

**Colonel Blotto's Top Secret Files:
Multi-Dimensional Iterative Reasoning in Action**

Ayala Arad

School of Economics, Tel Aviv University

Ariel Rubinstein

School of Economics, Tel Aviv University

Department of Economics, New York University

Abstract

We introduce a novel decision procedure involving multi-dimensional iterative reasoning, in which a player decides separately on the various features of his strategy using an iterative process. This type of strategic reasoning fits a range of complicated situations in which a player faces a large and non-ordered strategy space. In this paper, the procedure is used to explain the results of a large web-based experiment of a tournament version of the Colonel Blotto Game. The interpretation of the participants' choices as reflecting multi-dimensional iterative reasoning is supported by an analysis of their response times and the relation between the participants' behavior in this game and their choices in another game which triggers standard k-level reasoning. Finally, we reveal the most successful strategies in the tournament, which appear to reflect 2-3 levels of reasoning in the two main "dimensions".

Keywords: Colonel Blotto, Multi-dimensional iterative reasoning, Response Time

Many thanks to Eli Zvuluny who constructed and maintained the website gametheory.tau.ac.il through which the experiments were carried out, and to Yaniv Ben Ami and Hadar Binsky who provided research assistance in analyzing the data.

Our interest is in exploring the behavior of human beings in relatively complicated strategic situations. Most experimental research on strategic reasoning has focused on behavior in simple games in which the set of strategies is either small or naturally ordered. These games often trigger attractive decision rules such as playing the (pure strategy) Nash equilibrium, eliminating iteratively dominated strategies or implementing level- k reasoning. We wish to discover schemes of strategic reasoning used in situations in which it is hard to assess the strategy space due to its size and structure. In such situations, where there is no simple mental representation of the strategy space, the player is forced to formulate his strategy by choosing several features of a strategy. Furthermore, the player's payoff depends in a non-trivial way on the features chosen by the other player. A theory of choice in such situations could supply the instruments to construct new models of behavior in a variety of contexts.

In this paper we explore the experimental behavior of a large sample of participants in a variant of the renowned Colonel Blotto game. The game itself is interesting per se and has been applied to study a variety of economic and political situations. (Previous theoretical and experimental studies of the Blotto game will be discussed in Section 7.) The game also resembles other strategic scenarios such as allocation of funds between tasks, multi-characteristic product races and multi-object auctions. However, its main role in this paper is to serve as a platform for studying a novel strategic scheme that consists of reasoning about features, which can be relevant for analyzing other game situations as well.

1. The Colonel Blotto Game

Imagine you are a colonel in command of an army during wartime. You and the colonel of the enemy's army each command 120 troops. Your troops will engage the enemy in 6 battles on 6 separate battlefields.

It is the night before the battles and each of you must decide how to deploy your forces across the 6 battlefields. In the morning, you will win a battle if the number of troops you have assigned to a particular battlefield is higher than that assigned by your opponent. In the case that you have both allocated the same number of troops to a particular battlefield, the outcome of the battle will be a loss for both of you.

Your deployment of troops will face that of each of the other participants in a tournament. Your total score will be the number of battles you win against all the other participants.

How will you deploy your 120 troops?

Please pause for a second and try to devise a strategy. How would you play such a game?

The first possibility that comes to mind is the simple strategy of allocating the 120 troops evenly across the six battlefields. This instinctive strategy is likely to be chosen by other participants as well. If you choose this strategy, you will score no points against these participants. Furthermore, if some of the players will choose this strategy, and the marginal distribution of the other players' assignments in each battlefield will be symmetric around 20, the instinctive strategy will score less than 3 points on average. You speculate that winning the tournament requires scoring more than 3 points and thus choosing this strategy does not seem promising.

This may lead you to consider concentrating your troops in only some of the battlefields. But in how many? By allocating your troops evenly across five battlefields, you will win 5 battles and score 5 points against any opponent who chooses the instinctive strategy. However, what if the other players have the same thought and concentrate their troops in only five fields? In that case it might be better for you to assign a larger number of troops to only four battlefields, hoping to score 4 points against these players. It is not obvious where to stop this chain of arguments.

Suppose you have decided to deploy your troops in a certain number of battlefields and to abandon the others. Should you completely abandon those battlefields? Other players might also abandon them and therefore assigning even a small number of troops to these fields would score some easy victories. Other players are likely to go through a similar reasoning process and you need to decide on the exact number of troops to be assigned to the almost abandoned battlefields.

You might also doubt that the six fields are treated symmetrically. Is it possible that other players will systematically deploy more troops in some fields than in others? If so, you will gain more points by assigning larger masses to specific fields.

As you can see, the game is simple to describe but quite complicated to analyze since numerous considerations arise. At this point you are probably hoping that a game-theoretical analysis of the situation will provide some guidance in formulating a strategy. Calculating the equilibrium of this game, even if there are only two players, is extremely complicated. Even a best response for simple non-degenerated distributions of strategies (like the uniform distribution) is far from obvious. Nash equilibrium sometimes provides a reasonable prediction for games that are played by the same players many times.

However, it seems unlikely that the results will be close to the game's Nash equilibrium when each player participates in the tournament only once (see, for example, Holt and Roth (2004): "Experiments make clear that players often do not conform to equilibrium behavior when they first experience a game, even if it is a game in which behavior quickly converges to equilibrium as the players gain experience.")

Our goal in this paper is to advance our understanding of how people behave in a game of this type when playing for the first time. We will present the results of two experiments of the Blotto tournament, conducted with a huge number of participants, and will suggest a decision procedure that many of the participants appear to have followed.

A strategy in the Blotto game consists of an allocation of 120 troops across six battlefields. The number of possible strategies is around 250 million. A permutation is a set consisting of all strategies obtained by permuting a particular strategy (ignoring the labels of the battlefields). There are about 400 thousand permutations in the game. Indeed, thousands of strategies and almost a thousand permutations were actually chosen by the participants in our experiments. Thus, it is a challenge to identify and interpret common patterns of behavior in the data of such a game.

Unlike some other well-known simple games, it is hard to imagine a simple decision rule (such as successive elimination of dominated strategies) that could be used by a player in the Blotto tournament. It is also reasonable to assume that players do not adopt a process which involves finding a best response to a well-defined belief. This is because given the huge strategy space it is implausible that participants hold "single point" beliefs (which assume that most of the other participants will choose a particular strategy). If participants hold non-single point beliefs, then calculating the best response is enormously difficult, even if the belief's support consists of only a few strategies.

Our hypothesis is that the large size of the strategy space and its structure force a player to think in terms of *features* (or *dimensions*) of strategies rather than thinking about strategies per se. The paper presents a new decision procedure by which a player decides on each of the various features of his strategy separately and then integrates his choices in the various dimensions to formulate a strategy.

We discuss the decision procedure in Section 2. In Section 3 we describe the large-scale web-based experiment of the Blotto game which we carried out. In Section 4 we suggest that participants consider three particular features of a strategy and show that indeed many of the participants used our procedure and referred to these features. In Section 5 we further support the interpretation of choices in the Blotto game as reflecting multi-dimensional iterative reasoning by using additional data. Finally, in Section 6, we open Colonel Blotto's

“top secret files”, which reveal the salient strategies and the most successful strategies. We promise a surprise!

2. Multi-dimensional reasoning

We propose a new decision procedure that fits not only the Blotto game but also other games that are characterized by a rich structure. According to this procedure, a player has in mind some essential dimensions (or features) of a strategy. What these dimensions are depends on the structure of the game. The player makes a decision (chooses a value) for each dimension separately. He then picks a strategy that is consistent with the choices he made for each dimension.

The use of such a procedure is not uncommon in real life situations. Consider, for example, two competing manufacturers of fashion clothing, A and B, which produce a similar product. Each firm decides on the price and design of its product before it knows what the other firm has decided. In the previous year, both firms chose a price of \$12 and produced a similar design. In one of its meetings, A’s management decides to reduce the price to \$10 since it expects firm B to reduce its price to \$11 (A speculates that B anticipates that A will not be altering the price). Thus, in this dimension, A follows a decision rule in the spirit of level-k reasoning. In a separate meeting, A’s management decides to adopt a new and more modern design for the product since it expects firm B to stick to last year’s design. Given these two decisions, the planning department must come up with a new design that can be cheaply produced and has a modern look. The outcome might be a provocative design made out of cheap material. Of course, this outcome might not be a best response to firm B’s strategy.

Note that in this example a strategy is defined as a vector and players view the components of the vector as the dimensions for their strategic reasoning. However, even when a strategy is presented as a vector, players may have in mind dimensions that differ from those directly implied by the description of the strategy. As we will see later, this seems to be the case in the Blotto tournament.

We now give the procedure a more formal and detailed presentation. Let $\langle S, u \rangle$ be a symmetric game, where S is a set of strategies and $u(s, s')$ is the payoff (score) of playing the strategy s against a player who plays the strategy s' . We suggest that the choice of a strategy in complex games (i.e. when S is very large and lacks a clear ordering or a simple structure) is accomplished in two stages:

The editing stage

In this stage, a player performs a transformation of the game into a simplified framework in which he will later conduct his strategic deliberations. We propose that in this stage a player recognizes two elements:

(1) Dimensions.

The player has in mind a set of dimensions (features) of the game strategies. Dimension i must take a value from the set Z_i . For every dimension i , let $T_i(s)$ be a function that assigns to any $s \in S$ a unique element in Z_i . $T_i(s)$ is the value of the i 'th dimension of the strategy s .

Examples: (a) A new firm trying to compete with an established firm in the market might think of the following dimensions in deciding on its strategy: (i) product price, (ii) distribution area and (iii) advertising budget. A value for dimension (i), for instance, would be some non-negative number. (b) In a repeated game, one of the dimensions of a strategy might be the number of periods that a current action depends on and its value will be any non-negative integer. Another example could be the stationarity of the strategy and the value for this dimension might be "yes" or "no". (c) In a multi-object auction, one of the dimensions of a strategy might be the sum of all bids made by the player and its value will be any non-negative number. Another dimension could be the number of objects the player bids on.

(2) Proper response operator.

For each dimension i , a player has in mind a proper response function PR_i , which assigns to each value $z_i \in Z_i$ a value $z'_i \in Z_i$, with the interpretation that z'_i is beneficial against z_i . The proper response function represents a heuristic based on either an accurate calculation or a rough approximation.

Examples: (a) For any given $z \in Z_i$, a player considers the set $\{s | T_i(s) = z\}$, which is the set of all strategies for which the i 'th dimension takes the value z . He identifies the strategy s' which is a best response to the uniform distribution on this set and assigns $PR_i(z) = T_i(s')$.

(b) The player has in mind a representative strategy s with the value z in the i 'th dimension. He calculates the best response s' to s and assigns $PR_i(z) = T_i(s')$.

(c) $PR_i(z)$ is a heuristic which approximates a best response to the belief that the vast majority (but not all) of the other players will choose a strategy whose i 'th value is z .

The solution stage

The player is now able to implement a solution concept for the edited structure. The solution concept determines the values that a player will choose for each dimension and then the actual strategy he will play. The two elements of this stage are:

(3) *The method of choosing the value in each dimension.*

Examples: (a) A player has in mind an auxiliary game G_i in which each player chooses the value of the i 'th dimension. A player chooses a value that is a proper response to itself (a fixed point of the proper response operator).

(b) The player chooses the value that is a proper response to more values than any other.

(c) The player starts with an initial value $z_i(0)$ for the i 'th dimension. The initial value can be, for example, a salient value or the value of a salient strategy. He might end the process at this point and choose that value or alternatively he can recursively define $z_i(k) = PR_i(z_i(k-1))$. In this dimension, he is characterized by an integer k_i and chooses $z_i(k_i)$.

Note that a player might apply different dimensional decision rules in the various dimensions.

(4) *Picking a strategy.*

Once a player has chosen a vector of values z_i^* (one value for each dimension), he picks a strategy s such that $T_i(s) = z_i^*$ for all i . We will not make any particular assumptions regarding this element of the procedure.

Let us emphasize that the scheme described above is essentially based on heuristics. The application of such a scheme depends on the context. One may try to anchor the heuristics in a meta-optimization or to ground it on basic properties of human reasoning. However, we do not deal with the justification of the scheme here and make do with identifying traces of such a procedure in our experiment of the Blotto tournament.

Section 4 discusses the dimensions that participants in the Blotto experiments have in mind, the proper response operator they seem to use and their method of choosing a value for each dimension. We consider three dimensions of a strategy (element (1) in the above scheme). The proper response operator (element (2) in the above scheme) will be based on example 2(c). For the two central dimensions, the method of choosing a value (element (3) in the above scheme) will be based on iterative reasoning, as in example 3(c). In other words, a player starts with a salient value and might indeed choose it – this is step 0. Alternatively, he might implement step 1 of the reasoning in this dimension by choosing the value that is a proper response to the step 0 value. If he implements step 2 of the

reasoning, he chooses the value that is the proper response to the step 1 value and so on. Thus, we call our procedure multi-dimensional iterative reasoning.

Note both the similarity and differences between this procedure and the k -level reasoning approach (see Stahl and Wilson (1995)). A k -level model assumes that the population is partitioned into a collection of types, which differ in their depth of reasoning. Thus, a level-0 type is non-strategic and follows a simple decision rule. It is generally assumed that he randomizes uniformly (see Crawford and Irriberi (2007) and Arad (2009) for different specifications of the level-0 type, which take into account the instinctive attraction to salience). A k -level type, for any $k \geq 1$, best responds to the belief that all other players are level $k - 1$ types, or the belief that other players are level $k - 1$ or lower level types. In studies of other games, it was found that the most common types are level-1 and level-2 and it is rare to observe behavior consistent with a higher level. Our procedure differs from k -level reasoning in two major aspects: First, it relates to the features of the strategies rather than to the strategies themselves. This enables the level of reasoning to vary across dimensions. Second, it uses a proper response operator, which is only roughly related to the best response operator.

3. The Experiment

The platform used to conduct the experiment was the didactic website at gametheory.tau.ac.il. Each participant played our version of the Blotto game once. The participants came from two separate and distinct populations:

(i) **Classes:** Teachers of game theory courses occasionally assign their students virtual games and decision-theoretical problems from the site's bank of problems. The results obtained at the site are typically similar to those in laboratories experiments using monetary incentives (see Rubinstein (2007)).

Students were asked to participate in a tournament against their classmates; we will be reporting mainly the aggregated data of all the tournaments. The only incentive provided to the participants was that the three tournament winners in each class would have their names announced by the teacher. 4605 students had participated in the game. They belong to 129 groups in 25 countries (Argentina, Australia, Belgium, Brazil, Brunei, Canada, China, Denmark, Finland, France, Ireland, Israel, Mexico, Moldova, Netherlands, Norway, Slovakia, Spain, Switzerland, Taiwan, Thailand, Turkey, UK, USA and Vietnam).

(ii) **Calcalist:** "Calcalist" is a Hebrew business daily published in Israel. In the eve of Passover 2009, we invited Calcalist's readers to experience Game Theory by playing three

games posted on our website. The invitation was done through a newsletter, a link in a major news website, a link on "Calcalist on-line" and through Calcalist printed version in which the games were described. 1928 readers chose to participate. Prior to the Blotto tournament, the readers played two other new games: the Tennis Coach problem (Arad (2009)) and the "91-100" game (a variant of Arad and Rubinstein (2012a)), which will be described later in the paper. Both games naturally trigger k -level reasoning. It was promised to the readers that the names of the three tournament winners would appear in an article that will report the main findings.

Note that in both the Classes and the Calcalist experiment, participants were incentivized to try to win the tournament, i.e. to obtain the highest score among the participants. Our tournament differs from the standard two-player Blotto game in which each player is incentivized to maximize his expected score, although the games are quite similar both theoretically and psychologically. In particular, we have shown in a sequel paper (Arad and Rubinstein (2012b)) that when the number of players is very large, the equilibrium of any tournament of this type is close to a symmetric Nash equilibrium of the parallel two-player game. In the current study, we decided to use the tournament structure since we felt it would encourage daring strategies and strategic thinking.

Note also that even as a two-player game, our version is different from the classic variant of the Blotto game analyzed in the literature since it is not a zero-sum game (a tie in a battlefield is equivalent to a loss).

Given the huge number of possible strategies, as well as the variety of strategies actually chosen, thousands of independent observations are needed for a proper analysis of the considerations that arise in a player's mind. The use of web-based experimentation, which provided us with a sample of 6533 participants, overcame this hurdle. It also enabled us to confirm the robustness of the results through the sampling of two distinct populations: game theory students and readers of a financial newspaper, a less conventional group of participants for this kind of experiment. Despite the differences between the two populations, the results are very similar.

4. Multi-dimensional reasoning in the Blotto game

Before demonstrating how the multi-feature procedure can be plausibly applied in the Blotto game, it is worthwhile reviewing the difficulties involved in applying the standard k -level approach: We begin from the conjecture that the most prominent starting point for

iterative reasoning is the instinctive deployment of troops, i.e., 20 to each battlefield. We take this salient strategy to be the natural specification for level-0 behavior since it has the aesthetic features which make it the first strategy to come to mind (evidence for which is the particularly low response time associated with this strategy).

The difficulty arises when trying to specify the typical actions of higher-level types. Unlike some other simple games, here there are many best responses to the level-0 strategy, which makes the specification of the level-1 type unclear. Consequently, it is not reasonable to assign a single-point belief to the level-2 type. Holding a complex belief, which takes into account all the possible level-1 strategies, is not plausible either. Furthermore, the calculation of a best response to a non-degenerate distribution of level-1 strategies is very difficult even if the belief contains only a few strategies in its support.

Taking the level-0 to be a uniform randomization over the strategy space (that is, assigning equal probabilities to all 250 million strategies) is of no benefit since the calculation of the best response to this behavior is extremely difficult.

Thus, we find it more likely that a participant in the Blotto game employs the multi-dimensional reasoning process described in Section 2. In other words, he chooses his strategy after considering several dimensions and decides on the value of each dimension independently. For two of the dimensions, it is natural to assume that a player applies iterative reasoning (in the spirit of k -level reasoning) when choosing the value of the dimension.

We focus on three important features of a strategy: the number of reinforced battlefields, the unit digit in a single-field assignment and the order of the six single-field assignments. Our intuition (supported by comments we received from participants and readers) is that players consider the dimensions sequentially roughly in this order.

(a) The number of reinforced battlefields.

Choosing to reinforce $k = 0, 1, \dots, 5$ battlefields means that the participant has decided to strengthen his forces in k battlefields by assigning to each of them a large number of troops. We define the reinforcement of a battlefield as the assignment to it of more than 20 troops.

In this dimension we will adopt the proper response operation $PR_i(z)$ (described in 2(c), Section 2) which approximates a best response to the belief that the vast majority (but not all) of the other players will choose a strategy whose i 'th value is z . The method of choosing the value in this dimension will be based on an iterative process as described in 3(c) from Section 2. The description of the iterative process requires specifying the starting point and the proper response operator. The starting point of the iterative process in this

dimension is the reinforcement of 0 battlefields since the instinctive strategy, which assigns 20 troops to each battlefield, involves 0 reinforcements. This strategy is of course the only non-dominated strategy that involves 0 reinforcements.

As for the proper response function, we identify the reinforcement of 5 battlefields to be a proper response to the reinforcement of 0 battlefields and the reinforcement of 4 battlefields to be a proper response to the reinforcement of 5 battlefields. However, it is not clear that the reinforcement of 3 battlefields can be considered as the third step of reasoning. It is worthwhile elaborating on this point.

The first iterative step is to reinforce 5 battlefields. (The straightforward strategy of this kind involves deploying 24 troops in 5 of the battlefields.) If used against the step-0 strategy it will win 5 battles and thus score the maximum number of points possible in this game. This is the reason that reinforcing 5 battlefields is intuitively the optimal response to the step-0 strategy. Furthermore, if a player believes that a vast majority of participants, but not all, will reinforce 0 battlefields, then reinforcing 5 battlefields is a necessary condition for him to win the tournament. Given this belief, reinforcing 4 battlefields or less does not guarantee winning the tournament: if even one player in the tournament reinforces 5 battlefields, that player will score an average of almost 5 points, which is higher than the average score achieved by anyone who reinforces 4 battlefields or less.

As intuition suggests, the second step is the reinforcement of 4 battlefields. (The straightforward strategy of this kind involves deploying 30 troops in each of 4 battlefields.) The second-step type in this dimension believes that a vast majority of players reinforce 5 battlefields. By reinforcing 4 battlefields, the player expects to score about 4 points against the step-1 strategies, while the step-1 strategies will score around 3 points against each other. Thus, he will expect to win the tournament as long as the proportion of step-0 types is not greater than the proportion of step-1 types. Reinforcing less than 4 battlefields yields at most an average score of 3 points and is not successful given his beliefs.

An automatic continuation of the iterative process may lead to the thought that reinforcing 3 battlefields is a proper response to the reinforcement of 4 battlefields. Indeed, reinforcing 3 battlefields would generally yield a score of at least 3 points against step-2 strategies, whereas the average score of step-2 strategies against themselves is at most 3 (due to the possibility of ties, which are counted as losses). However, if in addition to step-2 strategies there are strategies of lower steps, a step-2 strategy will have the advantage of scoring about 4 points against these strategies. In such cases, reinforcing 3 battlefields may turn out to be inferior overall. Thus, the third iterative step is not clear-cut. In any case, the iterative chain stops here. Reinforcing less than 3 battlefields is not optimal

against strategies that involve reinforcing 3 or more battlefields.

Note that in the calculation of a proper response a player uses "an approximation of best response" to the lower step. He does so ignoring the other dimensions of the strategy and believing that if he reinforces less battlefields than his opponent, then he is likely to win in each of the reinforced battlefields. This argument makes sense since he has more resources for each reinforced battlefield. However, this is just an approximation since it does not take into account the possibility that the opponent's assignment in each reinforced battlefield can differ in size and some assignments may be very large.

We proceed into the analysis of the data. We will say that a participant has reinforced a battlefield if he assigned there more than 20 troops. This definition is somewhat arbitrary but it allows partitioning the strategies according to the number of reinforcements and captures the essence of concentrating troops in certain battlefields (for example the strategy 23-23-23-14-14 involves the reinforcement of 4 battlefields).

The following table presents the distribution of the number of reinforced battlefields for each population, as well as the median response time (MRT) for each category (measured in seconds). Note that when classifying the participants by level of reasoning we excluded a few dozen participants who did not assign all 120 troop.

# of fields with more than 20 troops		Classes			Calcalist		
		%	Score	MRT (σ)	%	Score	MRT (σ)
0	step 0	12%	2.24	113s (3.3s)	11%	2.04	79s (5.2s)
5	step 1	8%	3.15	182s (10.1s)	14%	2.92	135s (6.3s)
4	step 2	26%	3.08	192s (5.7s)	33%	3.05	143s (3.1s)
3	step 3 ?	22%	2.95	189s (6.3s)	22%	2.89	128s (4.7s)
2		19%	2.59	163s (4.1s)	13%	2.54	124s (6.3s)
1		12%	1.73	149s (6.3s)	7%	1.73	110s (7.9s)

Table 1

The data on response time supports our intuition regarding the structure of iterative reasoning in this dimension. The step-0 strategy is associated with exceptionally low response time, indicating that this choice is indeed instinctive. The step-1 and 2 strategies are associated with a relatively high response time, suggesting the use of a more complex deliberation process. The response time for the step-2 strategies is somewhat higher than that for step-1 strategies. The response time of strategies with 3 reinforced battlefields is

also high, suggesting that participants who made this choice were involved in a complex reasoning process as well. It is possible that these participants continued the iterative reasoning process intuitively, in an attempt to respond properly to step-2 strategies (though it is not clear that their choice is actually a proper response). Participants who decided to reinforce only one or two battlefields spent significantly less time on the decision, a hint that those bad choices were made hastily.

In Table 1 there is a difference between the Classes students and the Calcalist readers: the Calcalist readers tended to reinforce 4 or 5 fields relatively more often and to reinforce 1 or 2 fields less often. This might be because the Calcalist readers are more sophisticated. Alternatively, having participated in the two other games (the 91-100 game and the tennis coach problem) prior to the Blotto game may have triggered deeper iterated reasoning in this dimension among the Calcalist readers.

Note that the response time of the Calcalist readers was generally lower than that of the students. The difference may be explained by the fact that the Calcalist readers were presented with an Hebrew version of the game, which is much shorter than the English version presented to the students. Moreover, Hebrew is the Calcalist readers' mother tongue, whereas many of the students in the Classes group are not native speakers of English. In any case, we use the response time for comparisons within each group separately.

(b) The unit digit in single-field assignments

As in Dimension (a), we will adopt here the proper response operator described in 2(c) and the method of choosing a value described in 3(c) (Section 2).

A non-complicated and somewhat instinctive allocation of 120 troops across 6 battlefields involves single-field assignments that are multiples of 10 troops. Thus, we consider the use of the unit digit 0 in all battlefields as reflecting step 0 in this strategy's feature.

The most efficient way to win a battlefield is by assigning to that battlefield one troop more than the opponent. Thus, if a player suspects that a vast majority of the opponents' single-field assignments involve the use of a certain unit digit, using often a unit digit greater by one can be considered a proper response. (Note that the unit digits across the battlefields are not independent since they must sum up to a multiple of ten.) Of course, the assumption in the background is that the player's choice of tens digit will frequently be the same as his opponents'. Thus, the first iterative reasoning step would be the use of the unit digit 1 in some of the battlefields. The second step would be using the unit digit 2 and so

on. We doubt, though, that the choice of the unit digit 7, for example, is actually an outcome of 7 steps of reasoning. Even in simple games in which iterative reasoning is natural and straightforward, strategies that reflect more than 3 steps of reasoning were rarely observed in the literature (see a discussion in Arad and Rubinstein (2012a)).

Table 2 presents the distribution of unit digits in all the single-field assignments:

	Unit digit									
	0	1	2	3	4	5	6	7	8	9
Classes	62%	10%	3%	2%	4%	12%	1%	1%	2%	4%
Calcalist	56%	13%	5%	2%	5%	11%	1%	1%	2%	4%

Table 2

A majority of single-field assignments involved the unit digit 0 (almost all of them were either 0, 10, 20, 30 or 40). The unit digit 1 is heavily used (primarily in the choices 1 and 21). The choice of the unit digit 2 is less frequent (it appears primarily in the choices of 2 and 22).

We found that 30%-38% of the participants used the unit digits 1 or 2 in at least one single-battlefield assignment and half of them used those unit digits for at least three battlefields. The participants in this group spent significantly more time in deliberation than other participants: the MRT was 214s among the students and 153s among the Calcalist readers. In contrast, the MRT of those participants who chose only multiples of 10 (45% of the total participants) was 137s among the students and 106s among the Calcalist readers, reflecting this choice's instinctive nature. The MRT of the other participants who (i) did not use only multiples of 10 and (ii) did not use the unit digits 1 or 2 at all was somewhere in between (170s among the students and 127s among the Calcalist readers).

The unit digit 5 was the second/third most frequently used. This is the result of two patterns observed in the data: (i) One-fifth of the participants who used the unit digit 5 at least once, used the unit digit 1 five times. (ii) A large number of participants used only multiples of 5 (as in the strategy 25-25-25-25-10-10). The unit digit 4 was used more frequently than 6,7 or 8 due to the popularity of allocating 24 troops to 5 battlefields (step 1 in the reinforced battlefield dimension). The use of the unit digit 9 was clearly a corollary of the use of the unit digit 1.

In Table 2 we find again a difference between the students in the Classes group and the Calcalist readers: the Calcalist readers used the unit digits 1 and 2 more frequently and the

unit digit 0 less frequently. This difference is in the same spirit of the difference indicated above regarding the number of reinforced fields.

(C) *Order of single-fields assignments*

Once a player has chosen a particular partition of his 120 troops into 6 "divisions", he also needs to decide how to allocate the (perhaps) different-sized divisions among the 6 battlefields. Intuitive procedures of allocating the troops may treat the battlefields in a non symmetric way. For example, a player could allocate divisions successively, starting with allocating the strongest division to battlefield 1 and ending with allocating the weakest division to battlefield 6. Alternatively, he could concentrate the stronger divisions in the middle battlefields and the weaker in the edges (or the opposite).

Note that iterative reasoning is not as natural here as in the other two dimensions. Since there is more than one intuitive way to allocate the divisions, there is ambiguity regarding the step 0 value in this dimension. The definition of a proper response is also not as clear as in the previous dimensions. Suppose, for example, that a player believes that the other player is concentrating his troops primarily in the middle battlefields. One proper response would be to concentrate more troops in the middle battlefields and to assign a relatively small number of troops to the edges. Another plausible proper response would be to abandon the two central battlefields and assign more troops to all other battlefields in order to increase the chances of winning those battles.

Although an iterative process is less likely to be implemented in this dimension, the choice of value can still be an outcome of some other deliberation process. We do not identify a particular decision rule used in the order dimension, but we do trace in the data some interesting asymmetries that are presented below.

The six battlefields are numbered 1 to 6. Naturally, we focus on two types of symmetry:
directional: Are battlefields 1,2,3 treated identically to battlefields 6,5,4 respectively?
positional: Is the pair of battlefields in the center (3 and 4) treated the same as the pair of battlefields in the edges (1 and 6) and as the pair of battlefields in the mid-positions (2 and 5)?

It should be mentioned that the students and the Calcalist readers played versions of the game that differed in one framing detail: in the game played by the students, the 6 battlefields were arranged vertically, with battlefield 1 on top and battlefield 6 on the bottom. In the game played by the Calcalist readers, the battlefields were arranged horizontally with battlefield 1 on the left and battlefield 6 on the right. However, this did

not appear to have any effect on the results concerning this dimension.

The following graphs present the cumulative distribution of the number of troops assigned to each of the six battlefields. The graphs reveal that the marginal assignment of troops to each battlefield is far from being that induced by the uniform distribution over the interval $[0,40]$, which is the game-theoretic prediction for the related classical constant-sum Blotto game. Note the dramatic "jumps" around numbers like 20 and 30 and that only 10% (rather than 25%) of the choices in each field are above 30.




Figure 1: Cumulative distribution of single-field assignments in the six battlefields

The cumulative distributions are essentially ordered identically in the two populations by first-order stochastic domination: 3,4,2,5,1,6. For convenience, we also present the 33rd, 50th, 67th percentile points for each of the 6 battlefields:

Field	Classes						Calcalist					
	1	2	3	4	5	6	1	2	3	4	5	6
33rd percentile	10	15	20	20	15	8	10	20	20	20	18	10
Median	20	20	21	20	20	20	20	21	24	22	20	20
67th percentile	25	25	28	26	24	21	24	25	30	29	25	21

Table 3

Most noticeable is the low number of troops assigned to the 6th battlefield and the high number assigned to battlefields 3 and 4. This is in line with some other experimental results which demonstrate a tendency of people to avoid the edges and to concentrate resources on the center positions (see, for example, Rubinstein, Tversky and Heller (1996)).

There is almost no distinction between right and left. Battlefields 2 and 5 are treated almost symmetrically and the distributions for battlefields 3 and 4 are also very close. The fact that more troops are assigned to battlefield 1 than to battlefield 6 is the only directional asymmetry that was observed. A possible explanation is that a participant's instincts lead him to over-assign troops. Since allocations are often executed from the first battlefield to the last, battlefield 6 is treated as the "residual".

As can be seen in Table 4, some of the most popular single-field assignments appear in a non-symmetric way in the six fields. For example, the assignments 0 and 1 are twice as frequent in battlefields 1 and 6 as in battlefields 3 and 4 and the frequency of the assignment of 30 to each of the battlefields 3 and 4 is much higher than for the other pairs.

Field	Classes							Calcalist						
Assign.	%	1	2	3	4	5	6	%	1	2	3	4	5	6
0	14	19%	17%	11%	11%	16%	22%	16	19%	15%	8%	9%	14%	20%
1	4	7%	4%	3%	3%	4%	5%	4	7%	4%	2%	2%	5%	6%
30	15	14%	14%	19%	17%	12%	12%	15	13%	14%	21%	19%	14%	11%

Table 4

5. Support for the presence of multi-dimensional iterative reasoning

At this stage, we wish to introduce the "91-100 game". Calcalist readers (unlike the Classes group) played this game before playing the Blotto game. The fact that the same participants played both the 91-100 game and the Blotto game enables us to examine the relation between their observed behavior in the two games. This can help in evaluating our interpretation of participants' reasoning in the Blotto game. The data in this section is based solely on the Calcalist's participants.

5.1. The 91-100 game

Following is a description of the game as presented (in Hebrew) to the Calcalist readers:

You and another person are playing a game in which each player requests an amount of money. The amount must be an integer between 91 and 100 shekels. Each player will receive the amount he requests. A player will receive an additional amount of 100 shekels if he asks for exactly one shekel less than the other player.

What amount of money would you ask for?

In this game it is hard to think of more than one dimension for a strategy. We find the game in particular suitable for studying (one-dimensional) k -level thinking for four reasons:

(i) *The level-0 type specification is intuitively appealing*: The choice of 100 is a natural anchor for an iterative reasoning process. It is the instinctive choice when choosing a sum of money between 91 and 100 dollars (100 is the salient number in this set and "the more money the better"). Furthermore, the choice of 100 is not entirely naive: if a player does not want to take any risk or prefers to avoid competition, he might give up the attempt to win the additional 100 dollars and simply request the highest certain amount.

(ii) *Best-responding is easy*: Given the anchor 100, best-responding to any level- k action is very simple and leaves no room for errors.

(iii) *Robustness to the level-0 specification*: The type-1 action, i.e. choosing 99, is the unique best response to a wide range of reasonable beliefs including (a) all distributions in which 100 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) *Clear payoffs*: Unlike other games that trigger iterative reasoning (for example, the Centipede Game and the Traveler's Dilemma), in this game social preferences are not likely to be prominent. In particular, if a player believes that his opponent has chosen a number $n > 91$, choosing $n - 1$ will reward him with \$100 but not at the expense of the other player. Social preferences can appear in our game only in the extreme case in which a player would like the other player to get the bonus.

Assuming that all players wish to maximize the expected amount of shekels they receive, the game has a unique symmetric mixed-strategy Nash equilibrium (which yields an expected payoff of 100).

Action	91	92	93	94	95	96	97	98	99	100
Equilibrium	55%	9%	8%	7%	6%	5%	4%	3%	2%	1%
Experiment	18%	19%				10%	21%	18%	14%	

Table 5

As shown in Table 5, the behavior of participants is very far from the Nash equilibrium. Most notable is the low percentage of participants who chose 91 relative to the equilibrium prediction. The choice of 100, which in equilibrium appears only rarely, was chosen by

14% of the participants. The actions 97-98-99 that seems to exhibit 1-2-3 levels of reasoning were chosen by 49% of the participants, whereas in equilibrium they should have been chosen by only 9%. As noted above, higher levels of iterative reasoning are almost never observed in other studies of k -level reasoning. In our results, the actions 92-96 are also rare and appear much less often than expected by equilibrium.

In a parallel experiment of this game, about 160 students provided ex-post explanations for their choices. An analysis of their explanations suggests that the actions 92-96 are generally not an outcome of 4- to 8-level of reasoning but of simpler decision rules. It also validates the classification of 97-98-99 as the 1-2-3 levels of reasoning.

5.2. Association between behavior in "91-100" and "Blotto"

In this section, we seek support for our interpretation of some of the choices in the Blotto game as an outcome of multi-dimensional steps of reasoning. This is done by investigating the association between standard k -level iterative behavior in the 91-100 game and behavior in the Blotto game which seems to exhibit iterative reasoning in the various dimensions. More precisely, we examine the link between the choices 97-98-99 in the 91-100 game and either reinforcing 4 or 5 battlefields or using the unit digits 1 or 2 in the Blotto game. Given that the choices 97-98-99 clearly reflect 1-2-3 levels of reasoning in the 91-100 game (whereas the rest of the strategies are attributed to other decision rules), evidence of such an association will provide support for our intuition that these Blotto game choices emerge from iterative reasoning.

In analyzing the relation between the use of dimensional iterative reasoning in the Blotto game and k -level reasoning in the 91-100 game, we combined the 97-99 choices into one category. In our experience, the level of reasoning for a particular player is rarely the same in two different games not from the same class. Moreover, for any two players, the relative ordering of their levels is not stable between two games from different classes. This is supported by the findings in Georganas, Healy and Weber (2010). Thus, the most we can hope for is correlation between some use of level- k reasoning in the 91-100 game and iterative reasoning in a particular dimension in the Blotto game.

Table 6 shows the distributions of choices in the 91-100 game as a function of the participants' choice of the number of reinforced battlefields:

# of fields with more than 20 troops		Action in 91-100			
		91	92-96	97-99	100
0	step 0	25%	19%	35%	21%
5	step 1	17%	13%	56%	13%
4	step 2	16%	17%	55%	12%
3	Step 3 ?	16%	19%	50%	15%
2		20%	25%	42%	13%
1		16%	21%	42%	21%

Table 6

Participants who did not reinforce any of the battlefields (i.e. 0 reinforcements) tended to choose 91 and 100 more often and to choose 97-99 dramatically less often than the other participants. Of those who reinforced 4 or 5 battlefields, 55-56% chose 97-99 ($\sigma = 2\%$) whereas of those who reinforced 2 or less battlefields the proportion was only 42% ($\sigma = 2\%$). The behavior of participants who reinforced 3 battlefields resembled more that of the participants who reinforced 4 or 5 battlefields: 50% of them chose 97-99 ($\sigma = 2\%$). This finding supports our conjecture that strategies involving 3 battlefields reinforcements are the result of an intuitive continuation of the iterative reasoning process in this dimension.

Table 7 presents the mean number of battlefields with the unit digits 1 or 2 in the Blotto game as a function of the choices in the 91-100 game.

Action in 91-100	91	92-96	97-99	100
%	18%	18%	49%	15%
Blotto: Mean # of fields with digit 1 or 2 (σ)	0.82 (0.08)	0.97 (0.09)	1.20 (0.06)	1.03 (0.10)

Table 7

Table 7 demonstrates that participants who chose 97-99 in the 91-100 game tended to use the unit digits 1 or 2 significantly more often than the rest. Another way to see it: more than 57% of the participants who used the unit digits 1 and 2 at least once chose 97-99, while only 45% of those who did not use those digits chose 97-99.

A total of 47% of the Calcalist' participants reinforced 4 or 5 battlefields in the Blotto game and 38% of the participants used the unit digits 1 and 2 at least once. The choices of 25% of the participants exhibit iterated reasoning in both dimensions. We found that 59% ($\sigma = 2\%$) of them chose 97-99, whereas among those who reinforced fewer than 3

battlefields and did not use the digits 1 or 2, only 39% ($\sigma=2\%$) made those choices.

Incidentally, the choices of 97-99 are also associated with a high score in the Blotto game. Table 8 demonstrates this by comparing the 91-100 choices of the highest-performing 20% in the Blotto game with the choices of the rest of the participants. Of the Blotto game's top scoring players, 60% ($\sigma=2.5\%$) chose 97-99, whereas only 46% ($\sigma=1\%$) of the rest of the players chose these actions.

	91	92-96	97-99	100
Blotto's top 20%	10%	20%	60%	11%
the rest	20%	19%	46%	15%

Table 8

To summarize, we find that participants who exhibit 1-3 levels of reasoning in the 91-100 game tend to more often choose a strategy using 1 or 2 steps of reasoning in each of the two dimensions (i.e. the number of reinforced battlefields and the unit-digit numbers). This finding provides support for our interpretation of the reinforcement of five or four battlefields and the use of the unit-digit 1 and 2 as reflecting multi-dimensional iterative reasoning.

6. The Secret Files: Popularity and Success in the Blotto Game

The popular strategies

Nine of the 250 million strategies were chosen by at least one percent of the participants. The strategies and their average scores are presented in Table 9. The nine most popular strategies were together used by around 30% of the participants in each of the populations. Table 10 presents the permutations that were chosen by at least 2% of the participants in each of the populations. Apparently, there were eight permutations that were the most popular in both populations and all together were chosen by 41-45% of the participants. (A permutation's average score is calculated as the average for all participants whose strategies were within that permutation.)

#	Strategies						Classes		Calcalist	
	Battlefields						n=4605		n=1928	
	1	2	3	4	5	6	%	Score	%	Score
1	20	20	20	20	20	20	11.4%	2.33	11.1%	2.09
2	30	30	30	30	0	0	4.4%	2.87	4.6%	2.86
3	0	0	30	30	30	30	3.4%	2.97	4.6%	2.91
4	120	0	0	0	0	0	1.9%	0.98	1.1%	0.99
5	21	21	21	21	21	15	1.5%	3.19	3.3%	2.80
6	24	24	24	24	24	0	1.4%	3.08	1.6%	2.90
7	0	30	30	30	30	0	1.3%	2.93	1.2%	2.86
8	40	40	40	0	0	0	1.2%	2.76	1.6%	2.79
9	0	24	24	24	24	24	1.0%	3.16	1.9%	2.94

Table 9

	Permutation	Classes	Mean Score	Calcalist	Mean Score
		n=4605		n=1928	
1	20-20-20-20-20-20	11.4%	2.33	11.1%	2.09
2	30-30-30-30-0-0	11.2%	2.92	13.3%	2.90
3	120-0-0-0-0-0	4.7%	0.99	2.4%	0.99
4	40-40-40-0-0-0	4.1%	2.78	4.5%	2.81
5	30-30-20-20-10-10	2.9%	2.74	2.0%	2.66
6	24-24-24-24-24-0	2.8%	3.12	4.0%	2.94
7	30-30-29-29-1-1	2.2%	3.27	2.5%	3.18
8	21-21-21-21-21-15	2.0%	3.20	4.8%	2.81

Table 10

The most popular strategies and permutations were not very successful and scored well below the best-performing strategies in the tournament (which scored around 3.8 points on average).

Attentiveness of the participants

Participants were not forced to assign all the 120 troops across the battlefields. This was a device for checking their attentiveness. Among the students, only 5.4% of the participants

chose such a dominated strategy. Among Calcalist's readers, the proportion dropped to 3.0%.

One popular permutation, which may be the outcome of players not paying sufficient attention to the game, is the assignment of 120 to one of the battlefields. This permutation was chosen by 4.7% of the participants in the Classes and 2.4% among the Calcalist readers. The choice of the homogenous strategy is instinctive but does not necessarily imply that insufficient attention was paid to the game. It was chosen by 11% of our participants. To scale this fact, in the Avrahami and Kareev (2009)'s Blotto experiment (described below), which was carried out in a laboratory with monetary incentives, 25% of the participants chose the homogenous strategy in the first round of the game.

Relation to Nash Equilibrium

The distribution of scores in both the Classes and the Calcalist groups has a high variance and constitutes evidence that the recorded results are far from a Nash equilibrium of the tournament:

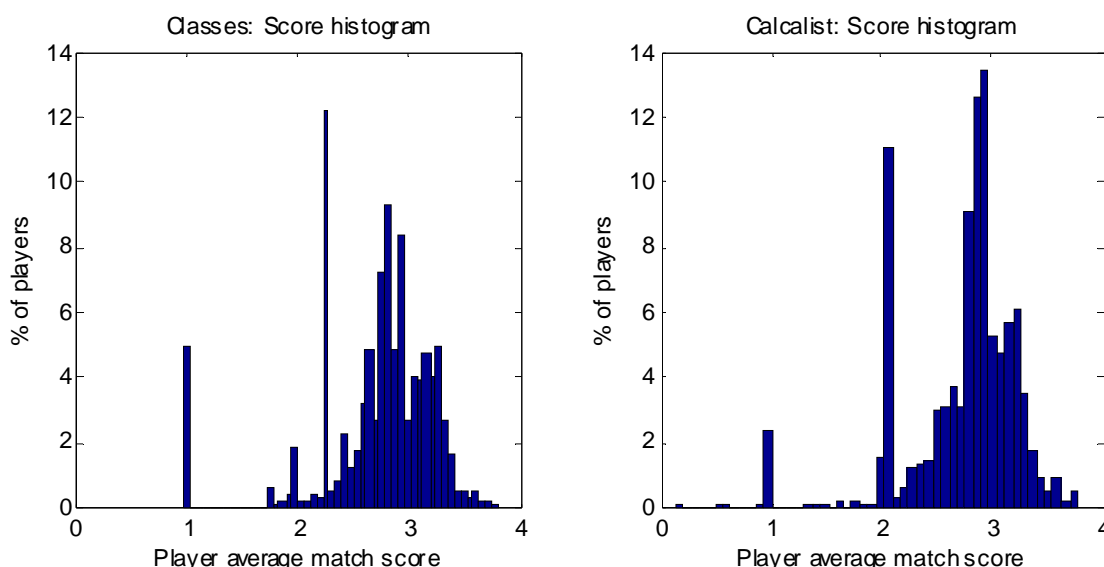


Figure 2: Distribution of average scores in the Classes and the Calcalist groups

The winning strategies

And now to the surprising result: The winning strategy in the Classes' grand tournament and in the Calcalist tournament was the same: 2-31-31-31-23-2 (and needless to say, was chosen by two different people...). Furthermore, there was also significant overlap in the lists of the top 10 strategies in the two tournaments (see Table 11). Four strategies are

common to both lists and, up to a permutation, 7 out of the top 10 strategies on the Classes list appear on the other list as well.

Classes' grand tournament								Calcalist tournament							
	1	2	3	4	5	6	Score		1	2	3	4	5	6	Score
1	2	31	31	31	23	2	3.83	1	2	31	31	31	23	2	3.77
2	3	31	31	31	21	3	3.80	2	2	32	31	31	22	2	3.76
3	3	31	3	31	31	21	3.76	3	2	23	31	31	31	2	3.75
4	1	31	31	31	25	1	3.76	4	1	1	32	32	32	22	3.72
5	2	27	31	31	27	2	3.75	5	1	1	31	31	31	25	3.71
6	2	31	23	31	31	2	3.74	6	2	27	31	31	27	2	3.71
7	1	1	31	31	31	25	3.73	7	2	31	1	31	31	24	3.70
8	2	21	32	32	2	31	3.72	8	1	31	31	31	25	1	3.70
9	1	1	31	31	25	31	3.71	9	1	25	31	31	31	1	3.69
10	1	31	31	25	31	1	3.69	10	1	1	34	31	31	22	3.69

Table 11

We were curious as to whether there is another strategy that was not chosen and could have done better than the others. We simulated 41,040 strategies in which all unit digits are 1,2 or 3 and found no such strategy.

The features of the winning strategy 2-31-31-31-23-2 are illuminated by the explanation provided by the Calcalist winner:

"In the first stage, I decided that I would "surrender" on two fronts, but not so easily. I thought that other people would decide to assign a few battalions to some of the fronts and perhaps would not deploy any battalions to other fronts. So I could win on an "abandoned" front at the inexpensive price of one battalion. Eventually, I decided to deploy two battalions on the weak fronts in order to overpower anyone who thought like me and placed one battalion on the weak fronts. It seems logical to me that the weak fronts would be on the edges. I was left with 116 battalions to allocate to four fronts, which is an average of 29 battalions per front. I decided to reinforce three of the four remaining fronts with two battalions - that is, to deploy 31 battalions - in order to defeat those who allocated the remaining battalions equally. In this way, I would also defeat those who allocated 30 battalions to each of the four central fronts."

Here are some of the features characterizing the ten leading strategies:

a) Two battlefields were essentially abandoned. In fact, all 30 leading strategies in the two tournaments used a low number of troops (1,2 or 3) in exactly two single-field assignments.

b) The most often almost abandoned battlefields are 1 and 6. This is profitable since these battlefields tended to be abandoned in the population much more than the middle ones.

c) Battlefields 2 and 5 were treated rather symmetrically (and, in particular, the strategy 2-23-31-31-31-2 does almost as well as the winning strategy).

d) 30+ troops are generally assigned to the middle battlefields. This is beneficial since the assignments to battlefields 3 and 4 tended to be the highest.

It is interesting to look at the winning strategies in the 11 tournaments of the largest classes, which contained at least 60 participants. In 4 of these classes (Argentine (2) and Canada (2)), a permutation of 1-35-1-31-31-21 was the winning strategy. In the other 4 (Switzerland (2), Thailand and Slovakia), a permutation of 31-1-31-1-31-25 was the winning strategy. In the remaining 3 large classes (in Switzerland and Argentina), the winning strategies were 31-31-31-21-3-3, 3-21-3-31-21-21 and 7-33-33-7-33-7. Note that 8 out of the 11 winning strategies belong to the same two permutations. All winning strategies in the 11 large classes, like the overall leading strategies (in the two grand tournaments), involved the reinforcement of 4 battlefields and the avoidance of multiples of ten. The winning strategies in the large classes performed well in the grand tournament as well. While the top 10 strategies in the grand tournament scored on average 3.7-3.83, the winning strategies in the large classes achieved an average score of 3.6-3.76 in the grand tournament.

The components of success

In order to evaluate the relative contribution of each of the three features of a strategy to a participant's score we ran a regression in which a participant's score is explained by three variables, corresponding to the three features discussed earlier:

A) $A = 1$ if the participant reinforces 5 fields, $A = 2$ if he reinforces 4 fields and $A = 0$ otherwise.

B) $B =$ the number of fields in which the participant used the unit digit 1 or 2.

C) $C = 1$ if the numbers of troops allocated to each of the extreme fields (1 and 6) were strictly less than the number allocated to each of the fields 2,3,4 and 5.

The results for the two populations are remarkably similar:

$$\text{ScoreClasses} = 2.41 + 0.28A + 0.10B + 0.20C$$

($R^2 = 0.32$, 95% confidence intervals for the four coefficients respectively: (2.40, 2.43), (0.26, 0.29), (0.09, 0.11) and (0.14, 0.25))

$$\text{ScoreCalcalist} = 2.49 + 0.24A + 0.10B + 0.16C$$

($R^2 = 0.40$, 95% confidence intervals for the four coefficients respectively: (2.46, 2.51), (0.23, 0.26), (0.09, 0.11) and (0.09, 0.23))

The large confidence interval for the coefficient of C is due to the small number of participants for whom $C = 1$ (6% among the students and 7% among the Calcalist readers). In contrast, $A > 0$ for 35%-47% of the participants and $B > 0$ for 27%-29% of the participants. Conditional on $A > 0$, the expected value of A is 1.7 – 1.8 and conditional on $B > 0$, the expected value of B is 2.7 – 2.8.

The results provide evidence that all three features play a significant role in explaining a player's success in the game. In assessing their relative importance, note that the variables have different domains (A is 0, 1 or 2, B is between 0 and 5 and C is 0 or 1) and thus the importance of a feature cannot be determined solely by comparing the absolute size of the coefficients. Overall, it seems that the variables A and B are more significant than C . Using the first/second step of reasoning in the number of reinforced fields dimension or using the unit digits 1 or 2 in two battlefields or more has a greater effect on the score than allocating the lowest single-field assignments to battlefields 1 and 6.

Time and Performance

Does an investment of more thought in the Blotto game translate into a better performance?

A standard regression affirms that response time contributes significantly and positively to performance in the Blotto game. Thus, for the Classes group we obtain the equation: $\text{score} = 2.07 + 0.12\ln(\text{response time})$ and for the Calcalist sample we obtain a similar equation: $\text{score} = 1.99 + 0.15\ln(\text{response time})$.

In order to better understand the correlation, we divided the two populations into ten deciles according to their response time. Figure 3 plots the average score and the two standard errors for each decile. We find that the mean score of the three bottom deciles is dramatically lower than that of the two top deciles. On the other hand, the average performance in the 4th-8th deciles is almost identical.




Figure 3: Mean score by response time deciles

Gender effects

Calcalist readers (but not the students) were asked to report their gender. Only 10% of the readers were females. Males did significantly better with an average score of 2.75 ($\sigma = 0.01$) as compared to 2.55 ($\sigma = 0.04$) for females. This is in spite of the fact that females spent more time on the game (females' median response time was 145s vs. 125s for the males).

Experience in Game Theory

We asked the Calcalist readers whether they had ever taken a course in Game Theory and one-sixth of them had. One might expect that a course in game theory would improve one's ability to play games; especially a synthetic game like the one experimented here. We found that taking a Game Theory course has only a marginal effect: The average score of Game Theory graduates was 2.78 ($\sigma = 0.03$) as compared to a close average score of 2.72 ($\sigma = 0.01$) for the others. In other words, a course in Game Theory may be entertaining, but there is no evidence that it helps much in playing games.

7. Bibliographic notes

The classic version of the game was suggested in Borel (1921). The equilibrium of its continuous version was characterized by Roberson (2006). The equilibrium of the discrete case, with B troops allocated to K battlefields, was characterized by Hart (2008). Both concluded that in an equilibrium, players treat the battlefields symmetrically and the marginal distribution of the troops in each battlefield is essentially uniform in the interval $[0, 2B/K]$. Our non-constant-sum version of the game is somewhat different from these versions since two players who tie on a particular battlefield do not split a point but rather get nothing.

The Blotto game has received widespread attention due to its interpretation within the political economics literature as a game between two presidential candidates who have to allocate their limited budgets to campaigns in the “battlefield” states (see Brams (1978)). Myerson (1993) suggested another interpretation of the Blotto game as a vote-buying game.

Only a few experiments of Blotto games have been conducted. Partington reports in his website (<http://www.amsta.leeds.ac.uk/~pmt6jrp/personal/blotto.html>) on a Blotto game tournament conducted in 1990. In his version, the participants had to allocate 100 troops across 10 battlefields. The winning strategy was 17-3-17-3-17-3-17-3-17-3.

Avrahami and Kareev (2009) report on an experiment of a "lottery version" of the constant-sum game. Each participant played 8 times in a row against a single player. In each round, once the two players have chosen their allocation of troops, one battlefield per player was randomly selected and the winner of the round was determined by comparing between the assignments in the two selected battlefields. This design prevents framing effects induced by the ordering of the battlefields. Among other things, the authors studied the case in which each player assigns 24 troops across 8 battlefields. In this case, the theory predicts that the marginal distribution of the assignment in each battlefield will be uniform in $[0, 6]$. In the vast majority of observations, 2-4 troops were assigned to each battlefield and a significant number of participants allocated the troops homogeneously (3 troops to each battlefield). For other recent experiments of variants of the game, see Chowdhury, Kovenock and Sheremeta (2009), Cinar and Göksel (2012), Horta-Vallve and Llorente-Saguer (2010) and Modzelewski, Stein and Yu (2009).

8. Conclusion

In the paper we focused on a version of the Blotto game. We consider the game to be an example of an interesting complicated game, in which the set of strategies is large and it is

difficult to calculate (or even to approximate) a best-response function. We have argued that in such games (and not only the Blotto game itself) it is natural to decide about each of the various features of strategies separately rather than about strategies per se. Often, the deliberation over a feature naturally involves an iterative process. The decision procedure we have proposed captures this sort of strategic reasoning.

The following are three classes of strategic situations in which using the procedure seems natural:

(i) Variations of the Blotto game. Each player has limited resources that he has to assign to various "tasks" (battlefields). The features of a strategy in such a context might be similar to those referred to in our discussion of the Blotto game.

(ii) Multi-object auctions. A set of objects is put up for auction. Players simultaneously place bids on each one. In this case, the features could be, say, the sum of the bids, the number of objects not to be bid on, and the number of objects on which to place high bids. The starting points of the iterative process in the various dimensions could be, say, "a sum of bids as determined by convention", "bidding on all objects" and "bidding very high on only half of the objects".

(iii) Product races. Consider two competing car producers who bring out a new model of executive car each year. Relevant features might be the timing of the launch, what generation to upgrade the technology to, investment in advertising and price. Last year's choices might serve as the initial values for the iterative process.

Note that we do not claim that people always use the multi-dimensional iterative reasoning scheme. All we argue is that in some games there is a significant group of people who choose their strategy based on a procedure of this type. People use a variety of decision procedures even in very simple situations and thus we do not think that one can hope for a universal procedure that would explain the behavior of all people in such games.

It is worthwhile discussing the difficulties in applying the multi-dimensional procedure to explain behavior in other games. The procedure has many degrees of freedom: the dimensions of a strategy, the starting point of the iterative process in each dimension and the proper-response operator. The identification of these elements is "ad-hockish" and might depend on the context and the framing of the situation. Future research may uncover principles for identifying these elements in any situation. However, we believe that at this stage, common sense could be used to point out these elements in particular strategic interactions of interest.

To conclude: to the best of our knowledge, this paper is the first attempt to define and identify in the data a process involving several non-inclusive forms of iterative reasoning.

We believe that our findings shed light on major considerations that arise in contexts similar to the Blotto game. Incorporating multi-dimensional reasoning in theoretical models and analyzing experimental data of other games in light of such a decision procedure are topics for future research.

References

- Arad, Ayala (2009). "The Tennis Coach problem: A Game-Theoretic and Experimental Study". *The B.E. Journal of Theoretical Economics* (contributions) (forthcoming).
- Arad, Ayala and Ariel Rubinstein (2012a). "The 11-20 Money Request Game: Evaluating the Upper Bound of k-Level Reasoning". *American Economic Review* (forthcoming).
- Arad, Ayala and Ariel Rubinstein (2012b). "Strategic Tournaments: How Does the Competition Influence Players' Behavior", mimeo.
- Avrahami, Judith and Yaakov Kareev (2009). "Do the Weak Stand a Chance? Distribution of Resources in a Competitive Environment". *Cognitive Science*, 33, 940-950.
- Borel Émile (1921). "La Théorie du Jeu et les Équations Intégrales à Noyau Symétrique". *Comptes Rendus Hebdomadaires des Séances l'Académie des Sciences* 173, 1304–1308. (Translated by L.J. Savage as "The Theory of Play and Integral Equations with Skew Symmetric Kernels". *Econometrica*, 21 (1953), 97–100.)
- Brams, Steven J. (1978). *The Presidential Election Game*. Yale University press.
- Chowdhury, Subhasish M., Dan Kovenock J. and Roman M. Sheremeta (2009). "An Experimental Investigation of Colonel Blotto Games". *Economic Theory* (forthcoming).
- Çinar Yetkin and Türkmen Göksel (2012). "An Experimental Analysis of Colonel Blotto Games: Decision Making Under Alternative Environments". *İktisat, İşletme ve Finans*, 27, 39-57.
- Crawford, Vincent and Nagore Iriberri (2007). "Fatal Attraction: Saliency, Naïveté, and Sophistication in Experimental "Hide-and-Seek" Games". *American Economic Review* 97, 1731-1750.
- Hart, Sergiu. (2008). "Discrete Colonel Blotto and General Lotto Games". *International Journal of Game Theory*, 36, 441–460.
- Holt, Charles A. and Alvin E. Roth (2004). "The Nash Equilibrium: A Perspective". *PNAS*, 12, 3999-4002
- Hortala-Vallve, Rafael and Aniol Llorente-Saguer (2010). "A Simple Mechanism for Resolving Conflict". *Games and Economic Behavior*, 70, 375-391.
- Modzelewski Kevin, Joel Stein and Julia Yu (2009). "An Experimental Study of Classic Colonel Blotto Games". MIT (mimeo).
- Myerson, Roger (1993). "Incentives to Cultivate Favored Minorities Under Alternative Electoral Systems". *The American Political Science Review*, 87, 856-869.
- Roberson, Brian (2006). "The Colonel Blotto Game". *Economic Theory* 29, 1–24.
- 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*, W.Guth et al. (editors), Springer-Verlag, 394-402.

Sotiris, Georganas, Paul J. Healy, and Roberto Weber (2010). "On the Persistence of Strategic Sophistication". Working Paper, University of Bonn.

Stahl, Dale O. and Paul W. Wilson (1995). "On Players Models of Other Players: Theory and Experimental Evidence". *Games and Economic Behavior*, 10, 218-54.