

Computing the Proportional Veto Core

Egor Ianovski¹ and Aleksei Y. Kondratev^{1,2}

¹ HSE University, International Laboratory of Game Theory and Decision Making, Russia

² Institute for Regional Economic Studies RAS, St Petersburg, Russia
george.ianovski@gmail.com, akondratev@hse.ru

Abstract

In social choice there often arises a conflict between the majority principle (the search for a candidate that is as good as possible for as many voters as possible), and the protection of minority rights (choosing a candidate that is not overly bad for particular individuals or groups). In a context where the latter is our main concern, veto-based rules – giving individuals or groups the ability to strike off certain candidates from the list – are a natural and effective way of ensuring that no minority is left with an outcome they find untenable. However, such rules often fail to be anonymous, or impose specific restrictions on the number of voters and candidates. These issues can be addressed by considering the proportional veto core – the solution to a cooperative game where every coalition is given the power to veto a number of candidates proportional to its size. However, the naïve algorithm for the veto core is exponential, and the only known rules for selecting from the veto core, with an arbitrary number of voters, violate either anonymity or neutrality. In this paper we present a polynomial time algorithm for computing the veto core and present a neutral and anonymous algorithm for selecting a candidate from it. We also show that a pessimist can manipulate the veto core in polynomial time.

Introduction

Suppose we have a society of 100 individuals who express the following preferences over alternatives a , b , and c :

- 60 individuals report $a \succ b \succ c$.
- 40 individuals report $b \succ c \succ a$.

Which alternative should the society adopt? It is reasonable to exclude c from the running, as it is Pareto dominated by b , but both a and b have good arguments to support them, and it is difficult to choose one over another without knowing what this society is and what they are electing. In a political context a lot can be said in favour of a . As a president, b 's initiatives would be constantly challenged by 60% of the public, and in a parliamentary system prime minister b would find it very difficult to achieve anything at all. On the other hand, if a , b , and c are meeting times, then presumably most people would be able to attend b , while with a the hall will be half-empty. As government budgets, a will write off the interests

of nearly half the population. As an arbitrated settlement to an armed dispute, a is unlikely to stop the violence.

Should we decide that in our context the minority matters, we need to study the profile from the bottom up. Rather than asking how many individuals would love to see a elected, the real issue is to how many such an outcome would be unacceptable. A line of research starting with Mueller (1978) considered voting rules that endowed groups or individuals with the power to unilaterally block certain outcomes from being elected, regardless of how desirable these outcomes may be to the rest of the population.

Mueller (1978) introduced the procedure of *voting by veto* to select a public good, with the explicit aim of protecting an individual from unfair treatment. The alternatives consist of one proposal from each voter, as well as the status quo. An order is formed over the voters and each, in turn, strikes off the alternative they like the least. Exactly one outcome is selected by this procedure, and clearly no voter will see their worst outcome elected. But this procedure is not anonymous, and requires that the number of alternatives be one greater than the voters.

Much of the subsequent work focused on the strategic behaviour of voters under this procedure. The sincere outcome of voting by veto and similar rules has been shown to correspond to a strong equilibrium solution (Peleg 1978; Moulin 1982), dominance solution (Moulin 1980), and maxmin behaviour (Moulin 1981b); truth-telling may not be an equilibrium, but among equally sophisticated voters the same alternative can be expected to be elected. However, these results depended on certain relationships between the number of candidates and voters, and these procedures all violated either anonymity or neutrality.

Moulin (1981a) extended the concept of voting by veto from individuals to coalitions, and studied the core of the resulting cooperative game. The veto core is anonymous, neutral, and is well defined for any number of alternatives and voters. Moulin showed that endowing each coalition with the power to veto a number of alternatives proportional to the coalition's size guarantees that the veto core is non-empty, and is the smallest possible core out of all such rules. Moulin also proposed *voting by veto tokens* for selecting a specific outcome from the veto core, however while the veto core as a concept is anonymous, voting by veto tokens returns to the problem of different voting orders leading to different out-

comes.

In the case of two voters, the veto core corresponds to the Pareto-optimal candidates ranked in the top half of both voters' ballots. It turns out that nearly every rule used in arbitrator selection selects from the veto core (Anbarci 1993; Barberà and Coelho 2017, 2020; Bloom and Cavanagh 1986; de Clippel, Eliaz, and Knight 2014; Erlich, Hazon, and Kraus 2018). Many of these rules are not anonymous, and it is not clear whether they can be extended to an arbitrary number of voters in a reasonable way.

In recent years, voting by veto has been rediscovered in a computational setting. Bouveret et al. (2017) introduce the Sequential Elimination rule, which is a generalisation of voting by veto to an arbitrary number of candidates. However their motivation is markedly different from Mueller (1978) – they emphasise the low communication complexity of the rule, rather than the guarantees it gives to minorities. Indeed, the main focus of the paper is finding an elimination order that best approximates the Borda winner – and the Borda rule selects the alternative that maximises the average rank, regardless of how many individuals end up with their worst outcome. Kondratev and Nesterov (2020) return to the original motivation, and study the veto core in the context of balancing the rights of minorities and majorities. Their main criticism of the rule is that it is unclear how the veto core is to be computed – the naïve algorithm is exponential.

Our Contribution

In this work we study the veto core from a computational perspective. Our main result is that a polynomial-time algorithm for the veto core exists. We also propose an anonymous and neutral rule for selecting from the veto core, and demonstrate that in the framework of Duggan-Schwartz (Duggan and Schwartz 1992; Taylor 2002) a pessimist can manipulate the veto core in polynomial time.

Preliminaries

We operate in the standard voting model. We have a set \mathcal{V} of n voters, a set \mathcal{C} of m candidates, and every voter is associated with a linear order over the candidates, which we term the voters' preferences, or ballots. We use \succ_i and P_i to denote the preference order of voter i . An n -tuple of preferences is called a profile. The notation P_{-i} denotes the preferences of all voters other than i . Thus (P_i, P_{-i}) is the profile P , and (P'_i, P_{-i}) is obtained from P by replacing voter i 's preference order with P'_i .

A voting rule is a function that maps a profile to a set of candidates: the tied election winners. A voting rule is *anonymous* if the names of the voters do not matter – for any permutation $\pi : \mathcal{V} \rightarrow \mathcal{V}$, we have $f(P) = f(\pi P)$. It is *neutral* if the names of the candidates do not matter – for any permutation $\pi : \mathcal{C} \rightarrow \mathcal{C}$, it holds that $f(\pi P) = \pi f(P)$.

Definition 1 (Moulin 1981a). A *veto function* is a mapping $v : 2^{\mathcal{V}} \rightarrow \mathbb{N}$. Intuitively, $v(T)$ is the number of candidates a coalition of voters T can veto. We call $v(T)$ the *veto power* of T .

A veto function is *anonymous* if $v(T)$ is a function of $|T|$ alone, i.e. only the size, not the composition, of T matters.

The *proportional veto function* is the anonymous veto function given by:

$$v(T) = \left\lceil m \frac{|T|}{n} \right\rceil - 1.$$

We say that a candidate c is *blocked* by a coalition T just if there exists a *blocking set* of candidates, B , such that:

$$\forall b \in B, \forall i \in T : b \succ_i c, \quad (1)$$

$$m - |B| \leq v(T) \quad (2)$$

Intuitively, condition (1) means that every voter in T considers every candidate in B to be better than c , and condition (2) means that the coalition T can guarantee that the winner will be among B by vetoing all the other candidates.

The set of all candidates that are not blocked with the proportional veto function is called the *veto core*. ■

For intuition as to why the veto function looks the way it does (rather than $\lfloor \frac{m|T|}{n} \rfloor$ or $\lfloor \frac{m|T|}{n} \rfloor - 1$) consider the case of the grand coalition and the singleton. We want the grand coalition to be able to veto $m - 1$ candidates, so a winner is left over, while $\lfloor \frac{mn}{n} \rfloor = m$; and a singleton, in the case of $m = n + 1$, should be able to veto at least one candidate, whereas $\lfloor \frac{n+1}{n} \rfloor - 1 = 0$. More formally, Moulin (1981a) has shown that proportional veto power occupies a distinguished position among the possible anonymous veto functions: any veto function that gives less veto power than proportional will have a larger core (in terms of set-inclusion), and any veto function that gives more power will sometimes have an empty core. The veto core is thus the smallest solution that is guaranteed to exist.

Theorem 2 (Moulin 1981a). *The veto core is non-empty.*

Note that the veto core is thus a function that maps a profile of preferences to a non-empty subset of candidates – in other words, it is a voting rule. The veto core is clearly anonymous and neutral. Other desirable properties it satisfies include Pareto efficiency, strongly monotonicity, homogeneity, and it never selects the candidates which are ranked last by more than $\frac{n}{m}$ voters. However, we shall see later that in general this rule is not single-valued.

Example 3. Consider a profile with five candidates and four voters with the following preferences:

- $e \succ_1 b \succ_1 c \succ_1 d \succ_1 a$.
- $b \succ_2 e \succ_2 c \succ_2 d \succ_2 a$.
- $d \succ_3 b \succ_3 e \succ_3 c \succ_3 a$.
- $a \succ_4 c \succ_4 d \succ_4 e \succ_4 b$.

In the case of $m = n + 1$ the veto power of a coalition of size k simplifies to $\lceil (n + 1) \frac{k}{n} \rceil - 1 = k$. In other words, k voters can veto exactly k candidates.

Candidate a is blocked by the singleton coalition $\{1\}$ (among others), with $B = \{b, c, d, e\}$; b is blocked by $\{4\}$, with $B = \{a, c, d, e\}$; d is not blocked by any singletons, but is blocked by $\{1, 2\}$, with $B = \{b, c, e\}$; c is blocked by the coalition $\{1, 2, 3\}$ with $B = \{b, e\}$. Thus the unique candidate in the veto core is e .

Now let us add a fifth voter:

- $c \succ_5 a \succ_5 d \succ_5 b \succ_5 e$.

With $m = n$, the veto power simplifies to $\lceil \frac{k}{n} \rceil - 1 = k - 1$. A singleton coalition can no longer block anything, but the coalition $\{1, 2\}$ can block a with $B = \{b, c, d, e\}$. However, no other candidate is blocked, so the veto core is $\{b, c, d, e\}$. ■

Computing the Veto Core

The veto core is a solution to a cooperative game, and a naïve approach would have us enumerate all possible coalitions. In this section we will show that it is possible to reduce this problem to detecting a biclique in a bipartite graph. To do this, we will need a lemma of Moulin (1981a), and a standard result of algorithmics. For completeness, we include proofs that were omitted in the original works.

Lemma 4 (Moulin 1981a). *Let T be a coalition of size k , and r, t be coefficients satisfying $rn = tm - \alpha$, where $\alpha = \gcd(m, n)$ is the greatest common divisor of m and n , and $t > \alpha n$. The veto power of T then satisfies:*

$$v(T) = \left\lceil m \frac{k}{n} \right\rceil - 1 = \left\lfloor \frac{rk}{t} \right\rfloor.$$

Proof. From $rn = tm - \alpha$ we obtain:

$$\frac{rk}{t} = \frac{km}{n} - \frac{k\alpha}{tn}.$$

Case one: $\frac{km}{n}$ is an integer. In this case:

$$\begin{aligned} \left\lfloor \frac{rk}{t} \right\rfloor &= \left\lfloor \frac{km}{n} - \frac{k\alpha}{tn} \right\rfloor, \\ &= \left\lfloor m \frac{k}{n} \right\rfloor - 1. \end{aligned}$$

Case two: $\frac{km}{n} = \lfloor \frac{km}{n} \rfloor + \epsilon$, $\epsilon \in \{\frac{1}{n}, \dots, \frac{n-1}{n}\}$. In this case $\lfloor m \frac{k}{n} \rfloor - 1 = \lfloor \frac{km}{n} \rfloor$, and:

$$\begin{aligned} \left\lfloor \frac{rk}{t} \right\rfloor &= \left\lfloor \frac{km}{n} - \frac{k\alpha}{tn} \right\rfloor, \\ &= \left\lfloor \frac{km}{n} \right\rfloor + \left\lfloor \epsilon - \frac{k\alpha}{tn} \right\rfloor, \\ &\geq \left\lfloor \frac{km}{n} \right\rfloor + \left\lfloor \epsilon - \frac{1}{n} \right\rfloor, \\ &\geq \left\lfloor \frac{km}{n} \right\rfloor. \end{aligned}$$

The equality follows because it is clear that:

$$\left\lfloor \frac{km}{n} - \frac{k\alpha}{tn} \right\rfloor \leq \left\lfloor \frac{km}{n} \right\rfloor. \quad \square$$

Proposition 5 (Garey and Johnson 1979). *Let G be a bipartite graph with vertices L on the left and R on the right, and $k > \max(|L|, |R|)$ an integer. The following are true:*

1. *If the largest biclique $K_{x,y} \subseteq G$ satisfies $x + y \geq k$, then $x + y = |L| + |R| - z$, where z is the size of the maximum matching in the bipartite complement of G , \bar{G} .*

2. *We can determine whether G contains a biclique with at least k vertices in polynomial time.*

Proof. Let $K_{x,y} \subseteq G$, with $x + y \geq k > \max(|L|, |R|)$, be the largest biclique in G . Observe that the vertices in $K_{x,y}$ form the maximum independent set in \bar{G} , and since $k > \max(|L|, |R|)$, this independent set has vertices on both sides of the graph. The complement of such a maximum independent set forms the smallest vertex cover, and is of size $|L| + |R| - x - y$. By König's theorem, the size of the smallest vertex cover is equal to z , the size of the maximum matching. Thus $z = |L| + |R| - x - y$, and $x + y = |L| + |R| - z$.

To determine whether a biclique of such size exists, construct \bar{G} , and find the size of the maximum matching, z . This is equal to the size of the smallest vertex cover, so $|L| + |R| - z$ is the size of the maximum independent set. Now verify that $|L| + |R| - z \geq k$. If so, then since $k > \max(|L|, |R|)$, this independent set spans both sides of the graph, and it corresponds to a biclique of size $|L| + |R| - z \geq k$ in G . If $|L| + |R| - z < k$ then, since every biclique in G corresponds to an independent set in \bar{G} , there can be no biclique of size k in G . □

The key insight in the proof of Theorem 8 is that we can reduce the search for a blocking coalition to the search for a biclique in a bipartite graph, and via Proposition 5 find that biclique in polynomial time.

Definition 6. The *blocking graph* for $c \in \mathcal{C}$ is a bipartite graph with rn vertices on the left and $t(m-1)$ vertices on the right, where r, t satisfy the conditions of Lemma 4. Every voter is associated with r vertices on the left, every candidate but c is associated with t vertices on the right, and a vertex associated with voter i is adjacent to a vertex associated with candidate d just if i prefers d to c . ■

Lemma 7. *The blocking graph for c contains a biclique of size tm if and only if c is blocked.*

Proof. Suppose that c is blocked by coalition T of size k . Since r, t satisfy the conditions of Lemma 4, the veto power of T is $\lfloor \frac{rk}{t} \rfloor$. This means there exists a B , $|B| = m - \lfloor \frac{rk}{t} \rfloor$, such that the voters in T prefer everything in B to c . Observe that the vertices associated with T and B form a biclique with $rk + tm - t \lfloor \frac{rk}{t} \rfloor$ vertices. Since $\frac{rk}{t} \geq \lfloor \frac{rk}{t} \rfloor$, it follows that $rk \geq t \lfloor \frac{rk}{t} \rfloor$ and the biclique has at least tm vertices.

Now suppose there exists a biclique $K_{x,y} \subseteq G$ with $x + y \geq tm$. Without loss of generality, we can assume $x = rk', y = tb'$ (since if voter i considers candidate d to be better than c , then all r vertices associated with i are adjacent to all t vertices associated with d). That is, there are k' voters who all agree that b' candidates are better than c . We must show that the coalition has enough veto power to force this outcome: $\lfloor \frac{rk'}{t} \rfloor \geq m - b'$. We know that:

$$\begin{aligned} rk' + tb' &\geq tm \\ rk' &\geq tm - tb' \\ \frac{rk'}{t} &\geq m - b'. \end{aligned}$$

Since $m - b'$ is an integer it follows that $\left\lfloor \frac{rk'}{t} \right\rfloor \geq m - b'$. \square

To see why duplicating the voters and candidates in the proportions dictated by Lemma 4 is necessary, contrast the cases of $m = n + 1$ and $m = 2n + 1$. In the first case, a coalition of size k needs to agree on a blocking set of size $m - k$, and we are searching for a biclique of size m – everything works, and no duplication is necessary. However, with $m = 2n + 1$, k voters need to agree on $m - 2k$ candidates – we are searching for a biclique of size $m - k$, with exactly k vertices on the left. This is a problem, because the argument in Proposition 5 only works if we only care about the total number of vertices in a biclique – the question “Does G contain a biclique of size x , with y vertices on the left” is NP-hard.

If we duplicate the voters and candidates, however, any biclique of size tm suffices – the size of the coalition is no longer relevant.

Theorem 8. *The veto core can be computed in $O(m \max(n^3, m^3))$ time.*

Proof. Using the Extended Euclidean algorithm we can find r', t' such that $r'n = t'm - \alpha$, where $\alpha = \gcd(m, n)$, and $|t'| \leq \frac{n}{\alpha}$ in polynomial time. It is easy to see that with $(r, t) = (r' + 3\alpha m, t' + 3\alpha n)$, we have $rn = tm - \alpha$ and $t > \alpha n$.

Lemma 7 together with Proposition 5 already gives us a polynomial time algorithm for the veto core – for every c , build the blocking graph, find the largest biclique, and add c to the core if the biclique is smaller than tm . However, since $t > \alpha n$, the blocking graph is of size $O(\alpha nm)$, and with a standard $O(E\sqrt{V})$ algorithm to find the maximum matching, this approach will find the veto core in $O(\alpha^2 m^3 n^2 \sqrt{\alpha mn})$ time – a heroic $O(n^{8.5})$ when $m = n = \alpha$. We can do better by reformulating the problem as one of maximum flow.

Let G be the blocking graph for c , and \bar{G} its bipartite complement. We wish to find the size of the maximum matching in \bar{G} without constructing the graph explicitly.

Define the *flow graph* of c to be the directed weighted graph with a source node S , a sink node T , a node for every voter, and a node for every candidate but c . There is an arc from S to each voter with capacity r , an arc from each candidate node to T with capacity t , and an arc of unbounded capacity between every voter node and every candidate node that the voter considers to be worse than c .

We will show that if \bar{G} has a matching of size F then the flow graph has a flow of at least F , and if the flow graph has a maximum flow of F then \bar{G} has a matching of size at least F . Combining the two, the size of the maximum matching in \bar{G} is equal to the maximum flow in the flow graph.

Suppose \bar{G} has a matching M of size F . For each edge (i', d') in M , where i' is a vertex of voter i and d' is a vertex of candidate d , increase the flow in (S, i) , (i, d) , and (d, T) by one. Since there are at most r and t edges in the matching incident on a vertex of i and d respectively, we have enough capacity in (S, i) and (d, T) to perform this operation for every edge in M , and thus construct a flow of size F .

Now suppose there is a maximum flow of F in the flow graph. By the integral flow theorem, we can assume that this flow is integral. For every unit of flow from i to d , match one vertex of i in \bar{G} with one vertex of d . Since the inflow of i is at most r , we have enough vertices of i in \bar{G} , and since the outflow of d is at most t , we have enough vertices of d .

Since we can solve the maximum flow problem in cubic time (Malhotra, Kumar, and Maheshwari 1978), we can check whether c is blocked in $O(\max(n^3, m^3))$ time, and compute the entire veto core in $O(m \max(n^3, m^3))$. \square

Selecting from the Veto Core

Any anonymous veto function with a non-empty core must include the veto core (Moulin 1981a); it is thus the smallest solution possible. In general, however, it is not a singleton. Simulations suggest that under the Impartial Culture assumption (drawing voters’ preferences from all $m!$ preference orders identically and independently) the veto core tends to be unreasonably large (Table 1). Indeed, in the case of a large number of voters and a small number of candidates, we can demonstrate that the veto core will contain half the candidates on average.

Proposition 9. *With the Impartial Culture assumption and fixed m , as $n \rightarrow \infty$ the following are true:*

1. *With probability one, the veto core consists of the candidates which are ranked last by less than $\frac{n}{m}$ voters;*
2. *The expected size of the veto core is half the candidates.*

Proof. Fix a set of candidates, B , $|B| = m - q$, $1 \leq q \leq m - 1$, and fix a candidate c not in B . We will show that as $n \rightarrow \infty$ the probability of the existence of a coalition blocking c with blocking set B is $\frac{1}{2}$ if $q = 1$, and 0 otherwise. Since $|B| = m - 1$ would imply that all members of the coalition rank c last, this will establish 1 and 2.

Let $p(q, m)$ be the probability that a voter believes every candidate in B is better than c . If $q = 1$, then c must be ranked last. There are $(m - 1)!$ possible preference orders the voter can have, so $p(1, m) = \frac{(m-1)!}{m!} = \frac{1}{m}$.

If $q > 1$, then the voter must rank c between positions $m - q + 1$ and m (any higher and the candidates in B could not possibly all fit above c). There are $(m - 1)!$ orders with c ranked last, and less than $(m - 1)!$ with c ranked in any of the other $q - 1$ positions. Thus there are more than $(m - 1)!$, and less than $q(m - 1)!$ preference orders that the voter could have. This gives us $\frac{1}{m} < p(q, m) < \frac{q}{m}$.

Recall that the veto power of a coalition of size k is $\left\lceil m \frac{k}{n} \right\rceil - 1$. Thus a coalition can veto q candidates if and only if the size of the coalition is strictly larger than $\frac{nq}{m}$. It follows that B is a blocking set for c if and only if more than $\frac{nq}{m}$ voters believe every candidate in B is better than c . Hence, by the Central Limit Theorem, the limit of the probability that B is a blocking set for c is:

$$1 - \lim_{n \rightarrow \infty} F \left(\frac{\frac{nq}{m} - np(q, m)}{\sqrt{np(q, m)(1 - p(q, m))}} \right),$$

which is equal to $\frac{1}{2}$ for $q = 1$ and 0 otherwise. \square

	$m = 2$	3	4	5	6	7	8	9	10	11	100	101	200
$n = 2$	0.75	0.39	0.41	0.28	0.28	0.22	0.23	0.20	0.20	0.17	0.04	0.04	0.02
3	0.50	0.69	0.35	0.33	0.40	0.28	0.27	0.31	0.25	0.24	0.09	0.09	0.06
4	0.68	0.45	0.66	0.35	0.37	0.35	0.42	0.30	0.32	0.31	0.15	0.15	0.11
5	0.50	0.42	0.47	0.66	0.36	0.37	0.36	0.39	0.44	0.31	0.21	0.20	0.17
6	0.65	0.63	0.48	0.50	0.67	0.35	0.39	0.41	0.40	0.41	0.26	0.26	0.22
7	0.50	0.48	0.41	0.47	0.54	0.68	0.35	0.39	0.41	0.39	0.29	0.29	0.27
8	0.63	0.44	0.64	0.46	0.52	0.58	0.70	0.36	0.40	0.41	0.32	0.32	0.30
9	0.50	0.63	0.53	0.41	0.52	0.54	0.60	0.71	0.36	0.40	0.33	0.33	0.32
10	0.62	0.51	0.51	0.65	0.47	0.50	0.57	0.63	0.71	0.36	0.37	0.34	0.35
11	0.50	0.45	0.44	0.58	0.42	0.49	0.54	0.60	0.64	0.71	0.35	0.35	0.34
100	0.54	0.53	0.55	0.57	0.49	0.54	0.52	0.57	0.58	0.57	0.73	0.37	0.55
101	0.50	0.50	0.53	0.53	0.48	0.52	0.50	0.55	0.57	0.56	0.73	0.74	0.55
200	0.52	0.49	0.54	0.54	0.52	0.50	0.55	0.54	0.56	0.54	0.68	0.41	0.73

Table 1: Proportion of candidates in the veto core, average of 1,000 IC profiles.

It is thus worthwhile to ask how a single candidate from the veto core can be selected. The first solution was proposed by Moulin (1981a):

Definition 10. Let r, t be coefficients satisfying the provisos of Lemma 4. Create r clones of every voter and t clones of every candidate. Let L be an order over the voter-clones.

Voting by veto tokens is the rule where the voter-clones successively veto their least preferred candidate-clone in the order of L until $\gcd(m, n)$ clones are left. The winners are all candidates who have at least one clone remaining. ■

While different ways of ordering the voter-clones may lead to different outcomes, such an outcome will always lie in the veto core.

Proposition 11 (Moulin 1981a). *Voting by veto tokens elects a candidate in the veto core.*

Proof. Suppose, for contradiction, that voting by veto tokens elects a candidate c , but c is blocked by voters T , $|T| = k$. This means there must exist a set of candidates B , $|B| = m - v(T)$, such that every voter in T believes every candidate in B is better than c . Let $W = \mathcal{C} \setminus B$, and W_i the set that voter $i \in T$ believes to be at least as bad as c . Observe that $W_i \subseteq W$. By Lemma 4, the number of clones of candidates in W is equal to $t|W| = tv(T) = t \lfloor \frac{rk}{t} \rfloor$.

Since a clone of c remains unvetoes by the end of the algorithm, then it means that each of the r clones of every voter $i \in T$ casts a veto against a clone of a candidate in $W_i \subseteq W$. Yet these rk vetoes have been insufficient to eliminate all the clones of $c \in W$ and hence $rk < t \lfloor \frac{rk}{t} \rfloor$, which is a contradiction. □

While voting by veto tokens is defined for all choices of n and m , it is inherently non-anonymous. This is the case of the majority of known procedures selecting from the veto core, and the few that are anonymous (the procedure of Peleg (1978), or the strong equilibria of antiplurality (Wilczynski 2018)) violate neutrality instead. If we are willing to admit a randomised procedure we could, of course, run voting by veto tokens with a random voting order; however it is possible to select from the veto core anonymously, neutrally, and

deterministically via a voting rule inspired by the probabilistic serial mechanism of Bogomolnaia and Moulin (2001):

Definition 12. *Veto by consumption* is the voting rule that is computed by an algorithm that has voters eat the candidates from the bottom of their order up. Every candidate starts with capacity 1, and is being eaten by the voters who rank it last. Each voter eats at speed 1.

The outcome can be computed as follows. In round k , let c_i be the capacity of candidate i and n_i the number of voters eating i . The round lasts until some candidate is fully eaten. To move to round $k + 1$, do the following:

1. Find an i which minimises c_i/n_i . Let r_k be this minimum ratio – this is the duration of the round.
2. Update all capacities, $c_j = c_j - r_k n_j$.
3. For all candidates who reached capacity 0, reallocate the voters eating these to their next worst candidates.

The last candidate to be eaten is the winner. In the case of two or more candidates being eaten simultaneously, a tie is declared among those candidates. ■

Observe that the duration of the algorithm is exactly $\frac{m}{n}$ units, and the time complexity is $O(mn)$. As a voting rule, veto by consumption is clearly anonymous, neutral, Pareto efficient, homogeneous, and never selects the candidates which are ranked last by more than $\frac{m}{m}$ voters. The results in Table 2 suggest that the number of winners tends to 1 as the number of voters increases.

Proposition 13. *Veto by consumption selects a candidate in the veto core.*

Proof. Suppose, for contradiction, that veto by consumption elects a candidate c , but c is blocked by voters T , $|T| = k$. This means there must exist a set of candidates B , $|B| = m - v(T) = m - \lceil \frac{mk}{n} \rceil + 1$, such that every voter in T believes every candidate in B is better than c . For voter $i \in T$, let W_i be the set of candidates i considers to be at least as bad as c . Observe that $W_i \subseteq \mathcal{C} \setminus B$.

Observe that as long as c is present, every agent will eat a candidate they consider to be at least as bad as c . Since we have assumed that c is eaten last, it follows that the k voters

in T eat only from $W = \bigcup W_i \subseteq \mathcal{C} \setminus B$ for the duration of the algorithm, and the capacity of W is $|W| \leq \lceil \frac{mk}{n} \rceil - 1 < \frac{mk}{n}$. However, it will take these k voters strictly less than $\frac{m}{n}$ units of time to eat all of W , but the algorithm runs for exactly $\frac{m}{n}$ units of time, which is a contradiction. \square

Given that we can compute the veto core in polynomial time, however, one could suggest any number of neutral and anonymous core selection algorithms – first compute the veto core, throw away all other candidates, and use your favourite voting rule to select a winner from what is left. Unfortunately, the veto core is vulnerable to a very simple form of agenda manipulation – padding the profile at the bottom with universally reviled candidates (Kondratev, Ianovski, and Nesterov 2019; Barberà and Coelho 2020). Indeed, suppose we add $(n - 1)m$ spoiler candidates at the bottom of every voter’s preference order. In the new profile k voters will need to agree on a blocking set of size $mn - \lceil \frac{mnk}{n} \rceil + 1 = mn - mk + 1$. Since there are only m original candidates, and they are universally preferred to the spoilers, it follows that only the grand coalition will be able to block anything, and the veto core will consist of all Pareto undominated candidates. And this change of agenda could be effected as easily as suggesting a group of English speakers choose a film not from a library of English-language films, but a library of all films.

Veto by consumption avoids this problem, as all the spoiler candidates are eaten before the voters move on to the real contenders.

Manipulating the Veto Core

The Gibbard-Satterthwaite framework of strategic voting presupposes a resolute voting rule, i.e. one that always outputs a single winner. As no anonymous rule can have this property, the standard approach in studying manipulation is to assume that ties are broken either lexicographically, in favour of the manipulator, or against the manipulator. Such an approach is reasonable if ties are a rare occurrence, but given the tendency of the veto core to be large, in our context such an approach would lead to a tangible deviation from neutrality or anonymity. Instead, we need to evaluate the manipulator’s preferences on the outcome sets directly. Duggan and Schwartz (1992) considered two kinds of preference extensions:¹ manipulation by an optimist, who evaluates a set by the best element, and a pessimist, who evaluates a set by the worst. While these extensions may not seem very imaginative, they are enough to derive an impossibility result – any rule that is immune to both types of manipulation and includes all singletons in its range is dictatorial.

Conceptually, manipulating the veto core in this framework is simple. The veto core consists of all candidates that are not blocked; therefore, an agent’s strategies are limited to either blocking a candidate they would not sincerely, or refusing to block a candidate they would have otherwise. In

the case of pessimistic manipulation, this gives us a polynomial time algorithm.

Lemma 14. *The top set of candidate c in order P_i is the set of all candidates ranked above c .*

If c is blocked in (P_i, P_{-i}) with top set X , and $X \subseteq Y$, then for any P'_i with top set Y , c is blocked in (P'_i, P_{-i}) .

Proof. Suppose c is blocked in (P_i, P_{-i}) with top set X . Let P'_i be an arbitrary order where the top set of c is Y , $X \subseteq Y$.

Let T be a coalition blocking c in (P_i, P_{-i}) with blocking set B . If $i \notin T$, then T is a blocking coalition regardless of what i votes, and remains a blocking coalition in (P'_i, P_{-i}) . If $i \in T$, then observe that $B \subseteq X \subseteq Y$, so i considers all the candidates in B to be better than c . The other voters’ preferences are unchanged, so they do too. Thus T is still a blocking coalition in (P'_i, P_{-i}) . \square

Theorem 15. *It can be determined whether the veto core is manipulable by a pessimist in polynomial time.*

Proof. Consider a profile P where c is the worst element in the veto core for the manipulator, who will be voter 1 for convenience. Call all the candidates 1 considers to be better than c the good candidates, all the candidates at least as bad (including c itself) the bad candidates.

The problem of pessimistic manipulation is thus finding a P'_1 such that every bad candidate in (P'_1, P_{-1}) is blocked. Call such a P'_1 a strategic vote.

Observe that as a consequence of Lemma 14, if c is blocked in (P'_1, P_{-1}) with top set X , then c is blocked in every (P_1^*, P_{-1}) with top set X . We can check whether this is the case in polynomial time by constructing an arbitrary preference order P_1^* with the candidates in X ranked first, then c , and running the veto core algorithm on (P_1^*, P_{-1}) . We will refer to this as *checking whether c is blocked with top set X* .

The algorithm functions by progressively filling the ballot from the top down. At step i of the algorithm, let G_i be the set of candidates already on the ballot. To initialise the algorithm, construct a partially filled ballot with all the good candidates ranked at the top in the same order as in P_1 . Thus G_0 is the set of good candidates. To advance from step i to $i + 1$ of the algorithm check if there exists a candidate $b \in \mathcal{C} \setminus G_i$ that is blocked with top set G_i . If so, rank b directly below the candidates in G_i and let $G_{i+1} = G_i \cup \{b\}$. If not, assert that no strategic vote exists.

First, observe that this algorithm will run for $O(m)$ steps, and each step will involve $O(m)$ calls to the algorithm computing veto core, so the algorithm is in P.

Next, we show that if the algorithm returns a vote it is indeed a strategic vote. To demonstrate this we will show that the algorithm has the invariant property that every candidate in G_i is either a good candidate or is blocked. At the beginning of the algorithm, G_0 consists solely of the good candidates. After step i , we add some candidate b who is blocked with top set G_i . We do not change the top set of any candidate in G_i , so those that were blocked at step i remain blocked at step $i + 1$. Thus if the algorithm returns a vote, all the candidates on the ballot are either good candidates or blocked.

¹Duggan and Schwartz (1992) use a more complicated definition in their paper, but it is equivalent to the pessimist/optimist definition found in Taylor (2002).

	$m = 2$	3	4	5	6	7	8	9	10	11	100	101	200
$n = 2$	1.5	1.167	1.42	1.22	1.38	1.24	1.37	1.25	1.35	1.26	1.31	1.3	1.31
3	1	1.6	1.04	1.02	1.41	1.04	1.03	1.33	1.04	1.03	1.04	1.04	1.04
4	1.38	1.24	1.58	1.13	1.28	1.2	1.36	1.15	1.23	1.18	1.13	1.12	1.12
5	1	1	1	1.42	1	1	1	1	1.21	1	1.02	1	1
6	1.31	1.36	1.21	1.15	1.3	1.08	1.1	1.14	1.09	1.08	1.03	1.03	1.02
7	1	1	1	1	1	1.18	1	1	1	1	1	1	1
8	1.27	1.08	1.24	1.08	1.11	1.06	1.11	1.04	1.05	1.03	1.01	1.01	1.01
9	1	1.22	1.01	1	1.08	1.03	1.02	1.06	1.01	1.01	1	1	1
10	1.25	1.14	1.11	1.13	1.05	1.03	1.02	1.01	1.04	1.01	1	1	1
11	1	1	1	1	1	1	1	1	1	1.02	1	1	1
100	1.08	1.01	1.01	1	1	1	1	1	1	1	1	1	1
101	1	1	1	1	1	1	1	1	1	1	1	1	1
200	1.06	1.01	1	1	1	1	1	1	1	1	1	1	1

Table 2: Number of veto by consumption winners, average of 1,000,000 IC profiles.

Finally, we show that if there exists a strategic vote P'_1 , the algorithm will find one.

We first claim that if there exists a strategic vote, there exists a strategic vote with all the good candidates ranked at the top of the ballot in the same order as in P_1 . Let P'_1 be a strategic vote where some good g is ranked behind a bad b . Obtain P'_1 by swapping g and b . We increase the top set of b , so if b was blocked it remains blocked. Thus P'_1 is still strategic, and we can repeat this operation until all the good candidates bubble to the top. Once we have a ballot where all the good candidates are on the top, their internal order does not matter since it will not change the top set of any bad candidate. There is thus no generality lost in assuming that the good candidates are ranked as in P_1 .

Now we proceed by contradiction. Suppose the algorithm fails to return a vote. This must mean that at step i of the algorithm none of the remaining bad candidates are blocked with top set G_i . By the above claim, we can assume that P'_1 ranks the set of good candidates G_0 at the top in the same order as P_1 and thus has bad candidates in the next i positions. Let X be the set of these i bad candidates, and $X' = G_i \setminus G_0$ be the set of candidates ranked in the corresponding positions by our algorithm. Observe that $X \neq X'$ – otherwise the algorithm could choose the same candidate in the i th step as in P'_1 . Thus there exists a $b \in X, b \notin X'$. Of all such b , choose one that is ranked highest in P'_1 . Let Y be the set of bad candidates (possibly empty) ranked above b in P'_1 . Since b is the highest ranked bad candidate in P'_1 that is not in X' , it must be that $Y \subseteq X'$. But that is a contradiction because if b is blocked with top set $G_0 \cup Y$ in P'_1 , it will certainly be blocked with top set $G_i = G_0 \cup X'$ in the algorithm. \square

In principle, it is also easy to check whether optimistic manipulation is possible – by Lemma 14, if it is possible for voter 1 to add some c to the core, then it must be possible to add c to the core by ranking c first, and the other candidates in any order. However, we have no examples of optimistic manipulation of the core, and suspect that it is not possible.

Future Directions

From inception, the concept of voting by veto was closely linked with implementation. Anbarci (2006) described the subgame perfect equilibria of the Alternate Strike scheme. It turns out that these equilibria correspond exactly to the veto by consumption winners in the case of two voters. Similarly, Laslier, Nunez, and Sanver (2020) proposed a mechanism which implements the two-voter veto core in Nash equilibrium. It is natural to ask whether something similar can be done with an arbitrary number of voters.

Veto by consumption is a novel voting rule, and hence the standard questions (axiomatisation, non-trivial properties, susceptibility to manipulation, etc) have not been asked about it. As an anonymous voting rule which selects from the veto core, it would be interesting to consider how well it functions in settings where the happiness of the minority matters, e.g. in group recommender systems (Masthoff 2015).

Acknowledgements

The section “Manipulating the Veto Core” is funded by Grant 20-71-00029 of the Russian Science Foundation. The section “Selecting from the Veto Core” is funded by Grant 20-01-00687 of the Russian Foundation for Basic Research. The section “Computing the Veto Core” is funded by the HSE University Basic Research Program.

We thank the referees for their helpful comments.

References

- Anbarci, N. 1993. Noncooperative foundations of the area monotonic solution. *The Quarterly Journal of Economics* 108(1): 245–258.
- Anbarci, N. 2006. Finite alternating-move arbitration schemes and the equal area solution. *Theory and Decision* 61(1): 21–50.
- Barberà, S.; and Coelho, D. 2017. Balancing the power to appoint officers. *Games and Economic Behavior* 101: 189–203.

- Barberà, S.; and Coelho, D. 2020. Compromising on compromise rules. Available at SSRN: <https://ssrn.com/abstract=3567049>.
- Bloom, D. E.; and Cavanagh, C. L. 1986. An analysis of the selection of arbitrators. *American Economic Review* 76(3): 408–422.
- Bogomolnaia, A.; and Moulin, H. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100(2): 295 – 328.
- Bouveret, S.; Chevaleyre, Y.; Durand, F.; and Lang, J. 2017. Voting by sequential elimination with few voters. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI'2017*, 128–134. AAAI Press.
- de Clippel, G.; Eliaz, K.; and Knight, B. 2014. On the selection of arbitrators. *American Economic Review* 104(11): 3434–58.
- Duggan, J.; and Schwartz, T. 1992. Strategic manipulability is inescapable: Gibbard-Satterthwaite without resoluteness. Working Papers 817, California Institute of Technology, Division of the Humanities and Social Sciences.
- Erlich, S.; Hazon, N.; and Kraus, S. 2018. Negotiation strategies for agents with ordinal preferences. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI'2018*, 210–218. AAAI Press.
- Garey, M. R.; and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York, NY, USA: W. H. Freeman & Co. ISBN 0716710447.
- Kondratev, A. Y.; Ianovski, E.; and Nesterov, A. S. 2019. How should we score athletes and candidates: Geometric scoring rules. *arXiv:1907.05082*.
- Kondratev, A. Y.; and Nesterov, A. S. 2020. Measuring majority power and veto power of voting rules. *Public Choice* 183: 187–210.
- Laslier, J.-F.; Nunez, M.; and Sanver, M. R. 2020. A solution to the two-person implementation problem. *halshs-02173504v3*.
- Malhotra, V.; Kumar, M.; and Maheshwari, S. 1978. An $O(|V|^3)$ algorithm for finding maximum flows in networks. *Information Processing Letters* 7(6): 277–278.
- Masthoff, J. 2015. Group Recommender Systems: Aggregation, Satisfaction and Group Attributes. In Ricci, F.; Rokach, L.; and Shapira, B., eds., *Recommender Systems Handbook*, 743–776. Boston, MA: Springer US.
- Moulin, H. 1980. Implementing efficient, anonymous and neutral social choice functions. *Journal of Mathematical Economics* 7(3): 249–269.
- Moulin, H. 1981a. The proportional veto principle. *The Review of Economic Studies* 48(3): 407–416.
- Moulin, H. 1981b. Prudence versus sophistication in voting strategy. *Journal of Economic Theory* 24(3): 398–412.
- Moulin, H. 1982. Voting with proportional veto power. *Econometrica* 50(1): 145–162.
- Mueller, D. C. 1978. Voting by veto. *Journal of Public Economics* 10(1): 57–75.
- Peleg, B. 1978. Consistent voting systems. *Econometrica* 46(1): 153–161.
- Taylor, A. D. 2002. The manipulability of voting systems. *The American Mathematical Monthly* 109(4): 321–337.
- Wilczynski, A. 2018. *Interaction among Agents via a Social Network in Computational Social Choice*. Ph.D. thesis, PSL Research University.