



Experience and insight in the Race game[☆]

Uri Gneezy^a, Aldo Rustichini^{b,*}, Alexander Vostroknutov^c

^a Rady School of Management, University of California, San Diego, United States

^b Department of Economics, University of Minnesota, 1925 4th Street South, 4-101 Hanson Hall, Minneapolis, MN 55455-0462, United States

^c Department of Economics, University of Maastricht, The Netherlands

ARTICLE INFO

Article history:

Received 18 May 2009

Received in revised form 19 March 2010

Accepted 12 April 2010

Available online 6 May 2010

JEL classification:

D830

C720

C900

Keywords:

Backward induction

Learning

Behavioral game theory

ABSTRACT

We study experimentally how subjects learn to plan ahead when they make sequential decisions. The task is the Race game. This game is played on a finite set of m possible positions occupied by a marker, which is initially in the first position. Two players alternate in the role of mover, and each one can move the marker forward by $1, 2, \dots, k$ places. The player who puts the marker in the final position wins.

Learning follows a similar pattern for all subjects. The experience of losses in early rounds induces them to switch the mode of analysis to *backward analysis*, which proceeds from the final position. The game has a simple dominant strategy, so we can calculate the frequency of errors made by subjects in each of the positions. The hypothesis that players follow a backward analysis gives precise predictions on the pattern of errors: for example that errors are more frequent, the further the position is from the end.

The experiment demonstrates that individuals are able to learn effective planning for future distant rewards, with a procedure of backward analysis. Their learning process may appear a pure insight, but is derived from evaluation of experience. Subjects are also able to transfer the knowledge they get from playing one game to a related game.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

We study experimentally whether people learn rational planning in a sequential choice problem. More important, we also study how they learn the optimal solution when they do, and whether they are able to use what they learned to understand optimal planning in similar problems. We are specifically interested in problems where actions that appear reasonable in the short run may have disastrous consequences in the future, and how people come to avoid those actions. The experiment is based on a simple game, the Race game, played on a finite set of positions, numbered 1 to m . The first player moves a marker from the initial position 1 by $1, 2, \dots, k$ places. No other move is allowed. Then the move goes to the other player, who can move the marker only forward by $1, 2, \dots, k$ positions. Players alternate in this fashion until one gets to the final position m and wins.

1.1. Do people learn to plan for future rewards?

Subjects in our experiment have to weigh whether to try and estimate from the start possible bad consequences of present actions, or wait until the problem is closer and easier to solve. The problem typically involves a real tradeoff. Worrying early

[☆] We thank audience at several seminars and conferences, and the insightful comments of two anonymous referees. The research was supported in part by the NSF grant SES 0924896 to AR.

* Corresponding author. Tel.: +1 612 625 4816; fax: +1 612 624 0209.
E-mail address: aldo.rustichini@gmail.com (A. Rustichini).

about future consequences of what we do presently may appear premature, and a seemingly natural intuition may lead to a rule of behavior opposite to careful planning: address a problem only when it is time to face it. This rule may be summarized by the ancient saying “I’ll cross that bridge when I come to it”. Such rule may not just be inspired by procrastination, but has instead an intuitive appeal on cognitive grounds: when the problem is close, its features are easier to appreciate, and a good solution easier to detect. Worrying about what to do too soon may just lead to unreasonable and ineffective fretting, before all the elements for an appropriate evaluation are available. Unfortunately, the sequence of short round actions taken initially may eventually lead to a situations in which all feasible outcomes are very bad. This is an important feature of real life economic problems. For example, if the bridge is retirement and crossing it means securing adequate financial cover after leaving the job, then the policy of waiting until the retirement is close is a recipe for disaster. Still, the policy of waiting to act until retirement is close seems to be the policy that is often used in real life financing of retirement (Yakoboski and Dickemper, 1997; Beeler and Lusardi, 2007). If even important decisions are not well planned, a crucial policy question is what may prevent people from adopting effective planning, and how they may learn more adequate ways to plan? Confronted with failure, people may adjust. Our experiment will test how many do that, and how fast.

1.2. *Where do insights originate?*

How do subjects take advantage of their experience to learn a better policy? Our experimental task has a unique solution, which is a dominant strategy. The strategy has an analytic form, simple to state but harder to see. Once presented to people who had not been able to discover it, the solution seems simple and clear, and subjects feel “they should have seen it”. People who do find it, may see it “in a flash”. The solution comes in an insight, defined as “the clear and sudden understanding of how to solve a problem” (Bowden et al., 2005). Two theories are provided to characterize the solution by insight in cognitive psychology. In one, solving a problem by insight is a unique cognitive process (with specific features, like relaxing self-imposed constraints, or chunk decomposition (Sternberg and Davidson, 1995; Knoblich et al., 1999). In another, insight solutions are achieved with the same cognitive mechanisms used in non-insight problems (Chronicle et al., 2004; MacGregor et al., 2001). The task we use allows us to study more closely how insight originates in the discovery of the solution. We propose here that insight is a flash which derives in a very natural way from the way people think about their previous experience.

1.3. *When and how do subjects learn to use backward analysis?*

The insight of subjects in our experiments takes a precise form: the switch from forward analysis to backward analysis. So understanding how and at what stage subjects solve a game is in our case a more specific question, namely how do subjects understand to proceed backward in the solution. We will see that this process is related but very different from the classical procedure of a backward induction solution of an extensive form game. This procedure starts by determining the optimal strategy in every decision node immediately preceding a final node, and replacing this final choice with the payoff induced by the optimal choice. This gives a new game to which the same procedure can be applied. The algorithm keeps rolling back until the initial decision stage is reached. In our game subjects will reason not on the extensive form of the game, which is extremely complex, but on the set of positions, which have a natural order.

1.4. *Do subjects transfer what they learn to different environments?*

The final question we address in this study is whether people transfer the learning they have acquired in a specific problem to a different one. This brings into the analysis a distinction between learning a strategy (that is an optimal way of behaving that is dependent on the general nature of the problem, but also on some specific parameters of it) and a rule (that is a more general procedure that applies to a class of problems). For example a person may acquire ability to play a bargaining game with two players, and may then find himself in a three players bargaining problem. The specific features of the strategy he has learned in the two players bargaining may not be directly useful, or might even be harmful, but the general approach to the strategic situation might very well be useful. Our experiment will allow a simple test of how able people are to transfer strategies and rules across strategic situations.

1.5. *Solving games*

There is a tradition of experimental testing of learning in games. A large literature documents deviations from the prescriptions of backward induction and subgame perfection. Typically, however, this evidence is based on experimental designs which confound different possible factors. First, in many cases deviations from rationality are confounded with distributional issues, making it difficult to identify their actual cause (e.g., in ultimatum games, see Güth et al., 1982). In this class of games the observed deviations from the subgame perfect equilibrium could be due to social preferences. Second, when the backward induction solution produces large losses of potential surplus (e.g., in the centipede game, see Rosenthal, 1982; Aumann, 1988) both players may be tempted not to follow backward induction, because its prescription results in a highly inefficient outcome. For this reason the game is not a good tool to study how people learn to reason backwards. Even a

player who understands the logic of the backward induction might not use it in the Centipede game if she thinks that the opponents are not doing the same (McKelvey and Palfrey, 1992).

Other games have been considered where efficiency issues are excluded: for example Fey et al. (1996) use as an experimental task a centipede game where the sum of the monetary payments is constant in every outcome. Still preferences over different allocations of the same surplus between the two player may matter: in their game the allocation of the constant amount becomes increasingly unequal as the game progresses, so players who have preferences over the joint distribution may consider the total welfare arising from a different allocation of the same amount differently.

A third concern in the literature studying backward induction is the complexity of the game. In the centipede game that we have discussed earlier, the game is very easy to analyze. Instead, in some decision problems the solution involves complex computations that players cannot make (or that are too costly to make relative to the benefit). As Simon (1957) noted, cognitive effort is costly, and thus it is not clear that the benefit of making the optimal decision is worth its cognitive cost.

Learning may help subjects and eventually lead to the optimal solution even when the game is complex: some empirical support for this can be found in Aymard and Serra (2001), who study the problem of optimal extraction from a renewable resource. They find that individual decisions are close to optimal, and individuals improve their performance with repetition. Their problem, like ours, has an analytic solution, and the authors can compare the behavior of subjects to its prescription. We try to use the solution to understand how subjects discover it. A minimal requirement for a study of whether and how subjects discover the solution to a game is that the theorist understands the solution and can compare the actual choices to the optimal choices.

The complexity of the game is important even if some of the players do understand the backward induction reasoning. Nagel (1995, 1998) used a Beauty Contest Game to estimate the degree of sophistication of players' actions (see also Crawford and Iriberry, 2007; Costa-Gomes et al., 2001; Costa-Gomes and Crawford, 2006). But in this class of games one needs to make an assumption about what others think. It is hard to disentangle the inferences the subject makes from the beliefs she has about the other players. Some evidence shows that experienced players adapt their level of sophistication to the players against whom they play (Slonim, 2005; Dufwenberg et al., 2005). For example, a Nim game is used by McKinney and Van Huyck (2007) to study the distribution of degrees of bounded rationality in a population of players.

1.6. The Race game

In this paper we study behavior in a game, called the Race game, that is immune to these problems. No fairness, distributional or efficiency issues arise, and the game does not involve high-level computations, just understanding the principle of reasoning back from the end. The game is zero-sum and the solution is in dominant strategies. This allows us to separate the reasoning of the players from the beliefs they have on the reasoning of others. The environment is simplified, and there is a cost and a benefit attached to this feature. The cost is clear: the implications for real life problems will have to be made carefully. The benefit, however, is that the optimal behavior is unambiguously defined, and the experiment can test how subjects learn the optimal solution. We use a game instead of a decision problem, because an essential part of the discovery of the solution is understanding that the behavior of others is not important if you hold a winning position.

Dufwenberg et al. (2010) also study a Race game. In their experiment subjects play $G(22, 2)$ ¹ and $G(7, 2)$, in the two possible orders. In the analysis of the results they emphasize the “epiphany” that some of the players experience, namely the sudden realization that the game may have a simple analytic solution (epiphany 1), which then prompts the search for such a solution. Epiphany is what we call here insight. Achieving the solution is a second epiphany, epiphany 2. The hypothesis they test is that subjects who play the simpler game first can achieve epiphany 1 earlier. We comment on the interpretation of the results in their and our experiment in the conclusions.

2. Learning in the Race game

2.1. Rules of the game

The game we use is a zero-sum, extensive form perfect information game, denoted by $G(m, k)$, where m and k are two integers with k smaller than m . Players alternate in moving. There is a state variable, which is a number between 1 and m . The initial state is 1. When a player moves, she can add to the current position any number between 1 and k . The player who gets to m first wins. Since this is a race to get to the position m first, we can call it the Race game. It is a combinatorial game, impartial, under the normal play rule. A *combinatorial game* is a zero-sum game with two players who alternate in moving, with a finite set of positions, and a set of feasible moves for each position, with no nature or random moves, and where some of the positions are end positions. The game is *impartial* if for every position the set of admissible moves is the same for both players. It has the *normal play rule* if the last player to move wins. The basic definitions and properties of these games are given and analyzed in Conway (2001).

¹ This is game G_{21} in their notation: their initial position is 0.

2.2. Winning strategy

A natural approach to solving such games is using *forward analysis*: You first consider the set of choices that you have, and where the different initial choices can take you. From the position where these lead, you can then consider what the other player would do. Then again you consider what you can do, and so on. This method is natural if you have no familiarity with the problem, because it proceeds from the current position, which is clearly defined, and gives you some familiarity with what lies ahead. As we are going to see, this is how most subjects in the experiment attack the problem. But in the Race game with m much larger than k the procedure quickly becomes very complex, and little useful information can be derived from it.

However, the game is finite, perfect information and zero-sum, with no ties, so there is a winning strategy for one of the two players. In $G(15, 3)$ the position is a winning position for the mover if and only if it is not 3, 7 or 11. The winning strategy is to move to the one among those three numbers that can be reached. In $G(17, 4)$ the position is a winning position for the mover if it is not 2, 7, 12, and again the winning strategy is to move to the one of these three numbers. All the other positions are losing positions, and what the player does there is irrelevant.

2.3. How the winning strategy can be discovered

Let us consider how this solution can be discovered in $G(15, 3)$: we claim that the process we outline provides a model for the learning process of subjects in the experiment. Consider a subject at his first try. By the rule of the game he will have eventually to move from one of the positions 11, 12, 13, or 14. At that stage, the game is easier to analyze: in these positions forward analysis provides the answer for the optimal choice and, more important, may give the key insight on how to solve the entire game.

In positions 12, 13 and 14 his choices is clearly no problem: he can reach the winning position 15 in one move. The really instructive position is position 11. If he looks forward to what follows after each of his feasible moves, he will realize that the opponent can win no matter what he does: Forward analysis easily provides the answer to what can be done in the subgame beginning at 11. Consider now the same subject in a later round. Again, by the rule of the game he will always have to move from one of the positions in $\{8, 9, 10, 11\}$. A subject that has gone through the previous chain of reasoning once, and finds himself in one of the winning positions 8, 9, 10, can now realize that he can put the opponent in position 11. Iteration of this process, which results from a combination of experience and analysis, can lead the subject to the discovery of the complete dominant strategy, or of important components of it, for example the strategy in positions closer to the final one.

2.4. Predictions

We can now formulate specific predictions on subjects' behavior in the experiment. Let us first define the variables we are going to measure. We classify the set of winning positions in three intervals separated by the three losing positions, numbered in the order in which they appear in the play. For example, in $G(15, 3)$ the first set is $\{1, 2\}$, the second is $\{4, 5, 6\}$, and $\{12, 13, 14\}$ is the fourth. In each of these sets the dominant strategy prescribes one and only one action: move to the position just to the right of the set: for example, move to 3 if in $\{1, 2\}$. Define error to be a deviation from the dominant strategy; obviously a player can make an error only when he moves from one of these subsets. The error rate is the empirical frequency of error; we can compute it for each of these subsets.

Our model predicts that the error rate on average over the rounds is decreasing in the distance from the final position, that is highest in the first set, and lowest in the last. This order on the frequency of errors is preserved over the experimental session: the frequency error rate is lower in the later stages of the game than in the early ones during the entire experimental session. We test these predictions in Section 4.1.

Learning the strategy in one game helps to understand the strategy in the other, similar game. The dominant strategy in $G(15, 3)$ is "move to 3, 7, or 11, whenever possible". This strategy is not useful in the game $G(17, 4)$; if applied leads to losing positions. Instead, the rule of discovering the strategy (proceed backwards from the end, which is a losing position by the rule of the game; go back by $k + 1$ steps, to find the next losing position, and continue to determine all losing positions at intervals of $k + 1$ positions) produces the winning strategy in every game. Subjects who discovered it through backward analysis should be able to transfer this understanding. Also, subjects who have a better and faster understanding of the strategy in the first game should be faster in understanding the strategy in the second game. We test these predictions in Section 4.2.

The second variable we can measure is the time to respond, that is the length of the time interval from the moment the opponent decides, to the moment of choice of the subject's next. The time to respond is a measure of the difficulty of the task the subject is facing. The time to think and respond follows the same pattern as the error rate: longer response times in positions further from the final position. Subjects may want to confirm the reasoning, so this pattern might persist over time. We test these predictions in Section 4.3.

3. Experimental design

Groups of subjects of different sizes in even number played in each experimental session. The sample for the entire study is 88 subjects, in 14 experimental sessions.

Table 1

Error rate in $G(15, 3)$, for the 72 subjects in the experimental sessions (excluding the talk-aloud). *Error* reports the mean of the mean error rate of the subject in the entire game. *Error 1* is the same mean but for the errors made when the subject was in position 1 or 2, *Error 2* for positions 4, 5 or 6; *Error 3* for positions 8, 9 or 10. *Obs* is the number of observations, *Std. err.* is the standard error.

	Obs	Mean	Std. err.	[95% Conf. int.]
Error	72	0.167	0.014	[0.138, 0.197]
Error 1	72	0.331	0.023	[0.284, 0.379]
Error 2	72	0.205	0.019	[0.165, 0.245]
Error 3	72	0.052	0.011	[0.029, 0.076]

In a first treatment we had a total of 72 subjects in 6 experimental sessions, with an average number of 12 subjects per session, ranging from a minimum of 6 to a maximum of 16. Participants were randomly and anonymously matched with another participant in their session; the pair was fixed for the entire session. Subjects played either one game or two games, one after the other. The first game was $G(15, 3)$, and it was played 20 times. The winner in the game was the player who won most rounds. The payment was 20 dollars to the winner, 0 to the loser, and 4 dollars in case of tie. In four of the six sessions after the first game, a second game $G(17, 4)$ was played 10 times. The players knew this description (the text of the instructions is reported in [Appendix 7.1](#)). At the end of the game players gave a short description of the strategy they had followed.

A second treatment was based on talk-aloud protocols. Subjects played the game $G(15, 3)$ for 10 rounds: but they also had to communicate to us their thoughts as they were contemplating what to do. The procedure is described in detail in Section 5. Each experimental session consisted of two players that never met. The payment was determined as in the previous treatment. A total of 16 subjects participated in this experiment. In both treatments the game was presented on computer with the z-Tree program ([Fischbacher, 2007](#)).²

3.1. Remarks on the design

Paying subjects when they win the majority of the games reduces the incentives for both subjects after one of the players has won 11 games. The cost is small: as we are going to see, one of the subject wins 11 games only late in the experimental session, when learning has typically already taken place. The benefit is large: subjects have a strong incentive to understand the optimal strategy because the returns are larger than they would be if we paid for each separate win. For the same reason, the total amount paid in case of a tie is only 8 dollars, instead of the 20 dollars paid when one of the subjects wins: this payment increases competition among the two subjects, and strengthens the incentive to look systematically for a way to win.

4. Results

We classify the type of error according to the position in which the error is made. So an error of type 1, or error 1 for short, is the error of a subject that in $G(15, 3)$ is moving from a position in the set $\{1, 2\}$ (and lands on the positions different than 3), and an error 4 is made by a subject moving from a position in the set $\{12, 13, 14\}$.

4.1. Error rate

For the $G(15, 3)$, [Table 1](#) shows the mean error rate, over subjects and rounds, in the three sets of winning positions.

The error rate in chance moves is 66 percent. The error rate among all subjects in the first move of the first round is 59 percent. [Table 1](#) shows large and significant differences in error rate. There are never any errors of type 4, in either of the two games, so the corresponding trivial statistic is never reported.

As predicted, the error rate is monotonic in the distance from the final position. The binomial confidence intervals show that the difference is significant. The next tables show that, also as predicted, the error rate decreases with the number of rounds, but the order is preserved: larger error rate for positions further from the final position. [Table 2](#) shows a large decrease between the first and the second half of the experiment.

The pairwise comparison between error frequencies gives further support to the claim that there is a significant difference among error rates. The Wilcoxon matched-pairs signed rank test of the hypothesis that the error rate is equal rejects the hypothesis at a p -value ≤ 0.0001 in all cases. For example for the entire experiment, error 1 is larger than error 2 ($z = 5.894$), error 1 is larger than error 3 ($z = 7.793$), error 2 is larger than error 3 ($z = 6.822$). The corresponding values for the early rounds (first 10) are $z = 6.822$, $z = 7.642$ and $z = 6.680$, respectively.

An additional test is provided by a panel data analysis of the choices of subjects over the rounds when they are in a winning position. The model used is the linear regression, to make the comparison with the results in [Table 1](#) easier. The

² Subjects were watching a screen displaying the 15 (or 17) positions, with the current position marked by an **x**. They would choose the number of steps from a set of options, and then confirm the choice with a click of the mouse. Once confirmed the move was final.

Table 2

Error rate $G(15, 3)$ early (first 10 rounds) and late (last 10 rounds) in the game. Variables are as described in Table 1.

	Obs	Mean	Std. err.	[95% Conf. int.]
Early rounds: from 1 to 10				
Error	72	0.280	0.021	[0.238, 0.323]
Error 1	72	0.518	0.030	[0.457, 0.580]
Error 2	72	0.339	0.030	[0.278, 0.399]
Error 3	72	0.090	0.021	[0.047, 0.133]
Late rounds: from 11 to 20				
Error	72	0.063	0.016	[0.029, 0.096]
Error 1	72	0.117	0.025	[0.066, 0.167]
Error 2	72	0.043	0.017	[0.008, 0.079]
Error 3	72	0.007	0.005	[-0.003, 0.017]

Table 3

Panel data analysis of error in game $G(15, 3)$.

	Mean error, b/se	Error early, b/se	Error late, b/se
Error 1	0.339 *** (0.017)	0.549 *** (0.029)	0.139 *** (0.014)
Error 2	0.193 *** (0.017)	0.352 *** (0.028)	0.025 * (0.014)
Error 3	0.042 ** (0.017)	0.083 *** (0.030)	0.005 (0.014)

* Significance at the 10 percent level.
 ** Significance at the 5 percent level.
 *** Significance at the 1 percent level.

variable indicating the interval in which the winning position is located (see Table 3) has an effect on error rate comparable to the one we have already determined.

The χ^2 test of the differences of the coefficients: $error1 = error2 : \chi^2 = 79.4, p < 0.0001$; $error1 = error2 : \chi^2 = 313.98, p < 0.0001$ and $error2 = error3 : \chi^2 = 81.87, p < 0.0001$.

In the last round of $G(15, 3)$ there is only one error over 145 total moves made by all 72 subjects from any winning position. So all the subjects learn the winning strategy after 20 rounds. This is confirmed by the brief statements that subjects wrote at the end of the experiment where they were asked to describe their strategy: in it, every subject described the optimal strategy as the strategy they were following. Clearly, subjects learn completely the dominant strategy in the course of the 20 rounds (for the $G(15, 3)$ game). Further insights into the process of discovery are provided by the analysis of the time pattern of the errors of different types, and by that of the response time in each position. Fig. 1 reports the average of the three different error rates for each round in the $G(15, 3)$.

In every round, the error rate for the first set of positions is larger than that for the second, and this in turn is larger than for the third and closest to the end. These data clearly indicate that subjects understand earlier the optimal policy in the later position, then in the intermediate one, and finally in the initial positions {1, 2}.

The later positions are easier: we claim that the knowledge of optimal policy in the late positions is used in the determination of the optimal policy in the early positions. An implication of this hypothesis is that each subject stops making




Fig. 1. Average error per round, $G(15, 3)$.

Table 4

Last round in which subjects make an error of type 1, 2, 3. *Last round 1* is the average value of the round number for the round in which subjects made an error of type 1 (that is an error in position 1 or 2); similarly for the others, for positions 4, 5 or 6, and 8, 9 and 10, respectively.

	Obs	Mean	Std. err.	[95% Conf. int.]
Last round 1	72	8.111	0.621	[6.871, 9.351]
Last round 2	72	4.722	0.485	[3.175, 5.691]
Last round 3	72	1.444	0.349	[0.747, 2.141]

Table 5

Error rate in $G(17, 4)$. Variables are as described in Table 1.

	Obs	Mean	Std. err.	[95% Conf. int.]
Error	52	0.152	0.030	[0.091, 0.212]
Error 1	52	0.264	0.048	[0.167, 0.360]
Error 2	52	0.213	0.044	[0.122, 0.303]
Error 3	52	0.026	0.012	[0.001, 0.052]

errors in the later positions earlier than she does in the later ones. Choice alone cannot tell precisely whether the subject has understood the strategy. The right choice may be done for the wrong reason, or by sheer chance. But the last round in which a subject makes an error in a specific position is a good indicator: certainly, he did not understand it earlier. Since the number of positions that he can move to is the same in all positions, doing the right thing by chance is equally likely in all positions, so the difference is a significant indicator. Table 4 reports the mean (over subjects) value of the last round in which subjects made an error of type 1, 2, and 3.

The average round for the last error of type 3 is less than 2: so on average subjects probably understood that the position 11 is a losing position after two moves. 85 percent made no error there after the 3rd round. For errors of type 1, the value is 8.11; so on average subjects probably understood the winning strategy half-way through the 20 rounds session. 72 percent made no error in the initial position after the 10th round.

4.2. Transfer of learning

A subset ($N = 52$) of the 72 subjects played 10 rounds of the $G(17, 4)$ game after the initial 20 rounds of $G(15, 3)$. A player has now four possible moves to choose from, and the dominant strategy prescribes a single action: so the probability of doing an error in a winning position for a random choice is 25 per cent. Still, the error rate is significantly lower than the one for $G(15, 3)$. Table 5 shows that the error rate from the first set of winning positions in the $G(17, 4)$ is 26.4 percent, compared to 33.3 in the $G(15, 3)$.

Since the set of feasible moves is now larger (four rather than three), the effective reduction in the error rate is even larger than indicated by these percentages.

Table 6 reports the error rate in first and second half of the rounds in $G(17, 4)$. The error rate is lower already in the early rounds, as compared to the one in the $G(15, 3)$ game (reported in Table 2).

In the later rounds (from 6 to 10) the error rate is zero for the third group of winning positions, and extremely low even for the initial ones. Subjects have learned the method of solution, and are able to apply it to the new situation. The classification of errors for the $G(17, 4)$ game is analogous to the one we gave for the $G(15, 3)$ game. There are no errors of any type after the sixth round in any play.

A subject-by-subject analysis of behavior in the two games supports the hypothesis that they are able to transfer learning from one game to the other. In Table 7 we consider two natural measures of understanding of the dominant strategy: the

Table 6

Error rate $G(17, 4)$ early (first 10 rounds) and late (last 10 rounds) in the game. Variables are as described in Table 1.

	Obs	Mean	Std. err.	[95% Conf. int.]
Early rounds: from 1 to 5				
Error	52	0.195	0.035	[0.126, 0.265]
Error 1	52	0.326	0.053	[0.220, 0.433]
Error 2	52	0.258	0.051	[0.154, 0.361]
Error 3	52	0.035	0.015	[0.004, 0.066]
Later rounds: from 6 to 10				
Error	52	0.064	0.033	[−0.002, 0.130]
Error 1	52	0.057	0.032	[−0.007, 0.123]
Error 2	52	0.038	0.026	[−0.015, 0.092]
Error 3	52	0.00	0.000	[0.000, 0.000]

Table 7

Transfer of learning: subject-by-subject analysis of last error in position 1 and mean error, in game G(15, 3) and G(17, 4) compared.

	Last error in G(17, 4), b/se	Mean error in G(17, 4), b/se
Last error in G(15, 3)	0.206*** (0.047)	
Mean error in G(15, 3)		0.613*** (0.216)
R ²	0.282	0.139
N	52	52

last round in which the subject made an error in the winning position 1 and the mean error in the game. The table reports, for each of these two measures, the regression of the value in the G(15, 3) and the G(17, 4) game.

As predicted, there is a strong and significant correlation between these performance measures for each subject in the two games.

4.3. Response time

An additional measure that can give insight into the solution process is the response time in each move: this is the length of the time interval between the moment in which the opponent moves, and the subject's move. The subject has all the time he wants to respond, so this is really a choice variable. Many factors potentially affect the length of the decision. One important component is the length of the analysis that the subject performs at that position. Another possible factor, the realization that the play is probably going to end with a loss, may stimulate the subject to search more accurately for a possible way out, inducing a longer reaction time in the losing positions.

The analysis of the response time shows that subjects think longer in the losing positions: precisely where the move they make is irrelevant, although at first they do not understand this clearly. These are the moments in which they seem to learn more about the game: they learn the solution because at the losing position they learn that no matter what they do, the opponent will win. Fig. 2 compares the average response time in different rounds, for the set of all positions and for the two separate groups of winning and losing positions. With the exception of the first round, the value for the winning positions is always below the average, and the opposite holds for the losing positions.

For the G(15, 3) the average response time in winning positions is 6.69, the one in losing positions is 9.22. In the non-parametric WMW test this difference is significant ($p < 0.00001$).

The hypothesis that the process of discovering the solution follows a backward inference direction predicts that the thought process should be longer in the initial stages when the game is in the final positions. In this case subjects have to (and can) understand the consequence of a certain final position. In contrast, in these early stages thinking in the initial positions is relatively unproductive: the conclusion is too far for subjects to derive anything useful from it. They should be aware of this, and they should spend relatively less time in thinking in these cases. In later stages, the conclusions reached in the earlier stages can then be used as an "input" in the inference about the value of being at some intermediate position. So in the later stages the response time should be longer in the intermediate positions than in the final ones, and longer than it was in the same positions in the initial stages.

Let us now consider the evidence, beginning with the first game (Fig. 3).

The response time in position 11 peaks in the initial round and then declines slowly. It is 6.9 s in the first 10 rounds, and 3.9 in the last 10. The time is short also for subjects that were not moving at 11 in the first round. We use a Wilcoxon–Mann–Whitney non-parametric test to study the hypothesis that the response time in the second round in position 11 is different for subjects who had been moving at position 11 in the first round, and those who were not. The difference is not significant (the p -value is 0.401). So both subjects seem to have made their inference when one of them found himself in position 11.




Fig. 2. Average response time in win and lose, G(15, 3).




Fig. 3. Average response time by round, $G(15, 3)$.




Fig. 4. Average response time by position after the right move is learned, $G(15, 3)$.

The response time in position 7 (3) initially increases from round 1 to 3 (4) to reach a peak at game 4 (5). Both times are higher than the average time. The time for position 3 remains below that in position 7, probably because at position 7 subjects realize that the game is closing, and they are making an effort to avoid losing.

We finally consider in more detail the relationship between the backward analysis that subjects are using and the well known procedure of backward induction solution of extensive form games. The backward induction algorithm replaces subgames with final nodes giving equilibrium payoffs in those subgames. If subjects were following this procedure we would expect that the response time in a given position, in the rounds after the right move was learned, is the same as the response time of a last move. An alternative hypothesis is that subjects, who have learned certain position, do not immediately cut the following subgame but rather trace their prediction forward each time they are in that position. In this case we would expect the response time in a position to be proportional to the distance of that position from the end of the game.

Fig. 4 shows the average response times in each position of $G(15, 3)$. These averages are taken across all subjects and for each subject the rounds that enter the averages are only those that follow the round at which the last mistake was made. Thus, these are the average response times of subjects who learned the right strategy.³

From positions 12, 13 and 14 to position 7 the average response time grows with the distance to the end of the game. At position 7 the average response time is 8.24. At positions 8, 9, 10 it is 6.59 and it is 4.00 at positions 12, 13, 14. The 95% confidence intervals are disjoint for positions 12, 13, 14; 11; and 8, 9, 10. This finding is consistent with the hypothesis that subjects check if their strategy is right each time they are in these positions.

5. Talk-aloud protocol

A way to reconstruct the thought process of the subjects is to ask them to state their thoughts as they are deciding. This method has a long tradition in the experimental analysis of decision making, although it is less common in experimental economics. The method was first proposed by Ericsson and Simon (1980), and is described in detail in Ericsson and Simon (1993).

³ No mistakes can be made in the positions 3, 7, and 11. For these positions the rounds chosen are the same as in the preceding set of positions. For example, the rounds averaged for position 11 are the same rounds as for the positions 8, 9, 10. The idea is that if subject understands what to do in 8, 9, 10 he should also understand what to do in 11.

Table 8

Times and rounds of discovery of the components of the optimal strategy in the protocol analysis. Experimental sessions were run in two separate consecutive months: the first column reports the group to which the session belongs. The column labeled “Position 11 (7, 3, respectively) time” reports the time in which the subject explicitly states that 11 (7, 3) is a losing position. The column labeled “Position 11 (7, 3) round” reports the corresponding number of the round. The column labeled “total time” reports the time length of the experimental session.

Pair	Subj.	Position 11		Position 7		Position 3		Total time
		Time	Round	Time	Round	Time	Round	
1	1	2.1	1	2.1	1	2.1	1	23.12
1	2	15.24	6	16.5	7	Never	Never	23.12
2	1	2.25	2	5.18	3	9.51	4	17.54
2	2	10.55	5	12.2	6	14.2	8	17.54
3	1	2.19	1	14.51	6	20.02	9	23.37
3	2	2.32	1	8.41	3	21.2	9	23.37
4	1	3.09	1	9.4	3	18.14	5	27.25
4	2	3.55	1	20.14	5	Never	Never	27.25
5	1	1.85	1	12.07	6	12.07	6	17.1
5	2	3.6	3	12.27	6	Never	Never	17.1
6	1	4.55	3	7.37	5	14.3	10	23.51
6	2	1.1	1	2.3	2	2.3	2	23.51
7	1	0.46	1	5.32	2	9.26	4	25.17
7	2	4.51	1	17.3	9	Never	Never	25.17
8	1	3.26	1	3.26	1	6.16	2	21.13
8	2	9.09	2	15.5	4	20.1	10	21.13

We run additional sessions with subjects playing, in a fixed pair, 10 rounds of the $G(15, 3)$ game. The design of the experiment was otherwise identical to the one described in Section 3 above. The instructions describing the rules of the game were read first. After that, the experimenter asked the subjects to talk aloud during the experiment. The specific instruction was to “say out loud what you say to yourself silently. Act as if you are alone in the room speaking to yourself.” The instructions explaining what the subjects had to do in the Talk-Aloud protocol are reported in [Appendix 7.2](#).⁴ The explanation was followed by two simple practice examples: multiplication of two integers and solution of an anagram. During the sessions, subjects who remained silent were reminded to report their thoughts. The statements of the subjects were recorded on digital voice recorders (Olympus DM-1), placed in full view of the subjects. Experimental sessions were run in different days, and one pair of subjects participated in each session. Subjects knew that there was a single other participant in each session. Overall, we have run 8 experimental sessions with 16 subjects in total.

Protocol analysis may of course affect the thought process. This issue is discussed at length in [Ericsson and Simon \(1980, 1993\)](#), see in particular the chapters 1 and 2 of their book. In our case a comparison of the performance of subjects in this treatment with those in the usual laboratory setup can settle the issue. The response time is obviously longer (these subjects have to talk before they move) in the protocol group: the average in this group is 16.49 s against 6.55 of the other subjects. The error rate, however, is not significantly different. The overall rate is 7.66 for the protocol group, versus the 7.56 for the rest (p -value in the WMW test: 0.9225). Detailed statistics for each single type of error gives the same result.

[Table 8](#) gives an estimate of the order in which the components of the optimal strategy are discovered. As the instructions indicate (see [Appendix 7.2](#)) subjects were asked to report regularly their thoughts, and in particular before the choice of a move they were asked to provide, if possible, a reason to justify their move, and to state explicitly when they thought they had lost the game, and why.

The first important result in the verbal reports is the regular use in the early stages of the forward mode of analysis of the problem. Many subjects (5 out of 16) choose the moves in the initial games after going through a complex list of possible sequences of choices of both players in the next moves. All the subjects except one (who discovered the optimal strategy before the first move, by analyzing (backwards) the game) make their first move explicitly stating that they want to see what will happen next. The analysis in this phase is characterized by an arbitrary elimination of some of the possible moves: the tree is too rich for the subjects to keep track of all the potential moves. It is interesting to note that a frequent rule for ignoring some of the possible developments of the game is wishful thinking: a subject considers as likely a move of the opponent that will make him win.

The second important result is the similarity of the motivation for a change in the mode of the analysis. In all cases, the experience that subjects gained when the game reached the position 11 spurred the insight that the game can be solved backward. The analysis is more explicit when the subject analyzing the situation is the one who has to move. However, both subjects understand that for any chosen move the opponent will be able to win. Some of the subjects understood the key role of position 11 from inspection. Interestingly, some of these subjects were still following a forward analysis analysis when, later, they found themselves in the initial positions.

The third result is the backward inference nature of the analysis. This pattern is confirmed in each instance. Subjects understand the implication of being at 11 the first time they face it, carry over this implication for the position 7 in the next

⁴ They follow closely the protocol described in [Ericsson and Simon \(1993\)](#), Appendix, pages 375–379.

round. This inference is explicit in the statements they make; as we are going to see in the next paragraph, subjects repeat each time the inference in its entire length, to check and confirm the accuracy of their reasoning. Each time a new piece of the inference is understood, they repeat the inference with the new step added. The inference for position 3 takes longer (two or three plays). The average length of the entire experiment (with 10 rounds) was less than 20 min. Most of the thinking is done in the first 4–5 rounds.

The fourth result is the observation that all subjects after they discover the dominant strategy feel the need to repeat, in every play, the detailed sequence of the possible moves that will follow. Upon moving to position 7 a subject will typically repeat the reason why, for example, the opponent will then be forced to move short of 11, allowing a successive move to 11, which will then be followed by a move to 12, 13 or 14 and then a final move. Subjects need to confirm the same fact several times over.

6. Conclusions

The results show that a process that proceeds backwards (Backward analysis) is the main cognitive procedure to which all subjects converge, as they learn the solution of this simple, finite, zero-sum combinatorial game. The speed of the convergence varies across subjects, but subjects converge to this common pattern.

Three results of our analysis support this conclusion. The first is that error rates in the successive positions decline at different rates: zero error rate for the position close to the end, slower for position one step away from the end, and even slower for positions more removed. The second result is that the response time in the position 11 is high in the initial round, and declines afterward. This is followed by a peak in the data for position 7 and then a later peak for those in position 3. The third result is the detailed analysis of the verbal reports, which confirms that subjects discover that position 11, 7, and 3 are losing positions in sequence, and in that order. They also use the conclusion reached for each of these positions to infer that the next one is also a losing position, so they use the fact that 11 is a losing position to conclude that 7 is a losing position.

What do these results tell us about planning in real life decisions, like retirement? An important aspect of the learning process in our experiment is that both the length of the time required to reach the correct conclusion, as well as the errors made are increasing in the length of the inference itself. The size of the difference is substantial: the length of time increases in a way which is proportional to the distance from the end. Learning the optimal move in the final positions is immediate, and error free. In all plays, for all subjects, in both games, there were no errors in this stage. Learning the optimal move for the second set (for example, of positions 8, 9 and 10 in the first game) takes a little longer, and a few mistakes occur in the initial stages, and so on. While in an abstract Backward induction procedure the elimination of the nodes closer to the end reduces the complexity of the inference in the next round, this does not seem to happen in the real thought process that subjects have to perform each time.

Why do subjects fail to do this simple inference? The evidence from the analysis of verbal reports indicates an important reason. Consider for example the case of a subject who already has some understanding that whoever is in position 11 is going to lose. Subjects need to examine the future sequence of moves from position 11, as a way of checking that they are not mistaken in concluding that the player who moves to 11 has won. This is a reasonable procedure: they are after all learning a game, and at each stage they want to make sure that something they have learned in the recent past has not changed the conclusion that they have reached previously (for example, that 11 is a winning position). Of course the fact that they are not completely sure of this conclusion prevents them from simply replacing (as the Backward Induction procedure would require) position 15 with position 11.

This uncertainty is even greater in decision problems like retirement, and it is a potentially important reason for the failure of effective planning in these domains. In real life problems there is no clear-cut, definite partial conclusion that can be substituted into the original problem. Having some income 10 years away from retirement is no guarantee of a precise income at retirement.

Dufwenberg et al. (2010) focus on the idea that learning the dominant strategy in complex games is possible only with the insight that the game has an analytic solution (epiphany 1), and this insight is derived from playing first a simpler one. We note that learning occurs in a natural way in complex games because even complex games have simpler games built in: for instance in $G(15, 3)$ a player must find himself either in the subgame beginning at position 11, which is equivalent to the simple $G(5, 3)$, or in a game beginning at 8, 9, or 10, which are equivalent to $G(8, 3)$, $G(7, 3)$ and $G(6, 3)$, respectively. Since in every round subjects play at some point a simple game, they can derive the insight from this experience. Since there are four different games possible, learning is harder. In fact our analysis suggests that the first and second epiphany are prepared by the experience in the previous plays of the same game. The first epiphany is prepared because the position 11 teaches the players that there are positions in which a player has lost no matter what he does (if the other plays optimally). The second epiphany is prepared because once that insight on position 11 is reached the player only needs to understand that he is in fact playing a Race to 11.

A last lesson is worth pointing out. There is increasing evidence (e.g., Banks and Oldfield, 2007; Lusardi and Mitchell, 2007; Benjamin et al., 2007; Dohmen et al., 2007; Burks et al., 2009) that differences in cognitive abilities beyond simple literacy and numeracy correlate significantly with the ability to plan effectively in financial and economic decisions. These papers provide a link between the abstract cognitive ability tested in our experiment and the real economic decisions we try to understand. It is interesting to note that Burks et al. (2009) use precisely the game examined in this paper as one of the measures of cognitive ability. This measure (called Hit 15 in the paper) is a better predictor of the behavior in economic

choices and economic performance than standard economic variables (like alternative income available) or demographic characteristics (like marriage status). For a measure of job performance, persistency in employment, this measure is a better predictor than other measures of cognitive skills, such as the IQ score measured with Raven's matrices. The results of this paper indicate how this effect may be produced. Neither a centipede game nor a beauty contest game seem to provide such test.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2010.04.005.

References

- Aumann, R., 1988. Preliminary notes on integrating irrationality into game theory. Mimeo. In: International Conference on Economic Theories of Politics, Haifa.
- Aymard, S., Serra, D., 2001. Do individuals use backward induction in dynamic optimization problems? An experimental investigation. *Economic Letters* 73, 287–292.
- Banks, J., Oldfield, Z., 2007. Understanding pensions: cognitive function, numerical ability and retirement saving. *Fiscal Studies* 28, 143–170.
- Beeler, J., Lusardi, A., 2007. Saving between cohorts: the role of planning. In: Madrian, B., Mitchell, O., Soldo, B. (Eds.), *Redefining Retirement: How Will Boomers Fare?* Oxford University Press, Oxford, pp. 271–295.
- Benjamin, D., Brown, S., Shapiro, J., 2007. Who is “behavioral”? Cognitive ability and anomalous preferences. Working paper, SSRN series. Available at: <http://ssrn.com/abstract=675264>.
- Bowden, E.M., Jung-Beeman, M., Fleck, J., Kounios, J., 2005. New approaches to demystifying insight. *Trends in Cognitive Sciences* 9, 322–328.
- Burks, S., Carpenter, J., Gooote, L., Rustichini, A., 2009. Cognitive skills affect economic preferences, social awareness, and job attachment. *Proceedings of the National Academy of Science* 106, 7745–7750.
- Chronicle, E.P., MacGregor, J.N., Ormerod, T.C., 2004. What makes an insight problem? The roles of heuristics, goal conception, and solution recoding in knowledge-lean problems. *Journal of Experimental Psychology: Learning, Memory and Cognition* 30, 14–27.
- Costa-Gomes, M.A., Crawford, V.P., Broseta, B., 2001. Cognition and behavior in normal-form games: an experimental study. *Econometrica* 69, 1193–1235.
- Costa-Gomes, M.A., Crawford, V.P., 2006. Cognition and behavior in two-person guessing games: an experimental study. *American Economic Review* 96, 1737–1768.
- Conway, J., 2001. *On Numbers and Games*. A.K. Peters Ltd., Wellesley, MA.
- Crawford, V.P., Iriberry, N., 2007. Fatal attraction: salience, naivete, and sophistication in experimental “hide-and-seek” games. *American Economic Review* 97, 1731–1750.
- Dohmen, T., Falk, A., Huffman, D., Sunde, U., 2007. Are risk aversion and impatience related to cognitive ability? IZA DP No. 2735.
- Dufwenberg, M., Lindqvist, T., Moore, E., 2005. Bubbles and experience: an experiment. *American Economic Review* 95, 1731–1737.
- Dufwenberg, M., Sundaram, R., Butler, D., 2010. Epiphany in the game of 21. *Journal of Economic Behavior and Organization* 75, 132–143.
- Ericsson, A., Simon, H., 1980. Verbal reports as data. *Psychological Review* 87, 215–251.
- Ericsson, A., Simon, H., 1993. *Protocol Analysis, Verbal Reports as Data*. MIT Press, Cambridge, MA.
- Fey, M., McKelvey, R., Palfrey, T., 1996. An experimental study of constant-sum centipede games. *International Journal of Game Theory* 25, 269–287.
- Fischbacher, U., 2007. Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10, 171–178.
- Guuuth, W., Schmittberger, R., Schwarze, B., 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3, 367–388.
- Knoblich, G., Ohlsson, S., Haider, H., Rhenius, D., 1999. Constraint relaxation and chunk decomposition in insight problem solving. *Journal of Experimental Psychology: Learning, Memory and Cognition* 25, 1534–1555.
- Lusardi, A., Mitchell, O., 2007. Baby boomer retirement security: the role of planning, financial literacy, and housing wealth. *Journal of Monetary Economics* 54, 205–224.
- MacGregor, J.N., Ormerod, T.C., Chronicle, E.P., 2001. Information processing and insight: a process model of performance on the nine-dot and related problems. *Journal of Experimental Psychology: Learning, Memory and Cognition* 27, 176–201.
- McKelvey, R., Palfrey, T., 1992. An experimental study of the centipede game. *Econometrica* 60, 803–836.
- McKinney, C.N., Van Huyck, J., 2007. Estimating bounded rationality and pricing performance uncertainty. *Journal of Economic Behavior and Organization* 62, 625–639.
- Nagel, R., 1995. Unraveling in guessing games: an experimental study. *American Economic Review* 85, 1313–1326.
- Nagel, R., 1998. A survey of research on experimental beauty-contest games. In: Budescu, D., Erev, I., Zwick, R. (Eds.), *Games and Human Behavior: Essays in Honor of Amnon Rapaport*. Lawrence Erlbaum Associates Inc., New Jersey, pp. 105–142.
- Rosenthal, R., 1982. Games of perfect information, predatory pricing, and the chain store paradox. *Journal of Economic Theory* 25, 92–100.
- Simon, H.A., 1957. *Models of Man*. Wiley, New York.
- Slonim, R., 2005. Competing against experienced and inexperienced players. *Experimental Economics* 8, 55–75.
- Sternberg, R.J., Davidson, J.E., 1995. *The Nature of the Insight*. MIT Press, Cambridge, MA.
- Yakoboski, P., Dickemper, J., 1997. Increased saving but little planning: Results of the 1997 retirement confidence survey. Employee Benefit Research Institute, Washington, DC, EBRI Issue Brief No. 191.