

Games and Phone Numbers: Do Short Term Memory Bounds Affect Strategic Behavior? *

Giovanna Devetag[†]

Department of Management and Computer Science

University of Trento

Massimo Warglien

Department of Business Economics and Management

Ca'Foscari University of Venezia

and

Cognitive Science Laboratory, Rovereto

JEL codes: C72, C92

Keywords: short term memory, information processing, iterated reasoning, heterogeneity

*Financial support from the Italian Ministry of Scientific and Technological Research (MURST) and the Cognitive Science Lab of Rovereto is gratefully acknowledged. We thank Marco Tecilla and Macrina Marchesin for precious help in the implementation and conduct of the experiments. The usual disclaimer applies.

[†]Corresponding author. Phone number: +39-0461-882109; fax: +39-0461-882124; email: devetag@cs.unitn.it.

Abstract

Research in experimental and behavioral game theory has revealed a substantial and persistent degree of heterogeneity in the strategic behavior of real individuals. While the prevailing theoretical explanations of the observed heterogeneity typically invoke underlying differences in beliefs among the population of players, we argue that a further source of heterogeneity may consist in the individuals' different ability to process information, of which short term memory capacity provides a measurable proxy. Research in cognitive psychology has shown that individuals typically differ in their short term memory capacity; furthermore, short term memory capacity provides a fundamental cognitive bottleneck to our ability to process information efficiently and hence seems correlated with performance in a variety of problem solving and reasoning tasks. In this paper we conduct experiments on a set of well-known games whose solution concepts require the application of some paradigmatic forms of strategic reasoning, such as iterated dominance and backward induction. We separately conduct standard short term memory tests on our subjects to detect the presence of a correlation between individuals' behavior in the games - here defined in terms of degrees of conformity to the standard game-theoretic prescriptions - and their short term memory score.

Our results show the presence of a significant and positive correlation between subjects' short term memory score and conformity to standard game-theoretic prescriptions in the games, thus confirming our hypothesis. While the robustness of our conjecture awaits to be confirmed by further data gathering in more interactive experimental settings, our preliminary results suggest a promising line of inquiry on the interconnections between information processing capacity and strategic behavior.

1 Introduction

Bounded rationality is about the limits of human information processing; it is not a generic deviation from the “olympic” rationality of neoclassical economics, but the result of intrinsic limitations in our cognitive architecture. It is resource-limited, not data-limited (Norman and Bobrow, 1975). In opposition to the widespread emphasis on the scarcity of information, Herbert Simon repeatedly argued that the most important problem is quite often abundance of information, given our limited processing capacities (Simon, 1978).

Within such view, a key research task is to single out the processing bottlenecks responsible for bounded rationality, to find parameters characterizing them, and to assess how such parameters affect cognitive performance. Simon criticized psychologists for directing too much efforts to hypothesis testing and too little attention to parameter evaluation (Simon, 1981). And he had little doubt that the most important constraint affecting human information processing is short term memory capacity.

The bottleneck (. . .) must lie in in the small amount of rapid-access storage (so-called short-term memory) available and the time required to move items from the limited short-term memory store to the large-scale long-term store” (Simon 1996, p.61)

Short term memory (STM) plays a central role in Simon’s view of the human information processing architecture. In his joint work with Allen Newell (Newell and Simon, 1972), he attributes to STM the key role of holding the input and output symbol structures of elementary information processes, thus being directly involved in all kind of processing activities.

While Simon himself dedicated some experimental research to STM (Simon, 1978 and 1989), his prototypical examples were from the domain of “toy” problem solving puzzles and learning tasks. There seems to be a surprising lack of research dedicated to the impact of STM capacity on economic interactive behavior. The recent resurgence of interest in bounded rationality among economists doesn’t fill this gap. Attempts at modeling “limited recall” and memory on deci-

sion making (Rubinstein, 1998) only address forgetting processes related to long term memory. Models trying to incorporate tradeoffs between computational complexity and rationality come closer by introducing explicitly the cost of information processing by finite automata (Abreu and Rubinstein, 1988). Complexity arises because automata have very limited working memory capacity. However, while in these models working memory capacity affects the execution of a strategy, it doesn't affect the thinking process by which a strategy is designed. Even more, such models typically assume that individuals can control complexity and have no constraints in making an accurate economic assessment of the costs and benefits of increasing complexity.

Experimental economists provide a more sympathetic view. A few experimentalists (e.g., Nagel 1995; Ho, Camerer and Weigelt, 1998) have suggested that individuals may differ in terms of the number of iterated steps of thinking they may perform. Since strategic reasoning in games usually implies some kind of recursive thinking, individual differences in the ability to perform the steps of such recursion may explain behavioral heterogeneity. While STM constraints are not usually mentioned (see Camerer, Ho and Chong, 2001 for an exception), they are natural candidates to explain difficulties in retaining the results of former reasoning steps while new steps are being performed.

The well-documented variance in individual STM capacity (Baddeley, 1990) suggests a potentially relevant source of heterogeneity in the cognitive and behavioral performance of individuals. While there is a growing attention to agents heterogeneity in game-theoretic modelling, it is usually assumed that such heterogeneity arises from differences in preferences, in the decision rules adopted and in beliefs about the behavior of others (e.g., Costa-Gomez, Crawford, and Broseta, 2001; Stahl and Wilson, 1995). We instead suggest that differences in information processing capacity may be a stable source of heterogeneity in interactive behavior.

In this paper we report some preliminary, exploratory results from a broader research project aimed at evaluating the impact of STM capacity on strategic behavior. In a related experiment (Devetag and Warglien, 2001), we have investigated a task in which subjects had to construct

a representation of relational structures isomorphic to the payoff order structures of some common normal form games. We found significant correlation between individual STM capacity and subjects' ability to construct correct representations of relational structures of increasing complexity - this hints at the fact that the individual ability at editing a correct mental representation of a game might be affected by STM constraints. In the experiments reported here, we investigate the relations between STM capacity and individual rationality in solving games that require different forms of iterated reasoning. We consider only games of perfect information, in which deviations from rationality cannot be attributed to a lack of information, but only to how information is encoded and processed. We find that STM capacity is significantly correlated to individual performance in such games.

The concept of a short term memory has been a subject of considerable controversy in the psychological literature (see Baddeley, 1990 for a classical discussion). However, while the existence of a separate, central psychological store of short term memory has been widely debated, there is a considerable consensus on the existence of STM phenomena revealing a limited capacity to store temporarily information while performing more complex cognitive tasks, and on the individual variability of such capacity. We will not deal here with the deeper psychological (and neurological: e.g., Cohen et al., 1997) issues raised by the former controversy - more pragmatically, we will only be looking at measures of STM capacity as a proxy of individual information processing bounds.

The paper is organized as follows: section 2 briefly describes the games that were presented to subjects; section 3 describes the experimental design; section 4 presents the results and finally section 5 offers some concluding remarks.

2 The Games

We selected three simple instances of games whose solution concepts require the application of some form of iterated reasoning. In all of the games, subjects had to interact with a hypothetical,

rational opponent who would always choose the payoff-maximizing strategy (motivations for such experimental design are given in section 3).

We briefly list and describe each of them in sequence:

1. Iterated Dominance

The following bimatrix illustrates the game:

Insert table 1 about here

In the game (see Warglien, Devetag and Legrenzi, 1998), the COLUMN player has a dominant strategy, C. If the ROW player realizes this, she should play B. The game is then solvable by one step of iterated dominance, and the unique equilibrium solution is (B,C). However, if the ROW player does not realize that her opponent has a dominant strategy, or if she is uncertain about the opponent's rationality, she may be tempted to choose A to assure herself a sure gain of 60. In a previous experiment on this game, 70% of the ROW players had chosen the 'safe' strategy A, against the 30% who had picked the correct dominance-solvable equilibrium solution Y. These data might be interpreted either as the result of excessive risk-aversion on the part of the ROW players, or, at least in part, by the fact that some of the ROW players simply did not realize that their opponent basically 'had no choice', because they did not focus on their opponent's available strategies and payoffs before deciding their own strategy. In other words, some of the players may not perform the single step of iterated reasoning that the game solution would require. In a de-biasing treatment (Warglien, Devetag and Legrenzi, 1998), subjects were explicitly asked to formulate a belief about the opponent's most likely move before deciding; in this treatment, the percentage of subjects picking the correct choice raised up to 70%, thus giving some support to the hypothesis that subjects in the control treatment did not focus spontaneously on their opponent's likely choice before deciding, but when prompted to do so, they were able to perform

on step of iterated reasoning correctly. Evidence on violations of even one or few steps of iterated dominance abounds in the experimental literature (see Camerer, 2000, ch. 5 for a review); the choice of this particular example of payoff matrix should eliminate all those instances in which violation of dominance occurs because of the existence of social preferences (like, for example, cooperation in the PD). Furthermore, telling subjects that their opponent is a payoff-maximizing algorithm should likewise eliminate choices caused by uncertainty in the opponent's level of rationality.

2. Common Knowledge: The “Dirty Faces” Game

This game is another example to test our subjects' performance in tasks requiring iterated steps of reasoning. The game takes its name from a story originally developed by Littlewood (1953), and it is often cited in discussions about common knowledge of rationality (e.g., Geanakoplos, 1994; Binmore and Brandenburger, 1990). In the experiment we used a modified version drawn by Fagin et al. (1995). In this version of the story, two children are playing with mud, and they both get “dirty faces”; each child can see the other child's dirty face but not his own. The baby-sitter arrives suddenly and announces: “At least one of you has a dirty face. Don't talk! You can only talk to answer this question: “Do you have a dirty face?” If you tell the truth, you won't be punished”. The two children both remain silent, and the baby-sitter then repeats the question. At this point, both children respond that they have a dirty face and avoid to be punished. In fact, both children can see the other's dirty face, and common knowledge that at least one of them has a dirty face results from the baby-sitter's announcement. Then each child is able to infer from the other child's silence that he himself has a dirty face (in fact, if he didn't have a dirty face, the other child would have immediately inferred that his own face was dirty after the baby-sitter's announcement and he would have confessed it). Restated in game-theoretic jargon, both players can be of one of two types, X and Y, and they both have an incentive to

take a certain action (telling the truth about their type to avoid punishment) as soon as they know their type with certainty. Common knowledge that at least one of the players is of type X assures to players that they can infer their type by observing the actions of others in the first round of play (confessing or remaining silent). The best outcome for both players is to confess in the second round. However, two requirements are necessary in order for this outcome to occur: common knowledge that at least one player is of type X, and common knowledge of rationality. Each player must hence perform a series of iterated reasoning steps to arrive at the solution.

Weber (2001) tested experimentally a modified version of the game, varying group size, the number of required steps of iterated reasoning, and allowing repeated play to test learning. He found that, while the majority of subjects conform to the theoretical prediction in the simplest case, the observed percentage of behavior consistent with theory sharply declines as more steps of reasoning are required. In addition, he found no group size effect and no significant evidence of learning even after nine rounds of play. In our experiments we used the hypothetical story above described, and gave subjects three possible answers, stating that only one of the three was the correct one.

3. Backward Induction: A Simple Game Tree

In the game (see fig. 1), our experimental subject is the first mover.

Insert fig. 1 about here

The rational strategy for player 1 is to choose STOP; in fact, the game is solvable through *backward induction*, by which players start at the end of the game tree, and ‘prune’ it by successively eliminating dominated moves. If YOU have a chance to move in the last round, you would choose A assuring a gain of 30 for YOU and 10 for the OTHER; however, the OTHER, knowing that, would choose STOP in round 2, assigning a payoff of 20 to herself and 10 to you. Hence, YOU should stop the game in the first round to gain of payoff of 20. The game’s unique

sub-game perfect equilibrium is in fact (STOP, STOP). However, if subjects do not perform the reasoning sequence correctly, they might be tempted to choose PASS to get the payoff of 30 in the last round. The game was never tested in previous experiments, although there is substantial evidence of heterogeneity and boundedly rational behavior in similar sequential games (e.g., Schotter, Weigelt and Wilson, 1994; Camerer, Johnson, Rymon and Sen, 2002).

3 Experimental Design

In this preliminary experimental study, subjects did not play the games against one another, but with a hypothetical, rational opponent that was programmed to always choose the payoff-maximizing strategy in each game. Since in each of the games there is only one optimal strategy for the opponent, we told subjects that they could imagine playing against a computer, but that in reality a payoff-maximizing algorithm would be applied to determine their opponents' choice, and hence their payoff would be determined with certainty once they had made their choice. The reasons of this choice of experimental strategy are twofold: first, we wanted our subjects to be motivated only by financial reward and not by other drives such as altruism, reciprocity, fairness, etc. In addition, we intended to rule out forms of sophisticated behavior induced by the belief that one's opponent may be irrational, revengeful, etc. Although proposing game-theoretic tasks in the form of hypothetical questions rather than as real interactive situations is an experimental design certainly more familiar to psychologists than to experimental economists, there is increasing evidence that in several cases results do not differ much among the two procedures (e.g., Rubinstein, 1999)¹.

A total of 67 subjects participated in the experiment. They were all undergraduate students who enrolled by responding to ads posted at the department buildings. The experiment was in form of a questionnaire and was conducted with paper and pencil at the Computable and

¹For a discussion on methodological differences between psychologists and experimental economists and their implications see Camerer (1996).

Experimental Economics Lab of the University of Trento. Before the experiment started, we conducted separately on each subject the standard Wechsler digit span test for short term memory capacity (e.g., Walsch and Betz, 1990). This simple test consists in asking subjects to repeat a series of digits which are to be read by the experimenter at the rate of one digit per second. The test is conducted sequentially on two independent sets of digit series of increasing length. For each set, the test stops when the participant fails to correctly repeat a given series. The subject's 'score' in each set is given by the length of the last series that was repeated correctly (so, for example, if a subject fails to correctly repeat a series of length 6, her score will be equal to 5). The subject's final score is then given by the higher among the two scores that were achieved in the sets. Although the score needs not directly reflect the number of 'short-term memory slots' available to an individual, it is generally assumed that higher scores correspond to a higher short-term memory capacity. After the test was taken, subjects received a sheet with instructions for the experiment (a translation of the instructions is available in Appendix). After reading the instructions, subjects were given a questionnaire with the three games mixed with other, unrelated questions. They were paid an amount of €3 for participation plus anything they could earn in the experiment for a total maximum of €25. The experiment lasted 45-50 minutes on average. Average earnings in the experiment were equal to €21.3. Subjects had no time limit to make their choice in the games.

4 Results

Our main research hypothesis is the existence of a correlation between a subject's score in the memory test and her conformity to theoretical predictions in the games.

Former attempts at eliciting behavioral strategies in games have revealed a bewildering variety of individual decision rules (e.g., Rapoport, Seale and Winter, 2000; Seale and Rapoport, 2000 among many). Notwithstanding such variety, we expect that, on average, STM constraints will affect subjects' ability to make optimal choices. Loosely speaking, there are two possibilities.

In one case, individuals may adopt resolution strategies which are too complex for their STM capacity; the resulting memory stress may be a source of crash in their execution, thus leading to errors. Alternatively, individuals may adapt (often by a process of trial-and-error) their decision rules to their own capacity constraints; such simplified decision rules may occasionally lead to correct choices, but on average will produce inaccurate responses. In both cases, it is reasonable to expect that more stringent STM constraints will increase the probability of errors. However, given the high variability in individual decision rules, we expect a high variance in responses in each single task, which may be absorbed as behavior in different tasks is aggregated².

Therefore, we have calculated the total number of correct responses in the three tasks for each subject (i.e., their cumulated score). Figure 2 reports the average cumulated score for each class of STM score (ranging from 5 to 8). The figure clearly shows a positive relation between average correct responses and STM capacity. We also computed the correlation coefficient between individual cumulated score and STM score. Table 2 reports the values of the two standard Kendall's τ_b and Spearman's ρ correlation coefficients.

Insert Fig. 2 about here

Insert table 2 about here

This result is encouraging: the table shows a positive and highly significant correlation between these two variables, suggesting that STM actually matters in determining average performance, notwithstanding the relevance of other sources of heterogeneity that cannot be captured

²In a discussion of problem solving in crypt-arithmetic tasks, Simon remarks that: "Different subjects do indeed apply different strategies ... But apart from the strategies, the only human characteristic that exhibits itself strongly in the crypt-arithmetic task is the limited size of short-term memory. Subjects get into trouble simply because they forget where they are, what assignments they have made previously, and what assumptions are implicit in assignments they have made conditionally." (Simon, 1981, pp. 62-63).

by our experimental design. Hence, the evidence found by Simon (1978; 1989) regarding the domain of individual problem-solving seems to apply also to interactive decision-making.

We then conducted the same analysis for each task separately. Figures 3, 4, and 5 report the frequencies of correct responses in each task as a function of the STM class. Two tasks (“dirty faces” and “iterated dominance”) show the expected monotonic relation, although with a flattening in correspondence of STM 6 and 7 in the “iterated dominance” case, while in the game-tree case, there is an unexpected peak in correspondence of the lowest STM value (notice, however, that the sample size of this STM class is the smallest one). We then computed nonparametric correlation coefficients between STM score and response in each task separately (see table 3).

Insert Fig. 3, 4 and 5 about here

Insert table 3 about here

As expected, results on individual correlations are weaker, with the only exception of the “dirty faces” game.

Some observations are noteworthy regarding the evaluation of our results: firstly, subjects in our experiments did not have any time constraints to make their choices, and the stimuli were constantly under their eyes. It is well known that these two conditions considerably dilute the effects of STM constraints (Baddeley, 1990). We chose this design, even though adverse to our hypothesis, because we wanted to disentangle the mere effects of remembering the task’s relevant information from that of remembering the intermediate steps of iterated reasoning required to make a decision, which is the main focus of our study. Secondly, the higher correlation observed between STM score and performance in the “dirty faces” task compared to the other tasks can be partly explained by our choice of an STM scoring test which is based on the phonological

loop' (Baddeley, 1990), measuring verbal STM performance. This particular measure neglects all visual components of STM. While the "dirty faces" task was presented to subjects in an entirely verbal form, the other two tasks were both supported by a graphical representation (i.e., a payoff matrix and a game-tree), that might have been used also as a support to the reasoning process. In our experiment, we didn't measure the visual STM capacity; we conjecture that a combined use of visual and auditory STM measures would provide a better insight into these phenomena.

5 Concluding Remarks

In this study we made a first attempt at exploring the link between short-term memory constraints and behavior in strategic settings. Our preliminary results seem to capture the existence of a relation, suggesting that short term memory capacity limits may indeed be a relevant source - though of course not the only one - of bounded rationality in games as well as in individual problem solving tasks, as Simon has extensively argued.

Further research should involve the use of protocol analysis to gain insight into the behavioral strategies that subjects adopt: in fact, an interesting research direction, in our view, should aim at identifying how STM constraints impact on the decision process, whether by simply affecting the probability of making 'mistakes' at some stage of the iterated reasoning process, or, more radically, by inducing subjects to the use of simplified rules of thumb as substitutes of more complex decision strategies. Even more radically, STM constraints may affect the mental representation that individuals construct of a strategic interaction situation: In a related experimental paper (Warglien, Devetag and Legrenzi, 1998) we suggest that individuals often use simplified representations of games that assimilate them to lotteries, allowing the application of 'naive' heuristics that neglect the most important characteristic of a strategic setting: the need to take into consideration the other players' motivations and preferences. Whether the switch to simplified representations is determined by underlying memory constraints is an open question.

These combined research directions could contribute to explain, at least in part, the behavioral

heterogeneity observed in experiments and which is usually accounted for in the literature by assuming a multiplicity of players' types. Types are usually assumed to differ by the beliefs they hold about the types of others, which in turn determine the decision rule adopted, so that quite often observed violations of equilibrium predictions are explained as best responses to such beliefs. It would be worthwhile to investigate whether and to what extent such differences in decision rules derive instead from an underlying computational bound.

6 Appendix: Instructions

Welcome! You are about to participate in an experiment on decision making; if you follow instructions carefully and make appropriate decisions, you will be able to earn an appreciable amount of money, which will be paid to you privately and in cash at the end of the experiment. In the following you will be presented with a series of interactive decision making situations, in which your potential payoff will depend on your choice as well as the choice of your opponent. These interactive situations will be called 'games'.

In all the games that will be presented to you, we ask you to imagine being paired with an opponent who is exclusively interested in trying to maximize his/her payoff, without being motivated by other concerns such as altruism, revenge, etc. besides, he/she imagines that you are behaving in the same way. If you think it's easier, you can imagine the other player as a computer; in reality, the moves of your opponent will be determined by an algorithm that maximizes the opponent's expected payoff in every game.

Furthermore, keep in mind that, while in common language the term "game" is usually used to denote purely competitive situations in which there is a winner and a loser, in the games presented here this will not necessarily be the case. In other words, there will be games in which both you and your opponent can earn less or more depending on the choices you make. Your earnings in the game will be expressed in a hypothetical currency which we call *schillings*. A schilling will be converted into euros at the rate of €0.12 per schilling. Your total earnings in

the games will be determined by the sum of your earnings in each game, and the corresponding amount will be paid to you privately and in cash at the end of the experiment. Upon request, you will be allowed to verify that your earnings have been calculated correctly (there is only one optimal choice for your opponent in every game, hence your earnings are determined with certainty once you have made your choice).

In all the game you will have a certain role which will be specified to you. The games will be presented to you in three possible forms: in verbal form, in form of a payoff matrix, and in form of a game-tree.

Games in form of a payoff matrix:

A payoff matrix is a convenient way to represent simultaneous games, i.e., games in which both you and your opponent must choose your actions without knowing what the other will choose before making your own choice. A payoff matrix is simply a table with a certain number of ROWS and COLUMNS. Usually, one of the players (the ROW player) is designated to choose one among the ROWS, while the other (the COLUMN player) can choose one among the COLUMNS. The cells inside the table report the earnings (in schillings) that both of you get in correspondence of every possible combination of your two decisions.

For example, in the following table the ROW player can choose among ROWS A and B, while the COLUMN player can choose among COLUMNS C and D. At the intersection between the row chosen by the row player and the column chosen by the column player one can read the payoffs that each of them gets; remember that the first number always indicates the row player's payoff, while the second number indicates the column player's payoff.

Insert table 4 about here

Remember that you can never choose directly the cell, but only the row or column. To help

your comprehension of a payoff matrix, you can imagine it analogous to a cartesian plan, in which one of the players can choose a value on the x -axis, and the other can choose a value on the y -axis. The intersection between the two identifies a point (i.e., the cell).

You will have the role of the ROW player. You will only have to read instructions carefully, think through the payoff table and then choose the row you prefer, imagining that the other player has identical information (except that he knows he's the column player), and knows that both of you have the same information.

Games in form of tree:

The tree-form is usually used to represent sequential games, i.e., games in which players move in sequence after having observed what the other has done. An example of a game tree is shown below:

In this game player 1 moves first, and can choose between two actions (the 'branches' of the tree), A and B. If she chooses A, the game stops and the two players receive the payoffs in parenthesis (the first number denotes player 1's payoff). If she chooses B, the it's player 2's turn to move, and she can choose between C and D. Here, like in the matrix case, you have to think carefully through the information contained in the game tree, and then choose the action you prefer, knowing that your opponent has the same information, and that she will observe your move before making hers.

You will have no time limit to make your choices. we ask you to think carefully and do the experiment in silence. If you have question, please raise your hand and the experimenter will come to you. Thank you for your collaboration.

References

- [1] Abreu D., & Rubinstein A. (1988). The structure of Nash equilibrium in repeated games with finite automata. *Econometrica* 56, 1259-1282.
- [2] Baddeley, A.D. (1990). *Human Memory: Theory and Practice*. Hillsdale, NJ: Erlbaum Associates.
- [3] Binmore, K. & Brandenburger, A.M. (1990). Common Knowledge and Game Theory. In K. Binmore (Ed.), *Essays on the Foundations of Game Theory*. Oxford: Basil Blackwell.
- Camerer, C. (2002). *Behavioral game theory: experiments on strategic interaction*. Princeton, NJ: Princeton University Press (in press).
- [4] Camerer, C. (1996). Rules for experimenting in psychology and economics, and why they differ. In W. Albers, W. Guth, P. Hammerstein, B. Moldovanu, and E. Van Damme (Eds.), *Understanding Strategic Interaction: Essays in honor of Reinhard Selten* (313-327). Berlin: Springer-Verlag.
- [5] Camerer, C, Ho, T., & Chong, J.K. (2001). Behavioral Game Theory: Thinking, Learning, and Teaching. Paper presented at the Nobel symposium.
- [6] Camerer, C., Johnson, E., Rymon, T., & Sen, S. (2002). Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining. *Journal of Economic Theory*, in press.
- [7] Cohen, J.D., Perlstein, W.M., Braver, T.S., Nystrom, L.E., Noll, D.C., Jonides, J., & Smith, E.E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature* 386, 604-608.
- [8] Costa-Gomez, M., Crawford, V., & Broseta, B. (2001). Cognition and Behavior in Normal-form Games. *Econometrica*, 69, 1193-1235.

- [9] Devetag, G., & Warglien, M. (2001). Representing Others' Preferences in Mixed Motive Games: Was Schelling Right?. Cognitive Science Laboratory (CSL) working paper, University of Trento.
- [10] Fagin, R., J.Y. Halpern, Moses, Y. & Vardi, M.Y. (1995). *Reasoning about knowledge*. Cambridge, Mass.: MIT Press.
- [11] Geanakoplos, J. (1994). Common Knowledge. In R.J. Aumann and S. Hart (Eds.), *Handbook of Game Theory*, Vol. 2. Amsterdam: Elsevier.
- [12] Ho, T., Camerer, C. & Weigelt, K. (1998). Iterated Dominance and Iterated Best Response in Experimental p-Beauty Contests. *American Economic Review* 88 (1998), 947-969.
- [13] Littlewood, J.E. (1953). *A Mathematician's Miscellany*. Meuthen & Co. Ltd.
- [14] Nagel, R. (1995). Unraveling in Guessing Games: An Experimental Study. *American Economic Review* 85, 1013-8026.
- [15] Newell A., & Simon H.A. (1972). *Human Problem Solving*. Englewood Cliffs (NJ): Prentice Hall.
- [16] Norman D.A., & Bobrow D.G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology*, 7, 44-74.
- [17] Rapoport, A, Seale, D. & Winter, E. (2000). An Experimental Study of Coordination and Learning in Iterated Two-Market Entry Games. *Economic Theory*, 16, 661 - 687.
- [18] Rubinstein, A. (1998). *Modeling Bounded Rationality*. Cambridge: MIT press.
- [19] Rubinstein, A. (1999). Pre and Post-class Problem Sets as a Didactic Device. *Games and Economic Behavior*, 29, 155-170.
- [20] Schotter, A., Weigelt, K., & Wilson, C. (1994). A Laboratory Investigation of Multiperson Rationality and Presentation Effects. *Games and Economic Behavior*, 6, 445-468.

- [21] Seale, D. A., & Rapoport, A. (2000). Elicitation of strategy profiles in large group coordination games. *Experimental Economics*, 3, 153 - 179.
- [22] Simon, H.A. (1978). Rationality as a process and as a product of thought. *American Economic Review*, 68, 1-16.
- [23] Simon, H.A. (1978). *Models of thought*, vol. 1. New Haven: Yale University Press.
- [24] Simon, H.A. (1981). *The Sciences of the Artificial*. Cambridge, Mass.: MIT Press.
- [25] Simon, H.A. (1989). *Models of thought*, vol. 2. New Haven: Yale University Press.
- [26] Stahl, D.O., & Wilson, P. (1994). Experimental Evidence on Players' Models of Other Players. *Journal of Economic Behavior and Organization*, 25, 309-327.
- [27] Stahl, D.O., & Wilson, P. (1995). On Players' Models of Other Players: A New Theory and Experimental Evidence. *Games and Economic Behavior*, 10, 218-254.
- [28] Walsh, W.B., & Betz, N. E. (1990). *Tests and Assessment*. Englewood Cliffs, NJ: Prentice Hall.
- [29] Warglien, M., Devetag, G., & Legrenzi, P. (1998). Mental Models and Naive Play in Normal Form Games. Universita' Ca'Foscari di Venezia working paper.
- [30] Weber, R.A. (2001). Behavior and Learning in the "Dirty Faces" Game. *Experimental Economics*, in press.

	C	D
A	60,20	60,10
B	80,20	10,10

Table 1: A game solvable by iterated dominance

Correlation coefficient	Significance (1-tailed)	
Kendall's τ_b	.248	.010
Spearman's ρ	.287	.009

Table 2: Correlation coefficients between individual STM score and cumulated score in the experiment

	It. Dom.		'Dirty Faces'		Game Tree	
	corr. coeff.	sig.	corr. coeff.	sig.	corr. coeff.	sig.
Kendall's τ_{ab}	.147	.098	.203	.037	.152	.091
Spearman's ρ	.159	.099	.219	.037	.164	.092

Table 3: Correlation coefficients between individual STM score and score in each task

	C	D
A	75,82	61,74
B	74,71	65,88

Table 4: An example of a payoff matrix




Fig. 1: a sequential game solvable by backward induction

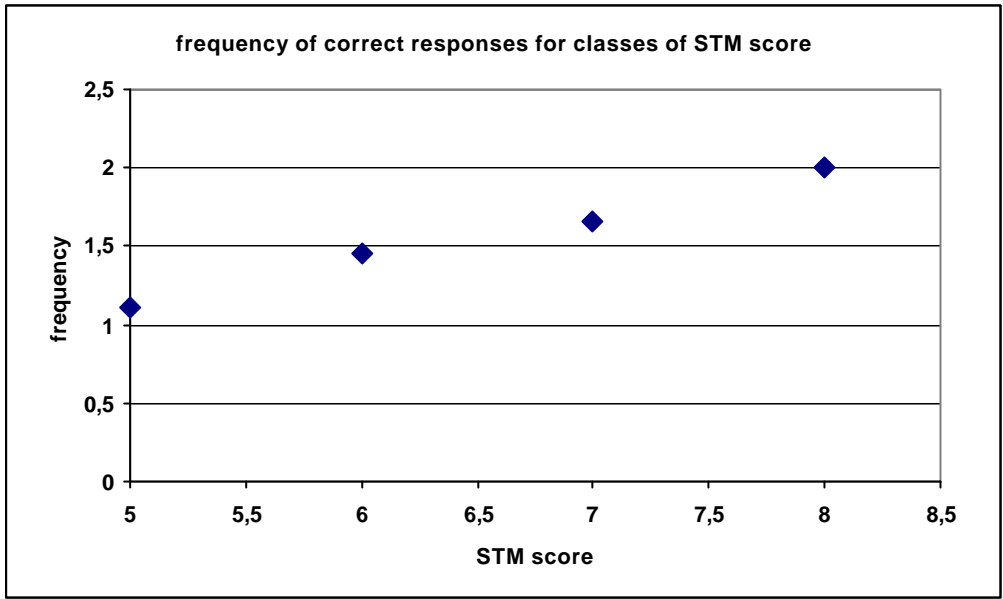


Fig. 2: frequencies of correct responses in the three tasks for classes of STM score

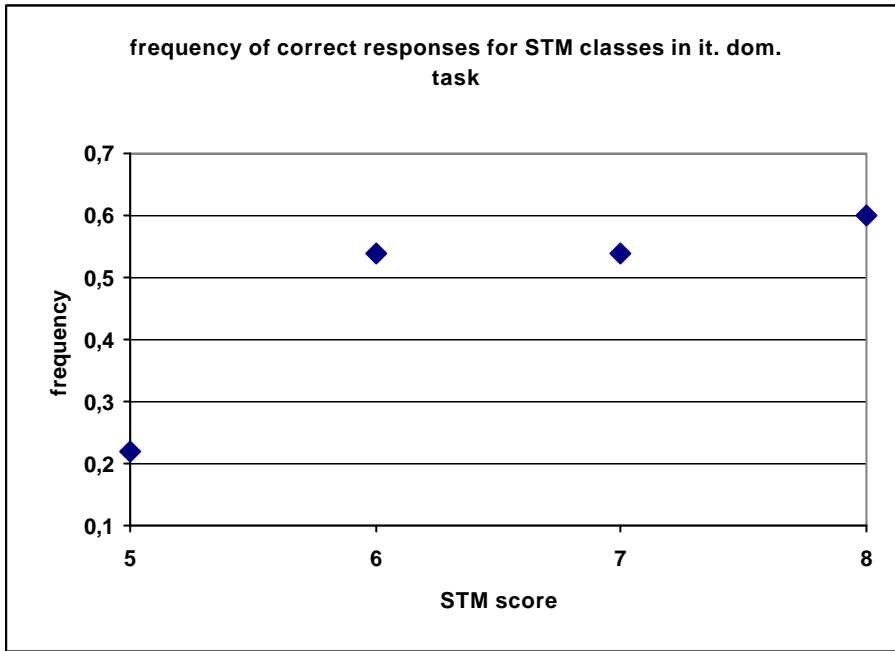


Fig. 3: frequency of correct responses for classes of STM score in the iter. dom. task

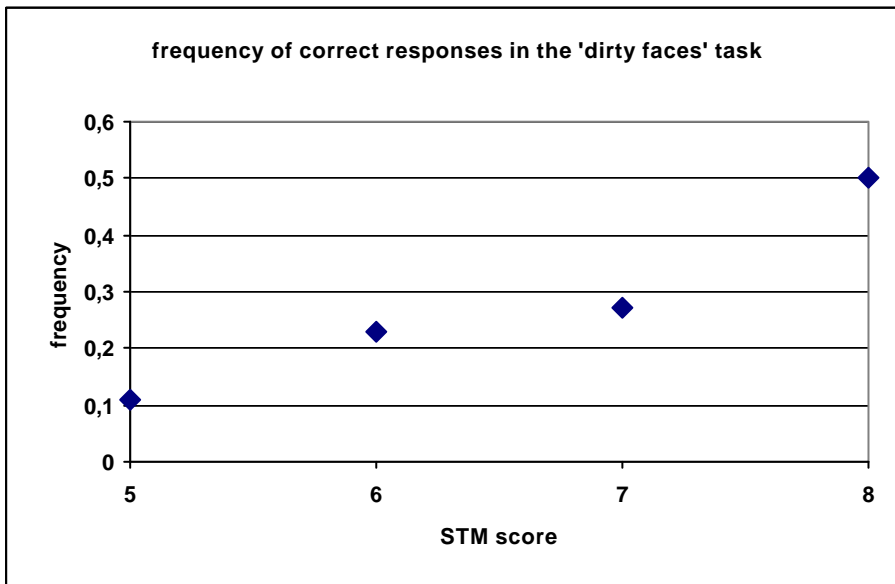


Fig. 4: frequency of correct responses for classes of STM score in the 'dirty faces' task

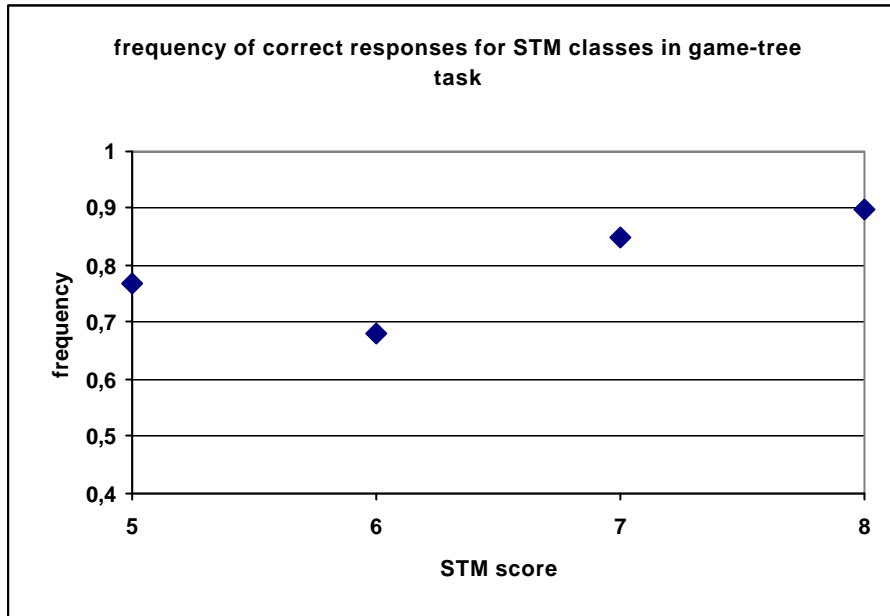


Fig. 5: frequency of correct responses for classes of STM score in the game-tree task