How Bad is Selfish Voting?

Simina Brânzei
Aarhus University

Ioannis Caragiannis
University of Patras

Jamie Morgenstern
Carnegie Mellon University

Ariel D. Procaccia
Carnegie Mellon University

Abstract

It is well known that strategic behavior in elections is essentially unavoidable; we therefore ask: how bad can the rational outcome be? We answer this question via the notion of the *price of anarchy*, using the scores of alternatives as a proxy for their quality and bounding the ratio between the score of the optimal alternative and the score of the winning alternative in Nash equilibrium. Specifically, we are interested in Nash equilibria that are obtained via sequences of rational strategic moves. Focusing on three common voting rules — plurality, veto, and Borda — we provide very positive results for plurality and very negative results for Borda, and place veto in the middle of this spectrum.

1 Introduction

Voting rules are designed to aggregate individual preferences into a socially desirable decision. However, the idea that the outcome of an election reflects the collective will of the agents is debatable, as agents can sometimes secure a better outcome for themselves by misreporting their true preferences. In fact, the Gibbard-Satterthwaite Theorem (Gibbard 1973; Satterthwaite 1975) implies that such situations are unavoidable: every "reasonable" voting rule is susceptible to manipulation. This is sad news: try as we might to design clever voting rules that faithfully select a desirable alternative, we will always end up basing our decision on the wrong collection of preferences!

Once we accept that selfish voting is inevitable, it is natural to ask: how bad can the resulting outcome be? To answer this question, we need several ingredients. First, the answer depends on the voting rule. For the next few paragraphs, let us focus on the plurality rule: each agent casts a single vote for its favorite alternative, and the alternative with most votes wins the election. Second, we need a way of quantifying the quality of the outcome. We will suppose that the quality of an alternative is its truthful *score*; in the case of plurality, this is simply the number of agents that honestly view it as the best alternative (we discuss this choice in Section 6).

Classic game theory predicts that selfish agents would possibly vote dishonestly but in a specific way. Indeed, their collective votes should form a *Nash equilibrium (NE)*: no agent can achieve a better outcome by unilaterally deviating, given the votes of other agents. We can now address the foregoing question using the well-known notion of *price of anarchy* (see, e.g., the paper by Roughgarden and Tardos (2002) whose title inspired our own): the worst-case ratio between the quality of the best outcome, and the quality of the equilibrium outcome. In our current example, this is the ratio between the plurality score of the truthful plurality winner, and the (truthful) plurality score of the equilibrium winner.

It is immediately apparent that for plurality this price is very high. To see why, let x be an alternative that everyone hates, and consider a situation where everyone mysteriously decides to vote for x. This collection of votes — known as a *preference profile* — is a NE: if any single agent changes its vote, x would still be the winner of the election, hence a single agent cannot improve the outcome by deviating. The truthful plurality score of the equilibrium winner is zero, whereas the truthful winner could be the favorite alternative of everyone.

Fortunately, such an equilibrium can only arise when all the agents simultaneously go mad; we would not expect to actually observe it in real elections (regardless of whether the agents are humans or software agents). We therefore need to refine our question, by refining our game-theoretic solution. Following Meir et al. (2010), we consider *sequences* of best responses, where agents iteratively change their votes to greedily obtain their best possible outcome at each step. Interestingly, this refinement is insufficient in and of itself, because a best response sequence that starts at the profile where everyone votes for the hated alternative x would end there (there are no unilateral improvements to be made). However, it is natural to suppose that the starting point is the truthful preference profile, which agents iteratively modify as they jockey for their preferred outcomes.

With all the ingredients in place, we can now formulate our research question:

> *Fixing a voting rule, what is the price of anarchy when the set of equilibria is restricted to NE that are obtained as the end of a best response sequence starting at the truthful profile?*

For conciseness we refer to this flavor of the price of anarchy as the *dynamic price of anarchy (DPoA)*. To the best of our knowledge we are the first to study the price of anarchy in the context of voting.

Our Results

We investigate three common voting rules, which belong to the family of *positional scoring rules*. The input to positional scoring rules is a collection of rankings of the alternatives, one per agent, which represents the preferences of the agents. A positional scoring rule is defined by a vector $\vec{s} = (s_1, \ldots, s_m)$, where m is the number of alternatives; each agent awards s_k to the alternative it ranks in the k-th position, and the alternative with the highest total score wins the election. Plurality is simply the positional scoring rule defined by the vector $(1, 0, \ldots, 0)$; we also study the *Borda* rule, which is defined by the vector $(m-1, m-2, \ldots, 0)$, and the *veto* rule, which is induced by $(1, \ldots, 1, 0)$.

For plurality we give a strongly positive answer to our research question. It turns out that in *any* profile on a best response sequence (and in particular in equilibrium), the truthful plurality score of the current winner cannot be lower than the truthful plurality score of the truthful plurality winner by more than one point! It follows that the DPoA of plurality is $1 + o(1)$.

Turning to veto, we show that for the case of three alternatives its DPoA is $1 + o(1)$, but when there are at least four alternatives its DPoA is $\Omega(m)$; we conjecture that the latter bound is tight. Finally, we place Borda on the negative end of the spectrum with a DPoA of $\Omega(n)$, which grows linearly with the number of agents.

Related Work

Meir et al. (2010) were the first to investigate best response dynamics in voting. They focus on the plurality rule, and on the question of whether best response dynamics are guaranteed to converge to equilibrium from any initial preference profile (although they do give special attention to the truthful profile as starting point). Even in the most natural setting — starting from the truthful profile, with unweighted agents and deterministic tie breaking — convergence is not guaranteed. However, their main result is that best response dynamics will converge under so-called *restricted* best response dynamics, which *a priori* rule out one of three types of best response moves. Two recent papers (Lev and Rosenschein 2012; Reyhani and Wilson 2012) follow up on the work of Meir et al. (2010), and — independently from each other — investigate veto and Borda (as well as a few other rules in passing). Both papers show that veto necessarily converges to NE under restricted best response dynamics, albeit using different arguments. In addition, both papers show that (even restricted) best response dynamics under Borda may not converge. In contrast to these three papers, we do not rule out any potential best response moves, as we are not worried about convergence issues; rather, we are asking: *if* best response dynamics converge to NE, how bad can it be?

There are rather few papers that discuss the quality of equilibria that are obtained via best response dynamics. The two that are most closely related to our work (Chekuri et al. 2007; Charikar et al. 2008) study a multicast game played by nodes in a rooted undirected graph. They show that best response dynamics lead to an improved (only polylogarithmic) price of anarchy. However, they require specialized initial conditions: the initial state must be obtained through a process where players arrive one by one and greedily choose a minimum-cost path. A major conceptual advantage of the voting setting is that there is a natural choice for the initial state: the truthful preference profile.

2 Model

Let $N = \{1, \ldots, n\}$ be the set of *agents*, and let A be the set of *alternatives*, $|A| = m$. The preferences of agent i are represented by a ranking $\succ_i$ of the alternatives. A *preference profile* is a vector $\vec{\succ} = (\succ_1, \ldots, \succ_n)$ that gives the preferences of all agents. A *voting rule* is a function that receives a preference profile as input, and outputs the winning alternative.

A *positional scoring rule* is represented by a vector $\vec{s} = (s_1, \ldots, s_m)$; each agent awards s_k points to the alternative it ranks in the k-th position, and the alternative with most points wins the election. We focus on the three prominent positional scoring rules: *plurality*, which is represented by the vector $(1, 0, \ldots, 0)$; *veto*, which is represented by the vector $(1, \ldots, 1, 0)$; and Borda, which is represented by the vector $(m-1, m-2, \ldots, 0)$. Plurality and veto have intuitive interpretations: under plurality, each agent votes for a single alternative (its highest-ranked alternative if it is voting truthfully); and under veto, each agent vetoes a single alternative (its lowest-ranked if it is voting truthfully).

The score of alternative $a \in A$ under preference profile $\vec{\succ}$ and positional scoring rule f is denoted $\mathtt{sc}_f(a, \vec{\succ})$. We make the common assumption that ties are broken according to a fixed ordering of the alternatives.

Best Response Dynamics

We consider an iterative process where the initial state is the truthful preference profile, and at every step an agent changes the preference profile by changing its own reported vote. Agents are assumed to be myopic in that an agent would only change its vote if the outcome under the new profile is preferred to the previous outcome according to its true preferences, i.e., only *improvement moves* are made. Formally, if o_t is the outcome after the t-th move by agent i, then $o_t \succ_i o_{t-1}$, where here $\succ_i$ represents that true preferences of agent i. A *best response (BR)* is a move by an agent i that achieves the best currently achievable outcome according to its true ranking $\succ_i$.

For example, assume that the voting rule is plurality, and let the true preferences of agent 1 be $a \succ_1 b \succ_1 c \succ_1 d$. Further, suppose that after $t-1$ moves agent 1 is casting its vote for a (truthfully), and that currently all alternatives are tied in terms of their plurality score. Finally, suppose that the tie is broken for d, that is, $o_{t-1} = d$. Note that agent 1 cannot change the outcome to a, as it is already giving a its vote. However, it can achieve $o_t = b$ or $o_t = c$. Both options would yield an improvement step, but since $b \succ_1 c$, casting a vote for b would be the best response.

A *Nash equilibrium* is simply a state where no agent has an improvement step. We are interested in Nash equilibria that are reached through a sequence of best responses starting at the truthful preference profile.

It will be useful to distinguish between different types of best responses. For plurality, an improvement move $b \to a$ — where an agent changes its vote from b to a — can belong to one of the three types listed below, depending on how it affects the winner.

Type 1 $o_t = a$ and $o_{t-1} \neq b$.

Type 2 $o_t \neq a$ and $o_{t-1} = b$.

Type 3 $o_t = a$ and $o_{t-1} = b$.

For veto, an improvement move $-a \to -b$ — where an agent changes its veto from a to b — can belong in one of the types listed below, depending on how it affects the winner.

Type 1 $o_t \neq a$ and $o_{t-1} = b$.

Type 2 $o_t = a$ and $o_{t-1} \neq b$.

Type 3 $o_t = a$ and $o_{t-1} = b$.

For both plurality and veto, convergence to equilibrium is not guaranteed when type 2 moves are allowed, hence previous papers on best response dynamics in voting (Meir et al. 2010; Lev and Rosenschein 2012; Reyhani and Wilson 2012) ruled such moves out. We do not make this assumption, but some of their results are nevertheless helpful in establishing ours.

In particular, for both plurality and veto, define the set W_t of *potential winners* after the t-th move in a best-response sequence as the set including the winner o_t as well as every alternative $a \in A$ that would become a winner by a type 3 move $o_t \to a$ (for plurality) or $-a \to -o_t$ (for veto), regardless of whether there is an agent that can perform this move at the next step $t+1$ or not.

Lemma 1 (Reyhani and Wilson 2012, Lemma 17)**.** *Let t, t' be steps of a BR sequence consisting only of moves of types 1 and 3 under plurality, with $t < t'$. Then $W_{t'} \subseteq W_t$.*

The following lemma is the analog of Lemma 1 for veto; notice that the containment is reversed.

Lemma 2 (Reyhani and Wilson 2012, Lemma 12)**.** *Let t, t' be steps of a BR sequence consisting only of moves of types 1 and 3 under veto, with $t < t'$. Then $W_t \subseteq W_{t'}$.*

Dynamic Price of Anarchy

For a truthful profile $\vec{\succ}$, denote by $\mathrm{NE}(\vec{\succ})$ the set of all Nash equilibria that are reachable from $\vec{\succ}$ via BR dynamics.

Definition The *dynamic price of anarchy (DPoA)* of a positional scoring rule f is

$$\mathrm{DPoA}(f) = \max_{\vec{\succ}} \max_{\vec{\succ}' \in \mathrm{NE}(\vec{\succ})} \frac{\mathtt{sc}_f(f(\vec{\succ}), \vec{\succ})}{\mathtt{sc}_f(f(\vec{\succ}'), \vec{\succ})}.$$

In words, the dynamic price of anarchy of f is the worst-case ratio between the maximum truthful score under the truthful profile, and the truthful score of the winner under a NE that is obtained via BR dynamics from the truthful profile.

We will also consider an *additive* version of the dynamic price of anarchy, which computes the worst-case difference between the two scores rather than the ratio.

3 Plurality

Plurality is the most commonly used and well-known voting rule. It is therefore encouraging that in the context of plurality our answer to the question "how bad is selfish voting" is "not bad at all!".

Theorem 3. *The additive DPoA of plurality is* 1.

Note that this immediately implies that the multiplicative DPoA is also extremely small. Specifically, since the score of the plurality winner is at least n/m, the theorem immediately implies a multiplicative bound of

$$\frac{\frac{n}{m}}{\frac{n}{m} - 1} \leq 1 + \frac{2m}{n},$$

where the inequality holds when $n \geq 2m$ (which is almost always the case).

In order to prove the theorem, we will establish a rather surprising fact: a best response sequence for plurality never includes type 2 and type 3 moves when starting from the truthful profile. Since the sequence consists only of type 1 moves, this implies that, at every step t, the alternative whose score increases is the one that had score either one point below or equal to the score of the winner prior to the step. Hence, no alternative that initially had a score two points below the score of the winner can ever win. The theorem will therefore immediately follow from the next lemma.

Lemma 4. *BR dynamics under plurality cannot contain type 2 and type 3 moves when starting from the truthful profile.*

Proof. Assume for contradiction that a BR sequence has type 2 or 3 steps and consider the first such step t in the sequence where some agent i removes its point from the winning alternative a, gives its point to alternative b, and alternative c (possibly different than b) becomes the winner. Clearly $c \succ_i a$ and $c \in W_{t-1}$.

Since agent i prefers c to a, its vote before step t is not truthful (it was voting for a, which is not its more preferred alternative). Hence, there is a step $t' < t$ in the sequence in which agent i makes a type 1 move (recall that t was the first move of type 2 or 3) by removing its point from an alternative z and moving it to alternative a, making a the winner. Clearly, alternative a is not the winner before this move. Also, since this is a type 1 move, z is not the winner before the move either. Observe that z can be c or some other alternative, but cannot be a.

We will show that c does not belong to $W_{t'}$. By Lemma 1, this contradicts the assertion $c \in W_{t-1}$ above and the lemma will follow. Denote by d the winner before step t'. We distinguish between two cases:

Case I: $z = c$. In this case, d beats c before step t', the score of c decreases by 1 after step t', and the score of a increases so that it defeats d. Hence, after step t', alternative c cannot simultaneously beat a and d via a single move, i.e., $c \notin W_{t'}$.

Case II: $z \neq c$. Observe that after step t', c cannot defeat d by increasing its score by 1 point. Indeed, since neither the score of d nor the score of c changes during step t', if this was the case, then agent i could remove its point from alternative z, give it to c and make it a winner during step t'. This would contradict the fact that the move at step t' is a

best-response move (recall that $c \succ_i a$). We again conclude that after step t', alternative c cannot simultaneously beat a and d via a single move, and therefore $c \notin W_{t'}$. □

4 Veto

When the number of alternatives is small, veto exhibits similar behavior to Borda.

Theorem 5. *The additive DPoA of veto with three alternatives is* 1.

For the case of three alternatives, the truthful veto winner must have a truthful score of $2n/3$. Therefore, for this case the theorem immediately implies a multiplicative DPoA of

$$\frac{\frac{2n}{3}}{\frac{2n}{3}-1} \leq 1 + \frac{3}{n}.$$

In order to prove the theorem, we will first show that type 2 moves never happen in BR sequences starting from the truthful profile. Hence, for the case of three alternatives, BR sequences are actually restricted BR sequences like the ones studied in (Meir et al. 2010; Reyhani and Wilson 2012; Lev and Rosenschein 2012).

Lemma 6. *BR dynamics under veto with three alternatives do not contain type 2 moves when starting from the truthful profile.*

Proof. Consider a profile with alternatives a, b, and c and assume without loss of generality that, initially, the set W_0 of potential winners contains at least alternatives a and b and alternative c is not the winner. Clearly, if W_0 had only one alternative, the lemma trivially holds since the initial profile would be an equilibrium.

So, assume that agent i is the first one that makes a type 2 move at step $t+1$. We distinguish between three cases depending on whether the agent vetoes the alternatives of W_0 before and after the move.

Case I: The move is $-a \to -c$. This means that $o_{t+1} = a$ and $o_t = b$; hence $a \succ_i b$. Since a is not the least preferred alternative of agent i, the agent must have performed a move $-x \to -a$ at some previous step t'.

We distinguish among four cases depending on the value of x and the type of move at step t'. If the move was $-b \to -a$ and of type 1, this means that $o_{t'} \neq b$, which implies that b was not in the set $W_{t'-1}$ of potential winners before that move. Since alternative b was initially in the set W_0 of potential winners, this contradicts the assumption that the first type 2 move happened at step $t+1$ (using Lemma 2). If the move was $-b \to -a$ and of type 3, we would have $o_{t'+1} = b$ and $o_{t'} = a$ which implies $b \succ_i a$, a contradiction. If the move was $-c \to -a$ and of type 1, we would have $o_{t'+1} = b$ and $o_{t'} = a$ which again implies $b \succ_i a$, a contradiction. If the move was $-c \to -a$ and of type 3, we would have $o_{t'+1} = c$ and $o_{t'} = a$, i.e., $c \succ_i a \succ_i b$. But then, since c is not the least preferred alternative of agent i, this means that there was a move $-y \to -c$ of type 1 or type 3 at some previous step $t'' < t'$. In any of these cases, we would have that $o_{t''-1} = c$ which means that the agent moved even though its most preferred alternative was the winner. So, case I is not possible.

Case II: The move is $-c \to -b$. This means that $o_{t+1} = c$ and $o_t = a$. Hence, $c \succ_i a$. Again, there must be a previous move $-x \to -c$ at some step t'.

We again distinguish between cases. If the move was of type 1, we would have $o_{t'+1} \neq x$; this implies that alternative $x \in \{a, b\}$ was not in the set $W_{t'}$ of potential winners before the move and (by Lemma 2) contradicts our assumption that the first type 2 move happens at step $t+1$. If the move was $-a \to -c$ and of type 3, we would have $o_{t'} = a$ and $o_{t'-1} = c$, i.e., $a \succ_i c$, a contradiction. Finally, if the move was $-b \to -c$ and of type 3, we would have $o_{t'} = b$ and $o_{t'-1} = c$, i.e., $b \succ_i c \succ_i a$. But then, since c is not the least preferred alternative of agent i, there must have been a move $-y \to -c$ of type 1 or type 3 at some previous step t''. Again, we would have that $o_{t''-1} = c$ which means that the agent moved even though its most preferred alternative was the winner. So, case II is not possible either.

Case III: The move is $-a \to -b$. This means that alternative c was the winner at step t. This may have happened by a move $-a \to -b$ (or $-b \to -a$) of type 1 or a move $-c \to -a$ (or $-c \to -b$) of type 3 of some agent j.

If the move was $-a \to -b$ (the argument for the case $-b \to -a$ is symmetric) and of type 1, we would have $o_t \neq a$ which means that $a \notin W_{t-1}$. This contradicts our assumption that the first type 2 move happens at step $t+1$.

The only other possible subcase is when the move is $-c \to -a$ (the argument for $-c \to -b$ is symmetric) and of type 3. In this subcase we would have $o_t = c$ and $o_{t-1} = a$, i.e., $c \succ_i a$. Therefore, there must have been a move $-x \to -c$ at a previous step t'. We consider several subcases. If it was a type 1 move, we would have that $o_t \neq x$ which means that alternative $x \in \{a, b\}$ was not in the set $W_{t'-1}$ of potential winners before the move; this contradicts the assumption that the first type 2 move happens at step $t+1$. If the move was $-a \to -c$ and of type 3, we would have $o_{t'} = a$ and $o_{t'-1} = c$, i.e., $a \succ_i c$, contradicting the assertion above. If the move was $-b \to -c$ and of type 3, we would have $o_{t'} = b$ and $o_{t'-1} = c$ and, hence, $b \succ_i c \succ_i a$. But then, there must have been a move $-y \to -b$ of type 1 or type 3 at some previous step t''. As in previous cases, we would have that $o_{t''-1} = b$ which means that the agent moved even though its most preferred alternative was already a winner. We conclude that case III is impossible and the lemma follows. □

We are now ready to prove Theorem 5. Note that, unlike in Theorem 3 where the DPoA bound holds for every step of the BR sequence, here we prove it specifically for the NE that is reached.

Proof of Theorem 5. Consider a profile $\vec{\succ}$ with three alternatives a, b, and c in which a is the winner under veto with $\mathtt{sc}(a, \vec{\succ}) = T$. For the sake of contradiction, assume that alternative c has $\mathtt{sc}(c, \vec{\succ}) \leq T-2$ and becomes a winner after a BR sequence of t^* moves that lead to a NE profile $\vec{\succ}'$. First observe that $\mathtt{sc}(b, \vec{\succ}) \geq T-1$. Indeed, if this was not the case (i.e., if the score difference between alternative a and

both of the other two alternatives was at least 2 in profile $\vec{\succ}$), then the first moving agent would need to make a type 3 move, and in particular would make its least preferred alternative a winner. So, clearly, both a and b initially belong to the set W_0 of potential winners.

It holds that $\mathtt{sc}(c, \vec{\succ}) < 2n/3$ because $2n/3$ is the average veto score and initially the score of c is strictly smaller than other scores. Because alternative c is the winner after step t^*, its score must be at least the average, and hence $\mathtt{sc}(c, \vec{\succ}') \geq 2n/3 > \mathtt{sc}(c, \vec{\succ})$. In other words, there exists an agent i that hates c but vetoes some alternative $x \in \{a, b\}$ after step t^*. Since $x \in W_0$, by Lemmas 2 and 6, we have that $x \in W_{t^*}$ as well. Hence, a type 3 move $-x \to -c$ could make x a winner. Since $x \succ_i c$, this contradicts the assumption that a NE has been reached after step t^*. □

The situation drastically changes when the number of alternatives is at least four. The arguments that yield Lemma 6 break down, and indeed type 2 moves are no longer impossible. We leverage this insight to establish a lower bound on the DPoA under veto that is linear in m.

Theorem 7. *When the number of alternatives is $m \geq 4$, the DPoA of veto is $\Omega(m)$.*

Before giving the proof, let us ponder a bit why it is nontrivial. We want the BR sequence to converge to an equilibrium where an alternative that is truthfully vetoed by many is the outcome. In order for this to be an equilibrium, none of the many agents who veto the equilibrium outcome should be able to topple the newly elected winner, even though all of these agents would prefer *any* other outcome; herein lies the difficulty.

Proof of Theorem 7. Our lower bound instance has a set A' of $m-2$ alternatives a_0, a_1, ..., a_{m-3} as well as two additional alternatives b and c. The tie-breaking ordering over the alternatives is $c > b > a_0 > ... > a_{m-3}$. There are $n = 4m^2 - 2m - 3$ agents that are partitioned into several sets. In particular:

- There is an agent C with preference $a_1 \succ a_2 \succ ... \succ a_{m-3} \succ a_0 \succ c \succ b$.
- For $i = 0, 1, ..., m-3$, there is a set of agents V_i, each with preference $a_i \succ a_{i-1} \succ ... \succ a_{i+2} \succ a_{i+1} \succ b \succ c$. Sets V_0 and V_1 consist of $2m$ agents while the remaining sets V_i have $2m+1$ agents each.
- There is a set X of $2m^2 - 5m - 2$ agents with preference $b \succ a_0 \succ ... \succ a_{m-3} \succ c$.
- There is an agent Y with preference $a_0 \succ b \succ c \succ a_1 \succ ... \succ a_{m-3}$.
- There is an agent Z with preference $c \succ b \succ a_0 \succ ... \succ a_{m-3}$.

These are the agents that move at least once in the BR sequence that will be described below. There are additional agents: two agents that veto a_i for every $i = 0, 1, ..., m-4$, $2m+3$ agents that veto b, and $2m+3$ agents that veto c. So, the alternatives in A' are tied with score $n-2 = 4m^2-2m-5$. The alternatives b and c have scores $4m^2 - 4m - 7$ and $4m$ respectively. Alternative a_0 is the winner. We denote the score of b by L. Observe that the differences between these scores are significant, i.e., $4m \ll L \ll n-2$.

In order to establish the lower bound (of at least $m-1$), we will identify a BR sequence of agent moves that lead to a NE profile in which c is the winner. The BR sequence consists of three rounds.

The first round starts with a move $-b \to -a_0$ by agent C. In this way, alternative a_1 (the most preferred alternative of agent C) becomes a winner. Then, for $j = 2, 3, ..., (2m+1)(m-2)-1$, a distinct agent from set V_i makes the move $-c \to -a_{i-1}$ where $i = j \bmod (m-2)$. This changes the winner from a_{i-1} to a_i which is the most preferred alternative of the deviating agent in V_i. Observe that the alternatives a_{i-1}, ..., a_{m-3} are tied with the highest score while the score of the alternatives a_0, ..., a_{i-2} (if any) is one point below. The last move in this subsequence is performed by an agent in V_{m-3}. After this step, alternative a_{m-3} is the winner and its score is one point above the score of b and the alternatives in $A' \setminus \{a_{m-3}\}$. The round ends with a move $-c \to -a_{m-3}$ by an agent in set X which makes alternative b win. At this point, alternative b and the alternatives in A' are tied with score $L+1$ (since the score of alternative b increased by 1 and the score of each of the alternatives in A' decreased by $2m+1$ during the round). During the first round, the score of alternative c increased by $2m^2 - 3m - 3$.

The second round consists of $m-3$ subrounds. For $i = m-3, m-4, ..., 2$, the subround i consists of $2m+1$ pairs of moves:

- A distinct agent from V_i plays $-a_{i-1} \to -a_i$ and makes its second most preferred alternative a_{i-1} win with a score of $L+2$. Since the agent cannot make its most preferred alternative a_i win, this is a best-response move. Note that the score of alternative a_i decreases during this step.
- An agent from X plays $-c \to -a_{i-1}$ to decrease the score of a_{i-1} to $L+1$ and make (its most preferred) alternative b win again.

At the end of subround i, the score of alternative a_i has decreased significantly to $L - 2m$ (i.e., it decreased by one point for each agent in V_i) while the score of c has increased by $2m+1$ (i.e., by one point for each agent in X that moved during the subround). The last subround consists of $2m+1$ pairs of moves as well. All pairs of moves besides the fourth one are as in the previous subrounds (using $i = 1$). The fourth pair of moves consists of the following two moves:

- Agent C plays $-a_1 \to -c$ and makes alternative a_0 win with a score of $L+2$. Note that, even though this agent prefers the alternatives in $A' \setminus \{a_0\}$ to a_0, these alternatives have significantly smaller score (namely, $L - 2m$) and cannot become winners.
- A distinct agent from X plays $-c \to -a_0$ and makes alternative b win again.

Overall, the score of alternative a_1 decreases significantly to $L - 2m + 1$ (i.e., it decreased by one point for each agent in V_1) while the score of c increased by $2m$ (i.e., by one point for each agent in X that moved during the subround

besides the one that moved in pair with agent C). Note that each agent from V_0 and X plays exactly once during the first two rounds while the agents in sets V_1, ..., V_{m-3} play twice during these rounds.

At the beginning of the third round, alternative b is tied with alternative a_0 with a score of $L+1$. The score of c has increased to L due to the moves during the first two rounds that removed vetos from c and put them on alternatives in A'. The score of all alternatives in $A' \setminus \{a_0\}$ is at most $L - 2m + 1$ and they cannot become winners during the (short) third round. Also, agent C vetoes alternative c while the vetos alternative b has are from those agents different than C that initially vetoed b. The third round consists of the following four moves:

- Agent Y plays $-a_{m-3} \rightarrow -b$. Now, a_0 becomes the winner with score $L+1$ while alternatives b and c are tied with score L. Clearly, this is a best-response move since a_0 is the most preferred alternative of agent Y.
- Agent Z plays $-a_{m-3} \rightarrow -a_0$. Now, the three alternatives a_0, b, and c are tied with score L and c is the winner. Again, this is a best-response move since c is the most preferred alternative of agent Z.
- Agent Y plays $-b \rightarrow -a_0$ and makes its second most preferred alternative b win. This is a best-response move since the agent cannot make its most preferred alternative a_0 a winner in any way. Now, the scores for alternatives b, c, and a_0 are $L+1$, L, and $L-1$.
- Finally, agent C plays $-c \rightarrow -b$ and makes c win with a score of $L+1$. Again, this is a best response move for agent C since a_0 (which C prefers to c) cannot become a winner via a move of agent C.

The profile reached after the last move is a NE. Since alternative c has priority over the alternatives in A' and each of these alternatives has score at least two points below the score of c, no move can make an alternative in A' win. Alternative b could become a winner only by a move of type $-b \rightarrow -c$ but this is not an option either since the agents that currently veto b have b as their least preferred alternative. This completes the proof of the theorem. □

We believe that for four alternatives the above bound is tight.

Conjecture 8. *When the number of alternatives is $m \geq 4$, the DPoA of veto is $\mathcal{O}(m)$.*

To justify this conjecture, suppose that a is the equilibrium outcome at time t, and b is the winner at time $t-1$. Further, suppose that almost all agents truthfully veto b. At time t, none of the agents who truthfully veto a can veto b, because otherwise they would be able to reverse the last move, which made a the winner, by making a $-b \rightarrow -a$ move (contradicting the equilibrium assumption). It follows that the score of b is extremely close to n (i.e., it is vetoed by few agents), and hence the score of a is close to n. Intuitively, it should be impossible for a to reach that high a score with so few agents that do not despise it. However, formalizing this argument turns out to be very challenging.

5 Borda

While our results for veto provide a mix of good and bad news, Borda is even more problematic from the point of view of selfish voting, as formalized by our last result.

Theorem 9. *When the number of alternatives is $m \geq 4$ and the number of agents is n, the DPoA of Borda is $\Omega(n)$.*

The proof of the theorem is by far our most intricate, and is omitted due to lack of space. It is worth clarifying though the relation between n and m; what we specifically prove is that for every $m \geq 4$ and infinitely many positive values of n, there is a profile with n agents and m alternatives and an initially truthful BR sequence of moves that leads to an equilibrium with ratio at least $\Omega(n)$. Typically the number of agents is much larger than the number of alternatives, hence we view this result as extremely negative, much more so than Theorem 7.

6 Discussion

So how bad is selfish voting? Our results suggest that the answers are "very good" under plurality, "not bad" under veto, and "very bad" under Borda. However, a caveat is in order. Our additive-1 upper bound for plurality and lower bounds of $\Omega(m)$ for veto and $\Omega(n)$ for Borda show a clear separation between plurality and the two other rules. However, we have not formally established a separation between veto and Borda. In particular, the veto upper bound of Theorem 5 covers the case of $m = 3$, while the Borda lower bound of Theorem 9 only applies to $m \geq 4$. And while Theorem 7 seems much less intimidating than Theorem 9, the former lower bound does not have matching upper bound. Nevertheless, we strongly believe that a separation does exist between veto and Borda, and indeed Conjecture 8 — together with Theorem 9 — would imply such a separation.

Perhaps the most debatable aspect of our model is our use of the score of an alternative as a proxy for its quality. Note though that this approach is not uncommon, as several papers aim to select alternatives whose score approximates the optimal score, to circumvent computational complexity (Caragiannis et al. 2012; 2013), obtain strategyproofness (Procaccia 2010), reduce communication (Service and Adams 2012), or deal with missing information (Lu and Boutilier 2011). The assumption underlying these papers is that the quality of an alternative is directly related to its score, hence the goal is to optimize score under constraints. Recent work by Boutilier et al. (2012) further supports our approach, by showing that the score of an alternative can equal its expected utility in settings where agents have exact utilities for alternatives, but only report a ranking; this is true for example with respect to Borda score when the prior over utilities is uniform on an interval.

Even if one objects to quantifying the quality of alternatives via their score, our results have a natural interpretation. For example, for large values of n and m our lower bounds for veto and Borda mean that under these rules an alternative that is ranked last by almost all of the agents can become a winner in equilibrium. On the other extreme, our additive-1 upper bounds clearly imply that only alternatives that are almost as desirable as the optimal alternative can be elected.

References

Boutilier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2012. Optimal social choice functions: A utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, 197–214.

Caragiannis, I.; Covey, J. A.; Feldman, M.; Homan, C. M.; Kaklamanis, C.; Karanikolas, N.; Procaccia, A. D.; and Rosenschein, J. S. 2012. On the approximability of Dodgson and Young elections. *Artificial Intelligence* 187–188:31–51.

Caragiannis, I.; Kaklamanis, C.; Karanikolas, N.; and Procaccia, A. D. 2013. Socially desirable approximations for Dodgson's voting rule. *ACM Transactions on Algorithms*. Forthcoming.

Charikar, M.; Karloff, H.; Mathieu, C.; Naor, J.; and Saks, M. 2008. Online multicast with egalitarian cost sharing. In *Proceedings of the 20th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 70–76.

Chekuri, C.; Chuzhoy, J.; Lewin-Eytan, L.; Naor, J.; and Orda, A. 2007. Noncooperative multicast and facility location games. *IEEE Journal on Selected Areas in Communications* 25:1193–1206.

Gibbard, A. 1973. Manipulation of voting schemes. *Econometrica* 41:587–602.

Lev, O., and Rosenschein, J. S. 2012. Convergence of iterative voting. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 611–618.

Lu, T., and Boutilier, C. 2011. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 287–293.

Meir, R.; Polukarov, M.; Rosenschein, J. S.; and Jennings, N. R. 2010. Convergence to equilibria in plurality voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 823–828.

Procaccia, A. D. 2010. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 836–841.

Reyhani, R., and Wilson, M. C. 2012. Best reply dynamics for scoring rules. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, 672–677.

Roughgarden, T., and Tardos, E. 2002. How bad is selfish routing? *Journal of the ACM* 49(2):236–259.

Satterthwaite, M. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10:187–217.

Service, T. C., and Adams, J. A. 2012. Communication complexity of approximating voting rules. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 593–602.

A Proof of Theorem 9

In order to prove the theorem we will first construct a profile R with a set of $m-2$ alternatives A and two additional alternatives e and z. The agents' preferences in R are such that the alternatives in A are initially tied with the highest Borda score while the score of alternative z is significantly smaller. Then, we will identify a BR sequence Σ of agent moves that starts from R and proceeds in rounds. In each round besides the last one, the score of the alternatives in A decreases while the score of z increases. The last round guarantees that alternative z eventually becomes a winner in such a way that no agent has any incentive to deviate from their last preference. Our construction uses alternative e as a placeholder.

We will first describe the construction (profile and sequence of moves) by using parameters k, t, s, and f. Then, we will appropriately set the values of these parameters so that the construction indeed yields the theorem. The alternatives in set A are denoted by a_0, a_1, ..., a_{m-3}. Alternative a_{j+1} has priority over alternative a_j; all alternatives in A have priority over alternative e. The priority of z is between alternatives a_{m-4} and a_{m-5}. The profile has $n = k(m-2)$ agents who are indexed with integers from 0 to $k(m-2)-1$ and have the following preferences:

- For $j = 0, ..., m-5, m-2, ..., k(m-2)-1-t$ (i.e., not including the values $m-4$ and $m-3$), agent j's preference is $a_j \succ a_{(j+1) \bmod (m-2)} \succ \dots \succ a_{(j-1) \bmod (m-2)} \succ e \succ z$.
- agent $m-4$ has the preference $z \succ e \succ a_{m-4} \succ a_{m-3} \succ \dots \succ a_{m-5}$.
- agent $m-3$ has the preference $z \succ e \succ a_{m-3} \succ a_0 \succ \dots \succ a_{m-4}$.
- For $j = k(m-2)-t, ..., k(m-2)-1$, agent j's preference is $e \succ a_j \succ a_{(j+1) \bmod (m-2)} \succ \dots \succ a_{(j-1) \bmod (m-2)} \succ z$.

The agents are partitioned into k *rows*. Row i (for $i = 0, ..., k-1$) consists of the agents $i(m-2), i(m-2)+1, ..., (i+1)(m-2)-1$. Rows different than row 0 that do not contain any of the last t agents are called *clean* rows. So, the parameters k and t are used in the definition of R.

We now define a sequence of moves Σ that consists of $s+f+2$ rounds: an *initial* round, s *slow* rounds, f *fast* rounds, and the *last* round. The initial round involves agents of row 0 who perform the following moves:

- For $j = 0, ..., m-5$, agent j changes its preference from $a_j \succ a_{(j+1) \bmod (m-2)} \succ \dots \succ a_{(j-1) \bmod (m-2)} \succ e \succ z$ to $a_j \succ z \succ a_{(j+1) \bmod (m-2)} \succ \dots, \succ a_{(j-1) \bmod (m-2)} \succ e$.
- agent $m-4$ changes its preference from $z \succ e \succ a_{m-4} \succ a_{m-3} \succ \dots \succ a_{m-5}$ to $a_{m-4} \succ e \succ a_{m-3} \succ \dots \succ a_{m-5} \succ z$.
- agent $m-3$ changes its preference from $z \succ e \succ a_{m-3} \succ a_0 \succ \dots \succ a_{m-4}$ to $a_{m-3} \succ e \succ a_0 \succ \dots \succ a_{m-4} \succ z$.

A slow round involves agents of a clean row, say i, and consists of the following moves:

- For $j = i(m-2), ..., (i+1)(m-2)-1$, agent j changes its preference from $a_j \succ a_{(j+1) \bmod (m-2)} \succ \dots \succ a_{(j-1) \bmod (m-2)} \succ e \succ z$ to $a_j \succ z \succ a_{(j+1) \bmod (m-2)} \succ \dots, \succ a_{(j-1) \bmod (m-2)} \succ e$.

A fast round involves agents of a clean row i as well and consists of the following moves:

- For $j = i(m-2), ..., (i+1)(m-2)-1$, agent j changes its preference from $a_j \succ a_{(j+1) \bmod (m-2)} \succ ... \succ a_{(j-1) \bmod (m-2)} \succ e \succ z$ to $a_j \succ z \succ e \succ a_{(j+1) \bmod (m-2)} \succ ..., \succ a_{(j-1) \bmod (m-2)}$.

Finally, the last round involves agents $m-4$, $m-3$ and $c = (k-1)(m-2)$.

- agent $m-4$ changes its preference from $a_{m-4} \succ e \succ a_{m-3} \succ ... \succ a_{m-5} \succ z$ to $z \succ e \succ a_{m-3} \succ ... \succ a_{m-5} \succ a_{m-4}$.
- agent c changes its preference from $e \succ a_0 \succ a_1 \succ ... \succ a_{m-3} \succ z$ to $a_{m-3} \succ a_0 \succ a_1 \succ ... \succ a_{m-4} \succ e \succ z$.
- agent $m-3$ changes its preference from $a_{m-3} \succ e \succ a_0 \succ ... \succ a_{m-4} \succ z$ to $z \succ e \succ a_0 \succ ... \succ a_{m-4} \succ a_{m-3}$.

Now that we have described the profile R and the sequence of moves Σ, we will show that the following six Conditions are sufficient, namely Σ is indeed a best response sequence that reaches an equilibrium in which alternative z is the winner.

1. There are at least seven slow rounds.
2. The $s+f+1$ first rows do not contain the last t agents.
3. agent $c = (k-1)(m-2)$ is one of the last t agents.
4. Initially, $\mathtt{sc}(a_i) - \mathtt{sc}(e) \geq 9$.
5. At the end of round $s+f$, the score of the alternatives in A is exactly $\frac{1}{m-1}$ higher than the score of z.
6. At the end of round $s+f$, the score of z is not smaller than the score of e.

Conditions 1 and 2 guarantee that slow and fast rounds indeed involve clean rows and agent c will indeed behave as described, respectively.

We will also show that Conditions 1, 4, 5, and 6 guarantee that Σ is a BR sequence. First, we claim that alternatives e and z are not winning alternatives during the first $s+f+1$ rounds. First, consider the score difference between a_{m-4} and e. By Condition 4, it is initially at least 9 and does not decrease until the beginning of the fast rounds (this implies that it is at least 9 when agent $m-4$ moves during the initial round). Then, it decreases but stays at least $\frac{1}{m-1}$ until the end of the last fast round. Similarly, the score difference between alternatives a_{m-4} and z decreases by at least $m-2-\frac{1}{m-1}$ during each of the (at least seven) slow and fast rounds. By Condition 5, this implies that it is at least $\frac{1}{m-1} + 7(m-2-\frac{1}{m-1}) \geq 12$ at the end of the initial round. It also implies that the score of z is significantly smaller than the score of a_{m-4} (by at least 9) during the whole initial round (observe that the score difference between a_{m-4} and z can increase by at most 3 during the initial round).

So, throughout the first $s+f+1$ rounds, the winning alternative belongs in A. We claim that the alternatives in A are tied at the beginning of each round. This is clearly true for the initial round since by pushing all alternatives of A by two position upwards in agents $m-4$ and $m-3$ and by one position upwards in the last t agents, the alternatives in A are uniformly distributed in the first $m-2$ positions in the preferences of all agents; hence, their scores are equal. For $j \geq 1$, before the move of player $i(m-2)+j$ in round i, the alternatives a_{j-1}, a_{j-2}, ..., a_0 are tied with the highest score, with alternatives $a_{m-3}, ..., a_j$ having score $\frac{1}{m-1}$ (for the initial round or slow rounds) or $\frac{2}{m-1}$ below (for the fast rounds). So, agent $i(m-2)+j$ promotes alternative a_j to the set of the alternatives with the highest score; after this move, the alternatives $a_j, a_{j-1}, ..., a_0$ are tied with the highest score and the tie is broken in favor of a_j. This is clearly a best-response move for all agents besides $m-4$ and $m-3$ since the alternative that becomes a winner is actually their most preferred one. For agent $m-4$ (and $m-3$), it is a best-response as well since a_{m-4} (respectively, a_{m-3}) is its most preferred alternative in A and the scores of its mostly preferred alternatives e and z is significantly smaller than the winning alternative when they move.

Let us examine now the last round assuming Conditions 5 and 6, namely that at the beginning of the round the alternatives in A are tied with some score T, alternative z has score exactly $T - \frac{1}{m-1}$, and alternative e has score at most $T - \frac{1}{m-1}$. In the first move, agent $m-4$ (that prefers z the most) makes z win by increasing its score by 1 (to $T+1-\frac{1}{m-1}$) and decrease the score of a_{m-4} by 1 (to $T-1$). This is clearly a best-response move. Next, agent c makes a_{m-3} win by increasing its score by $1-\frac{1}{m-1}$ to $T+1-\frac{1}{m-1}$ (a_{m-3} is now tied with z but the tie is broken in its favor) and decreasing the score of e by $1-\frac{1}{m-1}$ to at most $T-1-\frac{1}{m-1}$. Even though a_{m-3} was second last in the preference of agent c, agent c could neither increase the score of any alternative in $A \setminus \{a_{m-3}\}$ by more than $1-\frac{2}{m-1}$ nor decrease the score of z since she was ranked last in its preference. Hence, agent c's move is indeed a best-response move. The last move by agent $m-3$ makes its top choice z a winner by increasing its score by 1 and decreasing the score of a_{m-3} by 1. Now, the scores of the alternatives z, $A \setminus \{a_{m-4}, a_{m-3}\}$, a_{m-4}, a_{m-3}, and e are $T+2-\frac{1}{m-1}$, T, $T-\frac{1}{m-1}$, $T-\frac{1}{m-1}$, and at most $T-1-\frac{1}{m-1}$. Observe that any agent that prefers some other alternative x to z can decrease their score difference by at most $2-\frac{1}{m-1}$ (none among these agents has z in its top preference). Even though the score difference of z from the alternatives in $A \setminus \{a_{m-4}, a_{m-3}\}$ is exactly $2-\frac{1}{m-1}$, z has priority over all these alternatives. The score difference of z from alternatives a_{m-4}, a_{m-3}, and e is at least 2. Hence, no agent has any incentive to deviate from the current profile; we have reached an equilibrium.

We will now translate the conditions above into inequalities on the parameters. We will introduce a new integer parameter γ and set $t = s + 2f - \gamma(m-2)$; this only restricts our construction and will lead to simple inequalities. First observe that Condition 2 can be expressed as $(s+f+1)(m-2)+t \leq k(m-2)$. Using the definition of γ and rearranging, we have

$$s(m-1) + fm \quad \leq \quad (k-1+\gamma)(m-2) \qquad (1)$$

Condition 3 is $t \geq m-2$, which becomes

$$s + 2f \quad \geq \quad (1+\gamma)(m-2) \tag{2}$$

In order to express the remaining conditions, we will first compute the scores of the alternatives in profile R. Clearly, the score of alternative z is 2. Recall that, by pushing all alternatives of A by two position upwards in agents $m-4$ and $m-3$ and by one position upwards in the last t agents, the alternatives in A are uniformly distributed in the first $m-2$ positions in the preferences of all agents. Hence, their initial score is

$$\mathtt{sc}(a_i) = k(m-2)\left(\frac{1}{2} + \frac{1}{m-1}\right) - \frac{t+4}{m-1}.$$

Alternative e is ranked second last in all agents besides the two agents $m-4$ and $m-3$ who rank it second and the last t agents who rank it first. Hence, its initial score is

$$\mathtt{sc}(e) = \frac{k(m-2)}{m-1} - \frac{t+4}{m-1} + t + 2.$$

Using the expressions for the scores of alternatives a_i and e, Condition 4 becomes

$$2s + 4f \quad \leq \quad k(m-2) + 2\gamma(m-2) - 22 \tag{3}$$

Observe that the difference between the score of the alternatives in A and alternative z decreases by $m-6-\frac{3}{m-1}$ in the initial round, by $m-2-\frac{1}{m-1}$ in each slow round, and by $m-1-\frac{3}{m-1}$ in each fast round. Hence, using the notation $\mathtt{df}(a_i, z) = \mathtt{sc}(a_i) - \mathtt{sc}(z)$, Condition 5 can be expressed as

$$\mathtt{df}(a_i, z) - \left(m-6-\tfrac{3}{m-1}\right) - s\left(m-2-\tfrac{1}{m-1}\right) - f\left(m-1-\tfrac{3}{m-1}\right) = \tfrac{1}{m-1}$$

which can be seen to be equivalent to

$$2s(m-1) + 2fm \quad = \quad k(m+1) - 2m + 6 + 2\gamma \tag{4}$$

Also, the score difference between alternatives z and e increases by $m-6$ in the initial round, by $m-2$ in each slow round, and by $2-\frac{2}{m-1}$ in each fast round. Hence, Condition 6 can be expressed as

$$\mathtt{df}(z, e) + (m-6) + s\,(m-2) + f\left(2 - \frac{2}{m-1}\right) \geq 0$$

which is equivalent to

$$s(m-2) \quad \geq \quad k - \gamma(m-2) - m + 5. \tag{5}$$

We are ready to give specific values for the parameters k, s, f and t (and γ) that satisfy inequalities (1) – (5) for every $m \geq 4$ as well as Condition 1 ($s \geq 7$). We distinguish among three cases depending on the value of m:

- If $m \geq 6$, for every even integer $\alpha \geq 4$, we set $k = \alpha m - \alpha - 2$, $s = t = \frac{\alpha m + \alpha - 4}{2}$, and $f = 0$ (using $\gamma = 0$).
- If $m = 5$, for $k = 11, 15, 19, ...$, we set $s = \frac{3k-1}{4}$, $f = 0$, and $t = \frac{3k-13}{4}$ (using $\gamma = 1$).
- If $m = 4$, for every even integer $k \geq 16$, we set $s = k-8$, $f = 6$, and $t = 2$ (using $\gamma = 1 + k/2$).

It can be easily verified (see also (3)) that $t + 4 \leq k(m-2)$ and, hence, $\mathtt{sc}(a_i) \geq k(m-2)/2 = n \cdot \mathtt{sc}(z)/4$. The proof is complete. □

B Other Interesting Claims

Theorem 10. *The dynamic price of anarchy of veto of restricted initially-truthful BR sequences with $m \geq 4$ alternatives is exactly $\frac{m}{m-1}$.*

Proof. Denote by $V_t(a)$ the number of vetos for alternative a after the t-th move. Consider a sequence of T (type 1 or type 3) moves that reaches an equilibrium. Let c be the winner at equilibrium.

We first show that $V_T(c) \geq V_0(c)$. Assume otherwise; then, there is an agent i that vetoes alternative c in the initial truthful profile but vetoes some other alternative, say a, after step T. Since the final profile is an equilibrium, agent i cannot make alternative a win by changing its veto to c (even though he prefers a to c). This implies that $a \notin W_T$ and, by Lemma 2, a never belonged in the set of potential winners during the whole sequence of moves. Still, agent i vetoed it during some move. This could have happened only by a type 2 move which is not possible.

So, since c is the final winner, we have $n/m \geq V_T(c) \geq V_0(c)$. Hence the initial score of c is at least $n - n/m$ and n is an obvious upper bound on the initial score of the winner initially. The upper bound follows.

The tight lower bound can be proved using the following construction; we omit the formal proof due to lack of space. We have m alternatives $a, b, c, d(1), ..., d(m-3)$ and n agents (where n is a multiple of $m(m-3)$). The priorities are $c > a > b > d(1) > ... > d(m-3)$. The truthful profile has:

- a set $A(j)$ of $\frac{n}{m(m-3)}$ agents with preference $c \succ b \succ a \succ D(j) \succ d(j)$ for $j = 1, ..., m-3$,
- a set $B(j)$ of $\frac{n}{m(m-3)}$ agents with preference $c \succ a \succ b \succ D(j) \succ d(j)$ for $j = 1, ..., m-3$,
- a set F of n/m agents with preference $D \succ a \succ b \succ c$, and
- a set $G(j)$ of n/m agents with preference $c \succ a \succ b \succ D(j) \succ d(j)$ for $j = 1, ..., m-3$.

The scores are n for alternatives a and b, $n - n/m$ for c, and $\frac{(m^2-4m+1)n}{m(m-3)}$ for alternatives $d(1), ..., d(m-3)$. Now, the following sequence is a best-response one:

- At round $i = 1, ..., n/m(m-3)$
  - For $j = 1, ..., m-3$:
    - the i-th agent of set $A(j)$ changes its preference from $c \succ b \succ a \succ D(j) \succ d(j)$ to $c \succ b \succ D \succ a$. This is best-response since b (which is preferred to a becomes a winner and c cannot be made a winner in any way.
    - the i-th agent of set $B(j)$ changes its preference from $c \succ a \succ b \succ D(j) \succ d(j)$ to $c \succ a \succ D \succ b$. Again, this is best-response move.

□

Theorem 11. *Initially-truthful BR sequences of veto with (at least) four alternatives may not even have equilibria.*

Proof. We construct an instance with three agents 1, 2, and 3 and four alternatives with priorities $a > b > c > d$. The preferences in the initial profile are $b \succ_1 a \succ_1 c \succ_1 d$, $c \succ_2 a \succ_2 b \succ_2 d$, and $a \succ_3 b \succ_3 c \succ_3 d$. So, initially the scores of the alternatives a, b, c, and d are 3–3–3–0 (a is the winner). Consider the following cycle of moves:

1. agent 1 moves $-d \to -a$. Now the scores are 2–3–3–1 and b is the winner. This is obviously a best-response move since agent 1 prefers b the most.
2. agent 2 moves $-d \to -b$. The scores become 2–2–3–2 and c is the winner. Again, this is obviously a best-response move since agent 2 prefers c the most.
3. agent 1 moves $-a \to -b$. The scores become 3–1–3–2 and a is the winner. This is a best-response move since the agent cannot make its most preferred alternative b beat a and c simultaneously.
4. agent 2 moves $-b \to -a$. The scores become 2–2–3–2 and c is the winner. Again, obviously a best-response move.
5. agent 1 moves $-b \to -d$. The scores become 2–3–3–1 and b is the winner. Again, obviously a best-response move.
6. agent 2 moves $-a \to -d$. We have returned to the initial profile. This is a best-response move since agent 2 cannot make c win simultaneously a and b.

Let us now consider the incentives of agent 3 in each step. Clearly, he is happy with the winner a at the profiles before moves 1 and 4. At the profiles before moves 2 and 6, he cannot make its most preferred alternative a simultaneously beat b and c. Finally, at the profiles before moves 3 and 5, he cannot make a (or b) simultaneously beat c and d. So, the cycle above is the whole best-response dynamics that can be reached from the truthful profile. □