Finite Foresight in Chomp*

Zachary Breig[†] Hai Anh Pham[‡]

November 27, 2025

Abstract

We experimentally study sequential rationality in the Game of Chomp, a two-player win-lose game whose rules are simple but whose subgame-perfect Nash equilibria vary in complexity across initial states. Subjects frequently deviate from SPNE: the probability of making an SPNE-consistent move falls with the distance to the end of the game, and most can correctly implement SPNE only three to five moves ahead. Experience mainly improves play at intermediate distances. We also document a common opening heuristic that, although often theoretically suboptimal, is still associated with higher chances of winning.

JEL classification: C72, C73, C92, D90

Keywords: Chomp; experiment; sequential rationality; limited foresight

*The experiment discussed in this paper received ethical approval from UQ's BEL LNR review panel (project number 2024/HE001275).

[†]University of Queensland (z.breig@uq.edu.au).

[‡]University of Queensland (haianh.pham@uq.edu.au).

1 Introduction

Finite games of perfect information are among the simplest settings in which game theory delivers sharp predictions. Backward induction is a simple yet powerful solution concept that one might expect subjects to implement, yet real players may be unable to represent even moderately complex games as a tree or to carry out backward induction on large or intricate trees. Limited depth of reasoning, noisy best responses, and mis-learning can all undermine sequential rationality, and their impact is likely to grow with the game's complexity. This paper studies how much of the optimal play predicted by theory survives when players are boundedly rational, and how the resulting advantage varies across states within the same game.

In this paper, we study behavior in the Game of Chomp, a sequential, finite, perfect-information, two-player game introduced in Gale (1974). The game is played on an $n \times m$ rectangular grid of boxes. Players select one of the remaining boxes sequentially, and when any box is selected for removal, all those below and to the right of it are also removed. The player forced to remove the top-left box loses the game. The Game of Chomp has a first-mover advantage: under sequentially rational play, the first mover should always win. However, optimal play varies substantially depending on the initial dimensions of the game and does not, in general, follow a simple pattern. For this reason, the complexity and length of the game are easy to vary between matches without changing the fundamental ruleset, making it an ideal testing ground to measure how real-life participants deviate from sequential rationality as complexity varies.

This paper describes an experimental implementation of the Game of Chomp. We randomly match participants into groups of two to play the game. Roles (player 1 versus player 2) and initial game state (the dimensions of the grid) are randomized in each match. Participants are told that at the end of the experiment, one match would be chosen at random to be the match that determines payment. The winner of that match receives AUD$65, while the loser receives AUD$15.

Not surprisingly, participants are not perfectly consistent with sequential rationality. Even though the player chosen to move first has the power to win with certainty, they only win the game 51.72% of the time. However, this masks substantial heterogeneity across grids: Some initial game states exhibit player 1 win rates of upwards of 60%, while others have player 1 win rates that are lower than 40%.

We then investigate the determinants of deviations from sequential rationality. We find that a key determinant for whether a player makes a winning move is the distance to the end of the game according to the subgame perfect equilibrium path. This distance to the end of the game interacts with learning in an intuitive way: most learning (the increase in the likelihood of optimal play that is associated with experience) accumulates in game states that are of "intermediate" complexity – those that are 3-5 moves away from the end of the game. Finally, we show that subjects often chose the box $(2, 2)$ as the initial move of the game, despite it being theoretically suboptimal in non-square games. This may be in an effort to simplify the game, as many subjects appear to have learned to play optimally in the resultant "L-shaped" games.

The paper is organized as follows. Section 2 reviews the most relevant literature. Section 3 formally introduces the Game of Chomp, characterizes its subgame perfect Nash equilibria, and outlines our hypotheses. Section 4 details the experimental design. Section 5 presents the main findings on game outcomes and behavior in losing positions, and Section 6 investigates the underlying drivers of these behaviors. Section 7 offers concluding remarks.

2 Related Literature

There has been a substantial amount of experimental research into people's decisions in dynamic games with perfect information. Previous research has shown that deviations from the theoretical predictions likely stem from social preferences, limitations in backward induction, and complexity of the games.

Rosenthal (1981) introduced the Centipede Game, a dynamic game with perfect information, where two players take turns choosing to pass or stop at each node of the game. Continuing (usually) increases the total payoff for the group, but stopping gives the active player a higher proportion of the current total payoff. Generally, games are constructed so that in the subgame perfect Nash equilibrium (SPNE), subjects should choose to stop at the first node of the game. However, results from experimental studies of this game show that this does not match behavior. The Centipede Game was first studied experimentally in McKelvey and Palfrey (1992), which concludes that the failure to adhere to the SPNE is due to a mixture of altruism and decision noise. The existence of "altruists," who always choose "pass" at each node of the game, makes players more likely to continue the game if they believe their opponents are altruists. Fey et al. (1996) studies Centipede Games in which the total payoff does *not* increase with each pass, finding that Quantal Response Equilibrium fits behavior the best. Rapoport et al. (2003), in turn, compares Centipede Games under high and low stakes, finding that behavior is closer to SPNE and learning is stronger when stakes are high. Palacios-Huerta and Volij (2009) show that when professional chess players play Centipede Games against each other, outcomes are much closer to SPNE, whereas Levitt et al. (2011) find that even expert chess players who successfully solve a related "race to 100" game often fail to stop early in Centipede Games. In contrast to Centipede Games, a fully rational player acting as the first mover in the Game of Chomp does not need to form beliefs about the other player, because the game is zero-sum and has a binary outcome. Thus, any deviations from the SPNE can be attributed purely to the player's own ability to engage in backward induction.

Another strand of the experimental literature uses the "Race Game," which is a finite sequential zero-sum game of perfect information, to study how cognitive limitations and learning affect the implementation of backward induction. The game's rules can be described as follows: The state of the game at any point is represented as an integer m. Players take turns selecting a number $k \in \{1, \ldots, \bar{k}\}$, which is subtracted from the current state. The player who brings the state to zero wins the game. The Race Game has a simple winning strategy, and if the current state is a winning position, the player whose turn it is can guarantee victory regardless of the opponent's play (Dufwenberg et al., 2010; Gneezy et al., 2010). This strategy involves a relatively simple rule, depending only on m and $\bar{k}$. Both Dufwenberg et al. (2010) and Gneezy et al. (2010) focus on how learning this rule can transfer between different parameterizations (and perhaps subgames) of the game. Rampal (2025) uses the Race Game to compare behavioral game theoretic models, finding that the level-k model is the best fit for behavior in "short" games while the Limited Foresight Equilibrium of Rampal (2022) is better at explaining behavior in "longer" games. In contrast to the Race Game, the Game of Chomp does not *in general* admit a simple rule governing equilibrium play, so there is less reason to expect a single moment of "epiphany"' after which subjects play optimally.

The game of Nim, introduced in Bouton (1901), is a finite sequential zero-sum game of perfect information in which players alternately remove objects from a set of piles and the player who takes the last object wins. Bouton (1901) provides a complete solution, showing that every position can be classified as winning or losing, and provides a constructive procedure for optimal play. McKinney Jr and Van Huyck (2007) use Nim to show how increasing game complexity (measured by rank and related tree-based measures) reduces effective play: most subjects perform well only up to about rank 6, though a few can solve games up to rank 17. In follow-up work, McKinney Jr and Van Huyck (2013) find that learning mainly takes the form of "eureka" discovery of narrow heuristics, such as move-copying in two-row games, with little evidence that subjects acquire the full Bouton algorithm or more complex heuristics. In contrast to Nim, general Chomp positions do not admit a comparably simple global solution or a single dominant heuristic, so a Chomp experiment can shed light on how boundedly rational players search, form and transfer local heuristics in a structurally rich class of games.

Traditional game-theoretic models typically assume that players can carry out backward induction regardless of a game's complexity. Experimental work, however, shows that complexity matters: subjects perform better in simpler versions of the same strategic environment. Recent research has studied different facets of complexity. Oprea (2020) and Banovetz and Oprea (2023) show that subjects both find complex, multistep, and mentally implemented rules harder to follow and systematically choose simpler procedures because they are averse to implementing complex ones. Pycia and Troyan (2023) formalize complexity in terms of foresight, distinguishing between agents by how far they can plan into the future. Nagel and Saitto (2023) propose a measure of strategic complexity for mechanisms based on how many alternative actions or plans a player must compare in

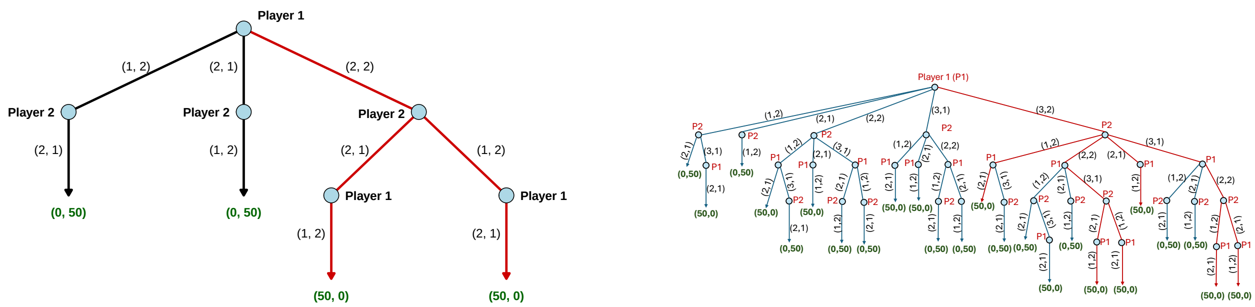

(a) (2×2) Game (b) (3×2) Game

Figure 1: Extensive form of the Game of Chomp, with its Subgame Perfect Equilibrium highlighted.

order to recognize a dominant strategy. The Game of Chomp provides a convenient testbed for these ideas: it has simple rules and direct winning strategies for square and $n \times 2$ grids, while variation in grid size and position allows us to vary both the number of available moves and the distance to the end of the game.

3 Theoretical Background

3.1 The Game of Chomp

The Game of Chomp, introduced by Gale (1974), is a dynamic game of complete information. The game's initial state is an $n \times m$ grid of boxes.[1] Two players take turns choosing from the remaining set of boxes. When a box is chosen, all boxes below and to the right of the chosen box are removed from the grid. Both players can always see the remaining boxes when they make their move. The player who is forced to remove the final (top-left) box is the loser.

Because Chomp is a finite game, it can be represented in extensive form. Figure 1 shows the game trees for grid sizes 2×2 and 3×2. The 2×2 game has only four terminal nodes. The extensive form of the 3×2 appears substantially more complex, with 24 terminal nodes. This complexity of the extensive form grows quickly with the dimensions of the game.

3.2 Winning Strategy for the First Player

Gale (1974) showed that the Game of Chomp is a first-mover advantage game: Player 1 can always guarantee a win with optimal play. Gale provides a non-constructive proof: Suppose Player 1 loses in the SPNE by choosing the initial move at the bottom right (n, m). Since choosing (n, m) leads to a loss, Player 2 has a best response, (a, b), that puts Player 1 into a losing position. However, Player 1

[1]Throughout the paper, we adopt the conventions that (1) an $n \times m$ game has n rows and m columns, (2) the first-mover is referred to as Player 1, and (3) the game is over when only a single box remains.

could have initially chosen (a, b), thus forcing Player 2 into a losing position. This is called "strategy stealing" (Berlekamp et al., 2004). Therefore, Player 1 can always win if he plays optimally. While there can be more than one optimal initial move in the Game of Chomp, the smallest grid size with more than one optimal initial move is 8×10 (Gale, 1974).

A standard method for identifying a winning strategy in an extensive-form game such as Chomp is to determine the SPNE through backward induction. This involves analyzing the final player's move: assuming sequential rationality, the player selects the action that maximizes their payoff if the game reaches that point. Given this choice, the preceding player then optimally responds, and so on. In Chomp, this process classifies every possible game state as either a winning or a losing position: a state is winning if at least one available move leads to a losing position, and losing if all moves lead to winning positions. Starting from the terminal 1×1 position, which is losing, each other state can be classified recursively in this way. However, because the number of possible states in a game of size $n \times m$ is $\binom{n+m}{n}$ (Zeilberger, 2001), the strategy space grows combinatorially, conceivably making it difficult for decision-makers to identify winning strategies in larger Chomp games.

While it is difficult in general to compute Chomp's SPNE, there are simple and straightforward winning strategies for square games and $n \times 2$ games, both of which are described in (Gale, 1974).

Square Games: For a square grid of $n \times n$, Player 1's optimal strategy is to initially choose the box at position (2,2). This results in an L-shaped game state with an equal number of boxes in each arm. Then, Player 1 mirrors all of Player 2's choices until the end of the game. For instance, on their next turn, Player 1 selects (1,a) or (a,1) if Player 2 selects (a,1) or (1,a), respectively. Player 1 repeats this process until she wins the game.

$\mathbf{n} \times \mathbf{2}$ Games: For grids that have either two columns or two rows, Player 1 wins by making the initial move at the bottom right box, which is $(n, 2)$ or $(2, m)$. This initial move results in an imbalanced game state, where the first column or row has one more box than the second. Regardless of Player 2's choices, Player 1 responds by selecting a box that maintains this imbalance. This guarantees a sure win for Player 1.

4 The Experiment

4.1 Experiment procedures

The experiment was conducted with two treatments: *Mixed Shapes* (MS) and *Rectangle Only* (RO). The MS treatment consisted of four sessions completed in August 2024, while the RO treatment consisted of two sessions completed in May 2025. All sessions took place in-person at the University of Queensland's Centre for Unified Behavioural and Economic Sciences Laboratory. The experiment was coded using oTree (Chen et al., 2016). Each session lasted 90 minutes, and no participant took part in more than one session.

In the experiment, subjects were provided with the game's instructions orally, on paper, and on their computer screen.[2] The instructions page provided an example grid size of 10×10 and explained how their choices changed the state of the game. Subjects could examine this by clicking on any box. Subjects were randomly rematched each round and continued playing new games of Chomp until the 65-minute session limit was reached, with the first match to finish after this cutoff serving as the final match of the experiment. Roles (Player 1 or Player 2, with Player 1 moving first) were randomly assigned in each round. The dimensions of the grid were also randomly selected for each group and each round, and the set from which they were selected differed according to the treatments that are discussed below. After the games were completed, subjects were asked to complete a short survey that included demographic questions, a cognitive reflection test (Frederick, 2005), and space to provide feedback on the experiment and their behavior.

Each Chomp game began with a grid of boxes. The top left box was yellow, and the others were white. The two players took turns selecting a box. Each time a box was selected, all the boxes below and to the right of it turned green. After the player clicked the "Submit" button, these boxes were removed from the grid. Each player had 30 seconds to make their selection. They could change their choice as many times as they wanted within the given time. After making their final decision, they had to click the "Submit" button to proceed to the next round.[3] Figure 2 shows an example of the choice page for Player 2 in a game with dimension 6×4 after Player 1 chose (4,3) in the first round. In each match, the winning player received AUD$65, while the losing player received AUD$15. At the end of this part, the computer randomly selected one match to count for the final payoff. Thus, the average payoff was AUD$40 for the 90-minute experiment.

Our experiment used a between-subjects design that varied the grid sizes subjects faced. In the *MS* ("mixed-shape") treatment, grids were randomly drawn from 2×2, 3×2, 4×2, 4×3, 5×3, 8×3, 4×4, 6×4, 6×5, and 6×6; among these, 2×2, 4×4, and 6×6 are square, while the rest are rectangular.[4] In the *RO* ("rectangular-only") treatment, grids were randomly drawn from 3×2, 4×2, 6×2, 4×3, 5×3, 6×3, 8×3, 5×4, 6×4, and 6×5.[5]

[2]Screenshots of all parts of the experiment can be found in Appendix B.

[3]If a player did not click "Submit" before time ran out, the computer randomly selected a box from all remaining boxes in the grid with equal chance. In our experiment, the timeout occurred in 351 out of 7425 turns. In order to avoid selection effects, our data analysis treats the random computer-generated choices as if they were made by the participant. Appendix Figure 9 shows how the proportion of turns that ended with a timeout is related to the number of available boxes.

[4]Here and throughout the remainder of the paper, we use the term "rectangular" to refer to games that are *not* square.

[5]As we discuss in Section 6.3, the most common initial move in the MS treatment is $(2, 2)$. Because this is the SPNE move in the square games in our design, we conjectured that subjects were learning that $(2, 2)$ works well on

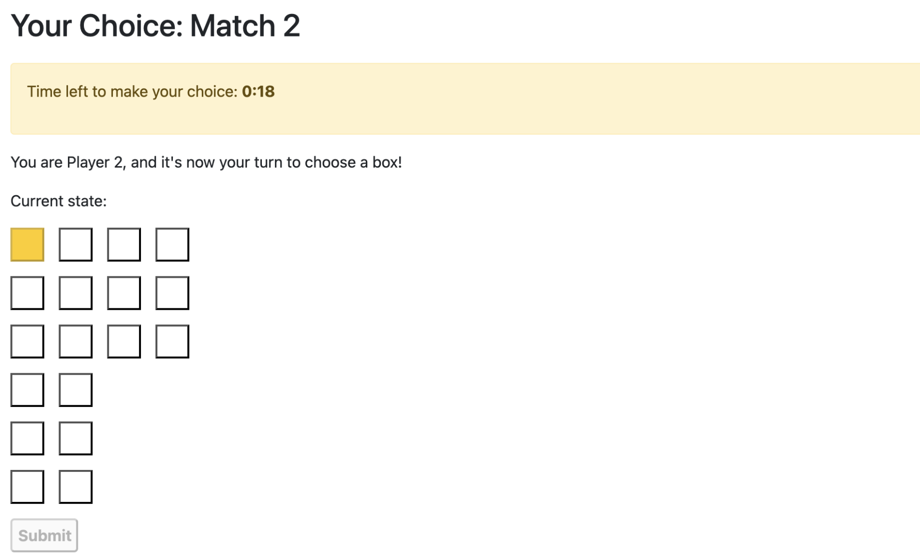

Figure 2: Choice page of Player 2

As the grid size increases, the complexity of SPNE and backward induction may also increase. The simplest versions of square games and $n \times 2$ games are 2×2 and 3×2, respectively. Grid sizes of 4×4 and 6×6 are examples of large dimensions for square games, while 4×2 and 6×2 are examples of $n \times 2$ games. The winning strategies in these games follow a simple rule, while optimal strategies in the rest of the rectangular games vary based on the dimensions of the grids. Also, all the grid sizes used in the experiment are smaller than 8×10. Therefore, each grid size has a unique optimal initial move.

4.2 Hypotheses

This subsection outlines the main hypotheses we test in our experimental framework. The rational benchmark delivers stark predictions in this setting: regardless of the current state of the game, players should always choose actions consistent with the SPNE, which in winning positions typically restricts play to a small subset of the many available options. Prior experimental work suggests that such behavior is rare. We therefore base our hypotheses on a combination of theoretical intuition and existing empirical findings, rather than on the SPNE benchmark alone.

For our first hypothesis, we consider how the game's initial state affects the likelihood that Player 1 will win the game. Under SPNE, the first mover should *always* win. However, this prediction relies on the prediction that Player 1 plays optimally even when not near the end of the game. A single

square boards and then misapplying this heuristic to rectangular boards. The *RO* treatment was therefore designed to eliminate all cases in which $(2, 2)$ is optimal. Nonetheless, $(2, 2)$ remained a common initial choice in the *RO* treatment.

error along the equilibrium path would put Player 2 into a winning position. For this reason, we expect that Player 1 is more likely to be able to win when the game is smaller and when the SPNE strategy is simpler.

Hypothesis 1. *"Player 1" is more likely to win "smaller" games, square games, and $N \times 2$ games.*

Our next hypothesis considers how players should choose when they are in a *losing* position. In these cases, SPNE gives no predictions, because all moves are expected to lead to a loss. However, Hypothesis 1 in combination with a recognition that players will make errors naturally leads to an intuitive strategy in these cases: a player in a losing position should remove fewer boxes in order to maintain the game's complexity and leave more chances for the other player to make a mistake.

Hypothesis 2. *Players tend to remove fewer boxes when they are in a losing position than in a winning position.*

Our next hypothesis considers how players' ability to implement SPNE varies with the state of the game. Generally, we expect that in more complex states subjects will be less able to play according to SPNE. We focus on two dimensions of complexity. First, a simple measure is the number of options available to the decision-maker. Second, implementing SPNE can require looking many moves ahead and evaluating multiple contingencies, so depth of reasoning may also be a constraint. When a subject is near the end of the game, fewer contingencies need to be considered and SPNE may be easier to implement. This is consistent with previous research, which finds that subjects are more likely to choose "take" near the end of the Centipede Game and to choose optimally near the end of the Race Game (McKelvey and Palfrey, 1992; Gneezy et al., 2010)

Hypothesis 3. *Players' moves are more likely to coincide with SPNE when there are fewer options available and when there are fewer steps remaining on the equilibrium path.*

Our final hypothesis is related to how behavior changes with experience. Previous work has tended to find that players' behavior becomes closer to SPNE as they gain experience (Dufwenberg et al., 2010; McKinney Jr and Van Huyck, 2013; Rampal, 2025). We expect similar results to hold for the Game of Chomp.

Hypothesis 4. *Players' moves are more likely to coincide with SPNE as they gain experience.*

5 Results

There were a total of 124 participants across the two treatments: 80 subjects completed 796 games across four sessions of *MS* and 44 subjects completed 486 games across two sessions of *RO*. Generally, our results are not substantially different between treatments. Subjects made choices 7425 times across 1282 games in the experiment, implying that the average number of turns per game was about 5.8. We report summary statistics for our demographic variables in Appendix Table 5.

5.1 Winner of the Game

This section considers the outcome of the game: who won, and under what conditions.

Result 1. Player 1 won roughly half of the time, and was more likely to win in "smaller" games and in square games. Player 1 did not win more often in $N \times 2$ games.

Result 1 shows that we partially reject Hypothesis 1. The evidence for Result 1 can be found in Figure 3, which shows a histogram of win rates for subjects in the role of Player 1, Figure 4, which shows win rates according to the initial game state and Appendix Table 6, which reports regressions of Player 1 winning on a game's initial characteristics.

Overall, Player 1 won in 663 out of 1282 games (a winning rate of 51.7%). Figure 3 shows that the winning rate varied widely by subject. Two subjects had winning rates that were above 0.9, and two *never* won when playing as Player 1. The histogram is left-skewed, with a peak winning proportion that falls between 0.6 and 0.65.[6]

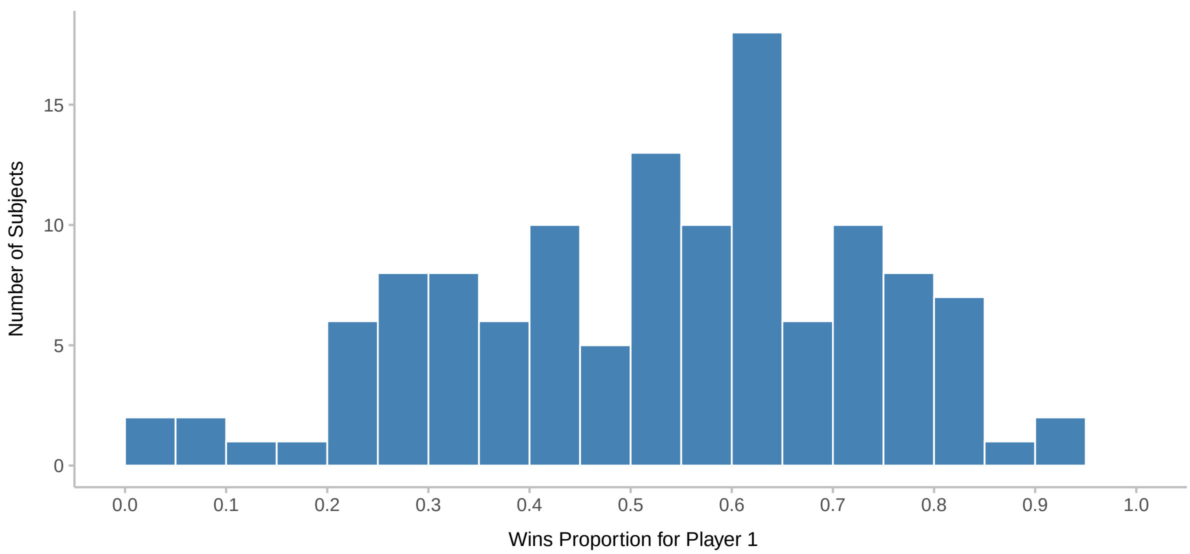

Figure 3: Histogram of proportion of winning as Player 1

We further analyze the winning proportions for Player 1 on different grid sizes and shapes. The left panel of Figure 4 shows the proportion of wins by Player 1 for different grid sizes, using observations from two separate treatments, while the center and right panels pool observations from both treatments to compare grid shape. Overall, subjects won more often with square games than with rectangular ones (roughly 70% vs below 50%) and less often with $N \times 2$ grids as compared to other shapes (around 49% vs above 52%). All square grid sizes have a winning proportion above

[6]Appendix Figure 10 shows the CDF for subject win rates split between the MS and RO treatments. The CDFs are almost identical.

0.5, with the highest winning rate of approximately 0.9 for grid size 2×2, whereas the figure for the majority of rectangular grid sizes is less than 0.5. In the MS treatment, the only rectangular games that have a winning proportion higher than 50% are the 3×2 and 6×4 games. In RO treatment, 3×2, 5×3, and 6×5 games have a winning proportion higher than 0.5.

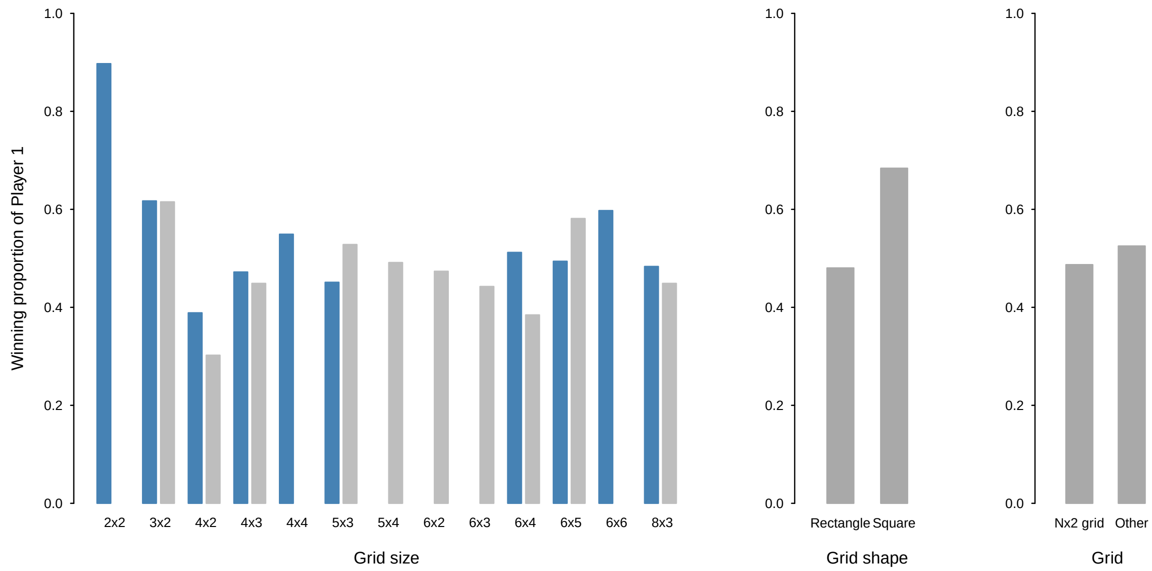

Figure 4: Winning proportion of Player 1 by grid size and shape. In the left panel, blue columns represent results from the *MS* treatment and grey from the *RO* treatment.

Column (1) of Appendix Table 6 reports a regression where the dependent variable is an indicator for Player 1 winning the game and the independent variables include characteristics of the initial game states. The regression shows that Player 1 wins square games roughly 20 percentage points more often than rectangular games, and that each additional available box at the start of the game is associated with a 0.44 percentage point drop in Player 1 winning. $N \times 2$ grids are associated with a 5 percentage point *decrease* in winning rates, but this relationship is not statistically significant. Column (2) explores the pathway of these effects by controlling for whether Player 1 played the SPNE move in their first turn. The coefficients on Square Game and Num. Available Boxes become near zero, while the coefficient on $N \times 2$ becomes negative and significant. This suggests that the apparent advantage of square and smaller boards operates mainly through their effect on the likelihood that Player 1 chooses the SPNE opening, but that conditional on that opening choice and other controls, grids with two columns are intrinsically less favorable to Player 1 than other shapes.

5.2 Moves Made from a Losing Position

All game states in Chomp can be classified as either winning or losing. In a winning state, a player who follows an SPNE strategy at the current and all future nodes is guaranteed to win. Because only a small subset of moves is optimal, SPNE provides sharp predictions for behavior in winning positions. By contrast, in a losing state SPNE implies that every available move ultimately leads to a loss and therefore does not distinguish among actions at that state—it does not make *any* prediction about what players will do. In this section, we study how players actually behave in these losing positions.

To test Hypothesis 2, which states that players in a losing position will remove fewer boxes, we need to account for several other features of the game's state. In particular, there are characteristics of the current position that are correlated with whether it is winning or losing and are also likely to affect how many boxes are removed. These include the current number of boxes in the grid, how close the player is to the end of the game under rational play, and how many boxes *should* be removed in the corresponding winning position according to SPNE.[7]

Result 2. Players do not remove significantly fewer boxes from losing positions than from comparable winning positions.

The evidence for Result 2, which offers little support for Hypothesis 2, is presented in Table 1. The regressions relate the number of boxes removed to characteristics of the current game state, controlling for "Match Number" (a proxy for experience), the current number of available boxes, and the number of moves left according to SPNE (discussed further in Section 6.1).

Column (1) of Table 1 shows that, controlling for other characteristics of the game state, players in a winning position remove about 0.1 more boxes than players in a losing position, on average. This difference is statistically significant but very small in magnitude: players must remove at least one box, and the average number of boxes removed is 2.88. When we additionally control for the number of boxes that *should* be removed in the corresponding winning position according to SPNE, the difference between winning and losing positions is no longer statistically significant.

[7] In cases where there are multiple SPNE moves, we define this variable as the maximum number of boxes that can be removed while remaining consistent with SPNE.

Dependent variable: Number of boxes removed		
	FE (1)	FE (2)
Winning Position	0.1007**	0.1041
	(0.0397)	(0.1006)
Winning Position x SPNE Boxes Removed		−0.0032
		(0.0880)
Num. Available Boxes	0.6526***	0.6556***
	(0.0271)	(0.0844)
Match Number	−0.0421***	−0.0421***
	(0.0056)	(0.0056)
SPNE Moves Left	−0.4217***	−0.4252***
	(0.0348)	(0.0888)
Observations	7,425	7,425
Adjusted R^2	0.7001	0.7001

Note: Standard errors are clustered at the subject level, and subject fixed effects are included. Significance: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 1: Regression on number of boxes removed

Although not our main focus, the coefficients on the other covariates are informative. Unsurprisingly, players remove more boxes when more boxes are available. Less obvious is that the negative and statistically significant coefficient on "Match Number" indicates that subjects remove fewer boxes as they gain experience. In addition, even holding the current number of boxes fixed, players remove fewer boxes when they are closer to the end of the game along the SPNE path. Overall, we do not find strong evidence in favor of Hypothesis 2.

6 Patterns and Determinants of Play

The results in Section 5 show that while subjects in our experiment do not behave exactly as theory predicts, there are clear patterns in how they choose, and these patterns depend on the state of the game that they face. In this section, we explore the drivers of behavior in the Game of Chomp.

6.1 Complexity and Limited Backward Induction

Most game-theoretic solution concepts, including SPNE, implicitly assume that decision makers face no constraints on their ability to compute equilibria. In practice, however, even computers cannot fully solve complex games such as chess or Go (Simon, 1955). Experimental evidence in McKinney Jr and Van Huyck (2007, 2013) similarly shows that game complexity limits subjects' ability to choose according to equilibrium. Thus, the complexity of the game is a natural candidate determinant of how subjects actually play. In this subsection, we investigate which features of the Game of Chomp make it difficult for players to implement backward induction.[8]

We focus on two sources of complexity: (i) the number of options available to the player and (ii) how far into the future the player must reason. The first is captured by the number of available boxes, since a player always chooses from the set of boxes that remain (except for the top-left box). To quantify the second, we define the variable "SPNE Moves Left," which measures the number of moves remaining along a particular SPNE path. More specifically, we use backward induction on the game tree, breaking indifference for the player in a losing position by selecting the option that pushes the end of the game furthest into the future, and for the player in a winning position by selecting the option that brings the end of the game closest. Using this rule, "SPNE Moves Left" is defined as the number of turns until the game ends.[9] Table 2 shows the values of these variables for each of the initial game states used in our experiment, and Appendix Figure 11 shows graphically how each measure is related to the likelihood of choosing in a way that is consistent with SPNE.

[8]A natural question is whether standard bounded-rationality models such as level-k, cognitive hierarchy, or quantal response equilibrium (QRE) can account for our findings (Nagel, 1995; Camerer et al., 2004; McKelvey and Palfrey, 1995). In a perfect-information win-lose game like Chomp, canonical level-k and cognitive hierarchy models would predict that any type above level 0 plays very close to SPNE, once she computes a best response, which conflicts with the large, depth-dependent errors we observe. Likewise, QRE treats deviations as payoff-sensitive noise that is roughly homogeneous across states, whereas in our data error rates depend strongly on the distance to the end of the game. These patterns are more naturally captured by our notion of limited foresight and complexity-driven computational constraints.

[9]By construction, all winning positions have an odd number of moves left, while all losing positions have an even number of moves left.

Grid Size	Moves Left	Available boxes
2x2	3	4
3x2	5	6
4x2	7	8
4x3	7	12
4x4	7	16
5x3	11	15
5x4	11	20
6x2	11	12
6x3	11	18
6x4	11	24
6x5	15	30
6x6	11	36
8x3	15	24

Table 2: Summary of available boxes and moves left for grid size.

Result 3. SPNE-consistent play falls sharply with the steps remaining on the equilibrium path and, conditional on that distance and the shape of the game, is essentially unrelated to the number of available options.

The evidence for Results 3, which partially contradicts Hypothesis 3, can be found in Table 3, which presents fixed-effects regressions where the dependent variable is an indicator for making a winning move.[10] We restrict the sample to choices made from winning positions, so the outcome equals one if the subject selects a move that guarantees they will be in a winning position when they next move and zero otherwise. The independent variables are characteristics of the game state the subject faces. All specifications include Match Number, our proxy for experience.[11] Our main focus in this analysis is on the variables "Num. Available Boxes" and "SPNE Moves Left," which capture different aspects of the complexity of the subject's current state.

[10]Implicitly, these regressions use the likelihood of making an SPNE-consistent move as a proxy for complexity. As an alternative proxy, we can consider the time subjects took to make each move, although each turn was capped at 30 seconds. Appendix Figure 12 shows that decision time increases with SPNE moves left and the number of available boxes for low values of these variables, but then levels off at about 20 seconds.

[11]The coefficient on Match Number is positive, indicating the intuitive result that as subjects gain experience they are more likely to make a winning move. We address learning more explicitly in Section 6.2.

Dependent variable: SPNE Consistent Move				
	FE (1)	FE (2)	FE (3)	FE (4)
Match Number	0.0042***	0.0056***	0.0057***	0.0056***
	(0.0011)	(0.0010)	(0.0010)	(0.0010)
Num. Available Boxes	−0.0373***		0.0054**	0.0028
	(0.0011)		(0.0022)	(0.0019)
SPNE Moves left		−0.0805***	−0.0898***	−0.0851***
		(0.0012)	(0.0038)	(0.0033)
Square Game State				0.0978***
				(0.0184)
Nx2 Game State				0.0017
				(0.0264)
Observations	5,555	5,555	5,555	5,555
Adjusted R^2	0.3489	0.4580	0.4593	0.4624

Note: Standard errors are clustered at the subject level, and subject fixed effects are included. Significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Effects of Complexity on SPNE Consistent Move

The negative and statistically significant coefficients for the number of available boxes in model 1 and for moves left in model 2 imply that subjects are less likely to make winning moves when the game is more complex. Model 3 presents a more surprising result: when controlling for the number of moves left to the end of the game, the relationship between the number of available boxes and the likelihood of making a winning move is *positive*, albeit small. However, this relationship disappears when we control for the "shape" of the game, as captured by dummies for the current grid being either a square or having two columns. Thus, in the Game of Chomp, the distance to the end of the game is a more relevant measure of complexity than the number of options available to the decision-maker.

The fact that the ability to implement the SPNE strategy depends strongly on the distance to the end of the game and only weakly on the number of available moves speaks directly to the complexity literature discussed in Section 2. In particular, our results line up closely with Pycia and Troyan (2023)'s emphasis on planning horizons and limited foresight. They further reinforce the idea, advanced by Rampal (2022) and Rampal (2025), that limitations to foresight should be explicitly incorporated into behavioral equilibrium concepts.

In order to further explore the prevalence of limited backward induction in our sample, we classify subjects according to the distance to the end of the game at which they are able to implement the SPNE strategy. Specifically, we say that a subject's Foresight Level[12] is k if they choose a winning move (i.e., a move that preserves a win according to SPNE) at least 70% of the time when the end of the game is at most k turns away, and less than 70% of the time when it is further away.[13] Thus, a subject with Foresight Level 1 tends to make a winning move when they can win the game in the current round, but not when they could guarantee themselves a win only by their next move.[14] In all cases, we restrict attention to choices made from winning positions.

Figure 5 shows the distribution of Foresight Levels in our sample.[15] The vast majority of subjects have Foresight Levels of either 3 or 5, suggesting that they can evaluate contingencies a few turns into the future, but not much further. Similar patterns appear in other dynamic games. In Nim, McKinney Jr and Van Huyck (2007) estimate a "rationality bound" and find that the average subject can reason effectively only up to about rank 6, with substantial heterogeneity across subjects. In the Game of 21 and the Race Game, Dufwenberg et al. (2010) and Gneezy et al. (2010) likewise report that behavior is close to equilibrium only a few moves from the end of the game and deteriorates further away. Our Foresight Levels, which cluster at three to five moves, are thus broadly in line with earlier evidence on finite planning horizons in dynamic games.

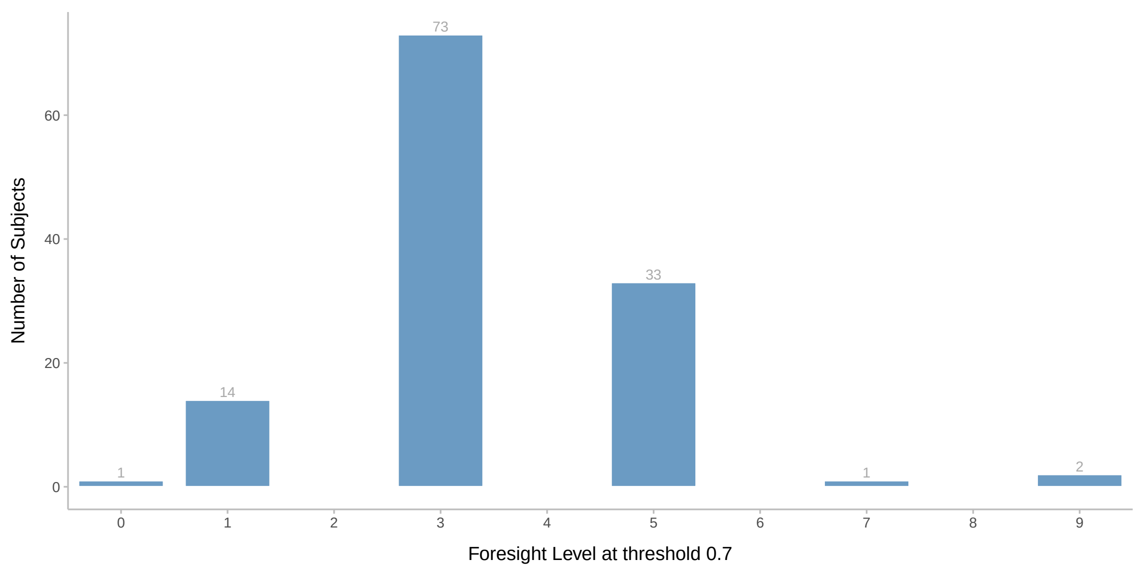

Figure 5: Subjects' Foresight Level at threshold 0.7

[12]We borrow this terminology from Rampal (2022).

[13]The cutoff of 70% was chosen arbitrarily. Appendix Figure 13 shows the classifications for cutoffs of 50% and 80%, respectively.

[14]A Foresight Level of 0 indicates that the subject chooses correctly less than 70% of the time even when they could win in the current round.

[15]Because of the structure of the Game of Chomp, Foresight Levels are necessarily odd: winning moves can only be made in game states that are an odd number of turns away from the end of the game according to SPNE.

Appendix Table 8 explores the correlation between survey responses and optimal behavior in our experiment. The dependent variable in column (1) is a binary indicator for making a choice consistent with SPNE (restricting to observations where the player is in a winning position), while the dependent variable in column (2) is the Foresight Level that we calculate for the subject. We find consistent results in both regressions. Conditional on all the covariates, being male is associated with a 5% increase in the likelihood of making a choice consistent with SPNE and a 0.67 higher Foresight Level. Each additional point on the CRT score is associated with a 5% increase in the likelihood of making a choice consistent with SPNE and a 0.30 higher Foresight Level. No other covariates have statistically significant coefficients.

6.2 Learning the Game of Chomp

In this section, we discuss how subjects' behavior changed as they gained more experience. Previous research has shown that, in many repeated games, subjects' play tends to move closer to Nash equilibrium as they gain experience (see, e.g., Van Huyck et al. (1991); Nagel (1995); Camerer and Hua Ho (1999)). Moreover, this learning extends to playing closer to subgame perfection (Gneezy et al., 2010; McKinney Jr and Van Huyck, 2013). The positive coefficient on Match Number in Table 3 is consistent with this: subjects make winning moves more often when they have more experience.

The Game of Chomp provides a unique and rich setting to study how decision-makers learn to play subgame perfect equilibrium. The game involves very simple game states, where the player is only a few moves away from winning, and complex game states for which even experts might have difficulty finding the optimal move.

Result 4. Subjects' consistency with SPNE increases with experience, and this improvement is concentrated in subgames of intermediate complexity.

We measure learning effects by using the indicator dependent variable "SPNE Consistent Move," and then binning game states by SPNE Moves Left, as shown in Table 4. Observations are restricted to moves at winning positions (and, thus, where the number of moves left until the end of the game is odd). The states of the game were categorized as having 3–5, 7–9, or more than 11 moves left to the end node of the game, with the corresponding dummy variables shown in the regression.[16] Interaction terms between the aforementioned dummy variables and the independent variable *Match Number* are also included. The baseline is when a move that wins the game is available in the current round.

[16] In Appendix Table 7, we reproduce these results using the number of boxes left instead of the number of moves left as the game state. The interpretation is largely the same: learning is concentrated in games of intermediate complexity.

Dependent variable: SPNE Consistent Move		
	MS	RO
Intercept	0.9716^{***}	0.9659^{***}
	(0.0117)	(0.0171)
3-5 Moves	-0.4051^{***}	-0.3736^{***}
	(0.0393)	(0.0412)
7–9 Moves	-0.7613^{***}	-0.7322^{***}
	(0.0488)	(0.0626)
More than 11 Moves	-0.8525^{***}	-0.9358^{***}
	(0.0283)	(0.0243)
Match Number	0.0011	0.0017
	(0.0009)	(0.0011)
3–5 Moves × Match Number	0.0147^{***}	0.0086^{***}
	(0.0029)	(0.0028)
7–9 Moves × Match Number	0.0050	0.0011
	(0.0037)	(0.0042)
More than 11 Moves × Match Number	−0.0018	−0.0010
	(0.0022)	(0.0015)
Observations	3461	2094
Adjusted R^2	0.4604	0.5206

Note: Standard errors are clustered at the subject level.
Significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 4: SPNE Consistent Move at different number of Moves Left

The patterns are similar across the two treatments. In both MS and RO, the coefficients on all distance-to-end dummies are negative and statistically significant at the 1% level, indicating that subjects are less likely to make a winning (SPNE-consistent) move when the game state is farther from the end node. When game states have 3–5 moves remaining, the probability of making a winning move falls by 0.41 in MS and 0.37 in RO relative to the baseline. When the game has 7–9 moves left, the probability decreases by over 0.70 in both treatments. For states with more than 11 moves left, the coefficients are even larger in magnitude (0.85 in MS and 0.94 in RO), reducing the probability of a winning move close to zero.

The coefficient on Match Number is positive but close to zero and not statistically significant, implying little detectable learning for states that are a single move away from the end of the game.[17] The interaction terms provide more insight into how learning varies across states, and they are generally consistent with Hypothesis 4. Subjects show some improvement in states that are a few moves from the end of the game: the interaction coefficients for states with 3–5 moves remaining are positive, and statistically significant at the 1% level. The coefficient implies that the probability of making an optimal choice increases by roughly 20 percentage points by the end of the experiment.

[17] This is not surprising, as the probability of making a winning move is already near one in these states.

By contrast, learning appears weaker and is statistically insignificant when there are 7–9 moves or more than 11 moves left.

6.3 L-shaped Games and the $(2, 2)$ Heuristic

This subsection reports an exploratory analysis of a pattern that emerged clearly in the data but that we did not foresee: the pervasive use of the move (2,2) as an opening and the role of the resulting L-shaped games.

Since all the grid sizes used in the experiment are smaller than 8×10, there is a unique optimal initial choice for each grid size. For instance, the optimal initial choice for a 2×2 grid is $(2, 2)$, and making this move guarantees Player 1 a sure win; for a 3×2 grid, the move $(3, 2)$ is optimal. As noted in Section 3, while $(2, 2)$ is the optimal first move for all square grids, it is *never* the optimal first move for *any* non-square grid.

Figure 6 shows players' first move in all grid sizes from the MS treatment.[18] The bottom (striped) portion of each bar represents the proportion of choices that were consistent with SPNE. The SPNE-consistent initial choice was only chosen higher than 50% of the time in 2×2 and 3×2 games. Despite being suboptimal in all non-square games, $(2, 2)$ is the most prevalent choice for all grid sizes other than 3×2 and 4×2. It is also the most popular initial move in the RO treatment, indicating that subjects started suboptimally and, according to SPNE, lost their initial advantage of being Player 1.

[18]The equivalent results for the RO treatment, which does not show substantive differences, can be found in Appendix Figure 14.

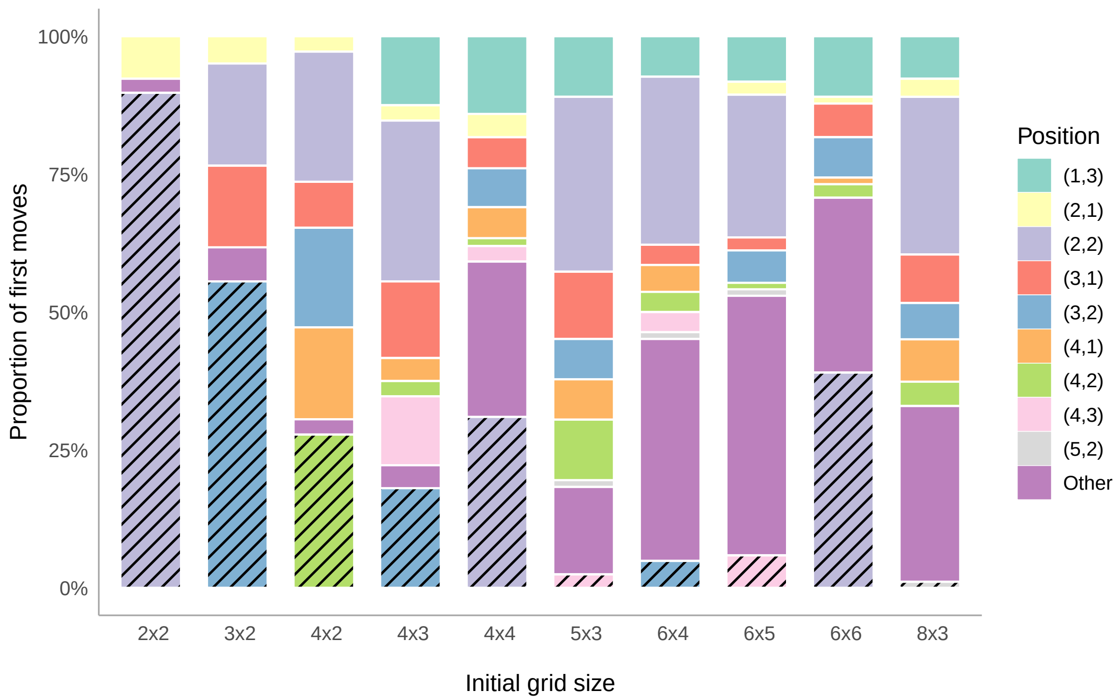

Figure 6: Proportion of initial choice by grid size in MS treatment

A more surprising result is that the rate of Player 1 winning when choosing $(2,2)$ as initial move is relatively high in both treatments, *even for rectangular games*. These rates are shown in Figure 7. In MS, the winning proportion is higher than 0.5 for most grid sizes; the only cases this does not hold are rectangular games with two columns, where the optimal initial choice is the bottom right box, as well as 5×3 and 8×3 games. In RO treatment, the exceptions are all games with two columns and the 6×3 game.

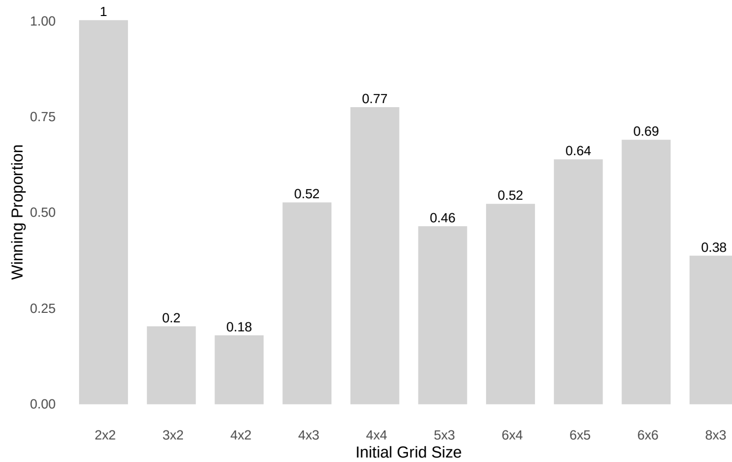

(a) MS treatment

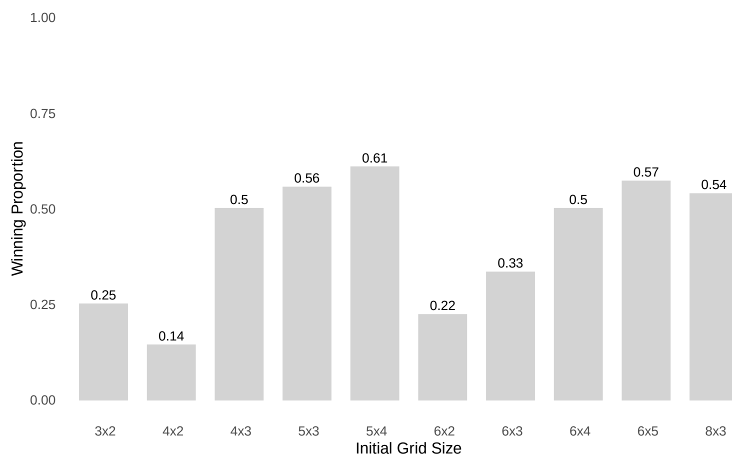

(b) RO treatment

Figure 7: Winning proportion when choosing initial choice $(2, 2)$ by grid sizes

We refer to the game state that results from choosing the box at (2,2) as an L-shaped game. Throughout the experiment, 35% of games eventually reached an L-shaped game. Because L-shaped games are both empirically prevalent and (relatively) strategically straightforward, we further study how subjects learned and performed in them. Figure 8 shows the proportion of winning moves made by 124 subjects in L-shaped games. Six subjects *never* made an optimal choice, while nine played perfectly. Roughly 60% of the subjects have a probability of making a winning move in L-shaped games of more than 0.5. This might explain why the winning proportions while choosing (2,2) as initial move are relatively high. Many subjects learned to play L-shaped games optimally.[19]

[19]Column (3) of Appendix Table 8 reports the results of a linear regression analysis of the probability of making a winning move in L-shaped games, based on the subjects' demographic information. Similar to the relationship between demographics and making optimal moves more generally, both the indicator for being male and the CRT score are positively related to the likelihood of choosing optimally in L-shaped games.

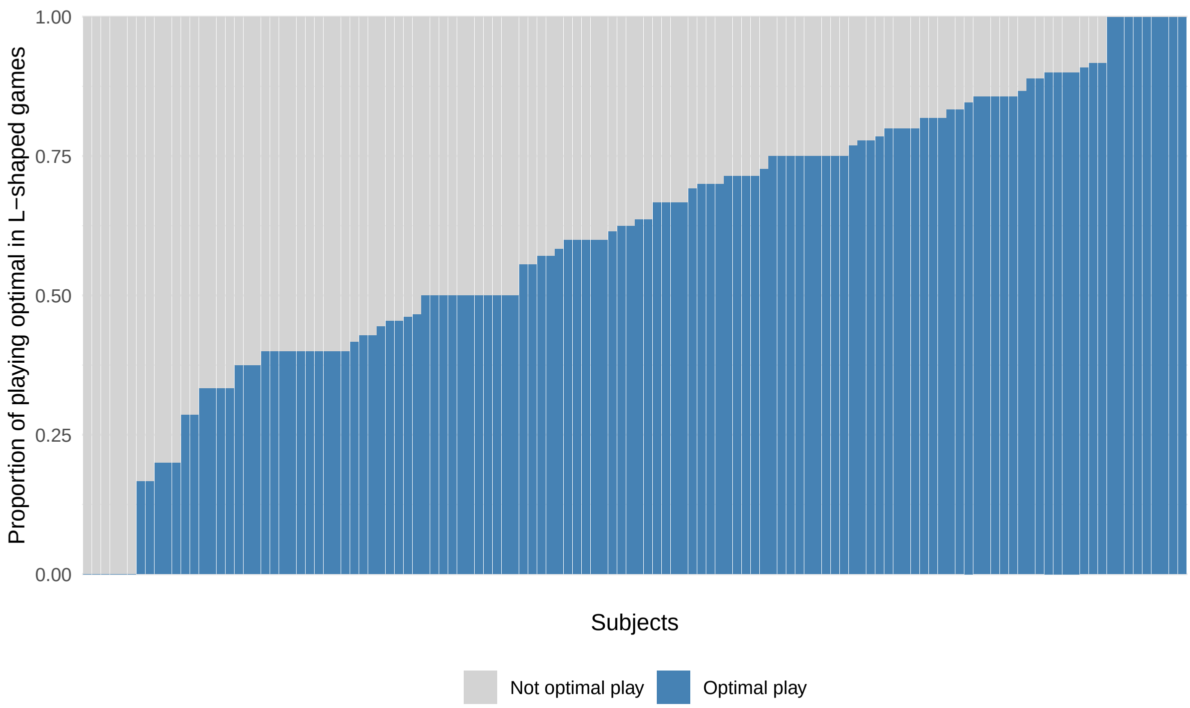

Figure 8: Probability of making a winning move in L-shaped games

Taken together, these findings suggest that many subjects adopt a simple $(2, 2)$ opening heuristic: they frequently choose (2,2) as an initial move, even when it is not SPNE-optimal, in order to transform diverse starting boards into a familiar class of L-shaped games. Because a large share of subjects learn to play these L-shaped subgames close to optimally, this heuristic remains profitable in practice and yields relatively high win rates for Player 1, even on rectangular boards where (2,2) is theoretically suboptimal. This pattern complements our earlier evidence on limited foresight: rather than computing optimal play from each initial state, subjects appear to reshape the game into strategically simpler subgames where they understand how to play optimally.

7 Conclusion

In this paper, we reported the results of an experiment implementing the Game of Chomp. Subjects played repeatedly with random partners and initial game states. They do not play according to SPNE: even when they could guarantee themselves a win, they often fail to choose an optimal move. We explored this behavior further and found that the ability to choose optimally decreases with the number of steps until the end of the game, but that subjects learn to play closer to optimally over the course of the experiment. Finally, we documented a simple heuristic that subjects appear to use to simplify the game, transforming a variety of boards into a familiar class of positions that they can play relatively well.

Our results have implications for behavioral and experimental game theory. At a basic level, they suggest that limited foresight is a key feature that needs to be incorporated into equilibrium concepts for dynamic games, supporting approaches such as that in Rampal (2022). More broadly, systematic deviations from SPNE should be taken into account both in dynamic mechanism design and when empirically analyzing dynamic games. Finally, our approach provides a simple tool for evaluating sequential rationality that is easy to implement and explain. We expect this to be useful in future work studying learning and the transfer of behavior across related dynamic environments.

References

Banovetz, J. and Oprea, R. (2023). Complexity and procedural choice. *American Economic Journal: Microeconomics*, 15(2):384–413.

Berlekamp, E. R., Conway, J. H., and Guy, R. K. (2004). *Winning ways for your mathematical plays, volume 4*. AK Peters/CRC Press.

Bouton, C. L. (1901). Nim, a game with a complete mathematical theory. *Annals of Mathematics*, 3(1/4):35–39.

Camerer, C. and Hua Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874.

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.

Dufwenberg, M., Sundaram, R., and Butler, D. J. (2010). Epiphany in the game of 21. *Journal of Economic Behavior & Organization*, 75(2):132–143.

Fey, M., McKelvey, R. D., and Palfrey, T. R. (1996). An experimental study of constant-sum centipede games. *International Journal of Game Theory*, 25(3):269–287.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4):25–42.

Gale, D. (1974). A curious nim-type game. *The American Mathematical Monthly*, 81(8):876–879.

Gneezy, U., Rustichini, A., and Vostroknutov, A. (2010). Experience and insight in the race game. *Journal of Economic Behavior & Organization*, 75(2):144–155.

Levitt, S. D., List, J. A., and Sadoff, S. E. (2011). Checkmate: Exploring backward induction among chess players. *American Economic Review*, 101(2):975–990.

McKelvey, R. D. and Palfrey, T. R. (1992). An experimental study of the centipede game. *Econometrica: Journal of the Econometric Society*, pages 803–836.

McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38.

McKinney Jr, C. N. and Van Huyck, J. B. (2007). Estimating bounded rationality and pricing performance uncertainty. *Journal of Economic Behavior & Organization*, 62(4):625–639.

McKinney Jr, C. N. and Van Huyck, J. B. (2013). Eureka learning: Heuristics and response time in perfect information games. *Games and Economic Behavior*, 79:223–232.

Nagel, L. and Saitto, R. (2023). A measure of complexity for strategy-proof mechanisms. *Working Paper*.

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):3913–3951.

Palacios-Huerta, I. and Volij, O. (2009). Field centipedes. *American Economic Review*, 99(4):1619–1635.

Pycia, M. and Troyan, P. (2023). A theory of simplicity in games and mechanism design. *Econometrica*, 91(4):1495–1526.

Rampal, J. (2022). Limited foresight equilibrium. *Games and Economic Behavior*, 132:166–188.

Rampal, J. (2025). Opponent's foresight and optimal choices. *Working Paper*.

Rapoport, A., Stein, W. E., Parco, J. E., and Nicholas, T. E. (2003). Equilibrium play and adaptive learning in a three-person centipede game. *Games and Economic Behavior*, 43(2):239–265.

Rosenthal, R. W. (1981). Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25(1):92–100.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118.

Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *The Quarterly Journal of Economics*, 106(3):885–910.

Zeilberger, D. (2001). Three-rowed chomp. *Advances in Applied Mathematics*, 26(2):168–179.

A Additional Empirical Results

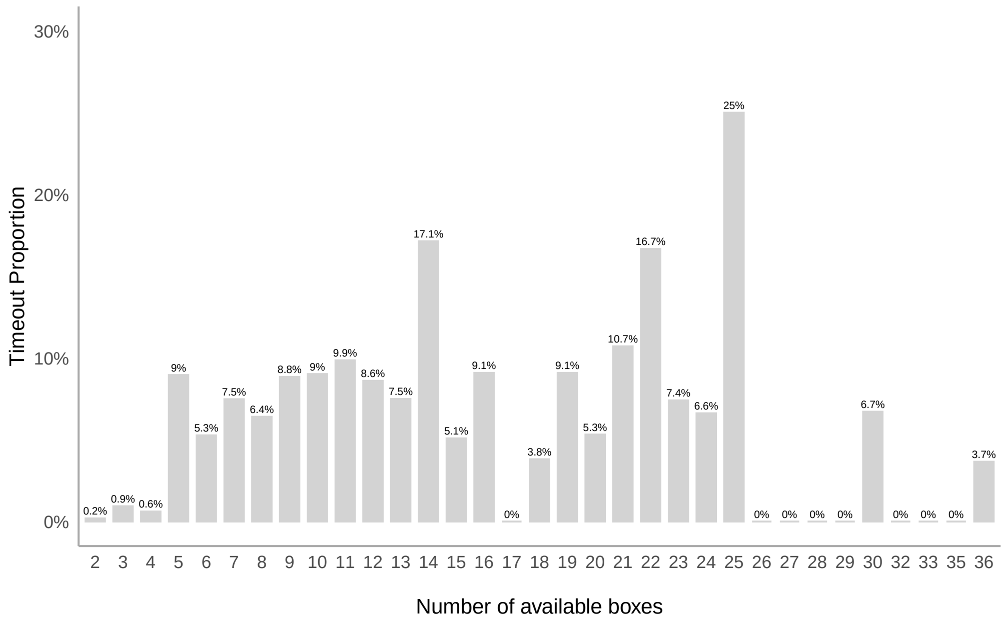

Figure 9: Timeouts by number of available boxes. This figure shows the proportion of timeouts at different number of available boxes. Timeouts are less likely to occur when game states have fewer than four boxes.

	Mean	Std. Dev.
CRT Score	1.40	1.17
Male	0.41	0.49
Age	23.27	4.87
English	0.20	0.40
Economics	0.33	0.47
Subjects	124.00	

Note: CRT Score is the number of correct answers on a Cognitive Reflection Test, ranging from 0 to 3. Male, English, and Economics are equal to one if the subjects report being male, speaking English as a first language, and majoring in Economics, respectively.

Table 5: Summary statistics

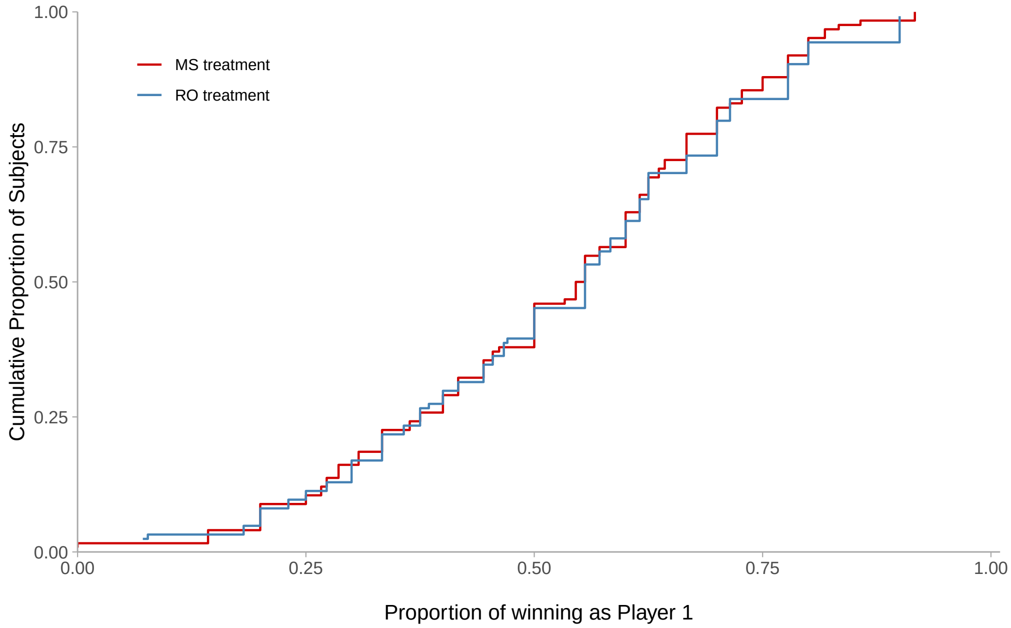

Figure 10: Proportion of winning as Player 1. This figure shows the distribution of subjects' win rates as Player 1 across the two treatments.

Dependent variable: Player 1 Winning		
	LPM (1)	LPM (2)
Intercept	0.5694***	0.4477***
	(0.0437)	(0.0452)
Square	0.1982***	0.0045
	(0.0363)	(0.0420)
Nx2 Grid	−0.0485	−0.1156***
	(0.0400)	(0.0375)
Num. Available Boxes	−0.0044**	0.0003
	(0.0018)	(0.0019)
First Move SPNE		0.4215***
		(0.0414)
Observations	1282	1282
Adjusted R^2	0.0265	0.1116

Note: Standard errors are clustered at the subject level.
Significance: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 6: Effects of initial game state on winning probability.

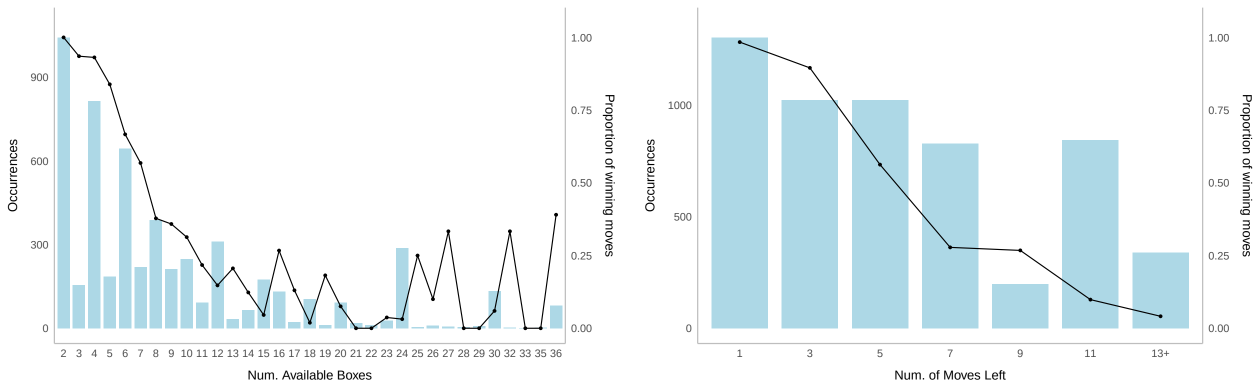

(a) Available Boxes (b) SPNE Moves Left

Figure 11: Relationship between complexity measures and making a winning move. The proportion of moves that are consistent with SPNE is generally decreasing with both complexity measures.

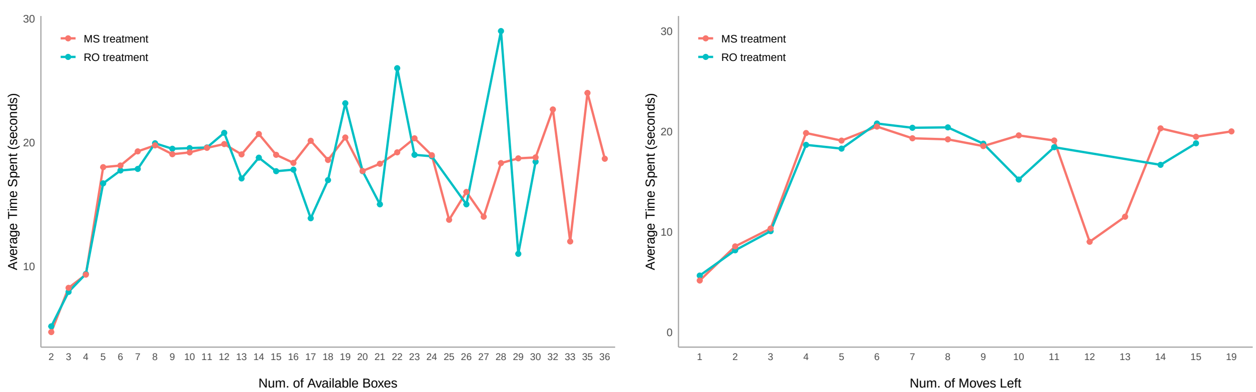

(a) Available Boxes (b) SPNE Moves Left

Figure 12: Average time spent by game state. This figure summarizes the average time spent at different game states, either by number of boxes (panel (a)) or moves left (panel (b)). Subjects spent less time when there are fewer than 4 boxes left and when there are 3 moves left until the terminal state of the game. On average, subjects spent around 20 seconds making a choice.

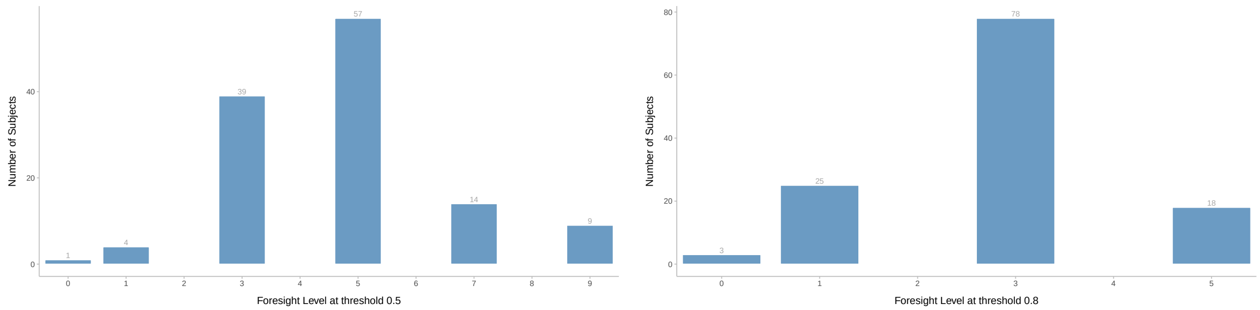

(a) Threshold of 0.5

(b) Threshold of 0.8

Figure 13: Foresight Levels. Panels (a) and (b) show the distribution of Foresight Levels (see text for definition) using thresholds of 0.5 and 0.8, respectively. With a 0.5 threshold, Foresight Levels range from 0 to 9; with a 0.8 threshold, they range from 1 to 5. One subject has Foresight Level 0 at the 0.5 threshold, meaning she chooses optimally less than 50% of the time even when the game is one move from the terminal state.

Dependent variable: SPNE Consistent Move		
	MS	RO
Intercept	0.9317***	0.9349***
	(0.0150)	(0.0194)
5–9 Boxes	−0.5425***	−0.4754***
	(0.0442)	(0.0501)
More than 10 Boxes	−0.7735***	−0.8543***
	(0.0281)	(0.0314)
Match Number	0.0036***	0.0023*
	(0.0010)	(0.0013)
5–9 Boxes × Match Number	0.0127***	0.0063*
	(0.0035)	(0.0033)
More than 10 Boxes × Match Number	−0.0031	−0.0000
	(0.0023)	(0.0021)
Observations	3461	2094
Adjusted R^2	0.4833	0.5103

Note: Standard errors are clustered at the subject level.
Significance: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 7: Learning by number of boxes left. This table reproduces the results from Table 4, using the number of boxes instead of number of moves left as the game's state. The results are consistent with those in the text, with learning concentrated in game states of intermediate complexity.

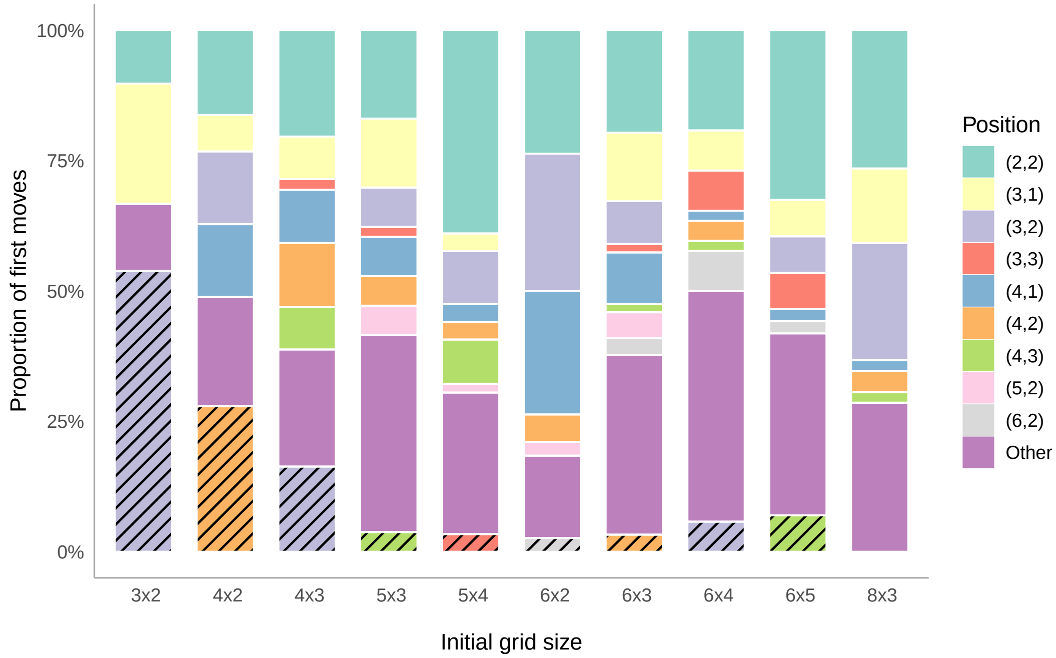

Figure 14: Proportion of initial choice by grid size in RO. Box (2,2) is the most popular initial choice in most grid sizes, excluding grids 3×2 and 4×2. For grid size 8×3, subjects never chose optimal choice (5,2) as their initial choice.

	(1)	(2)	(3)
Intercept	0.4861***	2.7141***	0.6264***
	(0.0664)	(0.8724)	(0.0882)
Male	0.0516**	0.6703**	0.0828**
	(0.0234)	(0.2700)	(0.0337)
Age	0.0000	0.0056	−0.0022
	(0.0028)	(0.0354)	(0.0039)
English	−0.0460	−0.4549	−0.0256
	(0.0279)	(0.2815)	(0.0433)
Economics	0.0044	−0.1128	0.0326
	(0.0270)	(0.2932)	(0.0387)
CRT Score	0.0453***	0.2989***	0.0670***
	(0.0088)	(0.1088)	(0.0132)
Observations	5555	124	1275
Adjusted R^2	0.0149	0.0913	0.0339

Note: Standard errors are clustered at the subject level.
(1) SPNE Consistent Move, (2) Foresight Level,
(3) SPNE Consistent Move in L-shaped Game.
Significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 8: Relationship between demographics and consistency with SPNE. Controlling for all covariates, CRT score and being male are positively and statistically significantly associated with all three consistency measures, while age, speaking English as a first language, and being an economics major are not significantly related to any of them.

B Screenshots from the experiment

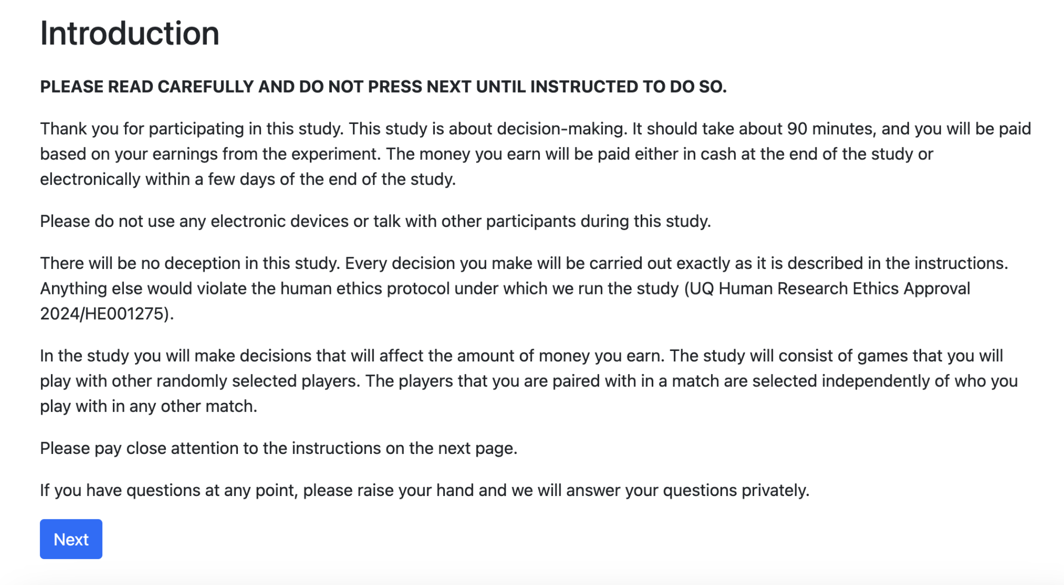

Introduction

PLEASE READ CAREFULLY AND DO NOT PRESS NEXT UNTIL INSTRUCTED TO DO SO.

Thank you for participating in this study. This study is about decision-making. It should take about 90 minutes, and you will be paid based on your earnings from the experiment. The money you earn will be paid either in cash at the end of the study or electronically within a few days of the end of the study.

Please do not use any electronic devices or talk with other participants during this study.

There will be no deception in this study. Every decision you make will be carried out exactly as it is described in the instructions. Anything else would violate the human ethics protocol under which we run the study (UQ Human Research Ethics Approval 2024/HE001275).

In the study you will make decisions that will affect the amount of money you earn. The study will consist of games that you will play with other randomly selected players. The players that you are paired with in a match are selected independently of who you play with in any other match.

Please pay close attention to the instructions on the next page.

If you have questions at any point, please raise your hand and we will answer your questions privately.

Next

Figure 15: Introduction page

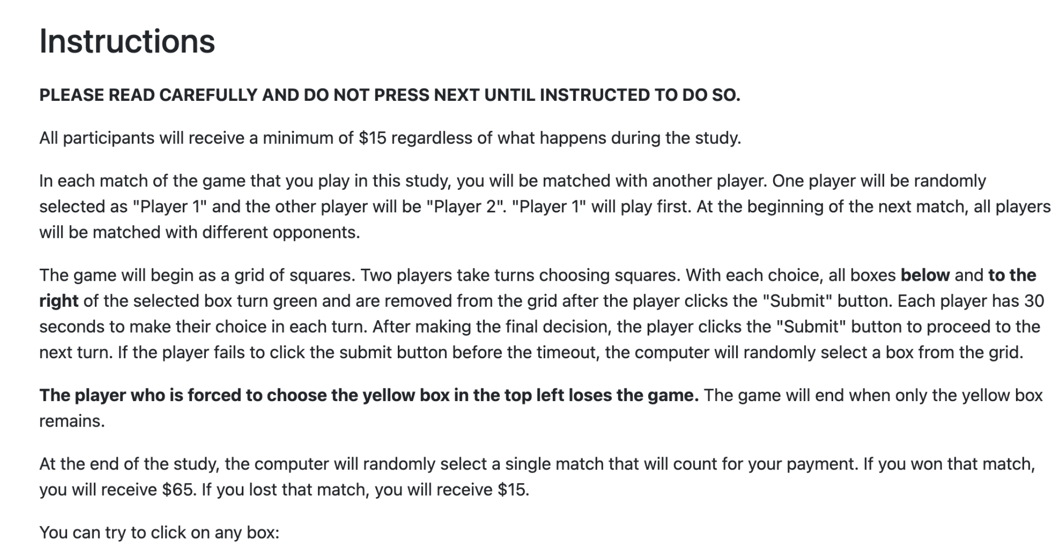

Instructions

PLEASE READ CAREFULLY AND DO NOT PRESS NEXT UNTIL INSTRUCTED TO DO SO.

All participants will receive a minimum of $15 regardless of what happens during the study.

In each match of the game that you play in this study, you will be matched with another player. One player will be randomly selected as "Player 1" and the other player will be "Player 2". "Player 1" will play first. At the beginning of the next match, all players will be matched with different opponents.

The game will begin as a grid of squares. Two players take turns choosing squares. With each choice, all boxes **below** and **to the right** of the selected box turn green and are removed from the grid after the player clicks the "Submit" button. Each player has 30 seconds to make their choice in each turn. After making the final decision, the player clicks the "Submit" button to proceed to the next turn. If the player fails to click the submit button before the timeout, the computer will randomly select a box from the grid.

The player who is forced to choose the yellow box in the top left loses the game. The game will end when only the yellow box remains.

At the end of the study, the computer will randomly select a single match that will count for your payment. If you won that match, you will receive $65. If you lost that match, you will receive $15.

You can try to click on any box:

Figure 16: Instruction page

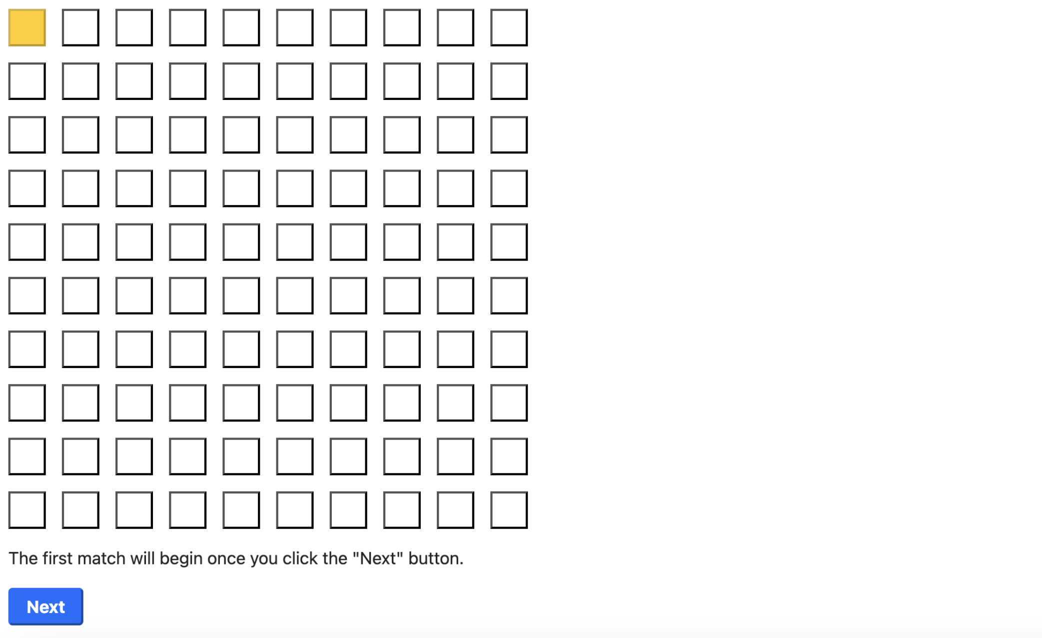

Figure 17: Instruction page (cont)

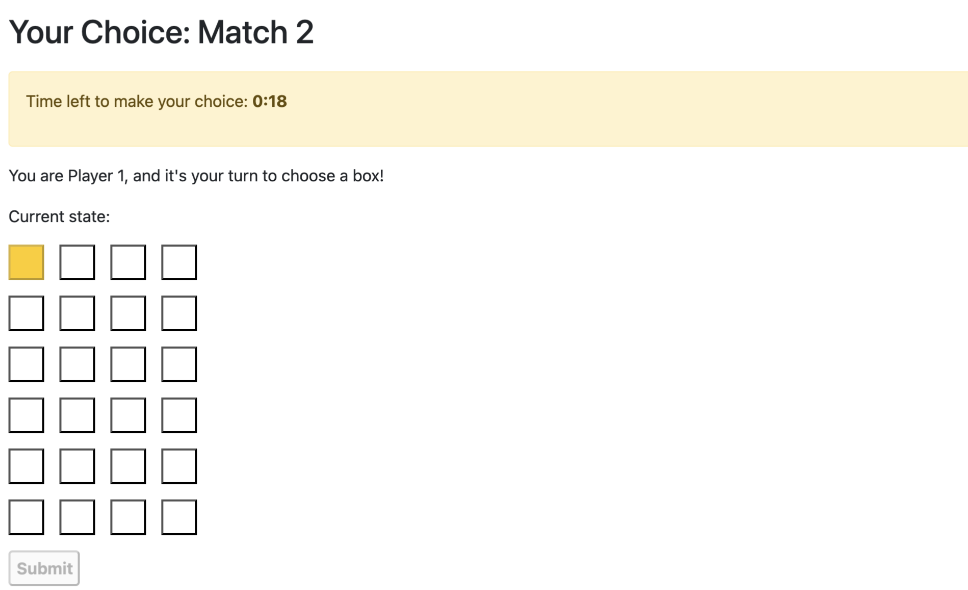

Figure 18: Choice page of Player 1

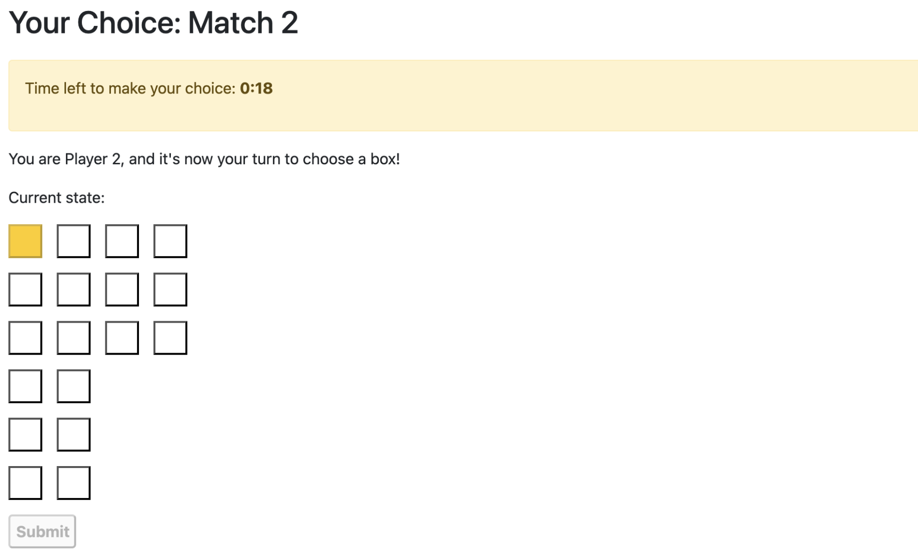

Figure 19: Choice page of Player 2

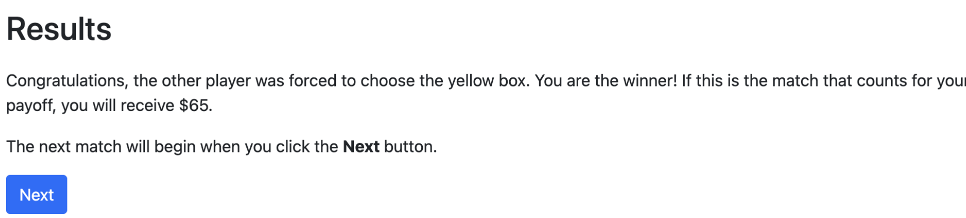

Figure 20: Result page for winner in each match

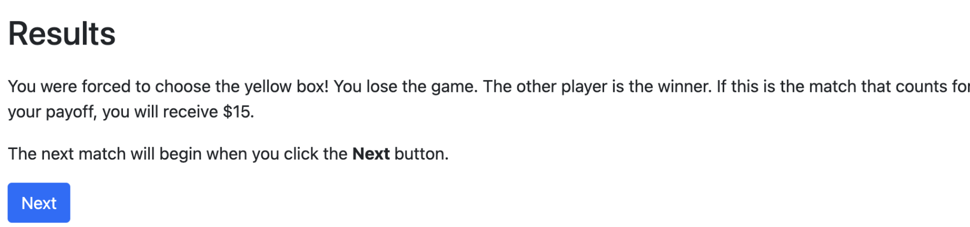

Figure 21: Result page for loser in each match

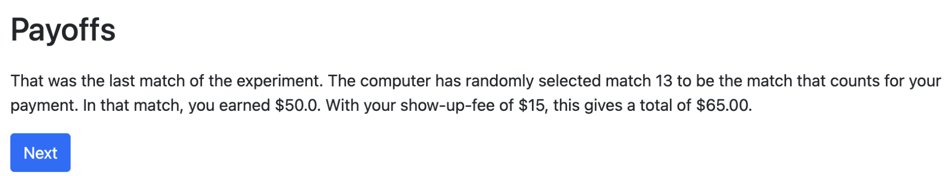

Figure 22: Payoff for winner in the experiment

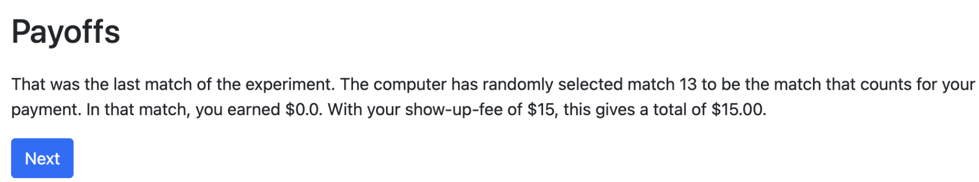

Figure 23: Payoff for loser in the experiment

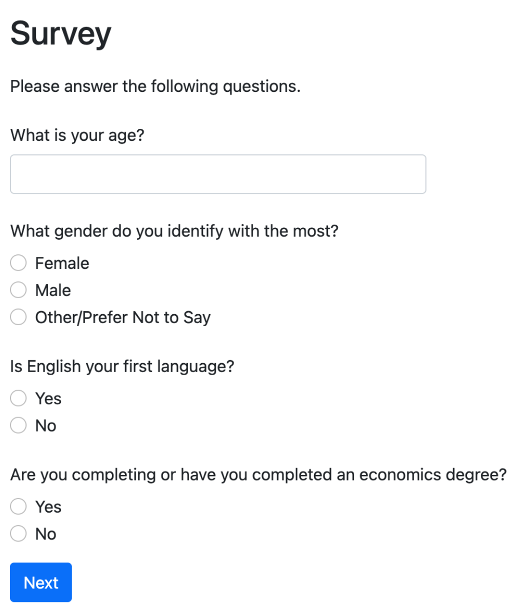

Figure 24: Demographics Page

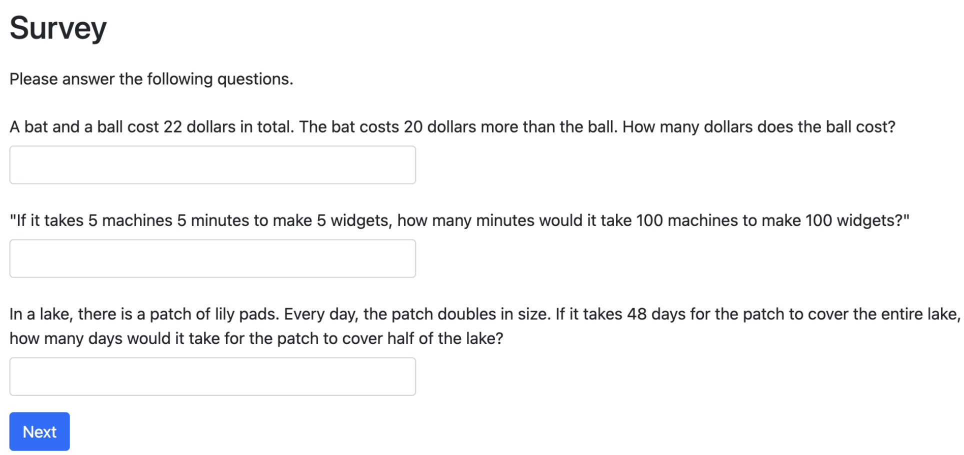

Figure 25: Cognitive Reflection Test

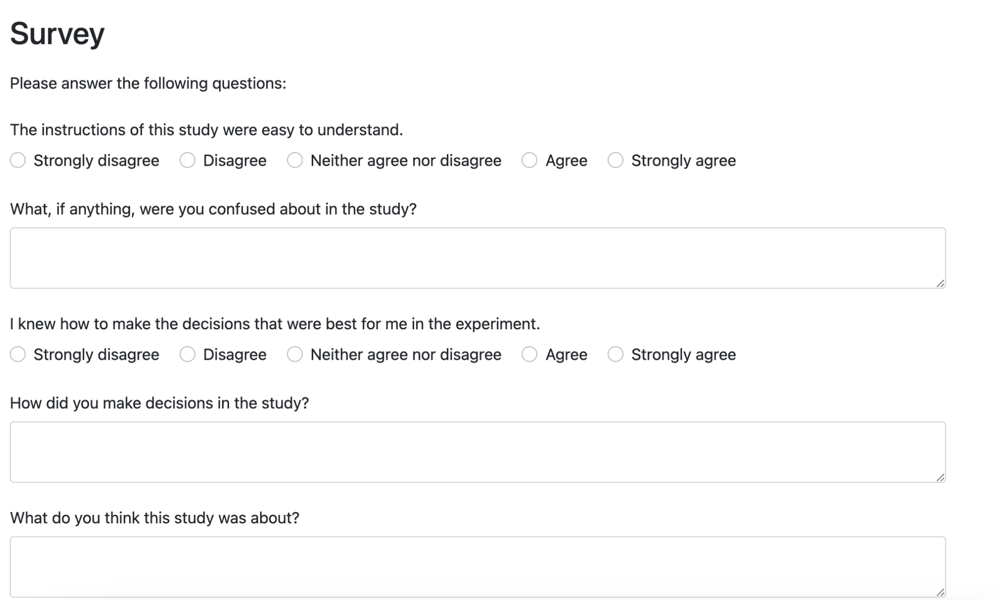

Figure 26: Experiment Feedback Page