

Analysis of Equilibria in Iterative Voting Schemes

Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis and Jeffrey S. Rosenschein

Abstract

Following recent analyses of iterative voting and its effects on plurality vote outcomes, we provide a general characterisation of the set of equilibria attainable by iterative plurality voting. We show that deciding whether a given profile is an iteratively reachable equilibrium is NP-complete; however, if *truth bias* is added, we show that it is possible to determine all equilibria in polynomial time. Furthermore, we fully characterise the set of iterative truth-biased equilibria. We then examine the model of *lazy voters*, in which a voter may choose to abstain from the election, showing that the iterative voting process in this case converges to a stable state. As in the case with *truth bias*, we show that it takes at most polynomial time to find stable states that are Nash equilibria.

1 Introduction

There are many aspects to coordination in multiagent systems that have engaged researchers in recent years, including questions related to the aggregation of multiple agents' preferences into a single system-wide choice. Researchers looking at group decision-making have explored the properties of voting schemes, which provide well-founded preference aggregation techniques that lend themselves to formal analysis; at times, voting can also provide intuitive and efficient prescriptive algorithms for resolving disagreements among participants.

Unfortunately, as the Gibbard-Satterthwaite theorem [10, 20] famously states, voting rules are susceptible to manipulation; under minor assumptions, for every voting rule there exists a set of voter preferences, in which at least one of the voters will be better off misreporting its preference. Given this negative result, there has been a solid body of research over the last two decades that has focused on the complexity of manipulation, as a potential barrier to the Gibbard-Satterthwaite theorem. In a different direction, and in lieu of having strategyproof voting rules (which is unattainable), another body of research has emerged on game-theoretic analysis of voting, initiated by Farquharson [7]. Viewing voters as strategic agents, it is then natural to examine the Nash equilibria of the underlying voting games, as a potential solution concept for preference aggregation scenarios.

However, Nash equilibria in this context, and without any further refinement, end up being a poor tool to predict voting behavior. For example, even if all voters rank the same candidate last, it is still a Nash equilibrium (in most reasonable voting rules) that all voters vote for this disliked candidate. More generally, there can be a very large number of equilibria in most voting games, and many of them are irrelevant to the effective analysis of the voting rule. Several methods have been proposed to handle this multitude of equilibria, and in this work, we focus on the following three:

- **Iterative voting:** Instead of examining all Nash equilibria, one can consider those equilibria that are “reachable” from certain initial voting profiles; that is, we can study sequences of moves, where during each move one of the players who is dissatisfied with the current result can change its vote, and achieve a better outcome for itself. This process is reminiscent of a group of friends trying to find a movie that they would all like to see, or a restaurant where they would all want to go to; dissatisfied with the

current result, they change their declared preference seeking to modify the outcome (see e.g., [16]). Several online services, such as Doodle, enable this kind of iterative preference aggregation.

- **Truth bias:** Another way to handle the large number of Nash equilibria is to add a small utility gain for voters who vote truthfully. This extra “bonus” should be small enough so that agents would still prefer to manipulate the election in cases where they can affect the outcome for their own benefit. Introducing this bias in the model dramatically reduces the number of equilibria by several orders of magnitude [21].
- **Voting with abstentions:** A different way to simulate real-world incentives is to add a small utility gain to voters who do not participate in the election, i.e., who abstain. The rationale is that coming to the election may incur a cost in time, effort, etc. Hence the abstention avoids this cost, but just as in truth bias, the agents would still have an incentive to manipulate if they can affect the outcome. Such voters are also referred to as *lazy voters* [4].

Contribution: We consider the above three techniques and their relationship, separately and combined, to Nash equilibria. We first examine iterative voting in a simple form, namely plurality voting with agents playing a best-response strategy. We find that even in this straightforward model (which is known to eventually converge [16]), checking whether a given profile is a reachable Nash equilibrium is NP-hard.

We then turn to the truth-bias approach, and combine it with iterative voting: the voting process stops only when it reaches a truth-biased Nash equilibrium. In this model, iterative voting converges to equilibria that are arguably more natural; truth bias reflects intuitive voter leanings, and eliminates many undesirable equilibria. We give a characterisation of the equilibrium profiles reachable under this model, and detail a polynomial-time algorithm for finding all such equilibria of a game.

Following that, we examine the lazy-voter approach, combined with iterative voting. We study the iterative process under the assumption that if a voter decides to abstain at some step, then he does not come back to the election at a later step. While convergence to a stable state is guaranteed under this model, the final state may not necessarily be a Nash equilibrium. However, we are still able to fully characterize the Nash equilibria that can be reached by this process.

1.1 Related Work

The iterative voting model we utilise is based on the one introduced in [16] and later expanded by additional researchers [15, 2, 12]. That model followed previous research into iterative and dynamic mechanisms, much of it summarised in [13], which also touched (lightly) on single-peaked preferences (as did [9], albeit not in a dynamic process). More recently, [1] examined another iterative process, using a type of plurality voting rule; however, it is quite different from our general iterative model. Furthermore, [5] used an iterative process that eliminates weakly dominated strategies (a requirement also used in the definition of equilibrium in [8]), and showed criteria for an election to result in a single winner via this process. While we do not deal with it in this paper, we note that there is also a line of research that discusses iterative processes in the context of limited information games, and studies the effects of lack of knowledge on the players’ strategies [3, 17]; and yet another line of work studies restricted dynamics in iterative processes, i.e., limitations on the allowed moves by the voters, e.g., [19, 11].

The notion of adding a truth bias to games was introduced (for a specific case) by [14], and was proposed for a specific voting rule (with limited results) by [6]. A more robust

model was suggested by [21], which introduced the general framework of the truth-bias equilibrium, and contained various empirical results when using plurality in truth-biased games. The theoretical side of that work was recently enhanced by [18].

The notion of lazy voting was studied in [4], as another way of eliminating some of the undesirable NE (Nash equilibria). The twist is that the utility function is changed so that voters have a slight preference for abstaining if they are not pivotal. However, they can still have an incentive to manipulate the election if they can affect the outcome. See [4] for more details.

Apart from iterative voting and truth bias, there have been many other attempts to escape the multitude of Nash equilibria in voting games. Some of these include a) introducing uncertainty, e.g., about how many voters support each candidate, as in [17]; b) changing the temporal structure of the game itself. [22] and [4] consider the case where agents vote publicly and one-at-a-time, and study subgame-perfect equilibria of these extensive-form games; In this paper, we do not address any of these approaches.

The rest of the paper is organised as follows. We start with definitions and notation in Section 2. Section 3 presents our first results, on equilibrium characterisation of iterative plurality voting, including our NP-completeness proof for a profile reachability decision. This is contrasted by the results of Section 4, which describe the properties of *truth-biased* iterative plurality voting, leading to a polynomial-time algorithm for equilibrium computation. We conclude in Section 6 with some final remarks, and a discussion of future work. While omitting many of this paper’s proofs due to space limitations, we provide key proofs and proof elements.

2 Definitions and Notation

We consider a set of m candidates $C = \{c_1, \dots, c_m\}$ and a set of n voters $V = \{1, \dots, n\}$. Each voter i has a *preference order* (i.e., a ranking) over C , which we denote by a_i . For notational convenience in comparing candidates, we will sometimes use \succ_i instead of a_i . When $c_k \succ_i c_j$ for some $c_i, c_j \in C$, we say that voter i prefers c_k to c_j .

At an election, each voter submits a preference order b_i , which does not necessarily coincide with a_i . We refer to b_i as the vote or ballot of voter i . The vector of submitted ballots $\mathbf{b} = (b_1, \dots, b_n)$ is called a *preference profile*. At a profile \mathbf{b} , voter i has voted truthfully if $b_i = a_i$. Any other vote from i will be referred to as a non-truthful vote. Similarly the vector $\mathbf{a} = (a_1, \dots, a_n)$ is the *truthful preference profile*, whereas any other profile is a non-truthful one.

Given a voter $i \in V$, and its vote b_i under a profile \mathbf{b} , we denote by $top(b_i)$ the top choice of the vote. A *voting rule* \mathcal{F} is a mapping that, given a preference profile \mathbf{b} over C , outputs a candidate $c \in C$; we write $c = \mathcal{F}(\mathbf{b})$. In this paper we will consider only Plurality under lexicographic tie-breaking. This is one of the most well-studied and widely-used voting rules. Under Plurality, each candidate is assigned a *score* equal to the number of ballots where it has appeared as the top choice. The winner of the election is then the candidate with the maximum score. I.e. the winner of the election under Plurality is the candidate who appears as the top choice in the maximum number of votes. In case of ties, we assume that tie-breaking is resolved by the linear order $c_1 \succ c_2 \succ \dots \succ c_m$.

Given \mathbf{b} , let s be the maximum score achieved by a candidate. We denote by $W(\mathbf{b})$ the set of tied candidates with score equal to s , i.e., all the potential winners before tie-breaking is applied. Also, let $H(\mathbf{b})$ be the set of candidates that receive $s - 1$ votes in \mathbf{b} , but would win a tie-break against any candidate in $W(\mathbf{b})$ (these are candidates who would need one extra vote to become a winner). These two sets play an important role in our analysis, as

they, together, define the “runner-ups” — the candidates that can win with an additional point. Finally, we denote by $\mathcal{F}(\mathbf{b})$ the winning candidate of a profile \mathbf{b} , and by $\text{sc}(c, \mathbf{b})$ the score of candidate c in \mathbf{b} .

2.1 Game theoretic considerations

In this work, we view elections as non-cooperative games. The standard way to do this is to associate a utility function u_i with every voter i , which is consistent with its true preference order. That is, we require that $u_i(c_k) \neq u_i(c_j)$ for every $i \in V$, $c_j, c_k \in C$, and also that $u_i(c_k) > u_i(c_j)$, if and only if $c_k \succ_i c_j$.

We study and compare three game-theoretic models. We refer to the first one as the *basic model* (following [16]) since it is the most standard approach. Under the basic model, we suppose that the payoff function of voter i when its real preference is a_i is:

$$p_i(a_i, \mathbf{b}, \mathcal{F}) = u_i(c_j), \text{ if } c_j = \mathcal{F}(\mathbf{b}),$$

where \mathbf{b} is the submitted profile.

The second model we consider is a variation of the first one, and we refer to it as the *truth-biased model*, following [21]. In this model, we suppose that voters have a slight preference for voting truthfully when they cannot unilaterally affect the outcome of the election. This bias is captured by inserting a small extra payoff, when the voter votes truthfully. This extra gain is small enough so that voters may still prefer to be non-truthful in cases where they *can* affect the outcome. If \mathbf{a} is the real profile and \mathbf{b} is the submitted one, then the payoff function of voter i is given by:

$$p_i(a_i, \mathbf{b}, \mathcal{F}) = \begin{cases} u_i(c_j), & \text{if } c_j = \mathcal{F}(\mathbf{b}) \wedge a_i \neq b_i, \\ u_i(c_j) + \epsilon, & \text{if } c_j = \mathcal{F}(\mathbf{b}) \wedge a_i = b_i. \end{cases} \quad (1)$$

The third model in this paper is also a variation of the first one. We refer to it as the *model with lazy voters*, following [4]. In this model, it is assumed that voters have a slight preference for not coming to the election if they are not pivotal. This bias is captured similarly to the truth bias in the previously described model variation. Let \perp denote the abstention ballot. If \mathbf{a} is the real profile and \mathbf{b} is the submitted one, then the payoff function of the voter i is given by:

$$p_i(a_i, \mathbf{b}, \mathcal{F}) = \begin{cases} u_i(c_j), & \text{if } c_j = \mathcal{F}(\mathbf{b}) \wedge b_i \neq \perp, \\ u_i(c_j) + \epsilon, & \text{if } c_j = \mathcal{F}(\mathbf{b}) \wedge b_i = \perp. \end{cases} \quad (2)$$

In the pathological case that the submitted profile is the vector (\perp, \dots, \perp) , we assume that no candidate is elected and each voter has a payoff of ϵ . This clearly cannot be realized as a stable state.

A Nash equilibrium in these games is a profile \mathbf{b} , where no voter can unilaterally improve its payoff, i.e., for every i and every vote b'_i , we have $p_i(a_i, \mathbf{b}, \mathcal{F}) \geq p_i(a_i, (b'_i, \mathbf{b}_{-i}), \mathcal{F})$ (\mathbf{b}_{-i} being the vector \mathbf{b} without player i 's vote).

2.1.1 Best Response Dynamics

Consider an election game, either in the basic or in the truth-biased model. We focus on an iterative process, where, starting from the truthful preference profile, voters can change their strategy by making improvement steps. An *improvement step* for a voter i at a profile \mathbf{b} is a switch to another strategy (i.e., vote) b'_i , leading to the profile (b'_i, \mathbf{b}_{-i}) , in which the payoff for i is strictly higher than before. A *best response* improvement step is one in which the voter changing its strategy achieves its currently best possible payoff.

We will focus only on best response steps, and we will refer to an *improvement path* as any sequence of submitted profiles such that each move from one to the next is a best

response step. We do not assume any fixed order by which voters update their strategies; we only assume that the voters start from the truthful profile (a natural starting point for such a process, as also argued in [2]) and then make their best response updates in an arbitrary order. The process in general can lead to different outcomes, or may not even converge. We are interested in studying the Nash equilibria of election games, that can be reached by best response improvement paths.

3 Basic model analysis

We begin by analysing the properties of Nash equilibria that are reachable under iterative voting in the basic model of plurality aggregation. First, recall the following property of best response improvement paths.

Lemma 1. *[Quoted from Lemma 4 in [2]] An improvement step can only take a vote for a non-winning candidate and transfer it to the winner of the newly formed voting profile.*

We can now consider the implications of voting iterations on Nash equilibrium profiles.

Definition 1. *Given a profile \mathbf{b} , we denote by $CS(\mathbf{b})$ (and refer to it as the chasing set of \mathbf{b}), the set: $CS(\mathbf{b}) = (W(\mathbf{b}) \cup H(\mathbf{b})) \setminus \{\mathcal{F}(\mathbf{b})\}$, i.e., the candidates who could become a winner if they were to receive one additional vote.*

The following fact which we use repeatedly in the sequel follows from the analysis in [2].

Fact 1. *Let \mathbf{b} be an equilibrium profile obtained by a sequence of improvement steps from the truthful profile \mathbf{a} . Then $\mathcal{F}(\mathbf{b}) \in W(\mathbf{a}) \cup H(\mathbf{a})$.*

Lemma 2. *Consider an equilibrium profile \mathbf{b} obtained by a sequence of improvement steps from the truthful profile \mathbf{a} . Then $CS(\mathbf{b}) \subseteq W(\mathbf{a}) \cup H(\mathbf{a})$, and if $\mathbf{b} \neq \mathbf{a}$, then $CS(\mathbf{b}) \neq \emptyset$.*

The properties identified in the previous lemmas help us formulate the following necessary conditions for reachable equilibria.

Theorem 1. *Let \mathbf{b} be an equilibrium profile obtained by a sequence of improvement steps from the truthful profile \mathbf{a} . Then the following holds.*

1. *For every voter i , $\mathcal{F}(\mathbf{b}) \succ_i c$, $\forall c \in CS(\mathbf{b}) \setminus \{top(b_i)\}$.*
2. *For every voter i such that $top(b_i) \in CS(\mathbf{b})$, $top(b_i) \succ_i c$, $\forall c \in CS(\mathbf{b}) \setminus \{top(b_i)\}$.*

Theorem 2. *Given a truthful profile \mathbf{a} and a profile \mathbf{b} , distinct from \mathbf{a} , it is NP-complete to decide if \mathbf{b} is reachable by iterative best-response updates, starting from \mathbf{a} .*

Proof Theorem 2: To show that the problem is in NP, it is enough to provide as a certificate the sequence of best-response updates that leads from profile \mathbf{a} to profile \mathbf{b} . One could then check that this is a valid sequence. It is important to note here that this is indeed a certificate of polynomial length. For this, it suffices to bound the number of possible best-response moves until we reach \mathbf{b} . The crucial observation here is that after any such move, the score of the currently winning candidate either increases or remains the same as the score of the previously winning candidate—it never decreases. This means that if \mathbf{b} is reachable, then it can be done in at most mn steps.

To prove NP-hardness, we provide a reduction from the Hitting Set (HS) problem, which is the following: we are given a set of ground elements $G = \{g_1, \dots, g_n\}$, and a family of subsets of G , $W = \{w_1, \dots, w_m\}$, $w_i \subseteq G$, $|w_i| = l_i$. We are also given a number $k \leq n$. The

Block 1	Block 2	Block 3	Block 4	Block 5
$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_m$	
$u_1 \dots u_k \dots u_1 \dots u_k$	$g_{1_{i_1}} \dots g_{m_{i_m}}$	$u_1 \dots u_k \dots u_k \dots u_1$	$t \dots t$	
$\vdots \ddots \vdots \vdots \ddots \vdots$	$w_1 \dots w_1 \dots w_m \dots w_m$	$\vdots \ddots \vdots \vdots \ddots \vdots$	arbitrary	stand-ins
$u_k \dots u_1 \dots u_1 \dots u_k$	$G \setminus g_1 \dots G \setminus g_1 \dots G \setminus g_{m_1} \dots G \setminus g_{m_{i_m}}$	$u_k \dots u_1 \dots u_k \dots u_1$	order	
$G \setminus g_1 \dots G \setminus g_n$	$t \dots t \dots t \dots t$	$G \setminus g_1 \dots G \setminus g_n$	over	
$t \dots t \dots t \dots t$		$w_1 \dots w_1 \dots w_1 \dots w_1$	all	
$D \cup W \dots D \cup W$	arbitrary order over all others	arbitrary order over all others	others	

Table 1: NP-Completeness proof profiles: Truthful profile. Recall, $|w_i| = l_i$.

Block 1	Block 2	Block 3	Block 4	Block 5
$t \dots t \dots t \dots t$	$w_1 \dots w_1 \dots w_m \dots w_m$	$w_1 \dots w_1 \dots w_1 \dots w_1$	$t \dots t \dots d_{ w_1 } \dots d_m$	as was
$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_{ w_1 -1} \dots t$	
$u_1 \dots u_k \dots u_1 \dots u_k$	$g_{1_{i_1}} \dots g_{m_{i_m}}$	$u_1 \dots u_k \dots u_k \dots u_1$		
$\vdots \ddots \vdots \vdots \ddots \vdots$		$\vdots \ddots \vdots \vdots \ddots \vdots$		
$u_k \dots u_1 \dots u_1 \dots u_k$	as was	$u_k \dots u_1 \dots u_k \dots u_1$	as was	
$G \setminus g_1 \dots G \setminus g_n$		$G \setminus g_1 \dots G \setminus g_n$		
$D \cup W \dots D \cup W$		as was		

Table 2: NP-Completeness proof profiles: Target profile. Recall, $|w_i| = l_i$.

decision problem is to ascertain that there is a hitting set $U \subset G$, so that $|U| \leq k$, and $\forall i \in [m]$, $U \cap w_i \neq \emptyset$. This is a well-known NP-complete problem.

We assume that we are given an instance that satisfies: 1) $|w_1| \geq |w_i|$, $\forall i \in [m]$, 2) $|w_1| \geq 3$, and 3) $m \geq n$ (we can always pad an instance by replicating a set to satisfy this). These three assumptions do not impact the complexity of the HS problem.

Given such an instance of the HS problem, we proceed by constructing an instance of our problem, i.e., a truthful profile \mathbf{a} , and a matching (non-truthful) profile \mathbf{b} , so that a sequence of iterative best-response updates going from \mathbf{a} to \mathbf{b} exists if and only if the HS instance has a solution.

Given an HS instance as above, we associate one candidate with each element of G and one candidate with each element of W . In addition, we introduce k candidates u_1, \dots, u_k , corresponding to the (up to) k elements of U . Finally, we also add a set D of m dummy candidates, $D = \{d_1, \dots, d_m\}$, and a special target candidate t . Overall, there are $n + 2m + k + 1$ candidates in our voting problem with the following lexicographic order for tie breaking: $d_1 \prec \dots \prec d_m \prec u_1 \prec \dots \prec u_k \prec w_1 \prec \dots \prec w_m \prec g_1 \prec \dots \prec g_n \prec t$. We slightly abuse the notation so that each w_j refers both to the set from the HS instance and the corresponding candidate in our instance—and similarly for the element candidates g_i .

We will now introduce five blocks of voters with preferences as depicted in Tables 1, 2. Notice that Block-1 and Block-3 contain nk voters each, Block-2 has $\sum |w_i| \leq mn$ voters, and Block-4 has exactly m voters. These cardinalities will be used in the later stages of the proof. Block-5 (the *stand-ins* in our Tables 1, 2) is incidental and its voters are only necessary to create an initial balance among the candidates in the truthful profile. These voters all have the order of preference $U \succ G \succ t \succ W \succ D$, broken for each voter only by shifting one particular candidate to be the top choice. We have as many votes in Block-5 as required to ensure that after counting all truthful votes in all the Blocks, each candidate receives exactly $2k + n$ in the truthful profile \mathbf{a} . This also means that given the tie-breaking rule, t is the winner in \mathbf{a} .

We will show that ascertaining reachability of the reported profile depicted in Figure 2, is equivalent to solving the original HS problem. First, observe that no voter in Block-2 can change its vote in \mathbf{a} to one that appears in \mathbf{b} (i.e., vote for some w_i) in a single deviation step, while following the best-response strategy. This is because initially the candidates in G have the same number of votes as those in W , but win in tie-breaking, hence the best response in the beginning is to vote for someone in G . To make \mathbf{b} reachable from \mathbf{a} , those candidates from G that *lock* the votes in Block-2 need to be eliminated from becoming winners, thus making the vote for w_i the best choice.

We show that the issue described above can be resolved so that the target profile \mathbf{b} is reachable from \mathbf{a} if and only if there is a solution to the HS problem. To see this, assume first that there is a solution U , with $|U| \leq k$, to the HS instance. Then, we can associate (possibly with replications in case that $|U| < k$) to each candidate u_i an element in G , say $g(u_i)$, so that for $\widehat{U} = \bigcup_{i \in [k]} \{g(u_i)\}$, it holds that $\forall j \in [m], \widehat{U} \cap w_j \neq \emptyset$.

Let each voter in Block-1, that has u_i $1 \leq i \leq k$, as his second choice and $g(u_i)$ as his first choice (there is exactly one such voter for each i), change his vote in sequence starting in the order from 1 to k . As a result of the best-response updates, each candidate u_i will in turn receive an additional vote, while $g(u_i)$ will lose at least one vote. In this new profile, the winner will be u_k , due to tie-breaking, with all candidates from U having $2k + n + 1$ votes, while those in the set $\widehat{U} \subseteq G$ will have at most $2k + n - 1$ votes. Denote this new profile by \mathbf{c} .

Now, in the profile \mathbf{c} , consider those voters in Block-2 that have $g(u_i)$ as their second choice, for each element $g(u_i) \in \widehat{U}$. Note that because U is a solution to the HS instance, this implies that for each sub-block of Block-2, having w_j as a third choice, with $1 \leq j \leq m$, there is a value of i and a voter from this sub-block where $g(u_i)$ is a second choice for this voter. For all these voters, it is not a best response to vote for their second choice $g(u_i)$, since in the previous round of updates all elements from \widehat{U} lost a vote and due to tie-breaking they cannot become a winner with a single step. Instead, the best response for these voters is to vote for w_j , thus “unlocking” the candidates of W . Let us choose one such voter for each w_j and let them change their vote in sequence. This changes profile \mathbf{c} to a profile \mathbf{d} , where all candidates in D can no longer become a winner with a single deviation, all candidates in U and in W have $2k + n + 1$ votes, and candidates in G have at most $2k + n$ votes each (some of them have $2k + n - 1$).

In the next round of updates, we will prevent all candidates in U and G from ever again becoming a possible winner by essentially “running” a competition between the candidates in W and the target candidate t . That is, we choose a voter sequence that will grant $W \cup \{t\}$, an ever increasing number of votes, eliminating any other candidate from becoming a best-response. The effect of this voting sequence will eventually be the emergence of the profile \mathbf{b} .

Let us first allow one voter from Block-4 to change his vote. The voter will naturally shift t to be his top-choice. This will give the target candidate t $2k + n + 1$ votes as well, completely preventing all candidates from G from ever becoming a winner, since all of them have less than $2k + n + 1$ votes and lose in tie-breaking to t . We can now cycle, repeatedly through all candidates from W selecting for each w_j a voter from Block-2 with w_j as his third (truthful) choice, which now is the best-response for that voter. At the end of each cycle iteration we will grant one more voter from Block-4 the possibility to change his vote. There will be at most $|w_1|$ such cycles. Notice that some cycles will be shorter in the sense that there will be some sub-blocks where all voters will have already voted for their corresponding candidate from W (since w_1 may have a strictly higher cardinality from the rest of the sets). Finishing this process we will have the voters from Block-2 and Block-4 vote as they are intended in the target profile \mathbf{b} . The voters from Block-1 and Block-3 will still be voting either for a candidate from U or a candidate from G , and the voters from

Block-5 remain as they are. In this intermediate voting profile, \mathbf{e} , candidates from W will have at most $2k + n + \sum |w_j|$ votes, and so will the target candidate t .

We can now reach profile \mathbf{b} from \mathbf{e} , if we alternate between allowing a voter from Block-1 and a voter from Block-3 (that still do not vote for w_1 or t) to change their votes. Notice that the best-response top-choice for these voters is indeed either w_1 or t . This will result in Block-1 and Block-3 transforming their votes into those prescribed by \mathbf{b} , completing the vote modification sequence from \mathbf{a} . Notice, additionally, that in \mathbf{b} the winning candidate is the target candidate t , with $2k + n + nk + \sum |w_j|$ votes and no voter can change the outcome, i.e., \mathbf{b} is an equilibrium.

Finally, for the other direction, assume that there is no solution to the underlying HS instance. It is then easy to check that there is no possible sequence of votes that can “unlock” at least one voter in Block-2 for each w_j . Hence, this makes the targeted profile \mathbf{b} unreachable. \square

4 Truth-Biased Iterative Voting Under Plurality

Having fully characterised the Nash equilibria in the iterative voting scheme under plurality, and in light of the negative result of Theorem 2, we now introduce the assumption of truth bias as described in Section 2. For the remainder of this section we will consider only truth-biased agents and investigate how this property affects the outcome of iterative voting. Once again, we begin our analysis by recalling some basic and already established properties of equilibrium profiles under truth bias, namely, the following lemma, which is proved in [18].

Lemma 3. *Suppose that $\mathbf{b} \neq \mathbf{a}$ is a non-truthful profile, which is a Nash equilibrium. Let $c = \mathcal{F}(\mathbf{b})$. Then all non-truthful votes in \mathbf{b} have c as the top candidate.*

Recall, $s = \text{sc}(\mathcal{F}(\mathbf{a}), \mathbf{a})$. The following lemma is an easy corollary of Lemma 2 in [18].

Lemma 4. *Suppose that $\mathbf{b} \neq \mathbf{a}$ is a non-truthful profile, which is a Nash equilibrium. Let $c = \mathcal{F}(\mathbf{b})$. Then $\text{sc}(c, \mathbf{b}) \leq s + 1$.*

We let $W_{\succ c}$ denote the set of all candidates $c_j \in W(\mathbf{a})$ such that $c_j \succ c$. $H_{\succ c}$ is defined similarly. It is easy to see that Fact 1 continues to hold in this model too. Hence, again the winner at an equilibrium $\mathbf{b} \neq \mathbf{a}$, belongs to the set $W(\mathbf{a}) \cup H(\mathbf{a})$. The next lemmas shed more light on each of the two possible cases (that $\mathcal{F}(\mathbf{b}) \in W(\mathbf{a})$ or $\mathcal{F}(\mathbf{b}) \in H(\mathbf{a})$) and they also highlight some important differences between the basic model and the truth-biased one.

Lemma 5. *Suppose that $\mathbf{b} \neq \mathbf{a}$ is an equilibrium profile, and that $c \in W(\mathbf{a}) \cup H(\mathbf{a})$ is the winner in \mathbf{b} . Then*

- *there is only one voter (say i) who submits a non-truthful vote in \mathbf{b} ;*
- *for every $c_j \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{\text{top}(a_i)\}$, $c \succ_i c_j$;*
- *If $c \in W(\mathbf{a})$, for every voter k and every $c_j \in W_{\succ c} \setminus \{\text{top}(a_k)\}$, $c \succ_k c_j$. If $c \in H(\mathbf{a})$, for every voter k and every $c_j \in W(\mathbf{a}) \cup H_{\succ c} \setminus \{\text{top}(a_k)\}$, $c \succ_k c_j$.*

These properties yield as direct corollaries the following:

Corollary 1. *If there exists a chain of iterative improvement steps that leads to the Nash equilibrium $\mathbf{b} \neq \mathbf{a}$ with winner c , then there exists such a chain consisting of only 1 step.*

As a consequence of the Corollary 1 we can obtain Algorithm 1 which finds all Nash equilibria achievable by the sequence of improvement steps with complexity of $\mathcal{O}(m^2n^2)$. Using the following more accurate description of iteratively achievable Nash equilibria we can construct an algorithm with even better running time.

Corollary 2. *There exists a chain of iterative improvement steps that leads to the Nash equilibrium $\mathbf{b} \neq \mathbf{a}$ with winner $c \in W(\mathbf{a})$ if and only if the following conditions hold.*

1. *There is at most one candidate $c_j \neq \mathcal{F}(\mathbf{a}) \in W_{\succ_c}$ for which there exists at least one voter i , with $c_j \succ_i c$ and $c_j \neq \text{top}(a_i)$.*
2. *If no such candidate c_j , as described above exists, then there exists a voter i such that $c \succ_i c_k$ for every $c_k \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{\text{top}(a_i)\}$ and $\text{top}(a_i) = \mathcal{F}(\mathbf{a})$. Otherwise, there exists a voter i with $\text{top}(a_i) = c_j$ and $c \succ_i c_k$ for every $c_k \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{c_j\}$.*

Corollary 3. *There exists a chain of iterative improvement steps that leads to the Nash equilibrium $\mathbf{b} \neq \mathbf{a}$ with winner $c \in H(\mathbf{a})$ if and only if the following conditions hold.*

1. *There exists at most one candidate $c_j \neq \mathcal{F}(\mathbf{a}) \in W(\mathbf{a}) \cup H_{\succ_c}$ such that there exists at least one voter i with $c_j \succ_i c$ and $c_j \neq \text{top}(a_i)$.*
2. *If no such candidate c_j , as described above exists, then there exists a voter i such that $c \succ_i c_k$ for every $c_k \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{\text{top}(a_i)\}$ and $\text{top}(a_i) = \mathcal{F}(\mathbf{a})$. Otherwise, there exists a voter i with $\text{top}(a_i) = c_j$ and $c \succ_i c_k$ for every $c_k \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{c_j\}$.*

Finally, we end this section with an algorithm for producing reachable equilibrium profiles.

Theorem 3. *Algorithm 1 finds all Nash equilibria with truth-bias, with complexity of $\mathcal{O}(mn)$.*

Proof. We shall first show why the algorithm only outputs equilibrium profiles, and then that there are no equilibria that it misses.

Exploiting Lemma 5, every equilibrium the algorithm finds is made of a voter that can change its vote to c , making it the winner, and, if all other voters remain the same, has no incentive to deviate to a different candidate. Furthermore, no other voters will deviate in retaliation — if there are, it means they can deviate to a candidate c' which can win over c and which they prefer over c , and these deviations are found by lines 11 and 25.

Now, suppose there is an equilibrium resulting in candidate c winning. According to the previously proven lemmas, $c \in W(\mathbf{a}) \cup H(\mathbf{a})$, and there is only one voter that will deviate. Therefore, that voter must be found with the algorithm's line 6 or 20. Since that equilibrium will only be eliminated due to finding a voter in lines 11 or 25, and such a voter will indeed destroy an equilibrium (as it will have an incentive to deviate), the algorithm will find all equilibria.

Since there are two nested loops, each counting through a subset of candidates and voters, complexity is $\mathcal{O}(mn)$, □

5 Lazy voters

We now turn to consider the lazy model, under iterative voting. In defining the iterative process, we feel that it is more natural to the spirit of lazy voting to make the following assumption: if at some point in the sequence of best responses a voter decides to abstain, then he never comes back to the election¹. Given this restriction, we can establish convergence, albeit not necessarily to a NE of the game.

¹If we allow voters who abstain to return to the election and vote for some candidate later on in the process, we run into the problem of not having convergence. This follows from the fact that Nash equilibria do not always exist under the lazy model [4].

Algorithm 1 Finding all truth-biased Nash equilibria

```
1:  $\forall c \in W(\mathbf{a}) \cup H(\mathbf{a}), eq[c] \leftarrow 0$   $\triangleright$  An array holding number of equilibria for every
   potential winning candidate
2:  $V, V'$  sets of all voters
3: for all  $c \in W(\mathbf{a})$  do
4:   for all  $v \in V$  do
5:     if  $v \in V', top(v) \neq \mathcal{F}(\mathbf{a})$  and  $top(v) \neq c$  then
6:       if  $c$  is the highest ranked candidate in  $W(\mathbf{a}) \cup H(\mathbf{a})$  excluding  $top(v)$  then
7:          $eq[c] \leftarrow eq[c] + 1$   $\triangleright$  This voter will deviate to make  $c$  win
8:          $V' \leftarrow V' \setminus v$   $\triangleright$  If this voter deviates for  $c$ , it will not deviate for any other
9:       end if
10:    end if
11:    if there is  $c' \in W_{\succ_c}(\mathbf{a}), c' \neq top(v)$   $c' \succ_v c$  then
12:       $eq[c] \leftarrow 0$   $\triangleright$  This is a blocking voter
13:      break  $\triangleright$  No point in examining this candidate further
14:    end if
15:  end for
16: end for

17: for all  $c \in H(\mathbf{a})$  do
18:   for all  $v \in V$  do
19:     if  $v \in V', top(v) \neq \mathcal{F}(\mathbf{a})$  and  $top(v) \neq c$  then
20:       if  $c$  is the highest ranked candidate in  $W(\mathbf{a}) \cup H(\mathbf{a})$  excluding  $top(v)$  then
21:          $eq[c] \leftarrow eq[c] + 1$   $\triangleright$  This voter will deviate to make  $c$  win
22:          $V' \leftarrow V' \setminus v$   $\triangleright$  If this voter deviates for  $c$ , it will not deviate for any other
23:       end if
24:     end if
25:     if there is  $c' \in W(\mathbf{a}) \cup H_{\succ_c}(\mathbf{a}), c' \neq top(v)$  and  $c' \succ_v c$  then
26:        $eq[c] \leftarrow 0$   $\triangleright$  This is a blocking voter
27:       break  $\triangleright$  No point in examining this candidate further
28:     end if
29:   end for
30: end for
```

Definition 2. After t steps of the iterative process, we let A^t be the set of active voters, i.e., the set of voters who have not chosen to abstain. Obviously $A^t \subseteq A^{t-1}$, for every t . We define a profile \mathbf{b} to be a stable state at time t , if no voter from A^t has an incentive to change his current vote.

Lemma 6. Every sequence of improvement steps converges, and the final state is a stable state.

Proof. Clearly if we have convergence, then it has to be to a stable state. To prove convergence, [16] showed that in the basic model, every sequence of best-response steps, beginning from *any* profile, always converges to a Nash equilibrium. In our model, every time there is an abstention, we can restrict ourselves to the profile of the remaining active voters as a non-truthful profile, from which we know convergence is guaranteed, if no further abstentions are made. Since there are at most $n - 1$ abstentions, this process is finite. \square

Observation 1. Since we have guaranteed convergence to stable states, obviously not all stable states are Nash equilibria, since Nash equilibria do not always exist in the lazy model

[4]. Furthermore, as in the truth-biased model, there are Nash equilibria which are not reachable using a sequence of improvement steps.

The following example demonstrates Observation 1.

Example 1. Figure 1 shows a game where no improvement path can lead to a Nash equilibrium. Since the tie-breaking rule is $c_1 \succ c_2 \succ c_3$, we can see that the profile where voter 1 votes his true preference, and the other two voters abstain is the only Nash equilibrium. Yet, from the truthful profile of Figure 1, the only available improvement move is for voter 1 to deviate to \perp and leave the election. Afterwards, voter 2 or 3 will also leave the election and the process converges to electing c_2 . Hence, we cannot converge to a Nash equilibrium.

1	2	3
c_1	c_2	c_2
c_3	c_3	c_3
c_2	c_1	c_1

Figure 1: A game without convergence to a NE.

We now proceed to characterize the reachable stable states.

Lemma 7. *If a profile \mathbf{b} is a stable state reachable by a sequence of improvement steps, then it consists of the truthful vote of exactly 1 voter, and an abstention by all other voters.*

Proof. The state described is obviously stable, as the absent voters cannot return to the election, and the result is the voter’s favourite option. Hence the active voter also does not have an incentive to change his vote or abstain.

Suppose a different state is stable: If there is only one active voter, it will always vote for its favourite candidate (as voters pursue a best-response strategy). If there is more than one active voter, and they are all voting for the same candidate, one of those voters would gain by abstaining, hence it would not be a stable state. And finally, if there are active voters voting for a losing candidate, then again these voters would gain by abstaining. \square

Theorem 4. *Suppose a profile \mathbf{b} is a reachable NE under the lazy model. Then \mathbf{b} satisfies the properties of Lemma 7 and if c is the top choice of the active voter in \mathbf{b} , it is ranked higher in the tie-breaking rule than all other candidates who are ranked above c by any other voter.*

Proof. The first condition is trivially true. As for the second, suppose it did not hold. Then, given Lemma 7, there is a voter who has abstained and who prefers a candidate other than c . This voter would have an incentive to return to the election and vote his true preference, i.e., \mathbf{b} cannot be a Nash equilibrium. \square

Theorem 5. *There exists a polynomial algorithm that finds all reachable Nash equilibria from the truthful state with lazy voters.*

Algorithm 2 is given a truthful profile \mathbf{a} and a truthful voter v with preferences $z \succ \dots$ and returns “yes” if there exists a path of best-response stages to the Nash equilibrium culminating in this truthful voter voting for z (according to Lemma 7, all others abstain). The algorithm basically goes over all options of reaching a NE, as long as the requirements of Theorem 4 are satisfied, first checking whether \mathbf{a} , the initial state, is a Nash equilibrium in the usual sense (without abstentions) or not, and finding the players which would enable the voter v to have a shot at being the only participant left in the game.

achieves them, for "no" we explain why.

We begin with the case where the starting state (\mathbf{a}) is a NE in the basic model, i.e., without abstentions. Note that in this case, as in all NEs of the basic case, all voters' best-response (potentially, excepting a voter for the winner) is to abstain.

Line 4 indicates the stage sequence where all other voters abstain (as this is a NE, they do not deviate to a different candidate) and then all but v abstain as well, until it's the only one left.

Line 6 indicates the stage sequence where all voters not voting for $\mathcal{F}(a)$ except v and the noted \tilde{c} voter abstain, and then all but 2 voters for $\mathcal{F}(a)$ abstain as well. Then the \tilde{c} supporter deviates to support z , and all other voters except v will now abstain.

Line 9 indicates the stage sequence where all voters not supporting $\mathcal{F}(a)$ or c except v and the voter supporting c over $\mathcal{F}(a)$ abstain, then $\mathcal{F}(a)$ supporters abstain until $|\text{sc}(c, \mathbf{a})|$ of them remain (or $|\text{sc}(c, \mathbf{a})| + 1$, depending on tie breaking rule). Then the z voter preferring c deviates to make c the winner, and then all voters abstain except v , and 2 c supporting voters (which include a z supporting one which deviated). This voter now reverts to z , making it the winner, and the voters except v abstain.

Line 11 indicates the stage sequence where all voters not supporting $\mathcal{F}(a)$ or c except v and the noted \tilde{c} voter abstain, then $\mathcal{F}(a)$ supporters abstain until $|\text{sc}(c, \mathbf{a})|$ remain (or $|\text{sc}(c, \mathbf{a})| + 1$, depending on tie breaking rules). Then the \tilde{c} voter deviates to c , making it the winner, and then our voter deviates to return $\mathcal{F}(a)$ to its victorious position (at this point, we know our voter prefers $\mathcal{F}(a)$ to others). Then all voters except ours abstain, and then our voter deviates back to z .

If the algorithm return no, this means all voters find that anyone they support over $\mathcal{F}(a)$ (that isn't their first preference) has only one point and loses to $\mathcal{F}(a)$ according to tie-breaking rule. This means no voter has any move except abstention, and hence, z can never become the winner.

We now turn to the case where the starting position is not a Nash equilibrium in the basic sense.

Line 14 indicates the stage sequence where we can simply have regular run of iterative plurality (without abstentions), when not allowing v to participate. We call the resulting state \mathbf{c} . z cannot become a winner (as it wasn't even a runner-up). If v prefers $\mathcal{F}(c)$ over any runner-ups, then we look at the run of iterative plurality, and replace the last deviation of a voter to $\mathcal{F}(c)$ with a deviation of v . Now, all voters except v can abstain. If v prefers a runner-up b , we let all voters for candidates that are not $b, \mathcal{F}(c)$ or voter v to abstain, and then let v deviate. Since v has deviated to the winner, all except v can abstain.

Line 17 indicates the stage sequence where if there is such a voter which voted for z , it deviates to b , and we can now follow the same sequence as for line 14 in the paragraph above, as z is no longer a runner-up. If there isn't (all of them prefer $\mathcal{F}(a)$), the noted voter deviates, making b the winner, and now one of the z voters deviates to $\mathcal{F}(a)$, making z no longer a runner-up, and again, we can now follow the same sequence as for line 14 in the paragraph above.

In line 19 we revert to the NE part, as the only voter wishing to deviate is v — without v this is a Nash equilibrium. However, in the case of 3 voters we can answer directly (line 20) — since v is the runner-up, the 2 other voters have voted for the winner and will not deviate.

Line 24 indicates the stage sequence where after reaching \mathbf{b} , all voters except v abstain.

Line 26 indicates the stage sequence where we pursue, after reaching \mathbf{b} a similar strategy to line 14. If v prefers $\mathcal{F}(b)$ over any runner-ups, then we look at the run of iterative plurality, and replace the last deviation of a voter to $\mathcal{F}(b)$ with a deviation of v . Now, all voters except v can abstain. If v prefers a runner-up b , we let all voters for candidates that are not $b, \mathcal{F}(b)$ or voter v to abstain, and then let v deviate. Since v has deviated to the winner, all except v can abstain.

Finally, in line 27 we are at a situation where, at most, only v wants to deviate, as in line 19. □

6 Conclusions and Future Work

The study of voting schemes has been challenged by the issue of manipulability, though some manipulations have been shown to be NP-hard. As an alternative game-theoretic approach to evaluating aggregation, it is possible to consider Nash equilibria as a solution concept for preference aggregation over strategic agents. This leads to the need for characterising and computing the resulting set of equilibria.

In this paper, we investigated characteristics of the recently proposed scheme of iterative voting. Though it has been previously shown that iterative plurality voting converges to an equilibrium, the set of equilibria has not been described. Here we complete this description.

This allows us, in turn, to consider the range of computational problems associated with NE reachability, from finding an arbitrary NE without any limitations on its properties to the problem of determining the iterative reachability of a specific NE profile. While the former is trivially polynomial (one simply has to let the iteration converge, as has been shown in [16]), we show that deciding on the reachability of a specific NE profile is NP-hard.

A multitude of NE property limitations may be considered between the two aforementioned extremes. For instance, one may seek an NE with a specific winner while remaining within the basic iterative plurality voting model. However, we find that more interesting NE features come from modifying the basic model by introducing voter bias. These modifications of the iterative voting model, in fact, have not been studied before, and we address this gap in our paper. Specifically, we consider the two most popular voter biases: truth-biased voters and lazy voters. Besides modifying the set of possible NE outcomes, voter bias also affects the complexity of finding a representative of the set. In particular, we show that the complete set of all iterative truth-biased plurality voting equilibria is computable in polynomial time.

Similarly, in the case of lazy voters (the most real-world voter model we have that incorporates abstention), it is also possible to find, in polynomial time, all the stable states of the iterative process that are Nash equilibria. In fact, we characterise the complete set of all stable states of the iterative process, though deciding on the reachability of an arbitrary *a priori* chosen stable state by the iterative process remains an open question.

7 Acknowledgments

This research has been co-financed by the European Union (European Social Fund — ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF) — Research Funding Program: THALES; and by ESRC grant RES- 000-22-2731; and by Microsoft Research through its PhD Scholarship Programme, Israel Science Foundation grant #1227/12, the Israel Ministry of Science and Technology — Knowledge Center in Machine Learning and Artificial Intelligence grant #3-9243, the Google Inter-University Center for Electronic Markets and Auctions, and the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI).

References

- [1] S. Airiau and U. Endriss. Iterated majority voting. In *Algorithmic Decision Theory*, volume 5783 of *LNCS*, pages 38–49. Springer, 2009.
- [2] S. Branzei, I. Caragiannis, J. Morgenstern, and A. D. Procaccia. How bad is selfish voting? In *AAAI*, 2013.
- [3] S. Chopra, E. Pacuit, and R. Parikh. Knowledge-theoretic properties of strategic voting. In *Logics in Artificial Intelligence*, volume 3229 of *LNCS*, pages 18–30. Springer, 2004.
- [4] Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *ACM EC*, pages 347–356, 2010.
- [5] A. Dhillon and B. Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46(1):55–75, 2004.
- [6] B. Dutta and J.-F. Laslier. Costless honesty in voting. in 10th International Meeting of the Society for Social Choice and Welfare, Moscow, 2010.
- [7] R. Farquharson. *Theory of Voting*. Yale University Press, 1969.
- [8] T. Feddersen and W. Pesendorfer. Voting behavior and information aggregation in elections with private information. *Econometrica*, 65(5):1029–1058, 1997.
- [9] T. J. Feddersen, I. Sened, and S. G. Wright. Rational voting and candidate entry under plurality rule. *American Journal of Political Science*, 34(4):1005–1016, 1990.
- [10] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–602, 1973.
- [11] U. Grandi, A. Loreggia, F. Rossi, K. B. Venable, and T. Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In *3rd International Conference on Algorithmic Decision Theory, ADT-2013*, pages 181–192, 2013.
- [12] N. S. Kukushkin. Acyclicity of improvements in finite game forms. *International Journal of Game Theory*, 40(1):147–177, 2011.
- [13] J.-J. Laffont. Incentives and the allocation of public goods. In *Handbook of Public Economics*, volume 2, chapter 10, pages 537–569. Elsevier, 1987.
- [14] J.-F. Laslier and J. W. Weibull. A strategy-proof condorcet jury theorem. *Scandinavian Journal of Economics*, 2012.
- [15] O. Lev and J. S. Rosenschein. Convergence of iterative voting. In *AAMAS*, volume 2, pages 611–618, 2012.
- [16] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria of plurality voting. In *AAAI*, pages 823–828, 2010.
- [17] R. B. Myerson and R. J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1993.
- [18] S. Obraztsova, E. Markakis, and D. R. M. Thompson. Plurality voting with truth-biased agents. In *SAGT*, 2013.
- [19] A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *AAMAS*, 2012.

- [20] M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- [21] D. R. M. Thompson, O. Lev, K. Leyton-Brown, and J. S. Rosenschein. Empirical aspects of plurality election equilibria. In *AAMAS*, 2013.
- [22] L. Xia and V. Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *AAAI*, pages 805–810, 2010.

Zinovi Rabinovich
Mobileye Vision Technologies Ltd.
Jerusalem, Israel
Email: zr@zinovi.net

Svetlana Obraztsova
National Technical University of Athens
Athens, Greece
Email: svetlana.obraztsova@gmail.com

Omer Lev
Hebrew University of Jerusalem
Jerusalem, Israel
Email: omerl@cs.huji.ac.il

Evangelos Markakis
Athens University of Economics and Business
Athens, Greece
Email: markakis@gmail.com

Jeffrey S. Rosenschein
Hebrew University of Jerusalem
Jerusalem, Israel
Email: jeff@cs.huji.ac.il