

Sequential Voting with Confirmation Network

Yakov Babichenko

Oren Dean

Moshe Tennenholtz

yakovbab@tx.technion.ac.il

orendean@campus.technion.ac.il

moshet@ie.technion.ac.il

Faculty of Industrial Engineering and Management
Technion—Israel Institute of Technology
Haifa, ISRAEL

We discuss voting scenarios in which the set of voters (agents) and the set of alternatives are the same; that is, voters select a single representative from among themselves. Such a scenario happens, for instance, when a committee selects a chairperson, or when peer researchers select a prize winner. Our model assumes that each voter either renders worthy (confirms) or unworthy any other agent. We further assume that the prime goal of each agent is to be selected himself. Only if that is not feasible, will he try to get one of those that he confirms selected. In this paper, we investigate the open-sequential voting system in the above model. We consider both plurality (where each voter has one vote) and approval (where a voter may vote for any subset). Our results show that it is possible to find scenarios in which the selected agent is much less popular than the optimal (most popular) agent. We prove, however, that in the case of approval voting, the ratio between their popularity is always bounded from above by 2. In the case of plurality voting, we show that there are cases in which some of the equilibria give an unbounded ratio, but there always exists at least one equilibrium with ratio 2 at most.

1 Introduction

Consider a committee voting to select a chairperson. Each committee member would like the honour of serving as chairperson himself. As a second best option he prefers one of several other members to win the position.¹ The committee members' preferences profile can be represented by a *confirmation network*, in which a direct edge indicates that the source of the edge confirms the target. In the confirmation network of Figure 1 member #5 is the *most popular* member — he is supported by three other members, while everyone else is supported by at most one.

As is the case in many committees, we assume the members are well-known to each other, and hence the confirmation network is known to everyone. The ballot is open and conducted sequentially. Since the committee members are always sited in the same places around the table, the voting order is prefixed and known. We consider two voting methods: plurality, in which each voter selects only one other member, and approval, in which members vote for any subset of the other members. In either voting method, a member is not allowed to vote for himself, but he is allowed to abstain. The member with the most votes is elected. Ties are broken by a predetermined and publicly known preference order.

Game-theoretically, we have a multi-stage game, describable as an *extensive-form game* — a tree with all possible voting-sequences, and an outcome at every leaf. The standard solution for this kind of game is a *subgame perfect equilibrium* (SPE). To find an SPE, we start with the last voter. For every voting history, we assume this voter will choose a ballot which gives him a best outcome (notice that

*This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement *n*^o 740435).

¹An equivalent situation arises when the committee is about to select a course of action (e.g., tenure a researcher) and each committee member is strongly associated with one of the alternatives.




Figure 1: Five committee members and their confirmations.

there may be more than one ‘best outcome’). Moving to the next-to-last voter, we know, for every voting history, how the last voter will respond to any of his ballots. Thus, we can find all his best votes and fix one of them for any sequence of voting history. We can continue this backward process until we select a best vote for all voters. Since voters are indifferent between their best votes, in general there can be many SPEs.

We exemplify the model and its complexity in two scenarios. Example 1 shows a case in which the most popular member is not elected in the unique SPE of plurality voting. The same scenario with approval voting leads to two different SPEs — in one of them the most popular member is the winner. In the network of Example 2 each member confirms at most one other member. Later (Proposition 4) we will see that under this condition the outcome is always ‘almost-optimal’. Nevertheless, Example 2 shows that in the case of approval the outcome is not trivial: one of the members manages to get a better result by voting for someone he does not confirm.

Example 1. Assume that in the network of Figure 1 the voting order is lexicographic, and so is the tie-breaking order.² We will show that in this case, we have a unique SPE for plurality voting, and a different unique SPE for approval voting. *If the voting method is plurality*, and voter #1 votes for #5, then #5 will be elected (even if voter #2 decides to vote for #1, voter #3 cannot get elected himself and he will vote for #5 which is his second best outcome). However, voter #1 has a better vote. He may abstain, thus forcing voter #2 to vote for him (if voter #2 votes for #5 then voter #3 may abstain and get elected with the help of the vote from #4 and the tie-breaking). Since #3 and #5 will not get more than one vote, voter #1 now wins by tie-breaking. Thus by abstaining #1 gets elected, even though #5 is the most popular member having the most confirmations. Now, *if approval is the voting method* and #1 abstains, member #2 may vote just for #1 and get him elected as before. He may also vote for both #1 and #5. In this case #3 still cannot get elected (due to the tie-breaking order) and will vote for his second-best outcome, namely #5. We see that if voter #1 abstains, voter #2 has two ‘best ballots’ (to vote for #1 or to both #1 and #5) that lead to an outcome he confirms (#1 or #5). Later, when we formalize the model (Section 1.3), we add a ‘truth-bias’ assumption which states that each voter prefers the vote which is closest to his true confirmation set. Under this assumption member #2 favours the vote {#1, #5} over just {#1}. In this case, #1 does not gain from abstaining; thus, using the ‘truth-bias’ assumption once more, we get that #1’s best vote is to be truthful (i.e. vote for #5). Everyone else will be truthful as well, and #5 will be elected.

Example 2. Figure 2 shows a network with four voters and at most one confirmation (outgoing edge)

²I.e., in case of a draw, the voter with lowest index wins.

for each voter. The voting order and the tie-breaking order are both lexicographic. Since voters #1 and #2 cannot get elected (as no one confirms them), if they have a vote which gets the one they confirm elected (#4 for voter #1 and #3 for voter #2), then this will be their best vote. Otherwise, they will just be truthful (again we assume they are truth-bias) and vote for the one they confirm. *In plurality voting* voter #1 has no vote which gets #4 elected, hence both #1 and #2 are truthful and #3 is elected after abstaining. However, *in approval voting*, #1 can vote for both #4 and #2. Since #2 precedes #3 in the tie-breaking order, #3 cannot be elected and will now vote for #4. Thus, in approval voting, #4 is elected.




Figure 2: Four committee members and their confirmations.

1.1 Related work

Voting systems and their limitations have been long studied as part of the broader field of social choice (see for example the classic book of Farquharson [10] and the more recent handbook, [15]). The classical voting model assumes that the sets of voters and alternatives are disjoint, and that each voter has a totally ordered preference over the alternatives. Sequential voting with this model has been studied before and showed to contain counter-intuitive ‘paradoxes’. Desmedt and Elkind [8] considered both simultaneous and sequential plurality voting. They showed that a sequential voting system with at least three alternatives is prone to strategic voting, which might lead to an unexpected outcome, such as a Condorcet winner who does not win the election. Conitzer and Xia [20] further exemplified this phenomenon in a wide range of sequential voting systems, characterized by their domination index. In [4] the authors showed even more extreme examples of strategic ballots in sequential voting systems.

A confirmation network as an underlying model for simultaneous voting has also been studied. Holzman and Moulin [13] took an axiomatic approach to show the possibilities and limitations of such electoral systems. The main requirement of such systems, in their paper, is that no voter will be able to manipulate the system to select him by delivering a dishonest, strategic ballot. Alon et al. [1] investigated the same model, and showed the impossibility of incentive-compatible (that is, ‘non-manipulable’), deterministic voting systems. They suggested a probabilistic system with a bounded ratio between the maximal in-degree and the expected in-degree of the elected agent. Further works with the same theme can be found in [6, 12, 3].

1.2 Our contribution

In this paper, we discuss for the first time sequential voting with the underlying model of a confirmation network, for both plurality and approval voting. We show the limitations of these voting systems by

demonstrating extreme cases in which an unpopular member is elected in an SPE. On the other hand, we prove upper bounds on the ratio between the maximal in-degree in the network and the in-degree of an SPE outcome.

1.3 The model

Let A be a set of agents. General social choice settings assume that each agent has ordinal preferences over all possible outcomes (i.e., an elected agent in our case). In such general settings it is arguable how to measure the quality of the elected agent. In this paper we restrict attention to a simplified setting where the ordinal preference of each agent has only three levels: each agent prefers himself, those he confirms are second, and those he does not confirm are last. We model the preferences of the agents using a directed graph $G(A, E)$ where the interpretation of $(x, y) \in E$ is that agent x confirms agent y . In this setting we have a natural measure of the quality of the elected agent: the number of incoming edges. As we saw in Examples 1 and 2, and will see in the results, even in this simplified setting the strategic analysis is quite involved.

The agents cast their votes sequentially and openly. We consider two voting rules: *plurality*, where each agent is allowed to vote for at most one other agent, and *approval*, in which each agent may vote for any subset of the other agents (abstentions are allowed). The winner of the ballot is the one who receives the most votes, under a predetermined tie-breaking order. We assume A is ordered; this order is used both as the voting order and the tie-breaking order.³ The utility of agent x from the outcome y is

$$\mathcal{U}_x(y) = \begin{cases} 1, & y = x \\ 1/2, & (x, y) \in E \\ 0, & \text{otherwise.} \end{cases}$$

There is nothing particular about this function; any three-level function will do. Actually, \mathcal{U} will not be explicitly used in the remainder of the paper.

We are interested in voting strategies that form a subgame perfect equilibrium (SPE). At least one SPE is guaranteed to exist [18]; however, if in some subgame more than one ‘best vote’ option is available to an agent, multiple SPEs exist, possibly with different outcomes. Such a situation can occur, for instance, when an agent does not confirm anyone and is also not confirmed by any other agent (i.e., the agent is an ‘isolated’ node). If many agents are isolated, and so indifferent to the outcome, they may each make an arbitrary vote and anyone may be elected. To avoid this problem, we take the same approach as in [9] (see also [14, 19, 17]). Namely, when an agent faces several best-votes, he will select the one which best reflects his true confirmations. In order to quantify the truth-bias assumption, we add the following bonus utility. Let $f(x)$ be the number of agents that x confirms and actually votes for, and let $g(x)$ be the number of agents he does not confirm and nevertheless votes for. Then his bonus utility is

$$\mathcal{B}_x = \varepsilon^2 f(x) - \varepsilon g(x),$$

where $0 < \varepsilon < 1/2n$.⁴ When the outcome is y , the actual utility of agent x is given by $\mathcal{U}_x(y) + \mathcal{B}_x$.

³It is only for convenience that we assume that the voting order and tie-breaking order are the same. The proofs of the upper bounds on the ratio (Theorems 7 and 8) do not use this assumption. The negative examples can be altered to give the same outcome for different tie-breaking orders (this is not to say that the same example works for all tie-breaking orders).

⁴The intuition behind this function is that a voter would rather not report some or all of his confirmations than vote for someone he does not confirm. The upper bound on ε is chosen so that the utility from the outcome always dominates the bonus utility.

1.4 Definitions and notations

We will use the following notations from graph theory. For $a \in A$ let $d(a) = d_{in}(a) := \#\{b \in A : (b, a) \in E\}$ be the *popularity* of a . We denote by $\Delta_{in}(G) := \max_{a \in A} d(a)$ the maximum in-degree in G . Similarly, $\Delta_{out}(G)$ is the maximum out-degree. An agent m is *most popular* if $d(m) = \Delta_{in}(G)$.

An agent is *achievable* if there is an SPE in which he is elected. We denote by $W \subseteq A$ the set of all achievable agents. Our goal in this paper is to find bounds on the difference and ratio between the highest/lowest popularity of achievable agents and the maximal popularity in the graph. Formally, for an achievable agent $w \in W$, let G_{-E_w} be the graph we get from G after removing all the out-edges of w . Let

$$\mathcal{D}(w) = \Delta_{in}(G_{-E_w}) - d(w), \quad \mathcal{R}(w) = \frac{\Delta_{in}(G_{-E_w})}{d(w)}$$

be the additive gap and the multiplicative ratio, respectively, between the highest popularity and the popularity of w . We have two justifications for defining these measures on G_{-E_w} and not directly on G . The first is philosophical: we do not want w 's own confirmations to influence the way he is measured.⁵ The second is mathematical elegance. We pay a small price in the definitions in order to get clearer theorems. It is obvious, though, that $\Delta_{in}(G_{-E_w}) \geq \Delta_{in}(G) - 1$; thus it makes little difference, especially for large values of Δ_{in} .

With a slight abuse of notation, we define for any graph G with plurality/approval voting, $\mathcal{D}(G) = \min_{w \in W} \mathcal{D}(w)$;⁶ for either plurality/approval let $\mathcal{D} = \sup_G \mathcal{D}(G)$.⁷ We shall promptly see that \mathcal{D} is unbounded for both plurality and approval voting. In order to give a better description of the limitations of the two voting methods, and to differentiate between them, we would like to chart the *asymptotic* bounds of the multiplicative ratio when $\mathcal{D} \rightarrow \infty$.^{8,9} To that end, we define for a graph G ,

$$\underline{\mathcal{R}}(G) = \min_{w \in W} \mathcal{R}(w), \quad \overline{\mathcal{R}}(G) = \max_{w \in W} \mathcal{R}(w),$$

the maximal/minimal multiplicative ratio between the most popular agent and the winners. For any positive integer k , we denote by \mathcal{G}_k the family of graphs with $\mathcal{D}(G) \geq k$, and define^{10,11}

$$\underline{\mathcal{R}} = \limsup_{k \rightarrow \infty} \underline{\mathcal{R}}(G), \quad \overline{\mathcal{R}} = \limsup_{k \rightarrow \infty} \overline{\mathcal{R}}(G).$$

1.5 Main results and paper organization

In Section 2, we show a sharp transition of the additive gap. In networks where each agent confirms at most one other agent (i.e., the maximum out-degree is one) there is a unique outcome, and $\mathcal{D}(G)$ is always zero (Proposition 4). However, already for networks where agents confirm at most two other

⁵This argument relates to the notion of ‘incentive compatibility’ which is central in [1] and [13].

⁶Our definition for $\mathcal{D}(G)$ uses the best winner (and not the worst). This strengthens our results for the additive gap which are all negative.

⁷For ease of notation we did not add a subscript to distinguish between \mathcal{D} of plurality and \mathcal{D} of approval. The results are the same for both anyway.

⁸Note that if we define \mathcal{R} without this asymptotic then it will be predominated by small graphs with low values of \mathcal{D} .

⁹Alternatively, we can take the asymptotic with respect to $\Delta_{in} \rightarrow \infty$.

¹⁰Again we do not have different notations for plurality and approval. It will be clear from the context to which of the two we refer.

¹¹ $\underline{\mathcal{R}}$ and $\overline{\mathcal{R}}$ are analogue to the ‘price of stability’ and the ‘price of anarchy’, respectively (see [16] Section 17.1.3).

agents (i.e., $\Delta_{out} = 2$), \mathcal{D} is unbounded (Proposition 6). In Section 3, we prove bounds on $\underline{\mathcal{R}}$ and $\overline{\mathcal{R}}$. For approval voting we show that $1.5 \leq \underline{\mathcal{R}} \leq \overline{\mathcal{R}} \leq 2$ (Theorem 7); whereas, for plurality voting we show that $\underline{\mathcal{R}} \leq 2$ and $\overline{\mathcal{R}} = \infty$ (Theorem 8). These results indicate that in worst-case scenarios approval voting succeeds in selecting more popular agents than plurality voting. In Section 4 we sketch a generalization of our results to k -approval voting. We wrap up with a discussion and open problems in Section 5.

2 Bounds on the additive gap

We start our discussion with the special case of graphs with a maximum out-degree of one (i.e. each agent confirms at most one other agent). In this case, we show that both plurality and approval voting have a unique winner in any SPE, and the winner is a most-popular agent or almost most-popular agent.¹² In approval voting, a vote of an agent to the set of his confirmed agents is called *truthful*. In plurality voting, a vote of an agent to one of his confirmed agents (or abstention if he does not confirm anyone) is called *truthful*. Our proof is based on the following observation, which is a consequence of our truth-bias assumption.

Observation 3. *In every SPE with outcome w , any agent who does not confirm w is truthful in the SPE path.*¹³

The reason is that the election of w is one of the worst outcomes for any agent who does not confirm w ; hence, being truthful is the only best vote for such an agent.

Proposition 4. *For the class of graphs with a maximum out-degree of one, both plurality and approval voting have a unique achievable outcome,¹⁴ and for both plurality and approval $\mathcal{D} = 0$.*

Before we prove this proposition, let us exemplify it in the scenario of Example 2. Notice that every node in the network of that example has either one out-edge or none, so the proposition applies. Indeed, we showed there that both plurality and approval have a unique SPE. In addition, in approval voting agent #4 is elected, and he is the most popular. In plurality, agent #3 is elected; notice, that if we remove his out-edge to #4, then he becomes one of the most popular agents.

Proof. We start by showing that the outcome is unique using backward induction. Given a subgame (i.e., a history of votes), if the agent which is about to vote has a vote which gets him elected, then this will be the outcome. Moreover, if he cannot get elected but he can get the one he confirms elected, then this would be the outcome (here we use the assumption that he confirms at most one agent). If he cannot get elected and cannot get the one he confirms elected, by Observation 3 his unique best action is to be truthful, and by induction, the outcome is determined uniquely.

Now fix an SPE. Let w be the winner of this SPE and let $m \neq w$ be one of the most popular agents. We denote by

$$C_W := \{a \in A : (a, w) \in E\}; C_M := \{a \in A : (a, m) \in E\};$$

the set of agents which confirm w and those which confirm m , respectively. By our assumption on the out-degree, $C_M \cap C_W = \emptyset$. Thus, by Observation 3 the agents in $C_M \setminus \{w\}$ are truthful. So m gets the votes of all those who confirm him, except perhaps the vote of w . Again by Observation 3, no agent in $A \setminus C_W$

¹²Meaning, that if we ignore his own confirmations, the elected agent is most popular.

¹³In fact, this simple observation holds even for a wider solution concept of Nash equilibria.

¹⁴Though the outcome might be different between plurality and approval; see Example 2.

votes for w , which means that w cannot get more votes than his in-degree. Since w is elected, we reach the conclusion that

$$|C_W| \geq |C_M| - \mathbb{1}_{(w,m) \in E},$$

and that is exactly the same as $\mathcal{D} = 0$. □

The proof of Proposition 4 can be generalized to subgames in which the remaining voters confirm at most one agent. Suppose we are in the middle of a voting process with graph G . Let $U \subset A$ be the voters who have not yet voted and let G' be the graph we get from G after removing the out-edges of vertices in $A \setminus U$. Let $\bar{s} = (s_1, \dots, s_n)$ be the current *scoring vector*.¹⁵ We define the *potential* of a vertex $a \in A$ in this subgame to be $\rho(a) = d_{in}(a, G') + s(a)$, where $d_{in}(a, G')$ is the in-degree of a in G' .¹⁶ Let $\mathcal{P} = \max_{a \in A} \rho(a)$.

Proposition 5. *Using the definitions above, if $\Delta_{out}(G') \leq 1$ then there is a unique SPE for the remaining voting process; if w is the outcome of this SPE and m is any agent with $\rho(m) = \mathcal{P}$, then*

$$\mathcal{P} - \rho(w) \leq \mathbb{1}_{(w,m) \in E(G')}.$$

We omit the proof which is very similar to that of Proposition 4. Proposition 5 will be used in the proof of Proposition 6.

In contrast to Proposition 4, we will now show that even for graphs with a maximum out-degree of two, \mathcal{D} is no longer bounded. In the proof, we will show a voting scenario in which voters who confirm both a very popular agent and a much less popular one are forced to vote only for the less popular.

Proposition 6. *For the class of graphs with a maximum out-degree of two, \mathcal{D} is unbounded, for both plurality and approval voting.*

Proof. We will build a series of graphs, $\{G_k\}$, such that $\forall k \geq 2$, $\Delta_{out}(G_k) = 2$, and the unique achievable outcome, for both plurality and approval voting, has popularity $\Delta_{in}(G_k) - k$. Figure 3 depicts the graph G_k . The agents in B and D are classified by their types (the number of agents in each type is denoted below its circle). The order of the agents starts with the agents in D by lexicographic order of their type, then agents in C by reverse lexicographic order and finally the agents in B (i.e., the order is $d_1, \dots, d_{k-1}, c_{k+1}, \dots, c_1, b_1, \dots, b_k$).¹⁷

Notice that by Observation 3, the winner in any SPE must be from C ; otherwise, we will have a winner which got no votes, while c_1 gets at least the votes of the agents in b_1 . Suppose we are in the subgame which starts right after the ballot of the last agent in D . Since all the remaining voters have at most one out-edge, according to Proposition 5 the winner must be an agent which will have the highest potential if he abstains. Since the tie-breaking order is the same as the voting order, the winner will be *the first* agent which will have the highest potential if he abstains.

Now, the agents in D confirm both c_{k+1} and one other agent. The point will be that the only best vote for these agents is to vote only for c_{k+1} . Before proving the general case, we demonstrate this phenomenon in the simplest case, when $k = 2$ (Figure 4). Here, if the agents of type d_1 give c_1 at least one vote (e.g., if one votes for c_3 and the other for $\{c_1, c_3\}$), then c_3 cannot be elected (since after abstaining his total votes will be at most two, and c_1 's potential is at least three). Therefore c_3 is truthful

¹⁵That is, s_i is the number of votes agent i received from the voters in $A \setminus U$.

¹⁶In other words, $\rho(a)$ is the maximum number of votes a can hope to reach when the voting is done.

¹⁷As explained in Footnote 3, we use the same order for voting and for tie-breaking.





Figure 3: The graph G_k . Agent c_1 has popularity $\frac{1}{2}k(k+1) + k - 1$ while the winner, c_{k+1} , has popularity $\frac{1}{2}k(k+1) - 1$.

and c_2 abstains and wins (he will have two votes from B and one from c_3 ; agent c_1 may also have three votes, but c_2 precedes him in the tie-breaking order). This result is unfavourable for the agents of d_1 . However, if the agents d_1 vote only for c_3 , then c_3 can now abstain; having the same potential as c_1 and c_2 , agent c_3 wins by the tie-breaking.

Turning to the general case, assume first that all the agents in d_i , $1 \leq i \leq k-1$ vote for c_i , and perhaps also for c_{k+1} (in case of approval voting). Since after the ballots of the agents in D , the agents c_3, \dots, c_{k+1} all have a lower potential than c_1 , while c_2 can abstain and have the highest potential, c_2 is the winner. This outcome is unfavourable for the voters of d_1 . We claim that a better vote for them is to vote only for c_{k+1} , since that leads to the election of c_{k+1} .¹⁸ Indeed, if now all the voters in d_i , $2 \leq i \leq k-1$ vote for c_i , then now c_1, c_4, \dots, c_{k+1} cannot be elected since c_2 will have a higher potential than theirs. However, c_3 can abstain and win by tie-breaking. Thus, the agents of d_2 are now dissatisfied. If they now all vote just for c_{k+1} the same reasoning continues and shows that now c_4 will be the winner unless all the agents of d_3 vote just for c_{k+1} . Eventually, if all the voters of D vote just for c_{k+1} and c_{k+1} abstains, he will get elected. The agents of D are all satisfied with this outcome, which shows that this is an equilibrium. Indeed, our reasoning shows that this is the only equilibrium for both plurality and approval voting. The difference between the popularity of the winner, c_{k+1} , and the most popular, c_1 , is k , which implies the claim of the proposition. \square

¹⁸To be more precise: each voter in d_1 considers the situation in his turn. If all the voters before him voted only for c_{k+1} then he sees the opportunity to make c_{k+1} elected. Otherwise, he cannot get a good result and is just truthful.

Figure 4: The graph G_2 .

3 Bounds on the multiplicative ratio

In Section 2 we showed that in general, as much as we can tell from the additive gap measure, both plurality and approval voting systems perform poorly. Notice, however, that in the series of graphs in the proof of Proposition 6, the maximum in-degree in the graph G_k is $\Delta_{in}(G_k) = \Theta(k^2)$. This means that the ratio between the maximum in-degree and the popularity of the unique achievable outcome is $1 + \Theta(k^{-1})$. This raises the question whether a constant fraction of popularity can be guaranteed in sequential voting. We shall see in Theorems 7 and 8, that the bounds of the multiplicative ratio are non-trivial and are quite different between plurality and approval voting.

Theorem 7. *In plurality voting, $\underline{\mathcal{R}} \leq 2$ and $\overline{\mathcal{R}}$ is unbounded.*



Proof. We shall first prove that $\overline{\mathcal{R}}$ is unbounded. We show a series of graphs and SPEs, such that the ratio between the most popular agent and the winner goes to infinity. To this end, consider the graph in Figure 5. Suppose the agents' order is: $d_1, d_2, d_3, c_3, c_2, c_1, b_1$. We claim that the following profile of strategies is an SPE.¹⁹

- Agent d_1 : always vote for b_1 .
- Agent d_2 : always vote for c_2 .
- Agent d_3 : if d_1 voted for c_1 , then vote for c_2 . Otherwise, vote for c_3 .
- The rest of the agents: be truthful (abstain).

To see that all the agents always act rationally, we start from the last voters and proceed backwards to the first. Agents c_1, c_2, c_3, b_1 confirm no one, thus, abstaining is always a best vote for them. Agent d_3 always gets an agent which he confirms elected, so his votes are best possible as well. Moving on to agent d_2 , he will get c_2 elected when d_1 votes for c_1 and that is a best outcome for him. On the other hand, if d_1 votes for b_1 , then d_2 knows that d_3 is about to vote for c_3 , so the result will be bad for him no matter how he votes. The best thing he can do is to vote for someone he confirms (like c_2). Lastly, agent d_1 is indifferent between voting for b_1 and c_1 because anyhow the elected will be someone he does not confirm (c_3 in the former case and c_2 in the latter). Thus, assuming that he votes for b_1 is legitimate.

This proves the existence of a graph and an SPE with a multiplicative ratio 3. Figure 6 shows the general case. Here, there is an SPE in which d_1, \dots, d_{k-1} vote for b_1, \dots, b_{k-1} , respectively; d_k then votes for c_k , who is elected. If d_i decides to vote for any of c_1, \dots, c_i , then d_{i+1}, \dots, d_k all vote for c_{i+1} , hence d_i gains nothing. This is an SPE with a ratio of k , which shows that $\overline{\mathcal{R}}$ is unbounded.

¹⁹Note that we only want to show that it is an SPE; we do not claim uniqueness here.

Figure 5: There is an SPE in which c_3 is elected.Figure 6: There is an SPE in which c_k is elected.

In order to prove that $\underline{\mathcal{R}} \leq 2$, we need to show that there is always an SPE in which the winner's in-degree is at least half of Δ_{in} . Let G be any graph, and let m be one of the most popular agents. Assume that every agent who confirms m would vote for him whenever it is one of his best votes. Fix an SPE with this assumption, and let $w \neq m$ be the winner. Notice first, that by Observation 3 w cannot get more than $d(w)$ votes, since anyone who does not confirm him would not vote for him. Let $C_{m,w}$ be the set of agents who confirm both m and w , and let $C_{m,\bar{w}}$ be the set of agents who confirm m and do not confirm w . By Observation 3 and our assumption on the SPE, all the agents in $C_{m,\bar{w}} \setminus \{w\}$ vote for m , which means that $|C_{m,\bar{w}} \setminus \{w\}| \leq d(w)$. In addition, $|C_{m,w}| \leq d(w)$. Hence, we get that m 's popularity in G_{-E_w} is at most $d(m) = |C_{m,w}| + |C_{m,\bar{w}} \setminus \{w\}| \leq 2d(w)$, and the claim follows. \square

In the next theorem, we prove finite bounds for both $\underline{\mathcal{R}}$ and $\overline{\mathcal{R}}$ in approval voting.

Theorem 8. In approval voting, $\frac{3}{2} \leq \underline{\mathcal{R}} \leq \overline{\mathcal{R}} \leq 2$.

Proof. The proof of the upper bound on $\overline{\mathcal{R}}$ is similar to the proof of the upper bound on $\underline{\mathcal{R}}$ in Theorem 7. This time we do not need to make any assumption on the SPE — by Observation 3 any voter who confirms the most popular agent and does not confirm the winner is voting for the most popular agent. The claim now follows in a similar way.

To show the lower bound, we construct a series of graphs $\{H_k\}_{k \geq 2}$ where $\Delta_{in}(H_k) = \Theta(k)$ and $\mathcal{R}(H_k) = 3/2$. In the graph H_k the agent m is the most popular, and there are four sets of additional agents:

- The k agents in $C = \{c_1, \dots, c_k\}$ are the only agents, besides m , with a positive in-degree. They all have precisely k in-edges less than m , and all confirm only m . We will show that c_1 is the winner in the unique SPE.
- The $k-1$ agents in $D = \{d_1, \dots, d_{k-1}\}$ are those who confirm m but are forced not to vote for him. For any $1 \leq i \leq k-1$, agent d_i confirms m and all the agents $\{c_j : j \leq i\}$.
- The $k-1$ agents in $E = \{e_2, \dots, e_k\}$ provide the threat which prevents agents of D from voting for m . Agent e_i , $2 \leq i \leq k$, confirms all the agents $\{c_j : j \leq i\}$.
- Finally, the set B contains agents of k different types which serve as ‘equalizers’ which ensure that all nodes in C have a popularity of $d(m) - k = \Delta_{in}(G_k) - k$. For $2 \leq i \leq k$ there are $2i - 3$ agents of type b_i and they only confirm c_i . In addition, there are $k-1$ agents of type b_m who confirm m .




Figure 7: The graph H_k . Agent m has a popularity lead of k over all other agents, yet agent c_1 is elected.

The general graph is represented in Figure 7. The agents in C, D and E are ordered alternately in lexicographic order: $c_1, d_1, e_2, c_2, d_2, e_3, \dots, c_{k-1}, d_{k-1}, e_k, c_k$; the agents in $B \cup \{m\}$ are ordered after them in arbitrary order.

We will prove by induction on k that the winner in the unique SPE is agent c_1 . Notice that the popularity of m in $G_{-E_{c_1}}$ is $3(k-1)$ while the popularity of c_1 is $2(k-1)$, which implies the claimed ratio. Our induction base is $k = 2$ (Figure 8). Here is a sketch of the unique SPE in this scenario. The voting

starts with c_1 who abstains. If d_1 and e_2 are both truthful, then c_2 will be truthful as well (since c_1 beats him anyhow), and m will be elected. As e_2 does not confirm this result, he will vote only for c_2 , who can now abstain and get elected. Agent d_1 , foreseeing this possibility, must vote only for c_1 . Everyone after d_1 will now be truthful and c_1, c_2 and m all end up with two votes, leading to the election of c_1 by tie-breaking. For a general k it is enough to show that the following holds after c_1 abstains:




Figure 8: The graph H_2 . Agent c_1 gets elected after abstaining.

- (1.) If d_1, e_2 are both truthful, then the winner is m .
- (2.) If d_1 is truthful and e_2 votes for c_2 , then c_2 is the winner.
- (3.) If d_1 votes for c_1 and e_2 is truthful, then c_1 wins.

Notice that in any case, agents in $B \cup D \cup E$ are not achievable since they have no in-edges.

Proof of (1.) After truthful votes from d_1 and e_2 , c_1 and c_2 have equal potentials. Moreover, every agent who confirms c_2 , except for b_2 , also confirms c_1 . Since in the tie-breaking order c_1 and c_2 are adjacent, it is not hard to see that due to the truthful bias assumption, in any SPE, any agent except b_2 must either include both c_1, c_2 or neither. This means that in any SPE, c_1 will precede c_2 by tie-braking, hence c_2 is not achievable. It is enough, therefore, to show that if c_2 and d_2 are truthful, then m is the only achievable agent. Indeed, suppose that there is an SPE in which the outcome is c_i , $3 \leq i \leq k$.²⁰ By Observation 3, all the agents in $(C \setminus \{c_1, c_2, c_i\}) \cup b_m$ will be truthful, which means that m will end up with at least $2k - 1$ votes,²¹ while the potential of c_i is $2k - 2$, which is a contradiction.

Proof of (2.) Notice that by an argument similar to that of (1), here c_1 is not achievable because c_2 has higher potential. We claim that the subgame which starts with the vote of c_2 is equivalent to the

²⁰Clearly, an agent from $D \cup E \cup B$ cannot get elected, since m already has three votes.

²¹Which are the votes of d_1, d_2 and $C \cup b_m \setminus \{c_1, c_i\}$.

voting (which starts with c_1) in the graph H_{k-1} . More precisely, we define a mapping $\varphi : A(H_k) \setminus (B \cup \{c_1, d_1, e_1\}) \rightarrow A(H_{k-1}) \setminus B$, by $\varphi(c_j) = c_{j-1}$, $\varphi(d_j) = d_{j-1}$, $\varphi(e_j) = e_{j-1}$, $\varphi(m) = m$. Since the agents of B vote last and support only one agent each, they are non-strategic.²² We thus treat these agents as part of the potential of the agents in C and ignore them in the description of an SPE. Now fix some SPE of the subgame which starts with the voting of c_2 , and denote by $V(a) \subseteq A(H_k)$ the ballot of agent a in this SPE. We claim that the profile of voting in H_{k-1} in which each agent $\varphi(a)$ vote for $\varphi(V(a))$ is an SPE. Indeed, the only difference between the two voting scenarios is that the potential of every agent $c \in C(H_k) \setminus \{c_1\}$ is higher by exactly two then the potential of $\varphi(c)$ in H_{k-1} .²³ Since the graph structure, voting order and tie-breaking are all the same, it is not possible that $\varphi(V(a))$ is not a best vote of agent $\varphi(a)$ in H_{k-1} . By the induction hypothesis, we know that the only SPE of H_{k-1} leads to the election of c_1 . Hence the only achievable agent in our subgame of H_k is c_2 , as claimed.

Proof of (3.) Just like in (1), c_2 is not achievable, hence c_2 is truthful. Again we get a subgame which is equivalent to H_{k-1} , only this time we ignore c_2 and map $\varphi(c_1) = c_1$. The rest of the claims are identical to the proof of (2). We get that the only achievable agent in this subgame is c_1 .

□

4 Generalizing to k -approval

The two voting methods we discussed (namely, plurality and approval) can be generalized to a k -approval voting method in which every voter is allowed to vote for *at most* k other agents.²⁴ So plurality is no more than 1-approval, and approval is the same as $(n - 1)$ -approval. The two propositions of the additive gap (Propositions 4 and 6) had a single proof for both plurality and approval, and it is not hard to see that it can be generalized for any k -approval. Likewise, the bound $\mathcal{R} \leq 2$ can be proved for any k -approval in a similar way to the proof in Theorem 7 (and see also the same argument in Theorem 8). We will now extend the second part of Theorem 7 to any k -approval with $k = o(n)$ by showing a series of graphs in which \mathcal{R} is unbounded.




Figure 9: There is an SPE in which c_3 is elected.

²²That is, their only best vote is to be truthful.

²³Here the potential includes the “sure votes” of agents in B .

²⁴We allow a voter to vote for less than k other agents and even abstain. Though this is not the standard definition of k -approval, since we allowed abstentions in plurality and approval, this is the correct definition to get the full range between plurality and approval with abstentions.

In the graph in Figure 9 each agent d_i , $i = 1, 2, 3$ confirms the agents c_j , $1 \leq j \leq i$. In addition, agent d_1 confirms k additional agents, denoted b_1 , and d_2, d_3 each confirms $k - 1$ additional agents, denoted b_2, b_3 , respectively. The agents' order is $d_1, d_2, d_3, c_3, c_2, c_1$ and then the rest. We describe an SPE in which c_3 is elected. In this SPE, if d_1 votes for any subset which includes c_1 , then d_2 and d_3 will both vote for a subset which includes c_2 , and c_2 is elected. Since the outcome is one of the worst d_1 can get, voting only for the k agents of b_1 is a best vote for him. In this case, d_3 decides to vote for c_3 and the $k - 1$ agents of b_3 , and c_3 is elected no matter how d_2 votes. The multiplicative ratio in this SPE is 3. It is not hard to see how to extend this example to get any ratio.

5 Discussion and open problems

Additive gap vs. multiplicative ratio

We have seen (Proposition 4) that in the special case where every agent confirms at most one other agent, the elected agent will be most-popular or almost most-popular. However, as soon as the maximum out-degree of the graph is higher than one, this is no longer the case. In fact, we have shown (Proposition 6 and the discussion in Section 4) that \mathcal{D} is unbounded even when $\Delta_{in}(G) = 2$, for any k -approval voting. We, therefore, come to the conclusion that the additive gap is not a sufficient measure for the quality of a voting method in this model. Thus we turned to the multiplicative ratio for a finer evaluation.

Indeed the multiplicative ratio gave us different bounds for plurality and approval voting. In the case of plurality, we have seen (Theorem 7) that even though there might be an SPE with a ‘bad’ outcome ($\overline{\mathcal{R}}$ is unbounded), for every graph, we are guaranteed to have an SPE with a ratio of 2, at most. Moreover, our proof explains how to distinguish this SPE from other SPEs: you just give a small extra incentive for those who confirm the most popular to actually vote for him. The case of approval voting is clearer. Here (Theorem 8) we have proved finite bounds for both $\underline{\mathcal{R}}$ and $\overline{\mathcal{R}}$.

Plurality vs. approval

In [7], Brams demonstrated the superiority of approval voting over plurality voting in simultaneous voting systems. We conclude from our results, that in our setting, plurality voting (and even k -approval voting for any $k = o(n)$) allows SPEs with unbounded multiplicative ratio, while in approval, this ratio, in any SPE, will be between 1.5 and 2. We cannot draw from our results a comparison of the ‘best outcome’. To achieve that, we need to bound $\underline{\mathcal{R}}$ from below for plurality voting.

Question 9. *In plurality voting, is it possible to construct a series of graphs, $\{G_k\}$, with $\Delta_{in}(G_k) \rightarrow \infty$ and $\underline{\mathcal{R}}(G_k) \geq \alpha$ for some $\alpha > 1$ and all k ?*

Notice that in the series of graphs of Proposition 6, when $\mathcal{D}(G_k) = k$, we have $\Delta_{in}(G_k) = \Theta(k^2)$. So in this particular example, $\underline{\mathcal{R}}(G_k) = \overline{\mathcal{R}}(G_k) = 1 + O(1/k)$. If there is a non-trivial (i.e. different than 1) bound for $\underline{\mathcal{R}}$ in plurality voting, then there is a series of graphs, $\{G_k\}$ such that $\mathcal{D}(G_k) = k$ and $\Delta_{in}(G_k) = \Theta(k)$ (this is exactly what we have shown for approval voting, when we proved the lower bound in Theorem 8). So a rephrase of the above question would be:

Question 10. *In plurality voting, is it possible to construct a series of graphs, $\{G_k\}$, with $\Delta_{in}(G_k) = \Theta(k)$ and such that $\mathcal{D}(G_k) \rightarrow \infty$?*

It is worthwhile to note here that the example giving the lower bound of Theorem 8 does not work for plurality. To see that, consider the graph H_2 (Figure 8). If c_1 abstains, then even if d_1 votes for c_1 , e_1

might opt to vote for c_2 , and as a result, c_2 will be elected. So, for the case of plurality voting, the graph H_2 has an SPE in which m is elected, and the proof fails.

For the approval voting method, we have proved both a lower and upper bound on $\underline{\mathcal{R}}, \overline{\mathcal{R}}$. Still, it could be nice to further narrow these bounds or even find the exact asymptotic values of $\underline{\mathcal{R}}, \overline{\mathcal{R}}$.

Question 11. *Can the bounds of Theorem 8 be narrowed down?*

k -approval and a threshold between plurality and approval

Finally, we have seen that k -approval has the same bounds as plurality for any $k = o(n)$. When $k = n - 1$ this voting method is precisely approval; and so a natural question is what can be said about the threshold function which separates k -approval from plurality.

Question 12. *Find a minimal function, $f(n)$, such that the voting method $f(n)$ -approval has a finite bound for $\overline{\mathcal{R}}$.*

References

- [1] Noga Alon, Felix Fischer, Ariel Procaccia & Moshe Tennenholtz (2011): *Sum of Us: Strategyproof Selection from the Selectors*. In: *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, TARK XIII, ACM, New York, NY, USA, pp. 101–110, doi:10.1145/2000378.2000390.
- [2] Alon Altman & Moshe Tennenholtz (2008): *Axiomatic Foundations for Ranking Systems*. *J. Artif. Int. Res.* 31(1), pp. 473–495, doi:10.1613/jair.2306.
- [3] Haris Aziz, Omer Lev, Nicholas Mattei, Jeffrey S. Rosenschein & Toby Walsh (2016): *Strategyproof Peer Selection: Mechanisms, Analyses, and Experiments*. In: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, AAAI Press, pp. 390–396. Available at <http://dl.acm.org/citation.cfm?id=3015812.3015872>.
- [4] Yakov Babichenko, Oren Dean & Moshe Tennenholtz (2018): *Paradoxes in Sequential Voting*. Available at <http://arxiv.org/abs/1807.03979>.
- [5] Marco Battaglini (2005): *Sequential voting with abstention*. *Games and Economic Behavior* 51(2), pp. 445 – 463, doi:10.1016/j.geb.2004.06.007.
- [6] Antje Bjelde, Felix Fischer & Max Klimm (2017): *Impartial Selection and the Power of Up to Two Choices*. *ACM Trans. Econ. Comput.* 5(4), pp. 21:1–21:20, doi:10.1145/3107922.
- [7] Steven J. Brams (2008): *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures*. Princeton University Press, doi:10.1515/9781400835591.
- [8] Yvo Desmedt & Edith Elkind (2010): *Equilibria of Plurality Voting with Abstentions*. In: *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10, ACM, New York, NY, USA, pp. 347–356, doi:10.1145/1807342.1807398.
- [9] Bhaskar Dutta & Arunava Sen (2012): *Nash implementation with partially honest individuals*. *Games and Economic Behavior* 74(1), pp. 154 – 169, doi:10.1016/j.geb.2011.07.006.
- [10] Robin Farquharson (1969): *Theory of voting*. Blackwell. Available at <https://books.google.co.il/books?id=fIMHAQAAMAAJ>.
- [11] Dan S. Felsenthal & Moshe Machover (2014): *Electoral Systems: Paradoxes, Assumptions, and Procedures (Studies in Choice and Welfare)*. Springer, doi:10.1007/978-3-642-20441-8.
- [12] Felix Fischer & Max Klimm (2014): *Optimal Impartial Selection*. In: *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, ACM, New York, NY, USA, pp. 803–820, doi:10.1145/2600057.2602836.

- [13] Ron Holzman & Herv Moulin (2013): *Impartial Nominations for a Prize*. *Econometrica* 81(1), pp. 173–196, doi:10.3982/ECTA10523.
- [14] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein & Nicholas R. Jennings (2010): *Convergence to Equilibria in Plurality Voting*. In: *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, AAAI'10, AAAI Press, pp. 823–828. Available at <http://dl.acm.org/citation.cfm?id=2898607.2898740>.
- [15] Herv Moulin (2016): *Handbook of Computational Social Choice*. Cambridge University Press, doi:10.1017/CBO9781107446984.
- [16] Noam Nisan, Tim Roughgarden, Éva Tardos & Vijay V. Vazirani (2007): *Algorithmic Game Theory*. Cambridge University Press, doi:10.1017/CBO9780511800481.
- [17] Svetlana Obraztsova, Evangelos Markakis & David R. M. Thompson (2013): *Plurality Voting with Truth-Biased Agents*. In Berthold Vöcking, editor: *Algorithmic Game Theory*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 26–37, doi:10.1007/978-3-642-41392-6_3.
- [18] Martin J. Osborne (2003): *An Introduction to Game Theory*. Oxford University Press.
- [19] David R.M. Thompson, Omer Lev, Kevin Leyton-Brown & Jeffrey Rosenschein (2013): *Empirical Analysis of Plurality Election Equilibria*. In: *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, pp. 391–398. Available at <http://dl.acm.org/citation.cfm?id=2484920.2484983>.
- [20] Lirong Xia & Vincent Conitzer (2010): *Stackelberg Voting Games: Computational Aspects and Paradoxes*. In: *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, AAAI'10, AAAI Press, pp. 921–926. Available at <http://dl.acm.org/citation.cfm?id=2898607.2898755>.
- [21] Lirong Xia, Jérôme Lang & Mingsheng Ying (2007): *Sequential Voting Rules and Multiple Elections Paradoxes*. In: *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge*, TARK '07, ACM, New York, NY, USA, pp. 279–288, doi:10.1145/1324249.1324286.