Opponent's Foresight and Optimal Choices

Jeevant Rampal*

September 19, 2020

Abstract

Using two experiments, this paper demonstrates that expert players of sequential-move games best respond to their opponents' backward-induction ability. In particular, I show that these experts take advantage of inexperienced opponents' weakness in backward induction. I find this when the expert is explicitly told that her opponent is inexperienced, but also when she infers the opponent's weakness from the opponent's preceding performance. I demonstrate that other-regarding preferences have no role in these findings. I find that a novel model of limited foresight and uncertainty about the opponent's foresight fits the data better than Level-k or Quantal Response models.

Keywords: Backward induction; Limited Foresight; Level-k; Quantal Response.
JEL Codes: D91, C91, D80, C72.

1 Introduction

Understanding optimal choices in sequential-move games or interactions is clearly of fundamental importance in economics, if not all social sciences. Ideally, each player should perform backward induction to take optimal decisions in sequential-move games. Backward induction entails calculating the optimal choice at each decision point by accounting for all possible future consequences of each possible current action. However, even with high stakes, decision-makers often choose sub-optimally when decisions have distant consequences, e.g. financial decisions while planning for retirement (Lusardi and Beeler (2007)). Laboratory experiments have identified several types of problems faced by novice players in backward induction (Johnson et al. (2002); Binmore et al. (2002)). Indeed, among other findings, Dufwenberg, Sundaram, and Butler (2010) show that it can take quite a bit of game-specific experience to understand backward induction for just that game.

*Indian Institute of Management Ahmedabad, Gujarat-380015, India. Email: jeevantr@iima.ac.in. This study was funded by the Decision Sciences Collaborative, and the Journal of Money Credit and Banking, Department of Economics, Ohio State University (OSU). OSU-IRB study number: 2015B0087. I thank James Peck and John Kagel for their advice and encouragement. I also thank Puja Bhattacharya, Jason Blevins, Mukesh Eswaran, PJ Healy, Arkady Konovalov, Dan Levin, Dilip Mookherjee, Kirby Neilsen, Antonio Penta, and Debraj Ray for their comments. Last, I thank seminar participants at the Econometric Society meetings, Midwest Economic Association, and Experimental Science Association conferences for their comments.

So situations where opponents have heterogeneous backward-induction expertise can arise simply due to the opponents being differently experienced with the sequential-move game in question, e.g. competing firms with differently experienced owners. Thus, two questions central to this paper arise: can we identify the effect of opponent's expertise in backward induction on experienced players' optimal strategy? Which theoretical model should we use to understand interactions among players with heterogeneous expertise in backward induction? This paper attempts to answer both.

I identify the effect of the opponent's backward-induction ability on optimal strategy, and I search for a model to explain this intuitive finding. The intuition is immediate: most of us would bet more money on winning a tic-tac-toe match if we knew that the opponent has never played it before. But surprisingly this intuition hasn't been identified empirically, possibly because it is hard to rule out other-regarding preferences as a confounding explanation of the relevant data. For example, in the Centipede game (Rosenthal (1981)) the belief that the opponent is altruistic/cooperative (e.g. see McKelvey and Palfrey (1992)) can have the same implication for one's optimal strategy as the belief that the opponent is weak in backward induction. This intuition has also not been explained in terms of a theoretical model. Clearly, any equilibrium concept that incorporates sequential rationality, including subgame perfect equilibrium (henceforth SPE), doesn't work since in these concepts each player assumes that the opponent performs backward induction flawlessly. Thus, one has to look to behavioral models to explain the intuition; however, the answer to "which model" is unclear from the literature.

I identify the effect of opponent's backward-induction (BI for short) ability on the optimal strategy by randomly pairing a mixed pool of experienced and inexperienced players for playing a specially designed winner-take-all game. I find that the experienced players switch strategies based on signals about the opponent's BI ability. In particular, I find that the signal indicating that the opponent is weak in BI causes experienced players to deviate from the unique sure-win strategy to try to win from a losing position, which entails a greater reward but is impossible unless the opponent makes a mistake in BI. I run two experiments, which use two different signals about the opponent's BI ability. In Experiment 1, the signal is information specifying if the randomly matched opponent is experienced or inexperienced. In Experiment 2, the signal is information specifying the opponent's choices and outcome in an immediately-preceding similar game. The use of winner-take-all games and choice-data analysis helps me rule out other-regarding preferences as a confounding explanation. The second aim of this paper is to find an explanation for the dependence of optimal strategy on the opponent's BI ability. For this, I perform an MLE comparison that shows that leading behavioral theory models have a lower likelihood with respect to the experimental data than a novel model that incorporates "limited foresight" and uncertainty about the opponent's foresight.

At the heart of the experimental designs in this paper is a modified *race game* I call *Avoid 9*, which is a two-player alternate-move game with perfect information.[1] At each move, the player

[1]Race games are "games of Nim." Race games have also been studied by Levitt, List, and Sadoff (2011), Dufwenberg, Sundaram, and Butler (2010) and Gneezy, Rustichini, and Vostroknutov (2010); the last two papers study how subjects learn to play these games.

moving there must remove 1, 2, or 3 items from an imaginary box containing 9 items. The game proceeds until one player has to remove the last (9^{th}) item; the player who removes the last item loses, while the other player wins. Note that there is a second-mover advantage: the second-mover can remove (4 minus the opponent's previous choice) at each move and guarantee herself a win. The winner gets a prize, the loser gets nothing. Thus, this game is a winner-take-all game with (weakly) dominant strategies. Notice that in the Avoid 9 game, no expert will deviate from the backward induction solution, no matter his beliefs about the opponent's strategy (due to dominant strategies) and for almost all conceivable other-regarding preferences he may have (due to the winner-take-all property).[2]

I make two design innovations to the Avoid 9 game to investigate if an expert may want to take advantage of an opponent's weakness in backward induction. First, to provide an *incentive to take advantage*, I make the prize for winning as the first-mover equal to 500 experimental currency units (ECU), and the prize for winning as the second-mover as 200 ECU. Second, to provide an *opportunity to take advantage*, I introduce a preceding stage to decide the first and second mover in the subsequent subgame. In this preceding stage, both players simultaneously report if they would like to be the first-mover or the second-mover. One of the players' decisions is selected at random and implemented. Henceforth, the Avoid 9 game will be understood as the game that includes the first/second mover decision stage and the different prizes from winning as the first or second mover.

The design of Experiment 1 is to "train" half the subjects, where training entails 12 rounds of experience of playing the Avoid 9 game with each other with stranger re-matching, then bring in the other half of subjects who are inexperienced, and randomly form pairs (again using stranger re-matching) each round among this combined pool of subjects to play the Avoid 9 game for eight rounds. In each round, each subject is told if her opponent is experienced or inexperienced. I find that the training of experienced subjects works; they choose the "sure-win" strategy against experienced opponents, i.e. they opt for the second-mover position and make no mistakes from winning positions.[3] The key finding from Experiment 1 is that when experienced subjects face an inexperienced opponent, they are significantly more likely to choose the first-mover position, i.e. actively risk a loss, to try to attain the higher prize of 500 by winning from a losing position, which is impossible unless the opponent makes a mistake as the second mover. Since the game is winner-take-all, which means cooperative outcomes are impossible, and experienced subjects never fail to convert a second-mover opponent's mistake into a win for themselves, I can confidently rule out any role for other-regarding preferences in the key finding described above.

In Experiment 2, conducted with different subjects and other minor changes, I replace giving the experienced subjects information about the opponent's experience-level with giving them information only about the opponent's moves and outcome in a preceding game. The preceding game is similar to the Avoid 9, but it is played against a computer. The preceding game is such that if a player loses it, then that is a clear signal of the player's weakness in the relevant back-

[2] Unless he values the opponent's payoff more than his own.

[3] This strategy is sure-win only if one's choice to be the second mover is selected in the first/second mover decision stage.

ward induction. I find that in Experiment 2, experienced subjects are more likely to opt for the first-mover position against an opponent who lost (rather than won) against the computer in the preceding game. Thus, I report the novel finding that experienced subjects infer and take advantage of the opponent's weakness in backward induction by only observing his/her past performance in a similar game. Since the related literature (including Experiment 1) only has evidence from experimental designs where players are informed about the opponent's expertise/experience-level by the experimenter, Experiment 2 extends the external validity of these findings to settings where only information about the opponent's preceding choices and outcome is available, and it is commonly known that there some inexperienced players present in the population.

Overall, the experimental results partly reconcile the findings of Palacios-Huerta and Volij (2009) and Levitt, List, and Sadoff (2011), henceforth P-HV and LLS. In the context of the centipede game, P-HV show that subjects switch strategies when facing an undergraduate student rather than a chess expert. P-HV conclude that this result is driven by the "assessment of opponent's rationality rather than altruism or other social preferences" (P-HV, page 1634). However, since LLS find little correlation between subjects' behavior in winner-take-all and centipede games, they "question the validity of using centipede games to draw inferences about backward induction" (page 988). Indeed, it is unclear how P-HV rule out the alternative explanation that subjects expected students to be more altruistic/cooperative than chess experts. In this paper, I rule out these alternative explanation by showing a P-HV-like result (optimal strategy depends on opponent's BI ability) in winner-take-all games that partially fit as games deemed "more appropriate tests of backward induction" by LLS (page 988), and where cooperative outcomes are impossible.[4]

The second aim of this paper is to identify a theoretical explanation for the experimental findings. To this end, I perform an MLE comparison between Level-k models (I test Ho and Su (2013) and Camerer, Ho, and Chong (2004)), the Agent Quantal Response Equilibrium (AQRE: McKelvey and Palfrey (1998)) model, and the novel Limited Foresight Equilibrium (LFE: Rampal (2020)) model. The results favor the LFE explanation: a high proportion of inexperienced players had limited foresight (i.e. they could only consider and understand a limited number of future stages at each move), and players with no foresight limitation (the experienced players) were aware of this heterogeneity in foresight within the subject population, and they best responded to it when they observed that the opponent was experienced/inexperienced (Experiment 1) or that the opponent had won/lost against the computer (Experiment 2).

As the notion of foresight is close to BI ability, the MLE results corroborate the claim that subjects were best responding to opponent's BI ability. By contrast, the Level-k and AQRE explanations would have been based on beliefs about the opponent's cognitive ability and payoff-weighted errors, respectively. The AQRE model has a worse fit since it provides no basis for strategic adjustment based on opponent's preceding choices or outcome, as required in Experiment 2. The key

[4]For appropriate tests, LLS use constant-sum winner-take-all games, where the backward induction strategy is also dominant. To give incentive and opportunity to take advantage of a weak opponent (as discussed above), the Avoid 9 game, while a winner-take-all game, is not constant-sum, and the backward induction strategy is dominant only after the first and second movers are decided.

challenge for Level-k models is the observed *pattern* of mistakes, where a mistake is a dominated choice, e.g. removing 3 items with 7 left. In both experiments mistakes reduced as items-left reduced, a pattern strongly indicative of limited foresight. No beliefs justify mistakes, which means only a Level-0 player can make mistakes, but the standard Level-0 doesn't make mistakes in the observed pattern.

Among other related literature, Mantovani (2016) is an independent working paper which uses a different modified version of a race game to demonstrate limited backward-induction ability, and contrast this with the Level-k model. In Mantovani (2016) the theoretical model and experimental design focus on limited backward-induction ability, but not heterogeneity therein, which is the focus in this paper. Le Coq and Sturluson (2012) and Dufwenberg, Lindqvist, and Moore (2005) study the effect of mixing experienced and inexperienced subjects in specific dynamic games, however they don't explore heterogeneity in backward-induction ability or its theoretical reasons.

The effects of combining experienced and inexperienced subjects in simultaneous move games have been studied extensively, for instance by Alaoui, Janezic, and Penta (2020), Alaoui and Penta (2016), Gill and Prowse (2016), and Agranov et al. (2012). Alaoui, Janezic, and Penta (2020) use a *tutorial* and *replacement* method to disentangle the effects of belief about the opponent's cognition and own cognition. It is worth noting that for the games I test, providing heterogeneous experience-levels induces heterogeneous sophistication, and tutorials aren't required. So, the present experimental setting is comparable to those "real-world" sequential-move games where past experience leads to better understanding of the relevant backward induction, players have heterogeneous experience-levels, and cues about the opponent's BI ability are available, like information about their experience or previous performance in similar situations.

2 Experiment 1 Design

Experiment 1 was conducted over 13 sessions with 154 subjects at The Ohio State University's experimental economics laboratory using the laboratory's subject pool. The sessions were conducted using zTree (Fischbacher (2007)). In each session, between 8 and 18 subjects participated. Each subject participated in only one session. A session lasted 62 minutes on average with an average payment of USD 14.45. Experiment 1 revolved around the *Avoid 9* game, whose rules are as follows.

The Avoid 9 game is an alternate-move two-player game. There are 9 items in an imaginary box. At every move, each player can choose to remove 1, 2, or 3 items from the box.[5] The player who removes the 9^{th} item loses, and his opponent wins. In the experiment, the Avoid 9 game had a preceding stage to decide the order of moves (henceforth, it is understood that the Avoid 9 game includes this stage). In this stage, both players simultaneously choose between "First Mover" and "Second Mover," i.e. if they want to be the first-mover or second-mover in the subsequent (sub)game. One of the players' decisions is selected at random and implemented. After it has been decided who the first-mover is, the two players choose alternately, until one wins. The

[5] If at a move the number of items left in the box is 2 or 1, then the maximum number of items that a player can remove at that move is 2 or 1, respectively.

players' first/second mover choices are not revealed, but all other prior choices (of number of items removed) in each round are common knowledge. If a player wins as the first-mover, his payoff is 500 ECU. If a player wins as the second-mover, his payoff is 200 ECU. If a player loses, he gets 50 ECU. The conversion rate used was 60 ECU = 1 USD.

The design of Experiment 1 was as follows. Each session comprised of two sub-sessions: first, *training sub-session*, and subsequently, *combined sub-session*. First, subjects were split into two types: **Experienced (Exp)** and **Inexperienced (Inexp)**, with Inexp subjects being at least 50 percent of the total subjects in any session.

1. The ***training sub-session*** was conducted as follows.

(a) **Exp subjects trained** (74 Exp subjects overall). **Exp** subjects were given the instructions about the Avoid 9 game while the Inexp subjects were not in the lab. Then, Exp subjects played 12 rounds of the Avoid 9 game among themselves, with stranger re-matching every round. Note that in this training sub-session, the designation "Experienced" is only a label since these subjects were completely new to the Avoid 9 game.

(b) **Inexp subjects** (80 Inexp subjects overall) **were kept uninformed.** They were not told about the Avoid 9 game in this first sub-session. The Inexp subjects were kept uninformed about the Avoid 9 game by making them play an unrelated bargaining game. It was announced to all Exp and Inexp subjects that the other set of subjects was playing an unrelated game.

2. In the second sub-session, i.e. the ***combined sub-session***, Exp and Inexp subjects were mixed and the Avoid 9 game was explained to all subjects (repetition for Exp subjects). It was announced to all subjects that Exp subjects had played 12 rounds of the Avoid 9 game in the first sub-session while Inexp subjects had played an unrelated game. Then, Exp and Inexp subjects played eight rounds (round numbers 13-20 of the session) of the Avoid 9 game together. The subjects were stranger re-matched into pairs before the beginning of each round. At all points during the Avoid 9 game, each player was told if his/her opponent was Exp (referred to as *type S* in the instructions) or Inexp (referred to as *type D* in the instructions).[6]

3. The earnings of a subject were: earnings in a randomly drawn round of the first sub-session plus earnings in a randomly drawn round of the second sub-session. At the end of the second sub-session, all earning details were displayed and payments were completed.

All aspects of the design were publicly announced to all subjects; only the rules of the Avoid 9 game and the bargaining game were withheld from the Inexp and Exp subjects, respectively, in the training sub-session.

In the Avoid 9 game, a player is said to be in a "winning position" if from that position (position means number of items left) she has a "sure-win" dominant strategy available that guarantees a win, regardless of the opponent's strategy. From a "losing position," any available action puts one's opponent in a winning position. The winning positions are $\{8, 7, 6, 4, 3, 2\}$, and the losing positions are $\{9, 5, 1\}$. Note that the second-mover has a sure-win strategy available: he can choose (4 minus

[6]While playing each Avoid 9 game, subjects moved sequentially after the first/second mover decision stage. Each subject at each move was given a clock with 45 seconds on it to remind them to move. The clock could only flash if the time taken was more than 45 seconds, and the game did not proceed without the subject's choice.

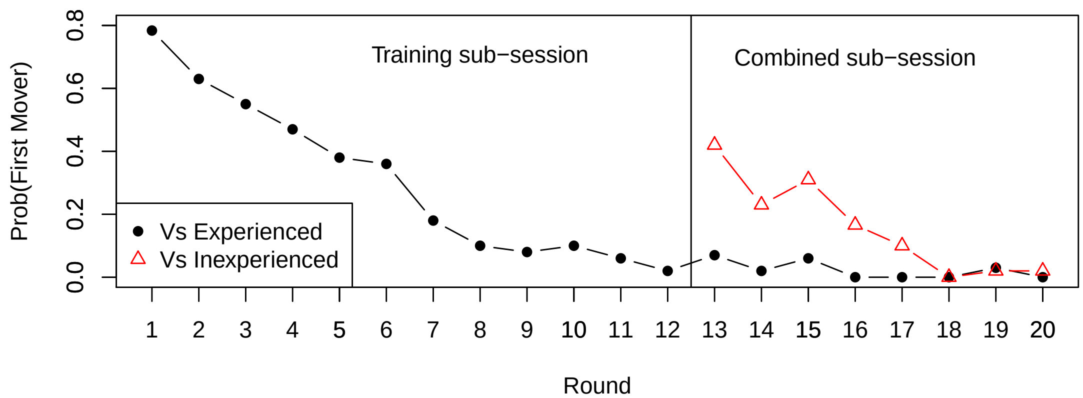

Figure 1. Experienced Subjects' Behavior Based on Opponent's Experience
Notes: The figure depicts the round-wise proportion of experienced subjects who chose "First Mover" when matched with an experienced opponent (black dots) and the proportion of experienced subjects who chose "First Mover" when matched with an inexperienced opponent (red triangles). The latter proportion is significantly higher, with a p-value < 0.015, for each of the rounds 13 through 16.

the opponent's previous choice) at each move to remove the 8^{th} item, forcing the opponent to remove the 9^{th} item. However, if the second-mover fails to put the opponent in a losing position after one of his moves, then the first-mover is in a winning position. The subgame perfect equilibrium (SPE) strategy for each player is: choose "Second Mover" (henceforth S) and then remove 3, 2, and 1 items from positions in $\{8, 4\}$, $\{7, 3\}$, and $\{6, 2\}$, respectively.[7] This strategy places no restriction on actions from a losing position. I label this strategy as the "perfect" strategy.

The perfect strategy is *not* a dominant strategy since there is an extra incentive to win as the first-mover. In particular, if a risk-neutral rational player believes with probability at least $\frac{1}{3}$ that his opponent will deviate from the perfect strategy as the second-mover, then choosing "First Mover" (henceforth F) is optimal. Let the subgame after the first and second movers are decided (but with 9 items still remaining) be labeled as $A9_{sub}$.[8] Note that in $A9_{sub}$, the perfect strategy *is* (weakly) dominant because the payoff from winning/losing is already decided, and the perfect strategy is a sure-win strategy from a winning position. I refer to a deviation from the perfect strategy in $A9_{sub}$ as a "mistake."

3 Experiment 1 Results

The central question in Experiment 1 is: in the combined sub-session, do Exp subjects infer their opponent's weakness in backward induction from the opponent's inexperience and take advantage? But first, I investigate the effectiveness of the training sub-session for Exp subjects.

[7] Regardless of the history of play in getting to that position.

[8] Strictly speaking, $A9_{sub}$ is not a subgame since there is imperfect information from the opponent's F/S decision not being revealed. I refer to $A9_{sub}$ as a subgame for labeling, and because the F/S choices don't matter once the order of moves is decided since each player then has a unique weakly-dominant strategy.

Result 1(a): *By the end of the training sub-session, almost all Exp subjects had understood the perfect strategy and their choices were as if they were convinced that a sufficiently high proportion of other Exp subjects understood the perfect strategy.*

Subjects can be tempted to choose F ("First Mover") until they gain understanding of the perfect strategy or until they are sufficiently convinced about their opponent's understanding of the perfect strategy. In the training sub-session (rounds 1-12) the proportion of Exp subjects who chose F declined significantly (p-value$<$ 0.01) between round 1 (78.4 percent) and round 12 (2.7 percent, which is insignificant). Further, there was a significant reduction (p-value$<$ 0.01) in the proportion of subjects who made a mistake as the second-mover: 36.5 percent in the first two rounds, to 5.4 percent in the last two rounds of the training sub-session.[9] Gneezy, Rustichini, and Vostroknutov (2010) find that in almost identical games (except for the F/S decision stage and different prizes), the experts understand the perfect strategy using backward analysis, which is a process similar to backward induction, except that the experts have to expend cognitive effort to check if the replacement of subgames with the appropriate SPE payoff profile is valid each time they approach a subgame. Thus, I interpret a demonstration of understanding the perfect strategy as a demonstration of understanding the relevant backward induction.

Result 1(b): *When experienced subjects chose "First Mover" in the combined sub-session, it was not due to any other-regarding preference, instead it was an attempt to attain the higher payoff which is possible only when the opponent makes a mistake.*

Result 1(b) is supported by the following fact. In the combined sub-session, if an experienced subject was the first mover and got into a winning position due to the opponent's mistake, then in 100 percent of such cases (15 in total) the experienced subject ensured that he/she won.[10]

Result 1(c): *Experienced subjects were significantly more likely to opt for the first-mover position when they were informed that the opponent was inexperienced, as opposed to experienced. Note that in opting for the first-mover position, the experienced subjects risk a loss to try for the higher payoff, attainable only if the opponent subsequently makes a mistake.*

In the combined sub-session (rounds 13-20), experienced subjects were significantly more likely to choose F against an inexperienced opponent than an experienced opponent. By Result 1(b), one can interpret the choice of F as an attempt to attain the higher payoff, which entails the risk of a loss in case the opponent plays perfectly. To see Result 1(c), notice in Figure 1 that the proportion of experienced subjects who chose F when matched with an inexperienced opponent

[9] The two p-values are obtained using Probits with choice F and "mistake" as the respective dependent variables and appropriate round dummies as independent variables. Errors are clustered at subject level.

[10] One can argue that the experienced subjects chose "First Mover" to selflessly maximize joint payoffs (i.e. to selflessly maximize efficiency). Such an explanation is not supported by the data since to maximize joint payoffs, the experienced players could have deliberately lost when selected as the second mover, giving 500 to the opponent and 50 to self; but in the combined sub-session, no experienced subject ever lost from a second-mover position.

Dependent Variable: Prob(First Mover)		
	Exp	Inexp
	(1)	(2)
Opponent is Inexp	1.18***(0.31)	0.12(0.13)
Round	-0.3***(0.07)	-0.43***(0.04)
Constant	2.65***(0.85)	6.2(0.67)
No. of Obs.	592	640
Pseudo R^2	0.3154	0.3273
Session Dummies	Yes	Yes

Table 1. Factors Influencing Probability of Choosing "First Mover"

Notes: * p <0.10, ** p <0.05, *** p <0.01. Numbers in parentheses are robust standard errors clustered at subject-level. For model (1), the dependent variable is Experienced subjects' choice of "First Mover" (dependent variable takes value 1) or "Second Mover" (dependent variable takes value 0). The sample is 74 Exp subjects across rounds 13-20, i.e. the combined sub-session. For model (2), the dependent variable is Inexp subjects' probability of choice of "First Mover" (takes value 1) or "Second Mover" (takes value 0). The sample is 80 Inexp subjects across rounds 13-20.

(red triangles) is greater than the proportion of experienced subjects who chose F when matched with an experienced opponent (black dots). Round-wise tests for difference in proportions show that this difference is significant for each of rounds 13-16 (p-value ≤ 0.015 for each round).

Next, consider the probit results for 74 Exp subjects and 80 Inexp subjects' F/S choice data from rounds 13-20, the combined sub-session. Model (1) in Table 1 reports the results from a probit estimation with the Exp subjects' choice of F/S in the Avoid 9 game as the dependent variable. The dependent variable took value 1 if an Exp subject chose F, and 0 if he chose S. So, a positive coefficient on an independent variable means that a higher value of the independent variable increased the probability with which the average player chose F. The independent variables are: (i) a dummy for if the opponent was Inexp (value 1) or Exp (value 0); (ii) round variable; (iii) session dummies. Table 1, model (2), marked with Inexp, reports the results from the same probit estimation done with the Inexp subjects' choice of F/S in the Avoid 9 game as the dependent variable. In both models, errors are clustered at the subject-level; clustering the errors at session-level doesn't change the sign or significance of any of the results in this paper.

Consider model (1). The highly significant coefficient on the dummy variable for "Opponent is Inexp" (p-value < 0.001) indicates that Exp subjects were significantly more likely to choose F when playing against an Inexp opponent than an Exp opponent. Thus on average, when experienced subjects are informed that the opponent is inexperienced, they infer that the opponent is weak in the required backward induction, which means that the opponent has a high chance of making a mistake. So the experienced subjects try to take advantage of this.

The highly significant negative coefficient of the round variable seems to be driven by an increase in Exp subjects' belief that the Inexp opponent understood the perfect strategy, at-least as the second-mover. This is reflected in the downward sloping red line (with triangles) in Figure 1.

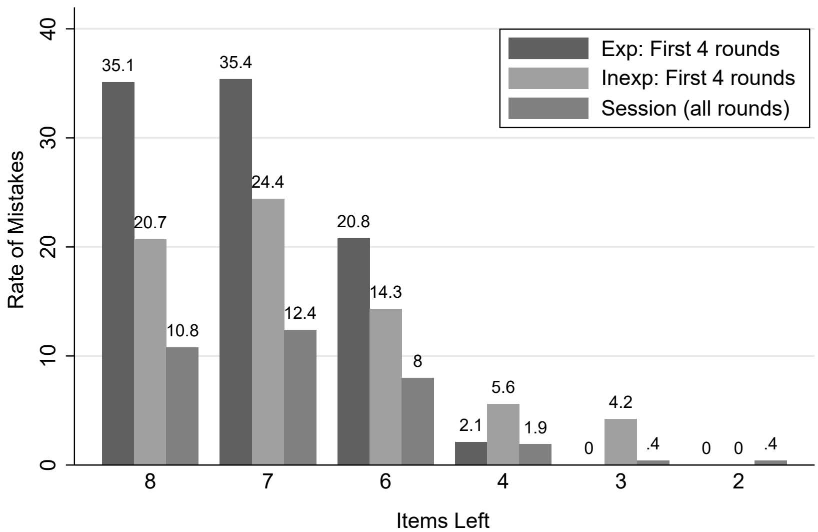

Figure 2. Rate of Mistakes by Items Left
Notes: For each winning position, the numbers show the percentage of actions that were mistakes.

Indeed, the data reveals that when facing Inexp opponents, it was optimal for a risk-neutral Exp subject to be the first-mover in the first two rounds of the combined sub-session (average earning as the selected first mover and second mover was 216 and 194, respectively), but over subsequent rounds, being the second-mover paid more.

Model (2) shows that Inexp subjects also increased the likelihood of choosing F when they faced an Inexp opponent, but this increase, measured by the coefficient on "Opponent is Inexp," was statistically insignificant. The round variable's highly significant negative coefficient (p-value$<$0.001) is consistent with Inexp subjects learning the perfect strategy with repetition.

Result 1(d): *The rate of mistakes declined significantly as item-left reduced. That is, as subjects got closer to the end of the Avoid 9 game, the rate of mistakes reduced.*

Figure 2 shows the observed rates of mistake by items left in the box for three parts of the data: the first four rounds of Exp (rounds 1-4 of the session) and Inexp (rounds 13-16 of the session) separately, and the overall session (training and combined sub-sessions together).[11] In each of these, the rate of mistakes is significantly more (p-values$<$0.01 in a Wilcoxon matched-pairs sign rank test) when 6-8 items are left, as compared to when 2-4 items are left. Overall, Result 1(d) replicates findings from all studies mentioned earlier that have tested similar race games.

Result 1(d) is a strong indicator of "limited foresight" among subjects, where the foresight of a player is defined as the number of subsequent stages that the player considers and understands from each move. In particular, a player with limited foresight cannot reason for the stages that lie beyond his limited foresight; he must do backward induction within his foresight horizon (see Section 6 and Rampal (2020) for more details). So if a player has limited foresight, then as she

[11] I don't report mistakes from losing positions because no action is a mistake at those positions.

gets closer to the end of the game (with fewer items left), her foresight extends to the end of the game, and she can understand the dominant strategy using backward induction and play perfectly. While Result 1(d) is an indicator of limited foresight, I reserve statements about *which model fits best* until a comparative MLE exercise (Section 6).

Result 1(e): *At losing positions, subjects were significantly more likely to remove 1 item rather than 2 or 3 items.*

With 9 items left or 5 items left - both losing positions where the perfect strategy is silent on which action should be chosen - subjects were significantly more likely to remove 1 item than the next most likely alternative (both p-values < 0.001 in tests of difference in proportions). This was also true for Exp subjects in the combined sub-session (see Table 5(a) in Appendix A). The predilection to remove 1 item was also present with 13 items left (in the Avoid 13 game in Experiment 2, discussed later: see Table 5(b) in Appendix A). All together, this appears to be evidence of the belief that keeping the opponent as far away from the end of the game as possible maximizes the chance that he will make a mistake. This intuition does work in the LFE model since it allows for uncertainty and beliefs about the opponent's foresight. With 9 items left (13 items, respectively), a player with foresight of 1-stage-ahead (3-stages-ahead) is more likely to make a mistake with 8 (12) items left rather than 6 or 7 (10 or 11). Thus, an expert with no foresight-limitation will choose to remove 1 item from 13 or 9 items left if he thinks there is a high enough chance that his opponent has such limited foresight. However, no model studied in this paper explains the tendency to remove 1 item when 5 are left. I discuss this more after the MLE exercise.

4 Experiment 2 Design

Experiment 2 is designed to investigate the following. First, in the presence of inexperienced subjects in the group, can experienced subjects infer their opponent's weakness in the relevant backward induction after observing preceding outcomes, without being explicitly informed about the opponent's experience level? And second, do experienced subjects take advantage of the inference that the opponent is weak in the relevant backward induction? In other words, Experiment 2 investigates whether the validity of Experiment 1 findings can be extended to scenarios where experts must infer the opponent's weakness in backward induction from their previous moves and/or outcome.

Experiment 2 comprised of 8 sessions. Each session contained between 8 and 18 subjects. A session lasted 64 minutes on average. The average payment made to the subjects was USD 14.80. Each experimental session used two games. First, the *"Computer 13"* game (C13 for short). Second, the *Avoid 13* game. The Avoid 13 game is exactly like the Avoid 9 game except there are 13 (not 9) items in the box. Otherwise, Avoid 13 is identical to Avoid 9: the first-mover/second-mover decision stage is identical; players move alternately, removing 1, 2, or 3 items until someone removes the last item and loses; payoffs are 50 ECU for a loss and 500 ECU (respectively 200 ECU) for winning as the first (second) mover.

The Computer 13 game The C13 (Computer 13) game is played by a human subject against his/her perfectly-playing computer. The subjects are told that "the computer plays perfectly to win."[12] There are two other key differences in C13 compared to Avoid 13: *first*, in the C13 game, *the human subject unilaterally decides who the first-mover will be - she or the computer.* After the first-mover is decided by the human subject, she and the computer move alternately, choosing 1, 2, or 3 items to be removed from an imaginary box that contains 13 items. The human player wins if the computer removes the 13^{th} item. The *second* key difference in the C13 game is that a player earns a flat 500 ECU for a win (the earning from a loss is 50 ECU as before). That is, *there is no extra-incentive to win as the first-mover.*[13] Therefore, it is a strictly dominant strategy to choose the second-mover position and then play "4 minus computer's previous choice" at each move. Thus, a win in C13 is a strong signal that a player understands its backward induction, and a loss in C13 is a clear signal that a player may have a weakness in the necessary backward induction.

Now I describe how the Experiment 2 design. Each round of Experiment 2 has two parts: play C13 against your computer, then play Avoid 13 with a randomly matched human opponent. The key design feature of Experiment 2 is: I do not inform Exp players about the opponent's experience-level (as is done in Experiment 1), instead, in each round, I only show them the move-history and outcome from opponent vs C13 of that round, just before Avoid 13 begins. Doing so, I want to investigate whether Exp players can infer the opponent's understanding of the relevant backward-induction by observing the opponent's play in C13, and if they use this inference to adjust their optimal F/S decision in Avoid 13.

The reader may be concerned about why Experiment 2 has two changes relative to Experiment 1: Avoid 9 is changed to Avoid 13, and rather than providing information about the opponent's experience level, information about the opponent's performance in C13 is provided. On this, it is worth noting that Experiment 2 is not meant to be compared to Experiment 1. That is, the aim is not to identify the effect of one of those changes. I use 13 items instead of 9 to further widen the gap between subjects who understood the backward induction and those who did not understand it, and to make this gap persist for more rounds. These features help in better distinguishing among competing theoretical models described later. Another reason for changing to 13 items was to give more confidence to Exp players that Inexp players who don't understand the backward induction will have a very low chance of playing the perfect strategy in C13. This was important since Exp players were relying only on their inference to identify the Inexp players from their performance in C13, and not an explicit signal from the experimenter.

Now, I describe the design of Experiment 2 in detail. In each session, each subject went through 2 sub-sessions. First, the subjects were split into two types: **Exp** (48 subjects overall) and **Inexp**

[12]The subjects were told that "The computer plays perfectly to win. That is, if there is a choice or a sequence of choices that the computer can make to win against you, the computer will make that choice or sequence of choices without making any mistake."

[13]Winning as the first mover is impossible in C13 since the computer plays perfectly. However, one needs to understand the backward induction to realize this.

(54 subjects overall), with Inexp subjects being at least 50 percent of the total subjects in each session.

1. In the ***first sub-session (training sub-session)***:

(i) **Exp subjects trained**. Each round of this sub-session comprised of two parts: C13 and Avoid 13.[14] At the start of each round Exp players were randomly stranger re-matched into pairs. Within each pair, in the first part of that round, both subjects separately played C13 against their respective perfectly-playing computer. After both players' respective C13 was complete, each of them was shown the other's moves and outcome in C13. Then, in the second part of the round, the two subjects played the Avoid 13 game with each other. In the training sub-session, Exp subjects played 8 rounds of this C13-Avoid 13 among themselves with stranger re-matching. The earning from a round was the sum of the earnings from its two parts. For C13, a flat 500 ECU for a win, or 50 ECU for a loss. For Avoid 13, 500 ECU (respectively, 200 ECU) for winning as the first-mover (second-mover), or 50 ECU for a loss. One round was randomly drawn as the round determining earning from the first sub-session.

(ii) **Inexp were kept uninformed.** They were not told about the Avoid 13 or C13 game in the first sub-session. The Inexp subjects were kept uninformed by making them play an unrelated bargaining game. It was announced to all Exp and Inexp subjects that the other set of subjects was playing an unrelated game.

2. In the ***second sub-session (combined sub-session)***: Exp and Inexp subjects were mixed and were randomly stranger re-matched into pairs before every round. They played 6 rounds as described in 1(i), i.e. in each round play the C13 with one's respective perfectly-playing computer and then play the Avoid 13 game with the human opponent (C13-Avoid 13), with two modifications: First, any pair of types, i.e., (Inexp, Inexp), (Exp, Exp), or (Exp, Inexp) was possible. Second, while the Exp subjects were shown the history of their opponent's moves and outcome versus his/her computer in C13, the Inexp subjects were *not shown* anything from the opponent's C13 play versus his/her computer. This was done to make the gap in understanding of the relevant backward induction (between Exp and Inexp) persist longer. The Exp subjects were told the fact that the Inexp subjects would not be shown the opponent's history or outcome of play in the C13 part of any round. Indeed, all aspects of the design of the combined sub-session were common knowledge to all subjects.

3. One round from the second sub-session was drawn at random for payment and added to the earning due from the first sub-session. The conversion rate used in the training and combined sub-sessions was 120 ECU = 1 USD. After the two sub-sessions, the payment due to the subjects was paid to them in cash.

The winning positions, or items-left-in-the-box, in the Avoid 13 and C13 games are {12, 11, 10, 8, 7, 6, 4, 3, 2} and losing positions are {13, 9, 5, 1}. In both games, the second-mover has a dominant sure-win strategy available: choose (4 minus the opponent's previous choice) at

[14] Exp players were playing C13 and Avoid 13 for the first few times in the training sub-session. Referring to them as Exp in the first sub-session is done only for labeling.

each move. However, if the second-mover fails to put the opponent in a losing position, then the opponent is in a winning position with a sure-win dominant strategy available.

Define the perfect strategy as: choose "Second Mover" (henceforth S) and then choose 3, 2, and 1 when the number of items left in the box is in $\{12, 8, 4\}$, $\{11, 7, 3\}$, and $\{10, 6, 2\}$, respectively. The perfect strategy puts no restriction on choices when 13, 9, or 5 items are left. The SPE prediction is: both players play perfectly in C13 and Avoid 13.

Note that playing the perfect strategy in the C13 part is strictly dominant (for the C13 part and the round) if the player is sure that the computer plays perfectly (which was communicated to the subjects), otherwise it is weakly dominant. To see why losing in C13 to influence the opponent's inference is not worthwhile, note that losing to the computer in the C13 part costs 450 ECU from round earnings, which is greater than any possible gain in the Avoid 13 part from misleading the opponent about one's understanding of the necessary backward induction.

Due to the extra incentive to win as the first-mover in Avoid 13, the perfect strategy in the Avoid 13 part is *not* a weakly dominant strategy. If a risk-neutral player believes with probability at least $\frac{1}{3}$ that his opponent will make a mistake as the second-mover, then choosing F is optimal for Avoid 13. Denote the subgame after the first and second-movers are decided in Avoid 13 as $A13_{sub}$. The perfect strategy *is* (weakly) dominant in $A13_{sub}$ because the payoff from winning/losing is already decided, and the perfect strategy is a sure-win strategy from a winning position.

5 Experiment 2 Results

The central question in Experiment 2 is: do Exp subjects infer their opponent's weakness in backward induction from the opponent's preceding performance in C13 (when he loses) and take advantage? Theoretically, the Exp subjects can also infer different degrees of weakness based on the exact position where the opponent makes a mistake in C13. But this may be expecting too detailed an analysis of an opponent's actions from subjects. Thus, I focus on the effect of the opponent's outcome in C13.[15] First, I investigate the effectiveness of the training sub-session for Exp subjects.

Result 2(a): *By the end of the training sub-session, approximately 73 percent of Exp subjects understood the necessary backward induction.*

Since playing perfectly in C13 is strictly dominant (weakly dominant if one is unsure about the computer's strategy), subjects' performance in C13 is a good measure of their understanding of the backward induction and thereby the dominant, sure-win perfect strategy. Recall that in C13, the subject unilaterally decides who the first mover is, thus playing perfectly guarantees a win, and is the only way to win. In the training sub-session (rounds 1-8), the proportion of Exp subjects who won in C13 against the computer increased significantly: from 17 percent in round 1 to 79 percent

[15] Indeed, controlling for opponent's outcome in C13, there is no evidence of any effect of the position at which the opponent makes a mistake.

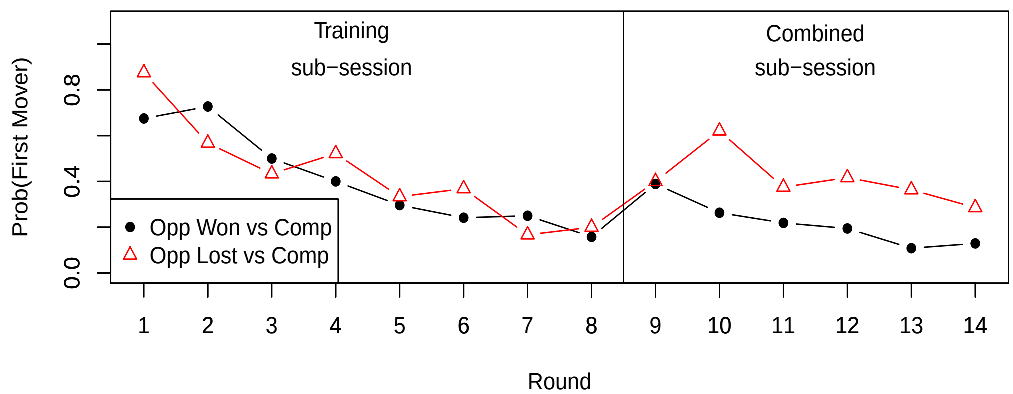

Figure 3. Experienced Subjects' Behavior in Avoid 13 by Opponent's Performance in C13
Notes: The figure depicts the round-wise proportion of experienced subjects who chose "First Mover" in Avoid 13 when matched with an opponent who won C13 (black dots) and the proportion of experienced subjects who chose "First Mover" when matched with an opponent who lost C13 (red triangles). The latter proportion is significantly higher, with a p-value < 0.05, individually for rounds 10 and 13, and a p-value<0.01 for the combined sub-session.

in round 8, a significant increase (p-value<0.01).[16] Notably, 72.9 percent of Exp subjects won both their last 2 attempts against the computer in the training sub-session.

Result 2(b): *In the combined sub-session, when experienced subjects chose "First Mover" in Avoid 13, it was not due to other-regarding preferences, instead it was an attempt to attain the higher payoff which is possible only when the opponent makes a mistake as the second mover.*

Result 2(b) is supported by the following facts. In the Avoid 13 part of the combined sub-session, if an experienced subject was the first mover and got into a winning position due to the opponent's mistake, then in 24 out of 27 such cases (89 percent) the experienced subject ensured that he/she won. Observing that the opponent lost against his/her computer in C13 didn't make the experienced players less ruthless - they converted 19 out of 22 opportunities of winning as the first mover due to such an opponent's mistake in Avoid 13.

Result 2(c): *In the combined sub-session, experienced subjects were significantly more likely to opt for the first-mover position in the Avoid 13 part of a round when they observed that the opponent lost to his/her computer in the C13 part of that round. In other words, experienced players could infer the opponent's weakness in the necessary backward induction if the opponent lost C13, and they tried to take advantage of this.*

Given Result 2(b), if an Exp subject chose F in the Avoid 13 part of the C13-Avoid 13 round, then I can conclude that she was choosing to risk a loss because she believed that there was a high enough chance that her opponent would make a mistake as the second-mover and hand her the

[16]The p-value is obtained using a probit with C13 win/loss as the dependent variable and appropriate round dummies as the independent variables; errors are clustered at subject level.

Dependent Variable: Prob(First Mover)				
	Exp	Exp	Exp	Inexp
	(1)	(2)	(3)	(4)
Opponent Lost C13	0.58***(0.2)	0.6***(0.21)	0.15(0.26)	
Round	-0.14***(0.05)	-0.17***(0.05)	-0.16***(0.05)	-0.39***(0.05)
Risk Aversion		0.14(0.53)	-0.18(0.54)	
Loss Aversion		-0.5**(0.22)	-0.67***(0.25)	
High			-1.32***(0.37)	
High*Opp Lost C13			0.23*(0.14)	
Constant	0.53(0.67)	2.04(1.32)	3.13***(1.41)	4.27***(0.69)
No. of Obs.	278	272	272	312
Pseudo R^2	0.1368	0.1940	0.2586	0.1765
Session Dummies	Yes	Yes	Yes	Yes

Table 2. Factors Influencing Probability of choosing “First Mover” in Avoid 13.
Notes. ***: two-tailed p-value<0.01. **: two-tailed p-value<0.05. *: *one-tailed* p-value<0.10. These results are for the combined sub-session (rounds 9-14) from 8 sessions of Experiment 2. The errors are clustered by subject and reported in parentheses. Probit (1) reports the results for 48 Exp subjects, while probits (2) and (3) report the results for 47 Exp subjects (risk aversion and loss aversion data for 1 Exp subject was lost). Probit (4) reports the results for 54 Inexp subjects. In probit (3), given the non-linearity of a probit model, the coefficient for the interaction term (High*Opp Lost C13), its standard error and its significance are stated at their respective mean level.

higher payoff of 500 ECU in Avoid 13. It is worth recalling that in each round of the combined sub-session, while subjects knew that more than half the subjects were Inexp, they had no information if the opponent was Exp or Inexp, so the Exp subjects had to infer the opponent’s weakness in backward induction from his/her performance in C13. Result 2(c) shows that Exp subjects did infer the opponent’s loss in C13 as such, and they used this inference to try to earn more money.

To see Result 2(c), note that in the combined sub-session (rounds 9-14), in Avoid 13, Exp subjects were more likely to choose F against an opponent who lost in C13 versus his computer, as opposed to an opponent who won his C13. Figure 3 shows that the Exp subjects whose opponent lost C13 (red triangles) chose F at a significantly higher rate than the Exp subjects whose opponent won C13 (black dots). The p-values of this difference in proportions are less than 0.05 for rounds 10 and 13. To derive results for the entire combined sub-session, I use the probit estimations below. I don’t observe any such significant difference in the training sub-session. The knowledge that one’s opponent is similarly experienced may have led the Exp subjects to not put too much weight on the opponent’s “current” performance in C13.[17]

Consider Table 2: models (1)-(3) have the Exp subject’s probability of choosing F in Avoid 13 as the dependent variable, while model (4) has the same for Inexp subjects. The positive and highly significant coefficient of “Opponent Lost C13” dummy in models (1) and (2) implies that, in the combined sub-session, if the average Exp subject observed that his opponent in Avoid 13 had lost

[17]Further, learning the necessary backward induction is one skill, learning how to benefit from the opponent’s lack of understanding is another skill. The latter might have taken additional practice to acquire, which may be why there is no effect of the opponent losing C13 in the F/S decision for Avoid 13 even in the first round of the combined sub-session (see round 9 in Figure 3).

against the computer in the C13 part of that round, then she was significantly more likely to choose F in the F/S decision stage of Avoid 13.[18] Indeed, over all rounds in the combined sub-session, if the opponent lost C13, then the average earning from being the first mover in Avoid 13 was more than the average earning from being the second mover.

The round variable is negative for all the models (1)-(4), which indicates that subjects were learning the perfect strategy due to repetitions, and that there was an increase in the belief that the opponent understands the relevant backward induction. Together, the two factors make choosing F a worse option.

In model (2) I also add the loss aversion and risk aversion parameters as measured using the DOSE (Wang et al. (2010)) estimation technique. At the very end of each session of Experiment 2, only the Exp subjects (not Inexp)[19] participated in a risky choice study to elicit these variables. The risk aversion coefficient's sign is positive, which is opposite of expected, but the coefficient is insignificant.[20] Loss aversion has a significantly negative effect on the probability that an Exp player chooses F.[21] In model (3) - described below - the risk aversion and loss aversion parameters have expected negative signs with the latter significant.

It is worth noting that the experienced subjects with a better understanding of the relevant backward induction drove Result 2(c), i.e. they were the ones more likely to take advantage of an opponent who appeared to be weak in the relevant backward induction. To verify this, I categorize an Exp subject as *High* (respectively, *Low*) if she won the C13 part in at least 3 (at most 2) of the last 4 rounds of her **training sub-session**. 31 Exp subjects are *high* and 17 are *low*. The *High* dummy variable indicates that the Exp subject displayed high expertise in the relevant backward induction during the training sub-session. I add the *High* variable and an interaction term, *High*Opponent Lost in C13*, to model (2) and estimate model (3) for the Avoid 13 data from Exp subjects. The coefficient of *High* is negative and significant (p-value of 0.003), which means that given that the opponent won the C13 part of the round, being the *High* Exp subject significantly reduced the probability of choosing F, which agrees with the definition of *High*. The coefficient of *High*Opponent Lost C13* is positive and significant (*one-tailed* p-value of 0.058). Note that in a probit model, the coefficient for the interaction term (*High*Opponent Lost C13*), its standard error and its significance vary for each observation based on the level of the dependent variable (predicted probability of choosing F). I find that the interaction effect is positive across all individuals, and the interaction is significant (*two-tailed* p-values<0.05) across 32 percent of the

[18]"Opponent Lost C13" dummy is not included in model (4) because Inexp subjects didn't observe any information about the opponent's C13.

[19]This is because the focus of this paper is on Exp players' behavior.

[20]For the Exp subject, conditional on his F/S choice being selected, choosing F in Avoid 13 entailed a gamble between winning (500 ECU) or losing (50 ECU) while choosing S guaranteed a win worth 200 ECU (unless the Exp subject made a mistake as the second-mover). So, the sign of the risk aversion coefficient should have been negative.

[21]The loss aversion term captures how much more a subject values losses compared to equivalent gains. If an Exp player uses the 200 ECU he could have received by winning from the second-mover position as a reference point (see Kőszegi and Rabin (2006) for a discussion on loss aversion with respect to reference points), and considers the reduction in earnings by 150 ECU due to choosing F a "loss," one would expect a significant decrease in the propensity to choose F due to a higher loss aversion term, as observed in models (2) and (3).

observations. Thus, *High* Exp type were more responsive than *Low* Exp type in choosing F to take advantage when facing an opponent who lost in the C13 part.

Results 2(d) and 2(e) approximately replicate the analogous findings from Experiment 1's Results 1(d) and 1(e), respectively. That is, in Experiment 2, I find the following: (i) the rate of mistakes declined significantly as the number of items left reduced; (ii) all subjects, including trained Exp subjects, were significantly more likely to remove 1 item from losing positions in Avoid 13 (see Table 5(b) in Appendix A).

6 Likelihood Comparison

There are three key findings to be explained from the combined sub-sessions of Experiments 1 and 2. First, players make mistakes in a specific pattern: mistakes reduce as items-left reduce (Results 1(d) and 2(d)). Further, experienced players try to take advantage of an opponent's weakness in backward induction when they are told that the opponent is inexperienced (Result 1(c)), but also when they infer this weakness from the opponent's preceding outcome (Result 2(c)).

In what follows, I compare three alternative theoretical models' likelihood with respect to the data: (i) Limited Foresight Equilibrium (LFE, Rampal (2020)), (ii) Level-k (here I test Ho and Su (2013), and Camerer, Ho, and Chong (2004)), (iii) Agent Quantal Response Equilibrium (AQRE: McKelvey and Palfrey (1998)).

6.1 Limited Foresight Equilibrium

In a sequential-move game, define foresight as the number of *subsequent* stages one can consider and understand, or - theoretically speaking - observe from a particular move. LFE models a scenario with two key components: first, players have different levels of foresight (where a foresight-level denotes a particular "type") in a sequential-move game; and second, players are uncertain about the opponent's foresight-level. Thus, the LFE setup converts a game of complete-and-perfect information into a game of incomplete information with observable actions where there is uncertainty regarding opponents' types.

For simplicity, I model *three possible foresight-levels (three types) for each player.* **First**, a type with foresight 0 (or type-0) is modeled as a player-type who cannot observe (i.e. understand) the game that follows after his immediate move. At each move, he observes only a "curtailed game," comprising only of the preceding stages, and the possible actions at the stage of his move. There is a "curtailed payoff" profile attached to each action at that move, where the payoff for each player is the $\frac{min+max}{2}$ of the payoffs possible for that player from that action in the original game setup by the experimenter.[22] The **second** player-type I model is one whose foresight is high enough to understand the backward-induction procedure once first and second movers are decided, but whose

[22] While the $\frac{min+max}{2}$ rule appears arbitrary, it works well in applications of LFE to other sequential-move games (Rampal (2020)), and furthermore this curtailment rule satisfies all of Ke (2019) axioms for curtailment rules, except that it satisfies "monotonicity" only weakly.

foresight is not high enough to understand the second-mover advantage when choosing between F/S (he also uses the same $\frac{min+max}{2}$ curtailment rule where his foresight ends). I denote this type as type-$(f-1)$ ("full minus one" foresight/type). **Third**, I model an ex-ante full-foresight type, labeled type-f, who has no foresight limitations. Beliefs about the opponent's foresight only matter for type-f. Denote type-f's belief (prior or updated) that the opponent is type-j by $\mu(\text{type}-j)$. Denote the three types' strategies in Avoid 9 and Computer13-Avoid13 as s_0, s_{f-1}, and s_f. These strategies can be summarized as follows (their derivation is in Appendix B):

- $s_0 \equiv$ *choose F and S with probability $\frac{1}{2}$ each; remove* 1, 2, 3 *items with equal probability if 5 or more items are left; play perfectly if 4 or less items are left.*

- $s_{f-1} \equiv$ *choose F in Avoid 9 or Avoid 13, choose F and S with probability $\frac{1}{2}$ each in Computer 13. Once order of moves is decided, play perfectly, randomizing uniformly from losing positions.*

- $s_f \equiv$ *in Avoid 9, choose F if prior $\mu(type-0) \geq 0.5$, otherwise choose S. In Computer 13, choose S. In Avoid 13, choose F if updated $\mu(type-0) \geq 0.375$, otherwise choose S. Once order of moves is decided, play perfectly, randomizing uniformly from losing positions.*

The belief, μ, of type-f can use an objective or a subjective prior (the latter is modeled as *LFE with subjective beliefs*, LFESB for short, in Rampal (2020)) (p_0, p_{f-1}, p_f) over the three types. Each prior generates an LFE strategy profile as per (s_0, s_{f-1}, s_f). Together, the prior and the LFE strategy profile generate a probability distribution over possible data observations in Experiments 1 and 2, where a data observation is a sequence of choices of a pair of subjects. For both experiments, I only consider the combined sub-session data. For each experiment, the MLE exercise estimates the prior from the data.

I use the more stringent objective formulation for Experiment 2 data where the experienced type-f player's prior belief is given by the (p_0, p_{f-1}, p_f) estimated using MLE. In Experiment 2, the type-f player updates this prior after observing the opponent's outcome in C13 given s_0, s_{f-1}, and s_f.[23] The objective formulation is also used for the sections of the data from Experiment 1 where Exp-Exp or Inexp-Inexp pairs played each other, for which I estimate $\mathbf{p_{Exp}} = (p_0, p_{f-1}, p_f)_{Exp}$ and $\mathbf{p_{Inexp}}$ respectively. For Experiment 1 data where Exp-Inexp pairs played each other, I allow for subjective beliefs about the distribution over types in the opposing group because Exp and Inexp players were both aware that they are playing a different subject group than their own. Thus, for this data I estimate separate $\mathbf{p^{EI}_{Exp}}$ and $\mathbf{p^{EI}_{Inexp}}$ for the Exp and Inexp players respectively, since I allow the Exp or Inexp type-f players (the only type for whom beliefs matter) to have subjective beliefs.

Observations that entail a mistake with 4 or less items left get probability 0 according to the LFE model since even the foresight 0 type doesn't make a mistake there. While there is no such observation in Experiment 2, there are some (1.4 percent) in Experiment 1. To account for the

[23]Recall that I am assuming that players only pay attention to the opponent's outcome in C13, not his moves.

0-probability case and make the likelihood function finite, I use the "uniform error rate" ϵ used by several Level-k models, e.g. Costa-Gomes, Crawford, and Broseta (2001), including the Ho and Su (2013) Level-k model (described next) that I compare with LFE. The uniform error rate works as follows. Let the number of possible outcomes be N (for example $N = 447$ in the Avoid 9 game). Then $\epsilon \in (0, \frac{1}{N})$ denotes the error probability that each of the N possible outcomes will occur. With remaining probability $[1 - \epsilon N]$, the model's prediction holds. Details about the LFE likelihood function are in Appendix C.

6.2 Level-k models

While the majority of Level-k models and their applications are for simultaneous-move games (Crawford, Costa-Gomes, and Iriberri (2013)), Ho and Su (2013) (henceforth HS) apply the Level-k model to sequential-move games. In the HS model, the Level-k strategy is a sequentially-rational best response to the Level-$(k-1)$ strategy, for $k = 1, 2, ...$, assuming that Level-0 randomizes uniformly at each move. Given these Level strategies, *rational agents* pick what they deem is the optimal level strategy given their *subjective belief* about their opponent's Level. Since players are assumed to be rational, the probability of Level-0 being chosen by any player is 0. So, a Level-0 player "only occurs in the minds of the higher-level players" (HS page 7). As a robustness check, I also test a dynamic version of the Camerer, Ho, and Chong (2004), henceforth CHC, model where Level-0 players do exist with positive probability. But, in the spirit of HS, I don't impose CHC's restriction on beliefs, i.e. I try a dynamic and more flexible version of the CHC model.

The higher level strategies in the Avoid 9 and C13-Avoid 13 games are as follows. The **Level-1** strategy is to choose F and subsequently play perfectly in Avoid 9 and Avoid 13; however, in C13, the Level-1 strategy is to choose S and then play perfectly because there is no extra incentive to win as the first-mover and also because subjects were informed that the computer plays perfectly.[24] (iii) The Level-2 or higher strategies are all identical: choose S and play perfectly in Avoid 9, and C13-Avoid 13 games. A prior distribution over levels ($Pr(L_0)$, $Pr(L_1)$, $Pr(L_{\geq 2})$) and the level strategies yield a distribution over outcomes. For each experiment, the MLE exercise estimates this prior from the data. The Level-k likelihood function is given in Appendix C.

Since the HS model doesn't allow for a Level-0 player, it fails to generate the high proportion of mistakes observed in the data (13.4 percent and 57 percent in Experiments 1 and 2 respectively) without an error term. As mentioned earlier, in applications, HS use the uniform error rate used in LFE. Another Level-k model designed for dynamic games is the Kawagoe and Takizawa (2012) model that uses a logit error structure (similar to the AQRE model studied here). I do not perform the MLE comparison with this model since the LFE model is more comparable with the HS model, because both use *error-less strategies* and the same uniform error rate structure. In particular, I don't use the logit error term for the LFE model because for the logit model it doesn't just matter

[24] Even if the Level-1 player doesn't believe the experimenter and expects the computer to be Level-0, the Level-0 computer has a $\frac{1}{27}$ chance of playing perfectly, which makes it strictly better for the Level-1 player to choose S and then play perfectly.

which action gives more expected payoff, the amount of difference also matters. Thus the logit prediction for the LFE model is not as robust to the "curtailment rule" used in LFE.

6.3 Agent Quantal Response Equilibrium

The Agent Quantal Response Equilibrium (AQRE) model of McKelvey and Palfrey (1998) extends their Quantal Response Equilibrium to extensive form games. An AQRE is a mixed-strategy behavioral strategy profile b, that (in their most commonly used logit-AQRE specification) depends on the parameter λ. In an AQRE b, each player at each move puts a certain probability weight on each action, where this weight is proportional to the expected payoff from that action (given b) relative to other actions. As λ increases, each player at each move puts more probability weight on the best response action. So, the parameter λ generates a certain AQRE b, which in turn implies a probability distribution over the possible outcomes.

In the MLE for Experiment 1 data, I estimate a *heterogeneous version of the AQRE model.* I.e. in Experiment 1, I estimate a separate λ_{Exp} for the Exp players, and a separate λ_{Inexp} for the Inexp players. To the best of my knowledge, this is the first application of such a model. In the MLE for Experiment 2, I estimate a single λ parameter because AQRE provides no framework for updating about the opponent's λ, and this is consistent with the MLE procedure for the other two models. Details about the AQRE likelihood function are in Appendix C.

6.4 MLE Results

Experiment 1 MLE Result- *For the first three rounds of the combined sub-session, pairwise Vuong (1989) tests yield the following likelihood rankings: LFE >AQRE (one-tailed p-value of 0.082) and LFE >HS (two-tailed p-value < 0.01).*[25]

Table 3 reports the likelihood and parameters of the three models discussed in Section 6.1-6.3. Following HS, I put 0 probability on $P(L_0)$ so that the distribution over levels is captured by $P(L_1)$ since $P(L_{\geq 2}) = 1 - P(L_1)$, and all levels 2 or above have an identical strategy. Note that for the Level-k model I estimate different parameters for Exp-Exp, Inexp-Inexp, and Exp and Inexp in the Exp-Inexp data. This is exactly the procedure for LFE, and for the same reason.

The LFE model has the highest likelihood overall and for pair-types other than (Inexp, Inexp), however the difference between the likelihood of the LFE model and the AQRE model is barely significant with a *one-tailed* p-value of 0.082. Further, both Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) also favor the LFE model.[26]

According to the MLE estimate of LFE, about 26 percent (0.31-0.05) Exp subjects switched from playing the perfect strategy to the s_{f-1} strategy when their opponent switched from Exp to

[25] I focus on the first three rounds because those are the rounds where there is a clear difference between the level of understanding of Inexp and Exp subjects, which is the focus of this paper. The LFE model remains the one with the highest likelihood if I take all the rounds of the combined sub-session into account.

[26] For the entire combined sub-session (rounds 13-20) the likelihoods are -2293, -2305, and -2325 for the LFE, AQRE, and HS models. The LFE model is still favored by the AIC and BIC criteria.

Model	Inexp Inexp	Inexp Exp	Exp Exp	Ln(L)
LFE	-291.8	-488.4	-144.1	-924.3
$P(0, f-1, f)$	$(0.3, 0.47, 0.23)$	Exp: $(.04, 0.31, 0.65)$ Inexp: $(0.33, 0.4, 0.27)$	$(0, 0.05, 0.95)$	
Error	1.5×10^{-4}	1×10^{-4}	1.24×10^{-19}	
Level-k	-300.2	-503.4	-144.1	-947.1
$P(L_1)$	0.65	Exp: 0.31; Inexp: 0.60	0.05	
Error	6.1×10^{-4}	3.4×10^{-4}	1.24×10^{-19}	
AQRE	-274.8	-526.1	-148.1	-949.1
λ	0.016	Inexp: 0.016; Exp: 0.024	0.024	

Table 3. MLE Results for Experiment 1
Notes: The table reports the log likelihood and parameters of the LFE, Level-k (HS), and AQRE models for the first three rounds of the combined sub-session of Experiment 1. The sample size is 231.

Inexp. Recall that in the (Exp vs Inexp) data, I am allowing Exp players to have subjective beliefs as per the LFESB model. Thus, under risk neutrality, the switch can be interpreted as follows: 26 percent of the ex-ante full-foresight type-f players' belief μ(opponent has 0 foresight) crossed the 0.5 threshold when they faced an Inexp rather than Exp opponent, which made it optimal for these type-f players to mimic the s_{f-1} strategy to take advantage of the opponent they believed had a more than 50 percent chance of having 0 foresight.

The key driver of the difference among the likelihoods of the LFE model and the Level-k (HS) model is that without the error term, the LFE model puts zero probability on 1.4 percent of the data, whereas the HS model puts zero probability on 13.4 percent of the data. This is because the HS model doesn't have a Level-0, which means HS puts 0 probability on any mistake, while LFE puts 0 probability only on mistakes with less than 4 items left.

For robustness, I estimate an "unconstrained" and dynamic version of the CHC model where players could choose any Level strategy, including Level-0 (who randomizes uniformly at every choice and therefore makes mistakes), without being constrained to having beliefs determined by the estimated distributions. With an overall likelihood of -941.2, this does provide an improvement over the HS model, but this likelihood is still significantly less than the likelihood of the LFE model (p-value$<$ 0.01 in a *Vuong (1989)* test).

The heterogeneous version of the AQRE model provides a slightly worse overall fit compared to LFE. I conjecture that the reason is that even with two parameters, λ_{Exp} and λ_{Inexp}, there are too many data patterns to match. For instance, λ_{Inexp} needs to be low enough to justify Exp player's F/S choice pattern (Result 1(c)), but it cannot be too low, otherwise the pattern of mistakes (Result 1(d)) isn't compatible.

Model	Log Likelihood	Parameters
LFE	-944.2	$Prob(0,\, f-1,\, f) = (0.12,\, 0.38,\, 0.5)$
Level-k	-1111.9	$\epsilon = 0.0078$. $P(L_1\vert$ Opponent Lost$) = 0.104$ $P(L_1) = 0.102 \approx P(L_1\vert$ Opponent Won$)$
AQRE	-1488.8	$\lambda = 0.0056$

Table 4. MLE Results for Experiment 2
Notes: The table reports the log likelihood and parameters of the LFE, Level-k (HS), and AQRE models for the combined sub-session. The sample size is 295.

Experiment 2 MLE

The total number of different sequences of moves possible in the C13-Avoid 13 round of Experiment 2 is more than 1 million. So to provide the HS uniform error term bite (each sequence must have the same minimum error probability of occurring), and for tractability, I divide the observations from Experiment 2 into 81 broad categories.[27] I categorize the observations on the following basis: (i) Exp/Inexp combination of a pair; (ii) the outcomes of a pair in their respective C13 (iii) F/S choices of a pair in Avoid 13; (iv) perfect/imperfect play by the selected first and second movers in Avoid 13. The details of this categorization are in Appendix D.

Experiment 2 MLE Result- *For the combined sub-session, pairwise Vuong (1989) tests yield the following likelihood rankings: LFE >AQRE; LFE >HS ; HS>AQRE (all p-values < 0.01).*

Table 4 states the MLE results and parameters for the data from the combined sub-session of Experiment 2 (rounds 9-14). The LFE model has the highest likelihood; further, both AIC and BIC criteria also favor the LFE model.

Note that the results in Table 4 are for the LFE model without an error term and with objective prior beliefs, which are updated after observing the opponent's win/loss in the C13 immediately preceding Avoid 13. For the estimated parameters, the LFE prediction is that if an Exp player's opponent loses C13, then the ex-ante full-foresight Exp player must choose "First Mover," and otherwise he must choose "Second Mover." This is clearly a stark prediction, yet, it fits the data better than the other models.

In Experiment 2 I attempt to incorporate the learning feature of the HS model wherein players' beliefs and thereby level distributions change across repetitions of a game. I don't do this in Experiment 1 since there I estimate four different prior distributions over levels for three repetitions with different opponent types across repetitions; I argue that this leaves only a minimal role for learning to change level distributions. However, in Experiment 2, one can argue that the two parts of the C13-Avoid 13 round are two repetitions of a very similar game and that the Exp subjects could have formed different subjective belief based on observing the opponent's play in C13. Therefore

[27]No such division was done for Experiment 1 as the number of possible outcomes was only 447.

in the results given in Table 4, I have also allowed for the Level ≥ 1 players to change their level in going from C13 to Avoid 13 based on the information about opponent's win/loss in C13. Given $P(L_0) = 0$ and that Levels 2 and above are identical, $P(L_1)$ captures the distribution for Inexp players who got no information after C13, but $P(L_1|\text{opponent lost})$ and $P(L_1|\text{opponent won})$ are allowed to be different from $P(L_1)$ for an Exp player when he observed that his opponent lost or won in C13, respectively.

As a different robustness check, in the HS model, I allow for learning across rounds, i.e., I allow for a separate distribution on levels (and separate additional parameters), in each of rounds 9 through 14. The log likelihood improves to -1074.3, which doesn't change the conclusions. As a last robustness check, I try the dynamic and flexible CHC model where I also allow for a non-zero proportion of players to be Level-0, who randomize uniformly at each choice. In this CHC model, I also allow for the adjustment of the strategy chosen by Level≥ 1 players between the C13 and Avoid 13 part of the rounds. The likelihood of this model is -969.9, which is also significantly lower than the LFE likelihood (two-tailed p-value= 0.079).

The standard HS model has low likelihood because it doesn't allow for subjects to lose in C13, or indeed make any mistake in C13- Avoid 13, which implies 0 probability on 57 percent of the data unless the uniform error term is used. Thus, the introduction of a Level-0 player using the CHC model helps a lot. But even this addition doesn't capture the pattern of mistakes made. While the Level-0 player has the same likelihood of mistake at each position, I observed that mistakes declined as items left reduced (Result 2(d)), a data feature partially captured by the foresight-0 type in LFE.

The AQRE model fares much worse than Level-k and LFE. I conjecture that this is because the AQRE model does not allow for the experienced player to update beliefs about the opponent's λ based on the opponent's outcome in C13. Thus, a heterogenous version of the AQRE model is not well defined for Experiment 2, and it may be futile even if attempted.

Discussion

Recall the peculiar but robust finding (Results 1(e) and 2(e)) that subjects had a predilection to remove 1 item from losing positions. A more elaborate version of the LFE model may have a partial explanation. If the ex-ante full-foresight type believed (objectively or subjectively) that there is some chance the opponent has a foresight of 3-stages-ahead or 1-stage-ahead, then removing 1 item becomes the best response from 13 items left or 9 items left, respectively.[28] However, nothing in LFE can explain the tendency to remove 1 item with 5 items left. Possibly, this indicates the need for a more "structural" approach to limited foresight in LFE, analogous to the Alaoui and Penta

[28]To see why the type-f player is more likely to remove 1 item with 9 items left, suppose his opponent can have foresight of 1-stage-ahead (type-1). From any given move, this type-1 player observes each of his actions, all actions following his actions, and then curtailed payoff profiles (constructed using the $\frac{min+max}{2}$ rule described earlier). The type-1 player can foresee that leaving 3 or 4 items after his move will cause a loss for sure, while he thinks a win or loss are both possible if he leaves 5 or more items after his move. Thus the type-1 player's probability of mistake when moving with 6, 7, and 8 items left is 0, $\frac{1}{2}$, and $\frac{2}{3}$, respectively. So, to maximize the chance that his type-1 opponent makes a mistake, the type-f player should remove 1 item when 9 items are left.

(2016) cost-benefit approach to Level-k.

7 Concluding Remarks

I demonstrate that experts best respond to the opponent's backward-induction ability; in particular, that experienced players can take advantage of the opponent's weakness in backward induction. I show that experienced players take advantage when they are told that the opponent is inexperienced (Experiment 1), but also when they must infer the opponent's weakness from his preceding choices and outcome, with no knowledge of his experience-level (Experiment 2). While these are intuitive findings, to my knowledge, this is the first paper to demonstrate these findings while clearly ruling out any possible role of other regarding preferences.

Somewhat surprisingly, these intuitive findings haven't been modeled theoretically in the literature. Indeed, in the MLE exercise, when I compare Level-k, AQRE, and LFE models, I find that the novel LFE model of limited foresight and uncertainty about the opponent's foresight provides the best data fit. While the limited foresight component explains the pattern of mistakes in the sequential-move games tested (the key challenge for Level-k models), the uncertainty component explains how experienced players may take advantage of opponents they believe have low foresight. Since the LFE allows belief updating (the key challenge for AQRE) about the opponent's foresight based on her preceding moves/outcomes, it helps explain the findings from Experiment 2.

While the experimental design of this paper shows how difference in backward-induction expertise can arise due to heterogeneity in past experience, there is a flip side. The external validity of the findings of this paper are limited to scenarios where experts had past experience with making mistakes in backward induction. The experienced players may well have realized that the inexperienced opponent, like them, would likely make mistakes in the first few rounds.

Broadly, the findings of this paper indicate that analysts and players of dynamic interactions should be mindful of heterogeneity in foresight among players, which means heterogeneity in backward-induction expertise. As shown here, this heterogeneity can arise simply from heterogeneity in experience with similar interactions, but it is easy to imagine other causes of this heterogeneity that are worth studying, e.g. heterogeneity in education. Thus a player's optimal strategy depends on her beliefs about other players' foresights. These beliefs can be based on knowledge about the other players' prior experience and/or dynamically updated based on their preceding choices and outcomes within the play of the dynamic interaction.

References

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

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

Alaoui, L. and A. Penta (2016). Endogenous depth of reasoning. *Review of Economic Studies 83*(4), 1297–1333.

Binmore, K., J. McCarthy, G. Ponti, L. Samuelson, and A. Shaked (2002). A backward induction experiment. *Journal of Economic Theory 104*(1), 48–88.

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

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

Crawford, V. P., M. A. Costa-Gomes, and N. Iriberri (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature 51*(1), 5–62.

Dufwenberg, M., T. Lindqvist, and E. Moore (2005). Bubbles and experience: An experiment. *American Economic Review 95*(5), 1731–1737.

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

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*(2), 171–178.

Gill, D. and V. L. Prowse (2016). Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy 124*(6), 1619–1676.

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

Ho, T.-H. and X. Su (2013). A dynamic level-k model in sequential games. *Management Science 59*(2), 452–469.

Johnson, E. J., C. Camerer, S. Sen, and T. Rymon (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory 104*(1), 16–47.

Kawagoe, T. and H. Takizawa (2012). Level-k analysis of experimental centipede games. *Journal Of Economic Behavior & Organization 82*(2), 548–566.

Ke, S. (2019). Boundedly rational backward induction. *Theoretical Economics 14*(1), 103–134.

Kőszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics 121*(4), 1133–1165.

Le Coq, C. and J. T. Sturluson (2012). Does opponents experience matter? experimental evidence from a quantity precommitment game. *Journal of Economic Behavior & Organization 84*(1), 265–277.

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

Lusardi, A. and J. Beeler (2007). Saving between cohorts: The role of planning. *Redefining Retirement: How Will Boomers Fare*, 271–295.

Mantovani, M. (2016). Limited foresight in sequential games: an experiment. *Working Paper*.

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

McKelvey, R. D. and T. R. Palfrey (1998). Quantal response equilibria for extensive form games. *Experimental Economics 1*(1), 9–41.

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

Rampal, J. (2020). Limited foresight equilibrium. *Working Paper*. Avalilable at https://sites.google.com/site/jeevantrampalecon/Research.

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

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica 57*(2), 307–333.

Wang, S. W., M. Filiba, and C. F. Camerer (2010). Dynamically optimized sequential experimentation (dose) for estimating economic preference parameters. *Working Paper*.

Appendix A: Tables

Table 5(a). Experiment 1: Avoid 9 data								
	Session				Exp in Combined Sub-session			
	Items Removed				Items Removed			
Items Left	1	2	3	p-value	1	2	3	p-value
9	41.1	29.3	29.6	<0.01	46.4	28.1	25.5	<0.01
5	53.5	24.2	22.3	<0.01	49	28	23	<0.01

Table 5(b). Experiment 2: Avoid 13 data								
	Session				Exp in Combined Sub-session			
	Items Removed				Items Removed			
Items Left	1	2	3	p-value	1	2	3	p-value
13	40.7	29.8	29.5	0.0004	41.1	34.9	24	0.3048
9	48	24.5	27.5	<0.01	43.4	29.2	27.4	0.0269
5	53.3	23.9	22.8	<0.01	53.3	23.9	22.8	0.0003

Table 5. Choices from Losing Positions
Notes: The figures are in percentages. The p-values are two tailed p-values comparing the proportion of times 1 was chosen, as compared to the next most chosen alternative. The table shows that from each losing position, subjects were significantly more likely to remove 1 item, than 2 or 3 items.

	Items Left		
	12-10	8-6	4-2
Exp: first 3 rounds	45.5	21.7	3.5
Inexp: first 3 rounds	34.8	12.2	0
Session	21	5	0.8

Table 6. Demonstrating Result 2(d): Rate of Mistakes by Items Left in Avoid 13
Notes: The figures are in percentage. For each of the three rows, Wilcoxon matched-pairs sign rank test reveals that the rate of mistakes is significantly more (p-values < 0.01) from positions with 12-10 items left as compared to positions with 8-6 items left, and from the latter compared to a positions with 4-2 items left.

Appendix B: LFE Strategy Profile

The type with foresight of 0 stages ahead. The type-0 player observes a curtailed version of the game being played from each move, where the game is curtailed immediately following his actions from that move. Type-0 assigns payoff profiles to each of his immediate actions in the curtailed game using a simple function on the payoff profiles of the original game (i.e., the game set up by the experimenter). The assigned payoff profiles are called "curtailed payoffs." As in the examples in Rampal (2020), I assume that this type-0 player follows the following simple rule. The curtailed payoff from action a at history h (where type-0 moves) is the $\frac{min+max}{2}$ of the payoffs possible in the original game after (h, a) is played. Assume that every type (not just type 0) considers his move as the "first move" in the simultaneous-move F/S decision stage. So, in Avoid 9, Avoid 13, and C13, at the F/S decision stage, the type-0 player's curtailment rule implies that the curtailed payoff from F or S is $\frac{50+500}{2}$.[29] So, the type-0 player is indifferent between the two options.

Throughout, I assume that all players or types uniformly randomize when indifferent among a set of actions that are all best responses. So, type-0 chooses F or S with probability $\frac{1}{2}$ each. Subsequently, from type-0's perspective, if 5 or more items are left, removing 1, 2 or 3 items yields the same curtailed payoff ($\frac{50+500}{2}$ if he is the first mover and $\frac{50+200}{2}$ if he is the second mover). This is because he thinks a win or loss are both possible, and since he doesn't understand how he and his opponent will choose after his current move. Once 4 or less items are left, type-0 observes that removing 1 less than the items left guarantees him the winner's payoff. So type-0 plays perfectly with 4 or less items left.

The type with foresight of $(f - 1)$. This corresponds to a foresight of 3-stages-ahead in the Avoid 9 game and a foresight of 5-stages-ahead in Avoid 13 and C13. The type-$(f - 1)$ analyzes the payoffs from the actions at the end of his foresight, where like the type-0 player, he uses the $\frac{min+max}{2}$ rule to assign a payoff profile to each of those actions. Thus, the type-$(f - 1)$ chooses F in Avoid 9 and Avoid 13 games because, conditional on his choice being selected, the max from the end of his foresight horizon is 500 if he chooses F, and 200 if he chooses S. The type-$(f - 1)$ randomizes with probability $\frac{1}{2}$ each among F/S in C13, because in C13, the max is 500 from choosing F or S. After the order of moves is decided, type-$(f - 1)$ plays perfectly since he observes the dominant strategy for the subgame.

The type with ex-ante full foresight The type-f understands the perfect strategy and the backward-induction procedure before the game begins. In this experiment, belief about the opponent's foresight matters only for type-f.[30] Let (p_0, p_{f-1}, p_f) be the distribution over the foresight-

[29] This is because from type-0's perspective, in Avoid 9 or Avoid 13, even after he chooses S, due to the "subsequent" moves of the other player and Nature, he may still end up as the selected first-mover.

[30] That uncertainty about the opponent's foresight doesn't matter for the type-0 player is obvious. To see that it doesn't affect the type-$(f - 1)$'s optimal strategy, recall that he understands the backward induction only after the first and second movers are decided, and beliefs matter only at the F/S decision stage. After order of moves is

levels (or types). In LFE, the type-f player best responds to the distribution (p_0, p_{f-1}, p_f) given s_f, s_{f-1}, and s_0. So, in Avoid 9 or Avoid 13, if the type-f player believes with high enough probability that his opponent is type-0 (the only type who makes a mistake from a winning second-mover position), then he chooses "First Mover" to maximize expected payoff. Otherwise, he chooses "Second Mover." Let $\mu(\text{type} - i)$ denote the type-f player's belief at the F/S decision stage that his opponent is type-i, where $i = 0$, $(f-1)$, or f. Assuming risk neutrality, in the Avoid-9 game, $\mu(\text{type}-0) \geq 0.5$ incentivizes the type-f player to choose "First Mover." In Avoid 13, this threshold is $\mu(\text{type} - 0) \geq 0.375$.

In Experiment 2, the experienced type-f player updates his beliefs after observing the opponent's outcome in C13. Note that here I am using the assumption that players only pay attention to the opponent's outcome in C13. The belief updating of the experienced type-f player in Experiment 2 is as follows. After observing that the opponent won in C13, the type-f's updated belief over (p_0, p_{f-1}, p_f) is given by:

$$\left(\frac{1}{18}\frac{p_0}{sumW}, \frac{1}{2}\frac{p_{f-1}}{sumW}, \frac{p_f}{sumW}\right); \text{ where } sumW = \frac{1}{18}p_0 + \frac{p_{f-1}}{2} + p_f. \tag{1}$$

And after observing that the opponent lost C13, the updated belief of the experienced type-f player is

$$\left(\frac{17}{18}\frac{p_0}{sumL}, \frac{1}{2}\frac{p_{f-1}}{sumL}, 0\right); \text{ where } sumL = \frac{17}{18}p_0 + \frac{p_{f-1}}{2}. \tag{2}$$

Appendix C: Likelihood Functions

Limited Foresight Equilibrium For each experiment, an observation is the observed choices of a pair of subjects in a round. Denote the observation i by o_i. Suppose o_i can be observed if the strategy profiles in $S^2(o_i) \ni (s_1, s_2)$ are played. Then the probability of o_i can be calculated using (s_0, s_{f-1}, s_f) and the prior distribution over the foresight levels, (p_0, p_{f-1}, p_f), as follows:

$$\begin{gathered} Pr(o_i|\mathbf{p}) = \\ \textstyle\sum_{(s_1,s_2)\in S^2(o_i)}[\sum_{(j,k)\in\{0,f-1,f\}^2} Pr(s_1|\text{type} - j,\, \mathbf{p})Pr(s_2|\text{type} - k,\, \mathbf{p})p_j p_k]. \end{gathered} \tag{3}$$

So, the probability of o_i according to LFE, given its parameters is:

$$Pr^{LFE}(o_i|\mathbf{p}, \epsilon) = \epsilon + (1 - \epsilon N)Pr(o_i|\mathbf{p}). \tag{4}$$

Suppose each observation o_i in Experiment 1 has frequency f_i in the data. Then the log likelihood function is given by:

$$LogL(LFE|\mathbf{p}, \epsilon) = \sum_i f_i.log(Pr^{LFE}(o_i|\mathbf{p}, \epsilon)). \tag{5}$$

decided, the perfect strategy is dominant.

In general, maximizing the likelihood using LFE implies searching over $(\mathbf{p}, \epsilon)$ to maximize (5). I use $\epsilon = 0$ in Experiment 2 because there is no 0 probability observation in the Experiment 2 data.

Level-k Like (4) and (5), the likelihood function for the HS model is calculated by first calculating the Level-k model's probability that outcome o_i occurs;

$$Pr^{L_k}(o_i|\mathbf{p}, \epsilon) = \epsilon + \\ (1 - \epsilon N) \textstyle\sum_{(s_1,s_2) \in S^2(o_i)} [\sum_{(j,k) \in \{0,1,2\}^2} Pr(s_1|L_j) Pr(s_2|L_k) P(L_j) P(L_k)], \tag{6}$$

and using this in the likelihood function as given by

$$LogL(Level_k|\mathbf{p}, \epsilon) = \sum_{i=1}^{N} f_i.log(Pr^{L_k}(o_i|\mathbf{p}, \epsilon)). \tag{7}$$

Maximizing the likelihood implies searching over $(\mathbf{p}, \epsilon)$ to maximize (7).

Agent Quantal Response Equilibrium Let o_i be an observed outcome in the data, and let f_i be its observed frequency. Then:

$$Pr^{AQRE}(o_i|\lambda) = \sum_{(s_1,s_2) \in S^2(o_i)} [Pr(s_1|b(\lambda)) Pr(s_2|b(\lambda))]. \tag{8}$$

So the likelihood function of the AQRE model is given by:

$$LogL(AQRE|\lambda) = \sum_{i=1}^{N} f_i.log(Pr^{AQRE}(o_i|\lambda)) \tag{9}$$

Appendix D: Data Categorization in Experiment 2

Recall that stage 1 of Avoid 13 is the F/S decision stage. Let I (Roman numeral 1) denote the selected first-mover in Avoid 13, and II denote the selected second-mover in Avoid 13. First categorize the data outcomes from $Avoid13_{sub}$ (moves in Avoid 13 after I and II are decided) into three categories: (a) I played an arbitrary action with 13 items left (stage-two), then II made a mistake in stage-three, then I left 9 or 5 items in the box (whichever was possible) in stage-four; (b) I played an arbitrary action with 13 items left (stage-two), then II played perfectly in stage three (left 9 items in the box); (c) I played an arbitrary action with 13 items left (stage-two), then II made a mistake in stage three, then I also made a mistake. While this ignores subsequent play following each of (a)-(c), the categories can be safely interpreted as (a) arbitrary-mistake-perfect; (b) arbitrary-perfect; and (c) arbitrary-mistake-mistake. In particular, assuming perfect play from just one "correct" move as the second mover is supported by the data. This is because in every case in which a player left 5 items in the box, two stages later he left 1 item and won the game. Further, in 97 percent of the data, if a player left 9 items after his move, two stages later, he left 5 (and then went on to win for sure).

Each of (a), (b), and (c) has three distinct cases possible for the choices in the F/S stage of Avoid 13. These three cases are: (1) in Avoid 13's F/S decision stage, I chose F and II chose F; (2) I chose F and II chose S; (3) I chose S and II chose S. For each of (1)-(3) in the F/S decision stage of Avoid 13, I have (a)-(c) in $Avoid13_{sub}$, which makes 9 categories of data outcomes from Avoid 13.

Experiment 2's combined sub-session also had two types of players playing together: Inexp and Exp. While the Exp subjects were informed about their opponent's play in C13 before the beginning of Avoid 13, the Inexp subjects were uninformed about their opponent's play in C13. Further, there were four possible outcomes for a pair of subjects from their respective play against their respective computer in the C13 part of the round: Win-Win, Win-Loss, Loss-Win, and Loss-Loss, where the first term is I's outcome in C13 and the second term is II's outcome in C13. Thus, I need to further broaden the categories of outcomes for Experiment 2.

I make a total of 81 categories for the data from experiment 2. The 9 categories above are repeated in 9 cases. Case (i): (Exp_W, $Inexp_W$) , which denotes that I is Exp and he won his C13 part, while II is Inexp and he also won his C13 part. Case (ii): ($Inexp_W$, Exp_W), with the order from (i) swapped. Case (iii): (Exp_W, $Inexp_L/Exp_L$), which denotes that I is Exp and he won his C13 part, while II lost his C13 part, and he can be either Inexp or Exp.[31] Case (iv): ($Inexp_L/Exp_L$, Exp_W) is the same as case (iii) with the order swapped. Proceeding similarly I make case (v): (Exp_W, Exp_W); (vi): ($Inexp_W$, $Inexp_W$); (vii): ($Inexp_W$, $Inexp_L/Exp_L$); (viii) ($Inexp_L/Exp_L$, $Inexp_W$); (ix) ($Inexp_L$, $Inexp_L$). For each of the cases (i)-(ix), the nine categories of Avoid 13 outcomes are repeated, making 81 total categories.

Appendix E: Subject Instructions (for online publication only)

Experiment 1

INSTRUCTIONS for Type D (Inexp) Subjects

This is an experiment in the economics of decision making. The Ohio State University has provided funds for conducting this research. The instructions are simple, and if you follow them carefully and make good decisions, you may earn a CONSIDERABLE AMOUNT OF MONEY which will be PAID TO YOU IN CASH at the end of the experiment.

Please pay careful attention to the instructions. If you have any questions at any time, raise your hand.

[31] I club the Exp_L and $Inexp_L$ in one category because the Exp player who loses in C13 cannot be type-f according to LFE. Thus, he has the same strategy in LFE, regardless of his beliefs. That is, in LFE, updating about the opponent's type within the round does not make a difference to $Exp_L's$ strategy, while $Inexp_L$ cannot learn by design. The Dynamic Level-k of HS doesn't allow for within round learning. As a robustness check I estimate a model where within round learning was possible for the players who chose Level-1 or higher in C13. But Level-1 or higher players should win in C13. Thus, the Exp players who lose their C13 cannot adjust their strategy based on the observation of the opponent's strategy, exactly like the Inexp player. The AQRE model does not account for any learning. So clubbing Exp_L and $Inexp_L$ in one category reduces the number of categories without affecting the result.

Do not communicate with other participants during the experiment. Put your phone away and make sure it is in silent mode or turned off.

Upon entry you were each assigned a unique subjectID, which you will be asked to enter several times. Please keep this carefully.

This is the second (first part was bargaining) of the two parts of the experiment. You will play 8 rounds of the AVOID Removing the 9th Item game described below. Your total payment will be a sum of your payment from the two parts of the experiment, and your show up fee ($5).

In each round of the AVOID Removing the 9th Item game:

1. Each of you will be asked to fill in your subjectID. Please do so accurately. Then each of you will be randomly assigned a participant whom you will interact with for that round. You will not know the identity of your assigned participant. I shall refer to him/her as "the other." Decisions made by you and the other in a particular round will affect how many experimental currency units you earn in that round. You will be told if the other is of type S or type D. Type S participants have already played 12 rounds of the AVOID Removing the 9th item game. Type D participants (like you) participated in a different, completely unrelated first part of the experiment, and are now playing the second part's AVOID removing the 9th item game with other type D and type S participants.

2. Once the "Avoid removing the 9th item" game begins, you will be asked to decide between being the first mover (you get to make the first choice) or the second mover (the other makes the first choice) in the subsequent task. The computer will choose one of you or the other with 50% chance each, and implement their first/second mover choice.

3. You and the other will take subsequent decisions alternately.

4. You and the other will alternately choose the number of items to remove from a box containing 9 items. You can only choose to remove 1, 2, or 3 items with any given choice. You can't choose to remove 0 items. Of course, you can't choose to remove more items than there are left in the box. For example: if 7 items have been removed before your move, you can only choose to remove either 1 or 2 items at your move. The round ends when the 9th item is removed:

    (a) The one who removes the 9th item from the box receives 50 experimental currency units (ECU) for the round.

    (b) If the other removes the 9th item, and you are the first mover, you receive 500 ECU.

    (c) If the other removes the 9th item, and you are the second mover, you receive 200 ECU.

5. Except your and the other's first-mover/second-mover choices, you will see all your past choices and all the past choices of the other at all points during the round.

You will play 8 rounds of this game in the second part. In each round, the other may be of type S or type D. The type of the other will be communicated to you at all points within a round.
Your earnings for the experiment: part one earning will be equal to your earning in a randomly selected round from part one. Part two earning will be equal to your earning in a randomly selected round from part two. Your total earnings will be a sum of your show-up fee and your earnings from each of the two parts of the experiment.

The conversion rate for ECU in the experiment is: 1USD=80ECU for the first part, and 1USD=60ECU for the second part. Your earnings will be rounded up to the nearest dollar.

INSTRUCTIONS for Type S (Exp) Subjects

This is an experiment in the economics of decision making. The Ohio State University has provided funds for conducting this research. The instructions are simple, and if you follow them carefully and make good decisions, you may earn a CONSIDERABLE AMOUNT OF MONEY which will be PAID TO YOU IN CASH at the end of the experiment.
Please pay careful attention to the instructions. If you have any questions at any time, raise your hand.
Do not communicate with other participants during the experiment. Put your phone away and make sure it is in silent mode or turned off.
Upon entry you were each assigned a unique subjectID, which you will be asked to enter several times. Please keep this carefully.
The experiment has two parts. In both parts, you will play several rounds of the AVOID Removing the 9th Item game described below. Your total payment for the experiment will be a sum of your payment from the two parts plus your show up fee ($5).

In each round of the AVOID Removing the 9th Item game:

1. Each of you will be asked to fill in your subjectID. Please do so accurately. Then each of you will be randomly assigned a participant whom you will interact with for that round. You will not know the identity of your assigned participant. I shall refer to him/her as “the other.” Decisions made by you and the other in a particular round will affect how many experimental currency units you earn in that round. You will be told if the other is of type S or type D. Type S will participate in several rounds of the same AVOID removing the 9th item game (described here) in both parts of the experiment. Type D will participate in a different, completely unrelated first part. After they finish the first part of the experiment, you will be informed. They will then join you in the second part of the experiment, i.e., they will also participate in several rounds of the AVOID removing the 9th item game in the second part, the same as you. You will be told when the type D subjects join and the second part begins.

2. Once the “Avoid removing the 9th item” game begins, you will be asked to decide between being the first mover (you get to make the first choice) or second mover (the other makes the

first choice) in the subsequent task. The computer will choose one of you or the other with 50% chance each, and implement their first/second mover choice.

3. You and the other will take subsequent decisions alternately.

4. You and the other will alternately choose the number of items to remove from a box containing 9 items. You can only choose to remove 1, 2, or 3 items with any given choice. You can't choose to remove 0 items. Of course, you can't choose to remove more items than there are left in the box. For example: if 7 items have been removed before your move, you can only choose to remove either 1 or 2 items at your move. The round ends when the 9th item is removed:

   (a) The one who removes the 9th item from the box receives 50 experimental currency units for the round.

   (b) If the other removes the 9th item, and you are the first mover, you receive 500 ECU.

   (c) If the other removes the 9th item, and you are the second mover, you receive 200 ECU.

5. Except your and the other's first-mover/second-mover choices, you will see all your past choices and all the past choices of the other at all points during the round.

You will play 12 rounds of this game in the first part and 8 rounds of this game in the second part. In the first part the other can only be of type S. In the second part, the other may be of type S or type D. The type of the other will be communicated to you at all points within a round. Part one earning will be equal to your earning in a randomly selected round from part one. Part two earning will be equal to your earning in a randomly selected round part two. Your total earnings will be a sum of your show-up fee and your earnings from each of the two parts of the experiment.

The conversion rate for ECU in the experiment is: 1USD=60ECU. Your earnings will be rounded up to the nearest dollar.

Experiment 2

INSTRUCTIONS for Type D (Inexp) Subjects

This is an experiment in the economics of decision making. The Ohio State University has provided funds for conducting this research. The instructions are simple, and if you follow them carefully and make good decisions, you may earn a CONSIDERABLE AMOUNT OF MONEY which will be PAID TO YOU IN CASH at the end of the experiment. If you have any questions at any time, raise your hand. Do not communicate with other participants during the experiment. Put your phone away and make sure it is in silent mode or turned off.

The experiment has Two Components. You already completed the first (bargaining) component. In this second component, I will be using a game called "Avoid Removing the 13th Item game." This game is played by two players. There is a box containing 13 items. The two players make choices alternately. Each player can choose to remove 1, 2 or 3 items from the box at their move. The players remove 1, 2 or 3 items alternately until 0 items are left in the box. One cannot remove 0 items or more than 3 items. Also, one cannot remove more items than are left in the box. The goal is to avoid removing the last/13th item. The player who makes his opponent remove the 13th item achieves the goal.

You are a type D subject. Type D subjects play a different bargaining game for the first component of the experiment. Type S subjects play the same round described below throughout the experiment, including the first component. In this component you will play 6 rounds of the round described below.

In every round of this experiment:

Each of you will be asked to fill in your subjectID. Please do so accurately. You will be assigned a random subject every round. You will interact with this subject for the whole round, and then be randomly re-matched. You will not know the identity of your assigned subject. I shall refer to him/her as "the other" or "the human opponent". Every round will have 2 parts.

Part 1 of a Round: Play the computer

You and your human opponent will separately play the AVOID Removing the 13th Item game (Avoid 13th for short) with the computer. *The computer plays perfectly to win. That is, if there is a choice or a sequence of choices that the computer can make to win against you, the computer will make that choice or sequence of choices without making any mistake.*

When you play the computer:

1. You will decide if you are the first mover or the second mover in the Avoid 13th game vs the computer.

2. Your choice for being the first or second mover vs the computer will be implemented. You will then commence playing the Avoid 13th game versus your computer as the first/second mover, as per your decision.

3. If you avoid removing the 13th item against the computer, you earn 500ECU from the first part of the round, regardless of your first mover/second mover decision. If you have to remove the 13th item, you earn 50ECU from the first part. 50ECU is the minimum payment for participating in each part of the round.

Part 2 of a Round: Play the human opponent

1. You and the human opponent will then play the AVOID 13th game with each other. Before the Avoid 13th game begins, both you and the human opponent will make the choice between being the first mover or second mover in the Avoid 13th game. One of your or the human opponent's first/second mover choice will be implemented with 50% chance each.

2. For the second part, if you AVOID Removing the 13th Item as the First mover, you earn 500ECU, while if you AVOID Removing the 13th Item as the Second Mover, you earn 200ECU. If you have to remove the 13th item, you earn 50ECU.

After the round ends, you will be randomly re-matched with another subject.

Your earnings for a round will be the sum of your earnings from the first and the second part of the round. To calculate your earnings from this second component, one round will be randomly selected from the 6 rounds you play.

Your earnings from the experiment will be the sum of your earnings from the two components (bargaining and Avoid 13th game) plus the show-up fee ($5). The conversion rate for ECU in this second component is: 1USD=120ECU. The conversion rate for ECU in the first component is: 1USD=120ECU. Your total earnings will be rounded up to the nearest dollar.

INSTRUCTIONS for Type S (Exp) Subjects

This is an experiment in the economics of decision making. The Ohio State University has provided funds for conducting this research. The instructions are simple, and if you follow them carefully and make good decisions, you may earn a CONSIDERABLE AMOUNT OF MONEY which will be PAID TO YOU IN CASH at the end of the experiment. If you have any questions at any time, raise your hand. Do not communicate with other participants during the experiment. Put your phone away and make sure it is in silent mode or turned off.

The experiment has Three Components. I explain the first two components here. The third component is unrelated and doesn't affect your earning from the first two components.

In the first two components, I will be using a game called "Avoid Removing the 13th Item game." This game is played by two players. There is a box containing 13 items. The two players make choices alternately. Each player can choose to remove 1, 2 or 3 items from the box at their move. The players remove 1, 2 or 3 items alternately until 0 items are left in the box. One cannot remove 0 items or more than 3 items. Also, one cannot remove more items than are left in the box. The goal is to avoid removing the last/13th item. The player who makes his opponent remove the 13th item achieves the goal.

Type D subjects play an unrelated bargaining game for the first component of the experiment. You are a type S subject. That means you will play the same round described below throughout the experiment.

In every round of this experiment:

Each of you will be asked to fill in your subjectID. Please do so accurately. You will be assigned a random subject every round. You will interact with this subject for the whole round, and then be randomly re-matched. I shall refer to him/her as "the other" or "the human opponent". Every round will have 2 parts.

Part 1 of a Round: Play the computer

You and your human opponent will separately play the AVOID Removing the 13th Item game (Avoid 13th for short) with the computer. *The computer plays perfectly to win. That is, if there is a choice or a sequence of choices that the computer can make to win against you, the computer will make that choice or sequence of choices without making any mistake.*

When you play the computer:

1. You will decide if you are the first mover or the second mover in the Avoid 13th game vs the computer.

2. You will be shown your human opponent's first mover/second mover decision in his/her interaction with his/her computer. You will then be asked about the type of your human opponent, i.e., your human opponent's type is ______(S or D). If you answer correctly, 100ECU (Experimental Currency Units) will be added to your earning from the round. Please note that this step is only for type S subjects. Type D subjects are never shown their opponent's moves vs the computer and never asked questions about their opponent's type.

3. Your choice for being the first or second mover vs the computer will be implemented. You will then commence playing the Avoid 13th game versus your computer as the first/second mover, as per your decision.

4. If you avoid removing the 13th item against the computer, you earn 500ECU from the first part of the round, regardless of your first mover/second mover decision. If you have to remove the 13th item, you earn 50ECU from the first part. 50ECU is the minimum payment for participating in each part of the round.

Part 2 of a Round: Play the human opponent

1. You will be shown the complete history of your human opponent's moves vs the computer and his/her resulting outcome in the first part of that round. You will again be asked about the type of your human opponent, i.e,. your human opponent's type is ______(S or D). A correct answer will add a further 100ECU to your earning from the round. (See the screenshot below). Please note that this step is only for type S subjects. Type D subjects are never shown their opponent's moves vs the computer and never asked questions about their opponent's type.

2. You and the human opponent will then play the AVOID 13th game with each other. Before the Avoid 13th game begins, both you and the human opponent will make the choice between being the first mover or second mover in the Avoid 13th game. One of your or the human opponent's first/second mover choice will be implemented with 50% chance each.

3. For the second part, if you AVOID Removing the 13th Item as the First mover, you earn 500ECU, while if you AVOID Removing the 13th Item as the Second Mover, you earn 200ECU. If you have to remove the 13th item, you earn 50ECU.

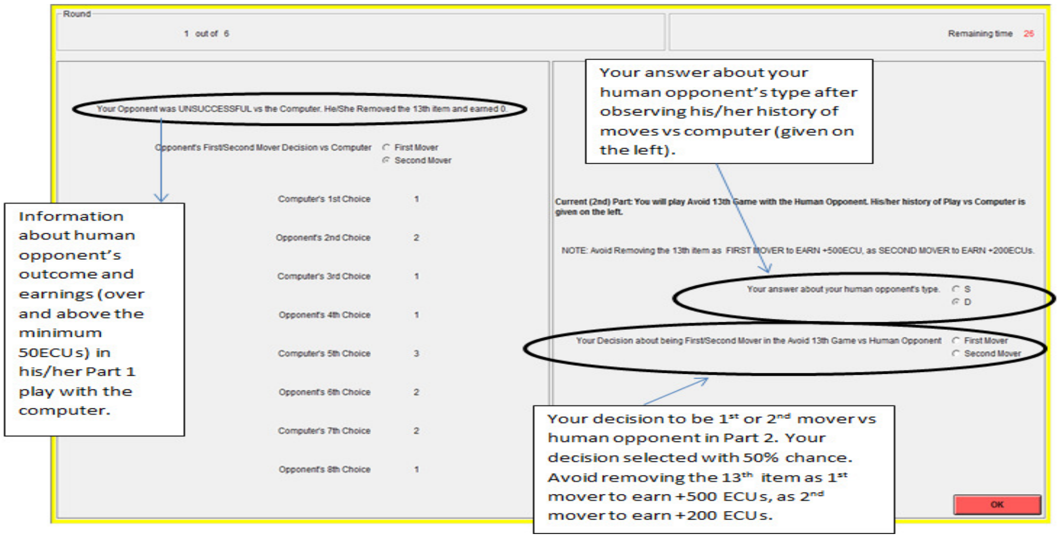

After the round ends, you will be randomly re-matched with another subject.

Your earnings for a round will be the sum of your earnings from the first and the second part of the round, including your earnings from the questions about your human opponent's type.
As stated above, this experiment will use three components:

First Component: Type S vs Type S. You will play 8 rounds (every round has two parts: play the computer, then play the human opponent) as described above. Every round your human opponent will be randomly redrawn from among the Type S subjects. That is, in the first component your opponent is always of type S. One of these 8 rounds will be randomly drawn to calculate your earning from this component.

Second Component: (Type S or D) vs (Type S or D). You will play 6 rounds as described above. Every round your opponent will be randomly redrawn from all the subjects in the experimental session. That is, in each round of the second component, your opponent may be of type S or type D. One of these 6 rounds will be randomly drawn to calculate your earning from this component.

Third Component: Risky choice study. This component is a short and unrelated study with 20 rounds. In each round you will be asked to make a choice among two options. According to your choice, you may end up losing some money earned in the previous two components. Press the 1 and 2 keys to make your choice, as explained below:

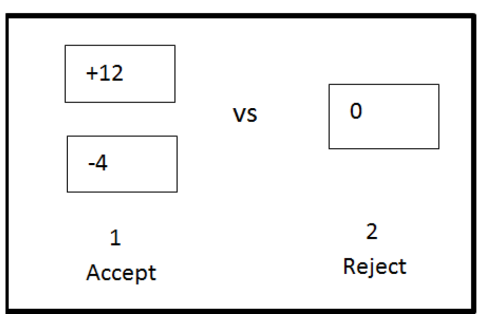

Option one, called ACCEPT (on the left of screen) consists of a possible reward and a loss. If you pick this option, the computer flips a fair digital coin (chances are 50-50). In case of heads, you earn additional money. In case of tails, you lose an amount that is subtracted from your previous earnings. To choose this option, press 1. Option two, called REJECT (on the right) is one reward. If you pick this option, you will get this reward for sure. To choose this option, press 2. At the end of this third component, one of these 20 rounds will be randomly selected by the computer. Your choice in that round will determine your earning in this third component. If you chose Accept in the selected round, your earning will be 20% of the amount drawn (reward or loss) from the computer's coin toss. If you chose Reject in the selected round, your earning will be 20% of the sure amount. Note: This last component will begin with a short practice with no payment. Your earnings will be the sum of your earnings from the three components plus the show-up fee ($5). The conversion rate for ECU in the first two components is: 1USD=120ECU. Your total earnings will be rounded up to the nearest dollar.