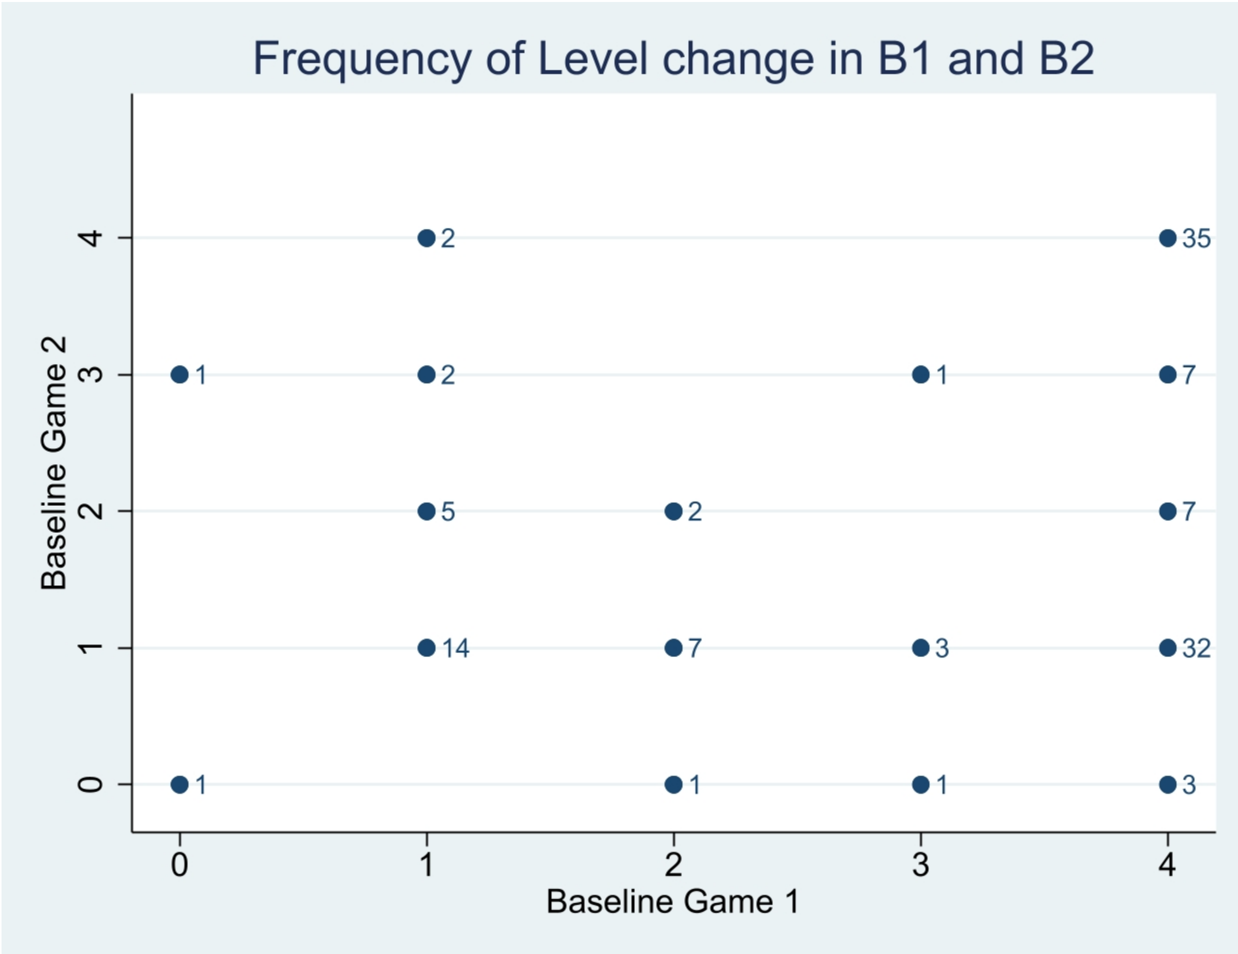

Figure 7: Inidivdual observed Lk in Game B1 and B2

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.

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. The following graphs illustrate how individual reasoning levels change with varying proportions of zero payoffs. The vertical axis represents the game with lower informational complexity. A clear trend emerges: as informational complexity increases, participants' reasoning levels tend to decline. This further supports the argument that game complexity systematically constrains higher-order rationality.

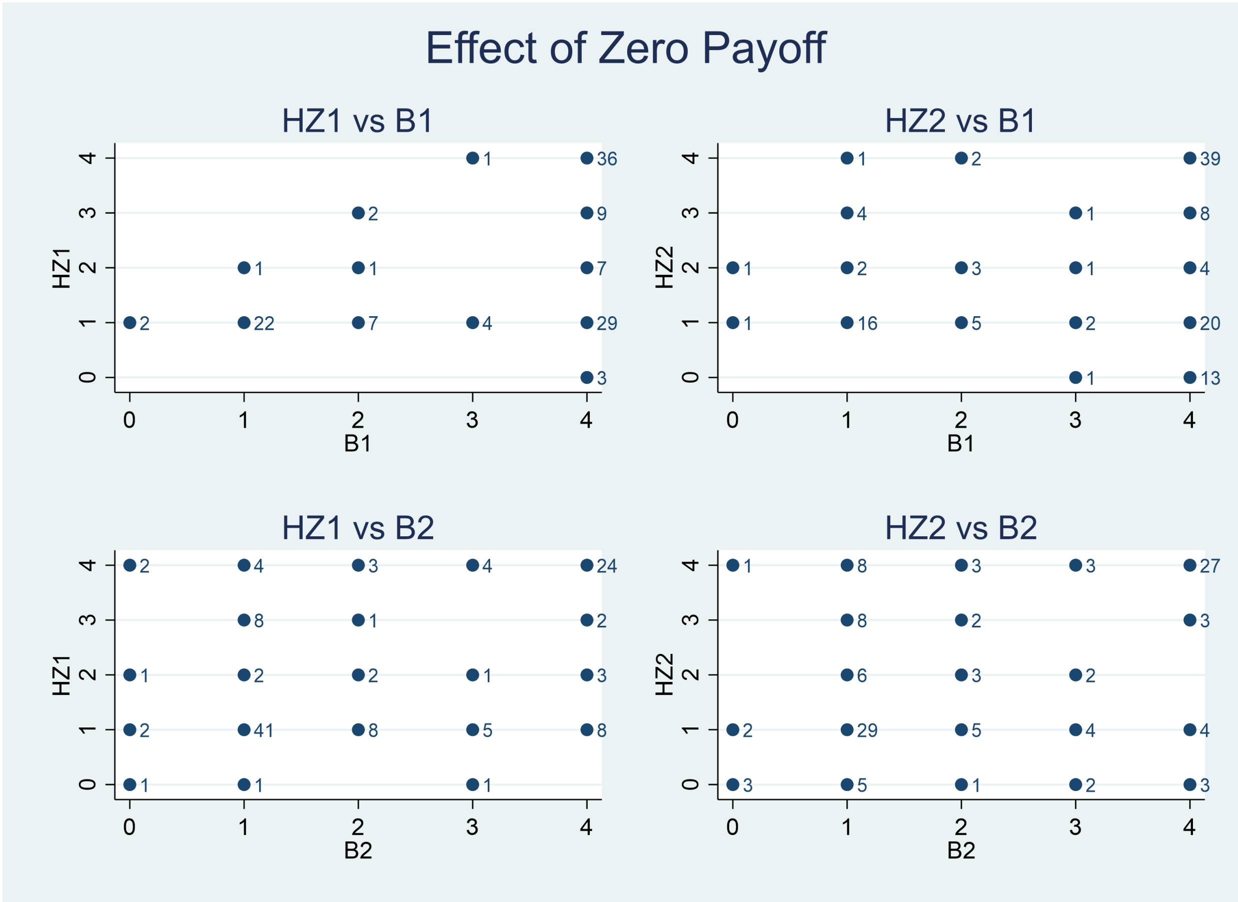

Figure 8: Individual Observed Lk in B1, B2, HZ1 and HZ2

In the upper pair of graphs, 47.58% of the data points fall below the 45-degree line, indicating that participants' reasoning levels generally declined when transitioning from the game with the highest proportion of zero payoffs (B1) to the Half-Zero games (HZ1 and HZ2). This drop is substantial—on average, 55.35% of Level 4 players in B1 exhibited lower reasoning levels in HZ1 and HZ2. Compared to the 58.3% drop in Level 4 players when moving from B1 to B2, the decline is of similar magnitude. However, as the proportion of zero payoffs continues to decrease, the effect on reasoning levels weakens. Specifically, the decline from the Half-Zero games back to B1 is much smaller, indicating a diminishing effect of informational complexity on revealed reasoning levels.

The lower pair of graphs further illustrate that as zero payoffs are eliminated, the distribution of reasoning levels becomes more variable. The decline in Level 4 players is less severe in the transition from B1 to the Half-Zero games compared to the transition from B1 to B2. For instance, when comparing HZ1 and B2, 27 participants experienced a decrease in reasoning level, while 29 participants saw an increase, suggesting a nearly balanced effect. A similar pattern is observed when comparing HZ2 and B2, though slightly more participants displayed a decrease (33) than an increase (29). These results suggest that reducing zero payoffs has a stronger impact when the initial proportion of zero payoffs is already low. However, once the proportion reaches a high level, the effect diminishes, reflecting diminishing returns on reasoning level adjustments.

Another key finding from this treatment is that the way in which zero payoffs are distributed does not affect strategic sophistication. In other words, how we lower informational complexity has no significant impact on the observed reasoning levels. This is evident from the figure above, where the distributions in the left-hand graphs closely resemble those on the right. Additionally, regression results confirm that the difference between HZ1 and HZ2 is not statistically significant ($p = 0.872$), reinforcing the conclusion that the allocation of zero payoffs does not influence strategic sophistication.

4.3 Structural Complexity

Result 3: Increasing structural complexity fosters strategic sophistication, and this effect intensifies as informational complexity rises.

This result is evident from the comparison of revealed reasoning levels in B2 with a 5×5 structure and B1 with a 3×3 structure. The following two figures summarize player behavior 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.

We identify two main factors driving this result. First is the Embedded Restriction of the Ring Game. For higher-level participants, it is natural to reveal a higher level of reasoning in more structurally complex games due to the inherent properties of the Ring Game. However, this is a limitation of the framework rather than an actual change in reasoning ability. In the Ring Game, any player whose reasoning level exceeds the equilibrium level is indistinguishable from an equilibrium-level player based on their choices. For example, in the 3×3 game, the equilibrium strategy can be reached by a Level 3 player. Therefore, any participant with a reasoning level higher than Level 3 will make the same decision as a Level 3 player, making it impossible to differentiate between them. As a result, in games with lower structural complexity, the highest

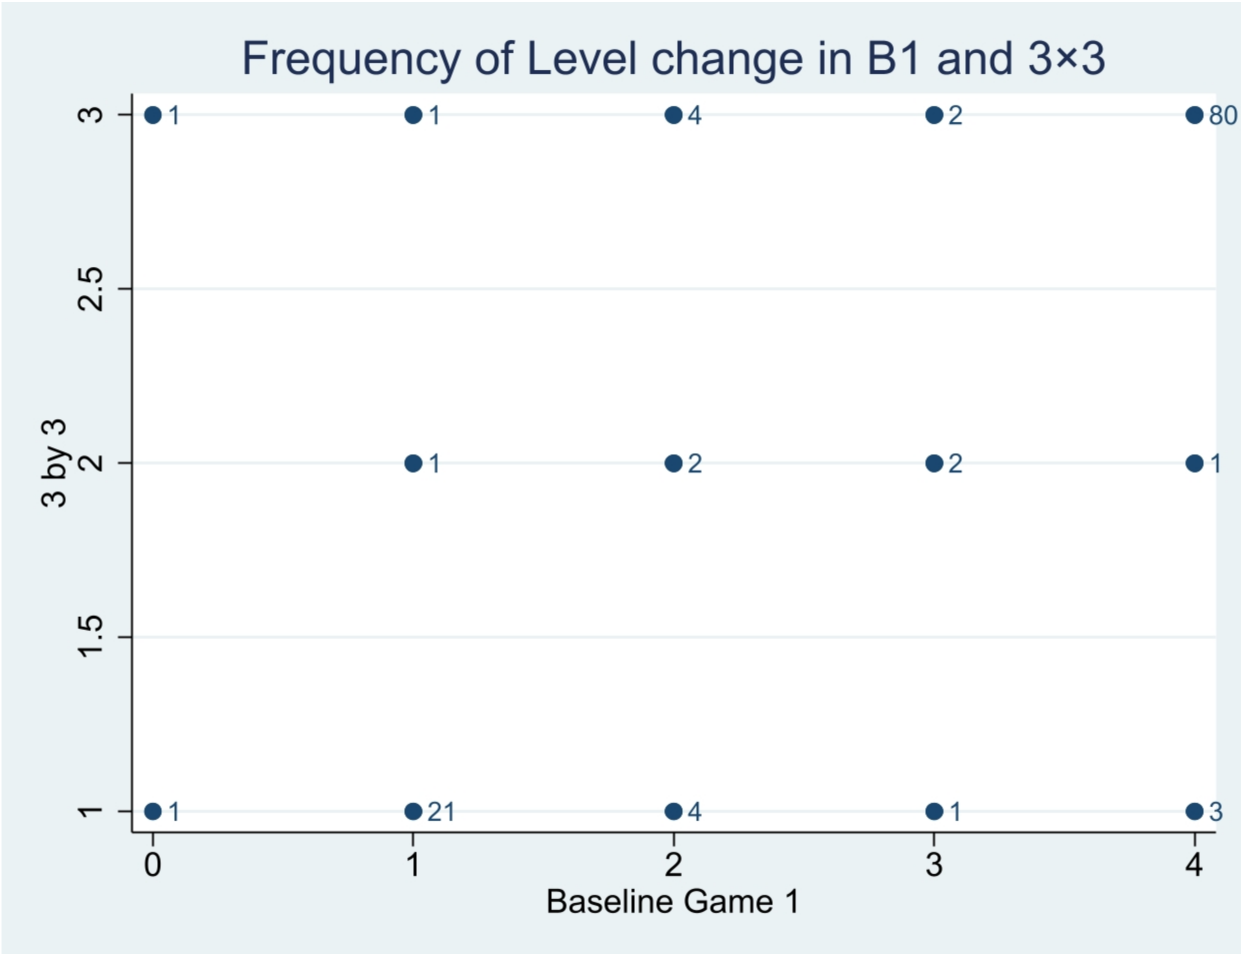

Figure 9: Individual Observed Lk in B1 and 3×3 Game

observed reasoning level is a mixture of equilibrium-level players and those with even higher actual reasoning levels.

This pattern is particularly pronounced in the comparison between B2 (Baseline Game 2) and the 3×3 game, where computational complexity is low. We find that 80 subjects were classified as Level 4 in B2 but could only be classified as Level 3 in the 3×3 game due to this limitation. A similar pattern emerges when comparing B2 with the 5×5 game, where computational complexity is high. In this case, 25% of participants were constrained by the game's structure—31 subjects were actually Level 5 or above but could only be identified as Level 4 in the other game. Thus, informational complexity amplifies the effect of structural complexity: when informational complexity is higher, the impact of increasing structural complexity becomes more pronounced.

Beyond the limitations of the Ring Game, we identify a second driving force behind the increase

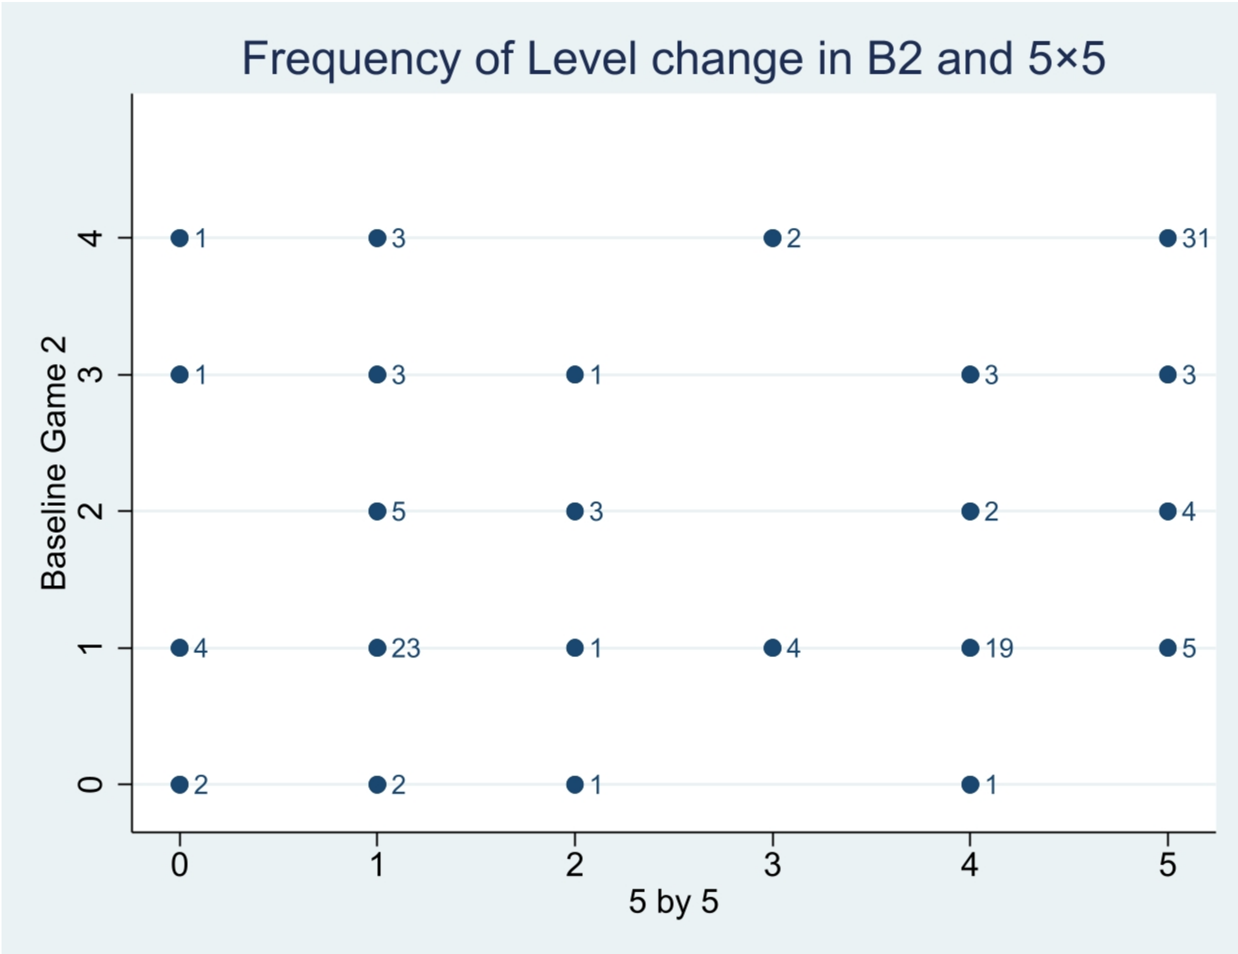

Figure 10: Individual Observed Lk in B2 and 5×5 Game

in revealed reasoning levels when structural complexity increases. While we cannot separately identify players who were genuinely displaying Level 4 behavior in B2 but moved to Level 5 in the 5×5 game, we can analyze players who had not yet reached the ability limit in the 4×4 game. Among these 87 players, more than half (51.7%) displayed an increase in their revealed reasoning level.

This suggests that beyond the structural limitations of the Ring Game, the true effect of increasing structural complexity is still significant. There are two possible explanations for this: Increased structural complexity makes additional information more salient, prompting participants to consider higher-order reasoning that they may not have engaged with in simpler games. Players often rely on heuristics when making decisions. The introduction of a more complex structure may cause these heuristics to fail, forcing participants to engage in deeper iterative reasoning.

4.4 Salience

The results indicate that the Salience of Dominant Strategy treatment exhibits a higher average reasoning level, suggesting that these two treatments have distinct effects on subjects' strategic sophistication. This finding is surprising, as it contradicts our initial hypothesis. We initially predicted that both interventions—Salience of Iteration and Salience of Dominant Strategy—would similarly enhance higher-order rationality. In the Salience of Iteration treatment, the intervention operates by restricting the information available in a player's own payoff matrix, thereby encouraging subjects to consider their opponents' payoff structures more carefully and engage in deeper iterative reasoning. However, the results suggest the opposite: rather than increasing attention to their opponents' payoffs, subjects in this treatment appear to focus less on them.

In contrast, the findings from the Salience of Dominant Strategy treatment align with our expectations. When the game's only dominant strategy is made salient, a greater proportion of subjects exhibit higher-level reasoning consistent with level-k behavior. This suggests that making the optimal choice more explicitly recognizable facilitates higher-order strategic thinking, while limiting available information may not necessarily induce deeper reasoning.

The impact of the Salience of Iteration (SI) varies between higher-level and lower-level players. For lower-level players, such as L1 players in B2, the SI treatment significantly increases their revealed reasoning level (44.6%). This suggests that for lower-level players, who perceive their own payoff structure as relatively certain—where any uncertainty stems solely from their opponent's choices—best responding becomes more straightforward, provided the opponent's choice is known. In such cases, these players may realize that they cannot rely exclusively on their own payoff matrix to make decisions, prompting them to engage in additional iterations of reasoning and, consequently, exhibit higher levels of strategic sophistication.

In contrast, for higher-level players, such as L4 players in B2, 51.35% (19/37) experience a decrease in their revealed reasoning level. A possible explanation is that when confronted with simpler payoff matrices, these players may perceive further iterative reasoning as unnecessary. Instead, they may determine their optimal choice directly from their own payoff matrix, without engaging in additional rounds of strategic reasoning. Since these opposing effects occur simultaneously—an increase in level for lower-level players and a decrease for higher-level players—the overall average

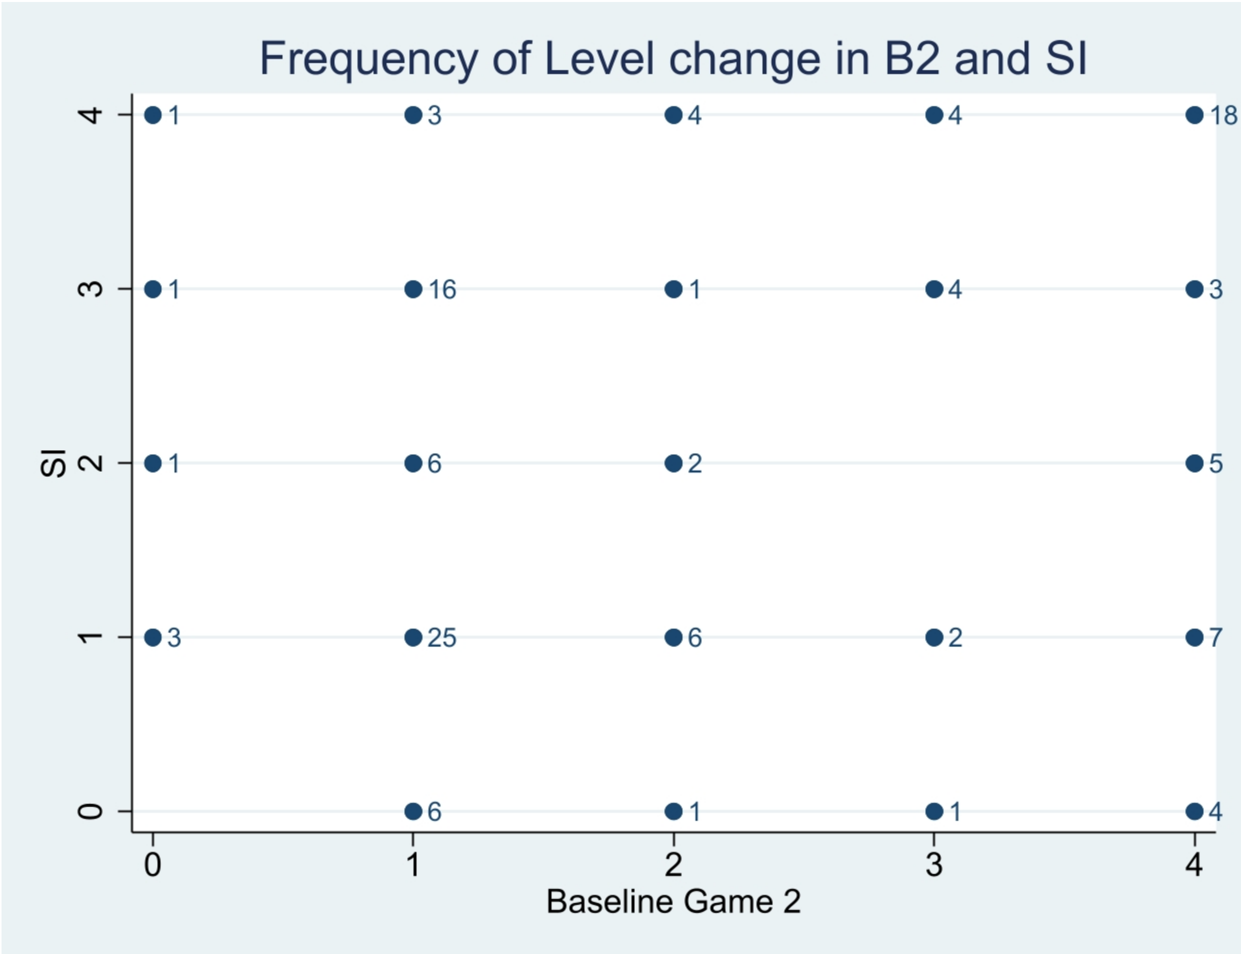

Figure 11: Individual Observed Lk in B2 and SI

reasoning level remains relatively stable. This highlights the nuanced ways in which players adopt different reasoning approaches, making the effects of iteration salience particularly intriguing.

In the Salience of Dominant Strategy treatment, we observe that making the dominant strategy more salient indeed enhances players' revealed reasoning levels. This effect is particularly pronounced among lower-level players, especially those classified as L0 and L1. These findings suggest that identifying the dominant strategy is genuinely challenging for some participants; however, once they recognize it, they become capable of engaging in additional rounds of iterative reasoning, leading to an increase in their strategic sophistication.

4.5 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 treatment dummies—Salience of Dominant Strategy (SD) and Salience of Iteration (SI)—which indicate whether the subject participated in these respective treatments. 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_4p are dummy variables capturing prior experience in game theory, with the baseline category being participants with the least exposure to game theory.

	(1)	(2)	(3)
Entropy	-0.389***	-0.389***	-0.389***
	(-7.15)	(-7.15)	(-7.15)
N	0.828***	0.827***	0.827***
	(8.09)	(8.08)	(8.08)
SD	0.268***	0.267***	0.267***
	(2.76)	(2.75)	(2.75)
SI	-0.111	-0.111	-0.112
	(-0.87)	(-0.88)	(-0.88)
CRT	0.157**	0.132**	0.127**
	(2.47)	(2.13)	(2.07)
G_2		0.157	0.166
		(0.85)	(0.90)
G_3		0.412**	0.410**
		(2.28)	(2.27)
G_4p		0.859***	0.830***
		(3.76)	(3.57)
age			0.0134
			(0.70)
gender_num			0.136
			(0.82)
reverse			-0.0546
			(-0.38)
Observations	1116	1116	1116

t statistics in parentheses

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

Table 4: Observed Level in all games.

The regression results reveal a strong negative relationship between Entropy and higher-order reasoning. The estimated coefficient for Entropy is -0.390 with a p-value below 0.001, indicating that as the informational complexity of a game increases, subjects become significantly less likely to engage in higher levels of reasoning. This finding aligns with the earlier descriptive results in the Informational Complexity section and is consistent with theories of bounded rationality, which suggest that increased cognitive demands constrain individuals' ability to process and respond optimally in strategic environments. The magnitude of this effect suggests that greater complexity shifts the distribution of reasoning levels downward, reinforcing the notion that cognitive constraints are a key determinant of strategic sophistication.

Similarly, CRT is positively and significantly associated with higher-order reasoning, with an estimated coefficient of 0.125 ($p = 0.039$). 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.

By contrast, the coefficient on Structural Complexity is positive and highly significant (0.828, $p<0.001$), indicating that when a game enables the identification of higher reasoning levels, subjects are more likely to engage in deeper strategic thought. This result underscores the role of game structure in shaping behaviour; rather than being an inherent trait, higher-order rationality is influenced by environmental features. The magnitude of this effect suggests that explicitly supporting and making higher-order reasoning discernible within the game significantly increases the likelihood of level-k reasoning. The treatment effects related to salience manipulations yield mixed results. The coefficient on the Salience of Dominant Strategy is positive and statistically significant at the 1% level (0.268, $p<0.01$), suggesting that making the dominant strategy more salient increases the likelihood of higher-order reasoning. This finding aligns with existing experimental evidence on the role of focal points in strategic decision-making. However, the effect size is smaller than that of Structural Complexity, indicating that while salience influences behaviour, structural game features exert a stronger impact.

In contrast, the coefficient on Salience of Iteration is negative but not statistically significant

(−0.110, p=0.39), suggesting that emphasizing iterative reasoning alone does not systematically affect the depth of strategic reasoning. The lack of statistical significance implies that without additional structural cues, subjects may not consistently adopt more sophisticated reasoning strategies.

The coefficients on the G_* variables capture the estimated effects of increasing levels of game theory experience on higher-order rationality (K). Since G_1 serves as the baseline group, all other coefficients are interpreted relative to this reference category. The results indicate a positive relationship between game theory experience and higher-order reasoning, with both the magnitude and statistical significance of the coefficients increasing as the experience level rises.

The coefficient on G_2 is positive (0.157) but not statistically significant at the 5% level, suggesting that individuals in this category do not exhibit significantly higher levels of reasoning relative to the baseline group. The lack of significance may indicate that at lower levels of experience, additional exposure to game theory does not necessarily translate into measurable improvements in strategic sophistication.

For G_3, the coefficient increases to 0.412 and is statistically significant at the 5% level ($p<0.05$), suggesting that individuals with moderate experience in game theory are more likely to engage in higher levels of reasoning than those in the baseline group. The significance of this result implies that formal exposure to game-theoretic concepts is associated with measurable improvements in strategic sophistication, though the effect remains relatively modest.

The most pronounced effect is observed for G_4p, which captures individuals with substantial game-theoretic experience [3] and yields a highly significant coefficient of 0.859. This finding suggests that extensive exposure to formal game-theoretic principles leads to substantial improvements in strategic sophistication. It reinforces the hypothesis that higher-order reasoning is not purely an innate cognitive ability but rather a skill that can be systematically developed through structured training in strategic reasoning. Moreover, the monotonic increase in the estimated effects across experience levels supports a learning-based explanation of strategic sophistication, wherein individuals progressively enhance their ability to engage in deeper levels of reasoning as they accumulate experience in game-theoretic settings.

These results provide strong evidence that both the complexity of strategic environments

[3] Self-reporting experience in game theory: G_4: know game theory and can use it, G_5: good at using it, we pool these two group as there are only 2 subjects reporting as G_5.

and the availability of identifiable reasoning structures influence human decision-making in games. The negative impact of entropy highlights the cognitive constraints individuals face in complex strategic settings, aligning with bounded rationality models that emphasize the limitations of human information processing in strategic interactions. Meanwhile, the positive effect of N suggests that higher-order reasoning is not solely an individual cognitive trait but is also shaped by the structure of the decision-making environment. The differing effects of salience treatments indicate that while highlighting dominant strategies can facilitate deeper reasoning, making iterative structures more salient does not necessarily induce higher levels of strategic thought.

4.6 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.

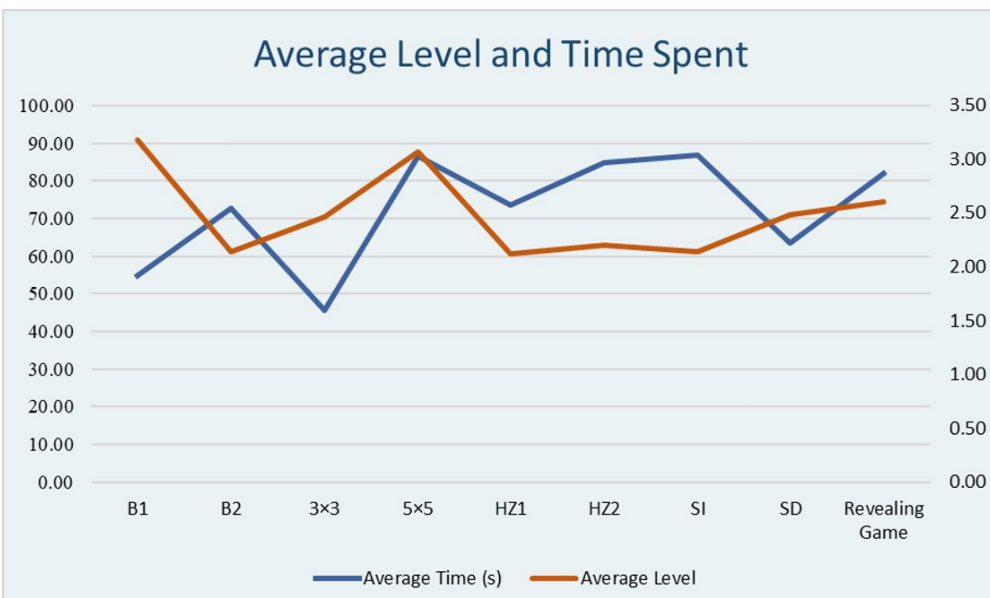

Figure 12: Average Level and Spent in each game

Figure 12 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.2, the coefficient on entropy is 12.34, indicating that as informational complexity increases, participants spend significantly more time making decisions. 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 (10.10***), 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 (-18.20***). 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, has no significant effect (5.235), implying that making iterative reasoning more prominent does not systematically alter response times. 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.

	(1)	(2)	(3)
Entropy	12.34***	12.34***	12.34***
	(7.09)	(7.08)	(7.07)
N	10.10***	10.10***	10.10***
	(5.18)	(5.17)	(5.17)
SD	-18.20***	-18.20***	-18.20***
	(-6.62)	(-6.61)	(-6.60)
SI	5.235	5.235	5.235
	(1.18)	(1.18)	(1.17)
CRT	1.231	1.545	1.161
	(0.70)	(0.90)	(0.74)
G_2		-4.505	-2.872
		(-0.87)	(-0.55)
G_3		-9.441**	-9.093**
		(-2.16)	(-2.10)
G_4p		-7.563	-4.758
		(-1.20)	(-0.84)
age			0.745
			(1.26)
gender			-6.108
			(-1.43)
Reverse Order			-9.519**
			(-2.51)
Observations	1116	1116	1116

t statistics in parentheses

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

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

When we look at the effect of Game Theory experience on time spent, the estimated coefficients for G_2, G_3, and G_4p 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 -9.441 indicates a significant reduction in decision time for participants with moderate game theory experience. In contrast, the estimates for G_2 (-4.505) and G_4p (-7.563) 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 experienced 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_4p 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.

4.7 Revealing Game

Figure 13 presents the heatmap of net highlight history in the Revealing Game, calculated as the number of times each cell was highlighted minus the number of times it was un-highlighted.

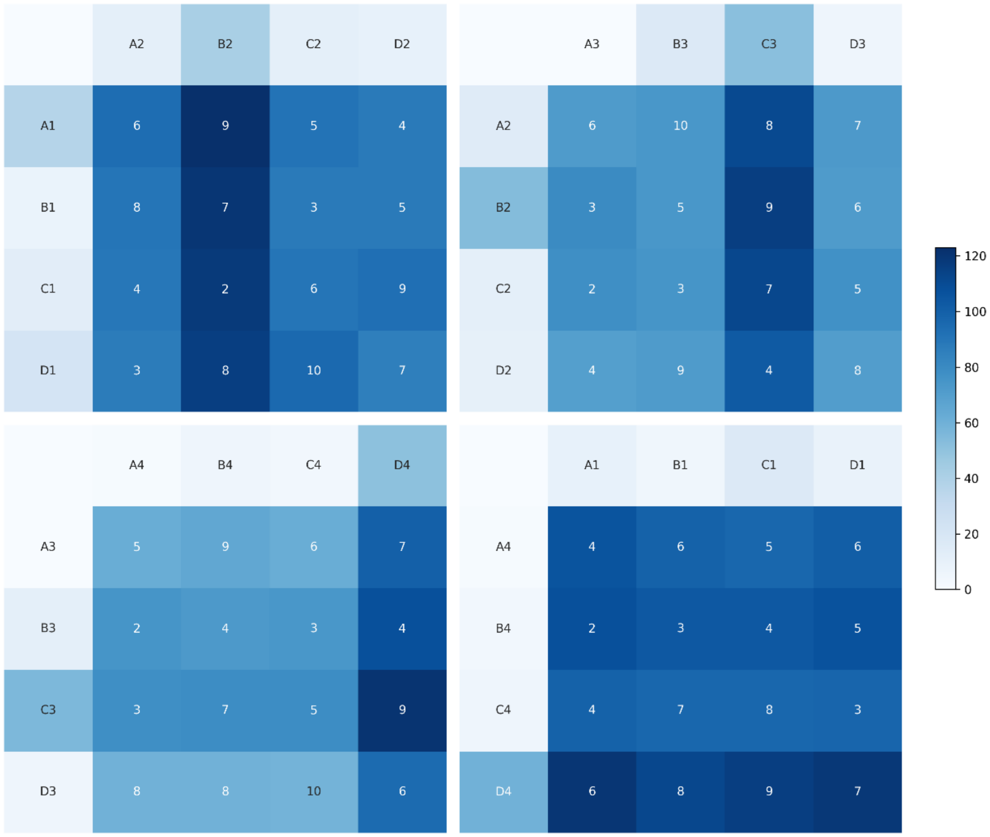

Figure 13: Heatmap of Net Highlight history [4]

This measure allows us to track the payoff tables each 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] The four tables correspond to the payoff matrices for Player 1 to Player 4 in the Revealing game, 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.

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 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 dominant 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.

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.