

The 11-20 Money Request Game: Evaluating the Upper Bound of k-Level Reasoning

Ayala Arad

School of Economics, Tel Aviv University

Ariel Rubinstein

University of Tel Aviv Cafes

School of Economics, Tel Aviv University

Department of Economics, New York University

Abstract

We study experimentally a two-player game which we find ideal for investigating k-level reasoning. Each player requests an amount of money between 11 and 20 shekels. He receives the amount that he requests and if he requests exactly one shekel less than the other player, he receives an additional 20 shekels. The best response function in this game is straightforward, the k-level strategies are invariant to the two prominent level-0 specifications (randomization or attraction to salience) and the situation calls for self-interest behavior only. Therefore, we propose the game as a test for the "upper bound" for the depth of k-level reasoning in a population. We support this conjecture by studying several variations of the game which manipulate the attractiveness of the level-0 strategy and the monetary cost of undercutting the other player.

Keywords: level-k thinking, salience, iterative reasoning

We wish to thank Eli Zvuluny for managing the website used in this study and to Aviad Shany for classifying the subjects' explanations. We are grateful to Anat Bracha, Vincent Crawford and Pedro Rey Biel for their valuable comments. We acknowledge financial support from the Center for Experimental Social Science at New York University and the Israeli Science Foundation, grant no. 259/08 and The Foerder Institute for Economic Research.

1.Introduction

When a symmetric two-player game is played only once, the concept of Nash equilibrium often fails (dramatically) to describe the experimental results. A major task of the experimental game-theoretic literature is to provide better tools for explaining behavior in such situations. Prominent in this literature is the concept of k -level reasoning, first introduced in Stahl and Wilson (1994,1995) and Nagel (1995).

A standard k -level model assumes that the population is partitioned into types, which differ in their depth of reasoning. A level-0 type is non-strategic and follows a simple decision rule. A level- k type (L_k), for any $k \geq 1$, behaves as if he best responds to the belief that all other players are level $k - 1$ types. Thus, a model is characterized by: (1) an L_0 behavior (which is the starting point for iterative reasoning) and (2) a distribution of types.

A typical study of k -level reasoning collects experimental data of a game and looks for the best fit to the data. That is, it finds an ex-post specification of an L_0 strategy and a distribution of k -level types that best explain the results in a statistical sense. Only a few papers have attempted to elicit subjects' actual levels of reasoning by analyzing other kinds of data in addition to the observed choice data. Examples include Costa-Gomes, Crawford and Broseta (2001) and Costa-Gomes and Crawford (2006) who examined data on information search behavior recorded using MouseLab; Bosch-Domenech, Montalvo, Nagel and Satorra (2002) and Arad (2009) who used ex-post explanations of their choices; and Burchardiy and Penczynski (2010) who analyzed subjects' arguments while attempting to convince their teammates to follow their advice. However, most of the studies simply searched for the best fit to the data, after having assumed the level-0 behavior (which is sometimes specified after having observed the data).

This research approach has also been applied to games in which the k -level procedure is not so natural, or in which the assumed L_0 strategy is not common knowledge. An example is Crawford and Iriberry (2007) which analyzed the hide-and-seek game introduced in Rubinstein, Tversky and Heller (1996). Crawford and Iriberry (2007) found that given their assumptions on the level-0 behavior, the distribution of types that best fits the data included many level-3 and level-4 types. In contrast, Burchardiy and Penczynski (2010)'s analysis of subjects' arguments in the same hide-and-seek game suggests that subjects do not have a common starting point for iterated reasoning (i.e. level-0) and only few subjects practice more than two steps of reasoning given their own starting point.

Not surprisingly, the distributions of k -level types differ between games. The fact that the

L_0 parameter cannot always be clearly identified and that the distribution of types is variable creates doubt whether one can use k -level reasoning to predict behavior in a new game.

In this paper, we study a very simple novel game, which is an ideal setting for applying a level- k procedure. We view the obtained distribution of k -level types in this game as an "upper bound" on the distributions of this population's actual level- k types in other one-shot games. This can be used (i) to narrow the predicted range of behavior in strategic situations that have not yet been studied and (ii) to examine whether a distribution of types found to best fit the data for a particular game is also reasonable as an actual explanation of subjects' behavior in that game.

Note, however, that we cannot expect the distribution of levels of reasoning to be identical across different populations. For example, students who have studied game theory are likely to apply higher levels of reasoning than students in the humanities. We therefore suggest the game as a tool for evaluating the upper bound on the k -level reasoning **given** a population.

Following is a description of the basic version of the game, which we will refer to as the **11-20 game**:

You and another student are playing a game in which each player requests an amount of money. The amount must be an integer between 11 and 20 shekels. Each player will receive the amount he requests. A player will receive an additional 20 shekels if he asks for exactly one shekel less than the other player.

What amount of money would you request?

There are five main aspects of the game that make it particularly suitable for studying k -level reasoning:

(i) *The level-0 type specification is intuitively appealing*: The choice of 20 is a natural anchor for an iterative reasoning process. It is the instinctive choice when choosing a sum of money between 11 and 20 shekels (20 is clearly the salient number in this set and "the more money the better").¹ Furthermore, the choice of 20 is not entirely naive: if a player does not want to take any risk or prefers to avoid strategic thinking, he might give up the attempt to win the additional 20 shekels and simply request the highest certain amount.

(ii) *Best-responding is straightforward*: Given the anchor 20, best-responding to any level- k action is very simple and leaves no room for error. The strategy 19 is the best

response to 20, 18 is the best response to 19 and so on (it is less straightforward that the strategy 20 is a best response to 11).

(iii) *Robustness to the level-0 specification*: The choice of 19, which is our level-1 strategy, is the unique best response to a wide range of reasonable beliefs on level-0 behavior including: (a) all distributions in which 20 is the most frequent choice and (b) the uniform distribution and a class of beliefs that are close to it. This makes the analysis robust to the specification of the level-0 behavior.

(iv) *Using k-level reasoning is very natural*: The game's payoffs are described using the best response function explicitly, a characteristic that triggers iterative reasoning. Moreover, it is hard to think of alternative decision rules for this game. There is no pure-strategy Nash equilibrium and the game lacks dominated strategies. In fact, all strategies are contained within the support of the unique Nash equilibrium (described later). The only other conceivable rules of behavior would be randomly choosing a strategy or arbitrarily guessing the other player's strategy and best-responding to it.

(v) *Clear payoffs*: Unlike some other games which also trigger iterative reasoning (such as Rosenthal (1981)'s Centipede Game and Basu (1984)'s Traveler's Dilemma), this game does not call for social preferences. In particular, if a player believes that his opponent will choose 20 (or some other amount), then requesting one shekel less will give him a bonus of 20 shekels but **not** at the expense of the other player.

The above characteristics make the use of k-level reasoning in this game very appealing. It is hard to think of situations in which a player will employ a higher k-level reasoning than in this game. By analyzing two additional versions of the game, we confirm that enhancing the anchor for iterative reasoning and eliminating the monetary cost of undercutting the other player does not increase the depth of reasoning. Thus, given a population, we nominate the distribution of k-level types in the 11-20 game as an upper bound for the distributions of types in other games.

2. The 11-20 game

Experimental Design: The subjects consisted of five classes of undergraduate economics students at Tel Aviv University (four classes of intermediate microeconomics and one class of the introductory course). The subjects had not studied game theory prior to the experiment (this topic is introduced only in a more advanced microeconomics course). Students were offered the opportunity to participate in a short experiment with monetary prizes during class time. All the students in each class decided to participate. Subjects

were asked to arrange themselves in the classroom so as not to be able to see their classmates' answers. Three different forms, which corresponded to three versions of the game, were randomly distributed among the subjects (one form to each subject). In this section, we report only on the results for the basic version of the game (described in the introduction), which was played by 108 subjects out a total of 207. The instructions for each game appeared on the forms; no additional instructions were provided. After all the subjects had made their decision, they were asked to write an explanation of their choice on the back of the form. They did not know in advance that they would be asked to do so. After all the forms had been collected, each one was randomly matched to another and each subject received his resulting payoff.

Results: Table 1 presents the unique Nash equilibrium distribution (assuming that players maximize expected payoff) and the actual distribution of choices in the experiment:

Action	11	12	13	14	15	16	17	18	19	20
Equilibrium					25%	25%	20%	15%	10%	5%
Results	4%	0%	3%	6%	1%	6%	32%	30%	12%	6%

Table 1: The 11-20 game (n=108)

The subjects' behavior cannot be explained by the Nash equilibrium. The strategy 20 was chosen by 6% of the subjects which is almost identical to the equilibrium prediction. However, only 7% chose the strategies 15 and 16 which is far from the equilibrium prediction of 50%. The vast majority of subjects (74%) chose the actions 17-18-19 which seem to exhibit 1-2-3 levels of reasoning respectively, whereas in equilibrium they should have been chosen by only 45% of subjects.

Every strategy in this game is by definition a level- k strategy for some k , but it does not necessarily mean that it actually reflects k -level reasoning. We interpret the choices in this game by analyzing the subjects' ex-post explanations of their strategies². The students' explanations of the choices of 19, 18 and 17 almost unanimously reflect one, two and three iterations, respectively (starting from the anchor of 20). Only 16% of the subjects chose one of the five strategies 12-16. We think that such a choice was seldom an outcome of 4-8 iterations. This intuition is supported by the subjects' explanations. There were no cases in which the choice of an action in the range 12-15 was explained as a 5-8-level of reasoning and only one subject described the strategy 16 as an outcome of the forth level of

reasoning. The strategy 11 could be an outcome of a different kind of reasoning based on an open ended interactive process. However, since only four subjects have chosen this strategy we defer the discussion on this strategy to another version of the game where it was more popular. The strategies 12-16 were frequently described as guesses ("It is just a guess"). Few explanations explicitly mentioned best-responding to a specific belief without describing the belief's origin. In contrast, almost all the subjects who chose a strategy in the range 17-19 explained their choices as an outcome of k -level reasoning ("I request 18 since many people believe that the majority requests 20 and thus request 19"). Such arguments were a little less frequent (77%) among subjects who chose 17. A typical explanation for the choice of 20 was "The chances to play against someone who requests one shekel more than me is very low. Therefore, I prefer to receive the highest sum for certain".

The proportion of subjects who go through a k -level reasoning process obviously depends on the characteristics of the game and the population. However, we speculate that in game experiments with similar monetary incentives and subjects (i.e. economics students prior to taking a course in game theory) the distribution of level- k types would be first-order "stochastic-dominated" by the one we identify here³. In other words, in other games we would **not** find: (i) a significant proportion of subjects who choose the level-4 or higher strategies, (ii) that significantly more than 80% of subjects⁴ choose the strategies corresponding to levels 1, 2 or 3, (iii) that significantly more than 32% of the population chooses the level-3 strategy and (iv) that significantly more than 62% chooses the level-2 or level-3 strategies. Needless to say, we are not suggesting that this study projects on the maximal number of steps of iterated elimination of dominated strategies in games where such a procedure is possible.

Studies of other games in the literature found that it is rare to explicitly observe subjects who practice more than three steps of iterative reasoning. However, one could think that it is due to the difficulty in using the level- k procedure in these games and in particular the complexity of the best-response function. In some other games, such as the Centipede Game and the Traveler's Dilemma, the behavior of many subjects is driven by social preferences which makes the use of iterative reasoning redundant.⁵

Our most striking finding is that despite the simplicity of the iterative reasoning process and the implausibility of social preferences considerations in our game, about three-quarters of the choices seem to reflect up to three levels of reasoning.

Note that the level of reasoning attributed to a player is the number of steps carried out from the starting point. It is possible that some subjects realized that the iterative process cycles and returned to the starting point of 20 on reaching strategy 11. In fact, this argument

appeared in only six explanations. The bottom line is that despite this understanding, they chose to conduct only up to three steps of reasoning from the starting strategy. The third version of the game which is reported on later, confirmed that this is not due to the "cost" of choosing a number lower than 17.

The statement we attribute to a level-1 type is "I think that [he is doing something]". The statement we attribute to a level-2 type is "I think that [he thinks that I am doing something]", in which the level-1 statement is embedded. Expressing explicitly the considerations of a level-4 type would involve three such embedments. Whether these considerations are implicit or explicit, it seems that behavior which is consistent with such considerations is unnatural. This may be rooted in a more fundamental phenomenon observed by psychologists. Kinderman, Dunbar and Bentall (1998) found that most subjects do not understand a sentence such as: "A thinks that B thinks that A thinks that B thinks that A is doing something", which is attributed to our level-4 type. (Almost all subjects understand sentences which we attribute to lower levels). They claim that this kind of sentence goes beyond the limit of reasoning normally used in real life and that most everyday situations probably do not require more than second-order intentionality. Literary scholars have also recently argued that "the zone of cognitive comfort" seems to be very limited. Zunshine (forthcoming) claims that using sentences with a higher number of embedments than in the sentence "I know that you think that he wants you to believe that she was angry at him" is rare and appears in the literature as a challenge to the reader. Thus, the scarcity of level-4 types seems to be in line with the findings in literature regarding the cognitive limit on the number of embedments and the rare use of high-order intentionality in everyday situations.

An alternative interpretation of the results: In this paper, we have adopted the definition of a level- k type as a player who best-responds to the belief that the other player is a level $k - 1$ type. Another version of the k -level model is that of the Cognitive Hierarchy model introduced by Camerer, Ho and Chong (2004), in which the level- k type best-responds to the conditional distribution of lower types and the distribution of types is assumed to be Poisson with parameter λ . Given the anchor of 20, the strategy 19 would be the choice of this model's level-1 type. Given the distribution of strategies in the experiment, the strategy 18 would be the best response to the conditional distribution of the lower types (who choose the strategies 19 and 20) and the strategy 17 would be the best response to the conditional distribution of 18, 19 and 20. Thus, the specification of levels 1-3 according to this model would be the same as before. The difference between the models emerges from the fact

that the strategy 17 is also a best response to the actual conditional distribution of the strategies 17 – 20. Thus, the strategy 17 could also be interpreted as the choice of any type higher than 3. We did not find any support in our subjects' explanations for the reasoning behind this alternative model. However, if we adopted this approach, we would look for the parameter λ that best fits the data (which would require additional assumptions on noise) and our study would suggest a tool for evaluating the upper limit of the parameter λ .

3. The Cycle Version

The following version of the game was designed to show that enhancing the "status" of the $L0$ action does not increase the depth of level- k reasoning:

You and another student are playing a game in which each player requests an amount of money. The amount must be an integer between 11 and 20 shekels. Each player will receive the amount of money he requests. A player will receive an additional amount of 20 shekels if:

(i) he asks for exactly one shekel less than the other player,

or

(ii) he asks for 20 shekels and the other player asks for 11 shekels.

What amount of money would you request?

This version was assigned randomly to 72 subjects from the same pool and the experimental design remains the same (as described in section 2). Table 2 presents the equilibrium distribution and the results. For the purpose of comparison, the results of the basic version are presented here again.

Action	11	12	13	14	15	16	17	18	19	20
Equilibrium					25%	25%	20%	15%	10%	5%
Cycle version	1%	1%	0%	1%	0%	4%	10%	22%	47%	13%
11-20 game	4%	0%	3%	6%	1%	6%	32%	30%	12%	6%

Table 2: The cycle version (n=72)

The cycle version differs from the basic version in one crucial detail. In the basic version, a player who chooses 20 could not receive the bonus whereas in this version he can (if the other player chooses 11). Thus, the choice of 20 is even more attractive and justifiable than

before. We hypothesized that this would strengthen the role of 20 as a prominent starting point for iterative reasoning and will alter the distribution of strategies.

In the results for this version, the strategy 20 is indeed more popular than in the basic version (13% vs. 6%, a difference which is significant at the 10% level). One might have expected that the use of level- k strategies (for $k > 0$) would increase due to the enhancement of the level-0 strategy. However, the proportion of subjects who used a strategy in the range 17-19 did not change significantly (it only increased from 74% to 79%). The explanations of the strategies in this range were very similar to those provided in the basic version. Level- k arguments appeared in all the explanations of the subjects who chose 18 and 19 and in 4 out of the 7 explanations of subjects who chose 17. On the other hand, now most of the choices in this range were of the strategy 19, which is the level-1 strategy. Furthermore, the difference between the conditional distributions of the low level- k types was significant. In other words, enhancing the $L0$ strategy induced subjects who used k -level reasoning to use fewer steps of iterative reasoning.

To summarize, the role of the strategy 20 as the $L0$ was enhanced in this version and indeed the proportion of subjects who chose the $L0$ strategy doubled (though it remained small). This could have no effect on the distribution of k -level types. We found that the proportion of subjects who chose the $L1$, $L2$ and $L3$ strategies altogether did not change. However, the proportion of subjects who chose the $L1$ strategy increased dramatically while the proportion of those who chose $L3$ decreased. It seems that many potential low level- k types, compared to the basic version, recognized the extra-justifiability of $L0$ and expected it to be chosen frequently.

We speculate that in other games, in which the $L0$ strategy is not as clear as in the basic version, we would not find levels of reasoning higher than those in the 11-20 game. It is more likely that the level- k decision rules will be less frequent in such games and the use of other decision rules (that are not based on iterative reasoning) will be more common.

4. The Costless Iterations Version

The following version of the game was designed to show that eliminating the cost of undercutting the other player does not increase the depth of level- k reasoning:

You and another student are playing a game in which each player chooses an integer in the range 11-20.

A player who chooses 20 will receive 20 shekels (regardless of the other player's choice).

A player who chooses any other number in this range will receive 3 shekels less than in

the case where he chooses 20. However, he can receive an additional amount of 20 shekels if he chooses a number that is lower by exactly one than the number chosen by the other player.

Which number would you choose?

Recall that in the basic version of the game, best-responding to 20 required giving up one "certain" shekel in the attempt to win an additional 20 shekels. In the same manner, performing another iteration and choosing 18 required giving up an additional shekel. One could suspect that the cost of additional iterations in the basic version is why subjects stop there at the level-3 strategy (17) and do not perform another iteration (which would require giving up an additional certain shekel). In the third version of the game, there is no cost to performing an additional iteration. The cost of choosing any integer in the range [11,19] is identical: instead of receiving 20 shekels for certain, each of these numbers guarantees only 17 shekels for certain. Therefore, when considering a strategy within the range [11,19], a player only considers the probability of winning the additional 20 shekels.

This version was randomly assigned to 27 subjects from the same subject pool of the two other versions. Table 3 presents the equilibrium predictions and the actual distribution of strategies in this version:

Action	11	12	13	14	15	16	17	18	19	20
Equilibrium				10%	15%	15%	15%	15%	15%	15%
Experiment	0%	0%	0%	8%	3.5%	3.5%	11%	26%	41%	7%

Table 3: The costless iterations version (n=27)

The distribution of choices is far from equilibrium, as in the other two versions, while it is very similar to that obtained in the cycle version (probably due to the framing which emphasizes the level-0 strategy). A vast majority of the subjects (85%) chose the strategies corresponding to L_0 , L_1 , L_2 or L_3 . According to the subjects' explanations, the choices of 17-19 (by 78% of the subjects) generally involved an iterative process, whereas the choices of 14-16 (by 15% of the subjects) were arbitrary guesses. Despite the small sample, it is reasonable to conclude that the cost of performing an additional iteration (i.e. losing another certain shekel) is not the reason that subjects do not perform more than three iterations in the other two versions of the game.

The results for another variation, which were obtained using a very large sample, are

reported in the next section and provide further support for this conclusion. In that variation, the cost of an iteration is negligible in comparison to the potential bonus. Nevertheless, subjects tended to adopt a very low level of iterative reasoning.

5. The 91-100 versions

In this section, we report on experiments involving variations of the basic version and the cycle version with the following differences:

- (i) The amounts of 11-20 shekels were replaced with 91-100 dollars (the bonus was changed from 20 shekels to 100 dollars),
- (ii) The experiments were conducted on-line, without monetary incentives, and
- (iii) The subjects were mostly undergraduate students in various countries who were studying game theory.

The didactic website at gametheory.tau.ac.il served as the experimental platform. The site is used by teachers of game theory courses to assign virtual games and decision-theoretical problems to their students. The results obtained at the site are typically similar to those in laboratory experiments with monetary incentives (see Rubinstein (2007)). Our subjects originated from Argentina, Canada, China, Columbia, Denmark, Italy, Korea, the Slovak Republic, Spain, Switzerland and the US.

Table 4 presents the mixed strategy equilibria distributions and the experimental results for each of the two 91-100 versions:

Version	91	92	93	94	95	96	97	98	99	100
Equilibrium										
Basic	55%	9%	8%	7%	6%	5%	4%	3%	2%	1%
Cycle	5.5%	14.5%	13.5%	12.5%	11.5%	10.5%	9.5%	8.5%	7.5%	6.5%
Experimental results										
Basic	26%	3%	3%	4%	6%	2%	6%	20%	17%	13%
Cycle	15%	4%	3%	1%	3%	3%	2%	11%	34%	24%

Table 4: The Basic version (n=666) and the Cycle version (n=186) of the 91-100 game

As in the experiments of the 11-20 versions, which were conducted with monetary incentives, the majority of subjects in both of the versions reported here are concentrated in

the level-0, level-1, level-2 and level-3 strategies (100, 99, 98, 97). However, there are three main differences:

(i) The proportion of subjects who chose the highest number is larger in the on-line 91-100 experiments. Our conjecture is that the strategy 100 is more salient in this range as is 20 in the range 11-20. It is also possible that the lack of monetary incentives increased the proportion of subjects who chose the instinctive action.

(ii) The proportion of subjects who chose the level-3 strategy was much smaller than in the 11-20 versions.

(iii) In both versions, a significantly higher proportion of subjects chose the lowest possible number (91) as compared to the proportion who chose 11 in the 11-20 versions. In this population, there were more subjects who are familiar with iterative elimination of dominated strategies and perhaps had been exposed to games such as the Traveler's Dilemma, in which the choice of the lowest number is the unique equilibrium. The choice of 91 may be an outcome of the similarity that they saw (mistakenly) between these games and the 91-100 games. We do not believe that a subject who chose 91 actually performed nine explicit iterations. It is more likely that he performed a few iterations and then continued the iterative process by induction to what he considered to be the end. Recall, however, that the best response to 91 is the strategy 100. We suspect that in the cycle version many subjects who practiced the open-ended process realized that it does not end at 91 and restarted the process from 100 (since in the cycle version the strategy 100 was described explicitly as a best response to 91). This is consistent with the lower frequency of the strategy 91 in the cycle version as compared to the basic version (15% vs. 26%).

We also conducted an on-line experiment (without monetary incentives) of a different version of the money request game in which a subject requests an amount of money in the range [\$1,\$100] and receives an additional \$100 if he requests exactly one dollar less than the other player. The subjects consisted of 135 students in various countries who had studied game theory. The results are presented in Table 5.

Action	1	2-90	91-95	96	97	98	99	100
Results	7%	6%	1%	1%	6%	22%	26%	30%

Table 5: Behavior in the 1-100 game (n=135)

The vast majority of the subjects (84%) chose one of the four strategies 97, 98, 99 and

100. Only 7% chose the unreasonable strategy 1. The rest (8%) were widely distributed within the range 2-96. These results provide support for our interpretation of choices in the range 92-96 in the basic 91-100 version as reflecting arbitrary guesses rather than high levels of iterative reasoning. If, in contrast to our intuition, the choices in the range 92-96 (made by 18% of the subjects) in the 91-100 game reflected between 4 and 8 levels of reasoning, then the strategies 92-96 should have been just as popular in the 1-100 game. However, only 2% of the subjects in the 1-100 version chose those strategies.

6. Conclusion

In this short paper, we have presented a simple game that provides an ideal test for the bounds on the depth of iterative reasoning in games. The game is easily understood and the k -level reasoning process is both intuitive and unambiguously specified. Nevertheless, subjects did not use more than three steps of reasoning. We conjecture that the distribution of level- k types in the 11-20 game can provide an upper bound on the distributions of types in other games. We do not argue that the distribution of types detected here is universal since it is likely to depend on the characteristics of the subjects. However, we do believe that the game reveals the upper distribution of k -level types within a population and can be used to identify the maximum depth of reasoning within the individuals level. To make our case more convincing, we studied several variations of the game and demonstrated that: (1) enhancing the $L0$ strategy does not increase the use of level- k strategies (for $k > 0$) overall and lowers the levels conditional of using these level- k strategies, and (2) subjects do not use more than 3 steps of reasoning even when the monetary cost of each level of reasoning (for $k > 0$) is the same.

References

Arad, Ayala (2009). "The Tennis Coach problem: A Game-Theoretic and Experimental Study", The Foerder Institute for Economic Research, Working Paper no 7-08.

Basu, Kaushik (1994). "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory". *American Economic Review*, 84, 391-395.

Branas-Garza, Pablo and Maria Paz Espinosa and Pedro Rey-Biel (2009). "Travelers Types" Barcelona Economics Working Paper Series Working Paper no 407.

Bosch-Domenech Antoni, Jose G. Montalvo, Rosmerie Nagel and Albert Satorra (2002). "One, Two, (Three), Infinity...: Newspaper and Lab Beauty-Contest Experiments", *American Economic Review*, 92, 1687-1701.

Burchardiy, Konrad B. and Stefan P. Penczynski (2010). "Out Of Your Mind: Estimating The Level-k Model" Working Paper.

Camerer, Colin F. Teck-Hua Ho and Juin-Kuan Chong (2004). "Cognitive Hierarchy Model of Games". *Quarterly Journal of Economics*, 119, 861-898.

Costa-Gomes, Miguel A., Vincent P. Crawford and Bruno Broseta (2001). "Cognition and Behavior in Normal-Form Games: an Experimental Study". *Econometrica*, 69, 1193–1235.

Costa-Gomes, Miguel A. and Vincent P. Crawford (2006). "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review*, 96, 1737–1768.

Crawford, Vincent and Nagore Iriberri (2007). "Fatal Attraction: Saliency, Naivete, and Sophistication in Experimental Hide-and-Seek Games", *American Economic Review* 97, 1731-1750.

Kinderman, Peter, Robin Dunbar, and Richard, P. Bentall (1998). "Theory-of-Mind Deficits and Casual Attributions", *British Journal of Psychology*, 89, 191-204.

Nagel, Rosmarin (1995). "Unraveling in Guessing Games: An Experimental Study". *American Economic Review* 85, 1313-1326.

Rubinstein, Ariel (2007). "Instinctive and Cognitive Reasoning: A Study of Response

Times". *Economic Journal*, 117, 1243-1259.

Rubinstein, Ariel, Amos Tversky and Dana Heller (1996). "Naive Strategies in Competitive Games". *Understanding Strategic Interaction - Essays in Honor of Reinhard Selten*, Eds: Albers, W. et al, Springer-Verlag, 394-402.

Rosenthal, Robert. (1981). "Games of Perfect Information, Predatory Pricing, and the Chain Store," *Journal of Economic Theory* 25, 92-100.

Stahl, Dale O. and Paul W.Wilson (1994). "Experimental Evidence on Players' Models of Other Players". *Journal of Economic Behavior and Organization*, 25, 309-327.

Stahl, Dale O. and Paul W.Wilson (1995). "On Players Models of Other Players: Theory and Experimental Evidence", *Games and Economic Behavior*, 10, 213-54.

Zunshine, Lisa (forthcoming). "1700-1775: Theory of Mind, Social Hierarchy, and the Emergence of Narrative Subjectivity." *The Emergence of Mind: Representations of Consciousness in Narrative Discourse in English, 700 – the Present*. Ed. David Herman. Lincoln: University of Nebraska Press.

Footnotes

1. The number 20 is also the focal point in the range 11 – 20 in the following sense: If two people play a different game, in which each chooses a number in this range and is awarded a prize only if both choose the same number, then clearly they would choose 20.

2. Subjects' explanations were classified by a research assistant who otherwise was not involved in the research. The research assistant chose the categories himself. The classification was straightforward and we have also confirmed that we agree with his analysis.

3. Economics students at Tel Aviv University are considered to be fairly sophisticated subjects and one would expect that a different population, with less analytic ability, would tend to use lower levels of iterative reasoning.

4. The upper bound of the 95% confidence interval.

5. Braoas-Garza, Espinosa and Rey-Biel (2009) study a version of the Traveler's Dilemma framed as firms' price competition. They analyzed subjects' explanations and indeed found that only 30% of the explanations included undercutting arguments, which are in the spirit of the level-k reasoning process (mostly level-1 and in a few cases level-2).