Anonymity-Proof Voting Rules

Vincent Conitzer

Departments of Computer Science and Economics
Duke University
Durham, NC, USA
conitzer@cs.duke.edu

Abstract. A *(randomized, anonymous) voting rule* maps any multiset of total orders (aka. *votes*) over a fixed set of alternatives to a probability distribution over these alternatives. A voting rule f is *false-name-proof* if no voter ever benefits from casting more than one vote. It is *anonymity-proof* if it satisfies voluntary participation and it is false-name-proof. We show that the class of anonymity-proof neutral voting rules consists exactly of the rules of the following form. With some probability $k_f \in [0, 1]$, the rule chooses an alternative uniformly at random. With probability $1 - k_f$, the rule first draws a pair of alternatives uniformly at random. If every vote prefers the same alternative between the two (and there is at least one vote), then the rule chooses that alternative. Otherwise, the rule flips a fair coin to decide between the two alternatives. We also show how the characterization changes if group strategy-proofness is added as a requirement.

1 Introduction

In many settings, a decision must be made on the basis of the preferences of multiple agents. Common examples include auctions and exchanges (where we must decide on an allocation of resources, as well as payments to be made or received by the agents) and elections (where we must decide on, say, one or more political representatives), but there are many other applications. A *(direct-revelation) mechanism* takes each agent's reported preferences as input, and produces a decision as output. An important issue is that self-interested agents will lie about their preferences if they perceive it to be to their advantage to do so. *Mechanism design* studies how to design mechanisms that produce good outcomes in spite of this. A key concept in mechanism design is that of *strategy-proofness*: a mechanism is strategy-proof if no agent can ever benefit from lying about her preferences. Strategy-proofness is roughly synonymous with *truthfulness* and *incentive compatibility*.[1] In mechanism design, attention is usually restricted to incentive compatible direct-revelation mechanisms. This is justified by a result known as the *revelation principle* [Gibbard, 1973; Green and Laffont, 1977; Myerson, 1979, 1981], which states (roughly) that, given that agents will misreport their preferences if

[1] To be more precise, strategy-proofness as the term is used here corresponds to *dominant-strategies* incentive compatibility. There are weaker notions of incentive compatibility, such as *Bayes-Nash* incentive compatibility, where *in expectation* over the other agents' preferences an agent is best off reporting her true preferences (assuming the others do so as well).

they perceive this to be to their benefit, anything that can be achieved by some mechanism can also be achieved by an incentive compatible direct-revelation mechanism.[2]

In mechanism design, the spaces of possible outcomes and preferences often display a great deal of structure, which facilitates the designer's job. For example, in auctions and exchanges, it is often assumed that agents can make and receive payments, that their utility is linear in this payment, and that the effect of the payment on utility is independent of the rest of the outcome. This enables, for example, Vickrey-Clarke-Groves mechanisms [Vickrey, 1961; Clarke, 1971; Groves, 1973], which always choose the efficient allocation. However, such structure is not always available: for example, in an election, payments can typically not be made. If we do not assume any structure on the agents' preferences, then agents can rank the possible outcomes (aka. *alternatives*) in any possible way. These general settings, in which each agent ranks all the alternatives, and the mechanism chooses an alternative based on these rankings, are commonly referred to as *voting* settings. The rankings are the *votes*, and the mechanism is usually called a *voting rule*.

The revelation principle applies to voting settings just as it does to any other mechanism design setting, so we should ask which rules are strategy-proof. Gibbard [1977] provides a complete characterization of strategy-proof voting rules that are allowed to use randomization. (This characterization generalizes the better-known, earlier Gibbard-Satterthwaite theorem [Gibbard, 1973; Satterthwaite, 1975].) He shows that any strategy-proof rule is a randomization over *unilateral* rules, in which only one vote affects the outcome, and *duple* rules, in which only two alternatives have a chance of winning. (Because the overall rule is a randomization over such rules, it can still be the case that every voter affects the probability with which an alternative is chosen, and that every alternative has a positive probability of winning. Hence, Gibbard's characterization is not universally seen as a negative result [Barbera, 1979a].) He also provides some additional conditions on these rules to obtain an exact characterization of the strategy-proof voting rules.

However, strategy-proofness is often not sufficient. In open, anonymous environments such as the Internet, an agent can manipulate the mechanism in other ways. For one, if an agent does not participate in the mechanism, then the party running the mechanism (aka. the *center*) is not even aware of her existence. Perhaps more significantly, an agent can open multiple accounts and participate in the mechanism multiple times under different identifiers—and the center cannot know which identifiers correspond to the same agent. This led to the concept of *false-name-proofness* [Yokoo *et al.*, 2004]. A mechanism is false-name-proof if an agent can never benefit from using multiple identifiers. Some positive and negative results on false-name-proofness have been obtained for combinatorial auctions and similar settings (*e.g.*, Yokoo *et al.* [2001]; Yokoo [2003]; Yokoo *et al.* [2004, 2006]; Rastegari *et al.* [2007]), but to our knowledge this concept has not yet been studied in voting settings.

In this paper, we define a (possibly randomized) voting rule to be *anonymity-proof* if it is false-name-proof, and it never hurts an agent to cast her (true) vote. Under the same

[2] To predict what will happen under a mechanism that is not incentive compatible, some solution concept from game theory must be used, and the version of incentive compatibility in the revelation principle depends on the choice of solution concept.

model as Gibbard [1977], we obtain a complete characterization of the anonymity-proof neutral voting rules. (A voting rule is *neutral* if it treats all alternatives symmetrically.) The proof is from first principles and (arguably) of reasonable length. The resulting class of voting rules is very limited (hence the result is mostly negative), but it does allow a modicum of responsiveness to the votes in cases where there is complete agreement among the voters on some pairs of alternatives. For example, in the special case where there are only two alternatives, the characterization tells us that if all votes prefer the same alternative, we can choose that alternative; but otherwise, we have to flip a fair coin to decide between them. This is in stark contrast to the case where we require only strategy-proofness, or even group strategy-proofness: for example, simply choosing the alternative that is preferred by more voters (the *majority* rule) is group strategy-proof.

1.1 Additional motivation

Our primary reason for studying false-name-proofness in general social choice (voting) settings is that these settings lie at the heart of mechanism design, and hence provide the most natural starting point for a thorough study of the concept of false-name-proofness. Nevertheless, perhaps surprisingly, anonymous voting is in fact a very real and growing phenomenon on the Internet. It may seem that anonymous elections are unlikely to result in outcomes that reflect society's preferences well (and, in fact, this paper can be seen as a commentary on just how unlikely this is). However, it appears that in practice, often, the party organizing the election has more interest in publicity than in a properly chosen outcome; moreover, the convenience of anonymous Internet voting appeals to the voters as well.

A very recent example of this phenomenon is the "New 7 Wonders of the World" election, a global election that was organized by businessman Bernard Weber to elect contemporary alternatives to the ancient wonders. Anyone could vote, either by phone or over the Internet; for the latter, an e-mail address was required. One could also buy additional votes (of course, using another e-mail address was a much cheaper alternative). In spite of various irregularities (including unreasonably large numbers of votes in some cases) and UNESCO distancing itself from the election, the election seems to have attained some legitimacy in the public's mind.

For better or worse, mechanisms such as these are going to feature increasingly prominently in our economy and social infrastructure. Hence, the theory of mechanism design must be extended so that it can provide guiding principles to maximize the efficiency and trustworthiness of such mechanisms. The sooner this happens, the fewer bad mechanisms will take hold.

Our results also apply to Internet rating systems in which anonymous reviewers rate products, sellers, *etc.* Here, the set of alternatives is the set of possible (final, aggregate) ratings. It should be noted that in this context, it makes sense for agents' preferences to be restricted: for example, it makes little sense for an agent to prefer high$\succ$low$\succ$medium for a product's final rating. Specifically, *single-peaked* preferences [Black, 1948] are a natural restriction in this domain; we will discuss such preferences in the conclusion.

2 Definitions

Let X, $|X| = m$, be the set of *alternatives* over which the voters are voting. A voter's preferences are given by a total order $\succ$ over the alternatives, together with a vector of utilities $\boldsymbol{u} = (u_1, \ldots, u_m)$ where u_i is the voter's utility for the alternative that she ranks ith. (It is required that $u_i > u_{i+1}$ for all $1 \le i \le m$.) Each voter seeks to maximize her expected utility. As in Gibbard [1977], voters only report a total order (ranking) of the alternatives (not their utilities); a reported ranking is called a *vote*. Again as in Gibbard [1977], we do not allow for indifferences (real or reported) between alternatives. We will use the notation $v = a_1 \succ \ldots \succ a_m$ for a vote. We will sometimes also use subsets in the order notation: for example, if $B = \{b_1, b_2, b_3\}$, then $a_1 \succ b_1 \succ b_2 \succ b_3 \succ a_2$ and $a_1 \succ b_3 \succ b_1 \succ b_2 \succ a_2$ are both *of the form* $a_1 \succ B \succ a_2$ (but, for instance, $a_1 \succ b_3 \succ b_2 \succ a_2 \succ b_1$ is not of this form). A *voting rule* f takes a multiset[3] of votes V as input, and chooses the winning alternative based on these votes (possibly using randomization). Let $P_f(V, a)$ denote the probability with which f chooses a given votes V; the function P_f defines the rule f. A voting rule is *neutral* if it treats all alternatives symmetrically—that is, if π is a permutation of the alternatives, then $P_f(\pi(V), \pi(a)) = P_f(V, a)$ (where $\pi(V)$ is the multiset that results from replacing each alternative a by $\pi(a)$ in each vote in V). In fact, the following weaker definition of neutrality will also suffice for our purposes: if a subset B of the alternatives is symmetric in V (that is, for any permutation π for which $\pi(a) = a$ for all $a \in X - B$, $\pi(V) = V$), then $P_f(V, b_1) = P_f(V, b_2)$ for all $b_1, b_2 \in B$. We are only interested in neutral voting rules.[4]

Definition 1 *A voting rule f is* false-name-proof *if for any multiset of votes V, for any $v \in V, v = a_1 \succ \ldots \succ a_m$, for any decreasing $\boldsymbol{u} = (u_1, \ldots, u_m)$, and for any multiset of votes V', we have $\sum_{j=1}^{m} P_f(V, a_j)u_j \ge \sum_{j=1}^{m} P_f(V \cup V', a_j)u_j$. That is, the voter corresponding to v cannot increase her expected utility by additionally casting votes V'.*

It should be noted that under this definition, a voter who uses false names is assumed to cast at least one vote representing her true preferences. This only weakens the requirement. All of the rules in the characterization result of this paper are also false-name-proof in the stronger sense where none of the votes cast by the false-name voter are required to represent her true preferences. Hence, the characterization remains the same if this stronger requirement is used.

Definition 2 *A voting rule f* satisfies participation *if for any multiset of votes V, for any $v \in V, v = a_1 \succ \ldots \succ a_m$, for any decreasing $\boldsymbol{u} = (u_1, \ldots, u_m)$, we have $\sum_{j=1}^{m} P_f(V, a_j)u_j \ge \sum_{j=1}^{m} P_f(V - \{v\}, a_j)u_j$. That is, the voter corresponding to v cannot increase her expected utility by not casting her vote.*

[3] This is implicitly assuming that every vote is treated equally; anything else would seem unreasonable in open, anonymous environments. Rules that treat every vote equally are commonly called *anonymous*; this is not to be confused with the definition of anonymity-proofness.

[4] Sometimes rules that are not neutral are of interest, for example if one alternative is the incumbent and should be treated specially; but in most settings, only neutral rules are of interest.

Definition 3 *A voting rule is* anonymity-proof *if it is false-name-proof and it satisfies participation.*

Anonymity-proofness does not directly mention strategy-proofness. Thus, it may appear that even if a rule is anonymity-proof, it is possible that a voter can benefit from misreporting her preferences. However, all of the rules in the characterization result of this paper are also strategy-proof (this is implied by the fact that they satisfy the stronger version of false-name-proofness). Hence, the characterization remains the same if strategy-proofness is added as a requirement.

3 The characterization of anonymity-proof rules

In this section, we prove the main result. Showing that all the rules in the proposed class are anonymity-proof is not difficult; most of the proof consists of showing that all rules that are anonymity-proof are in the class. We prove the latter part using a sequence of six lemmas. Assuming the rule is anonymity-proof, these lemmas demonstrate how to transform any multiset of votes to a particular multiset of only two votes, without affecting one given alternative's probability of winning; and they demonstrate that this alternative's probability of winning in those two votes is as the theorem states.

The first lemma is a fundamental building block of the proof. It states that if we add a vote that agrees with an existing vote on the top k and bottom $l-k$ alternatives, then the probability of winning for each of those alternatives does not change.

Lemma 1 *Consider a multiset of votes V, and suppose that for some $v \in V$, v is of the form $a_1 \succ \ldots \succ a_k \succ B \succ a_{k+1} \succ \ldots \succ a_l$. (Please note that l is equal to m only if B is empty.) Let v' (not necessarily in V) be another vote of the form $a_1 \succ \ldots \succ a_k \succ B \succ a_{k+1} \succ \ldots \succ a_l$ (that is, it is identical to v except for the internal ordering of B). Then, if f is anonymity-proof, for any $1 \leq i \leq l$, $P_f(V, a_i) = P_f(V \cup \{v'\}, a_i)$.*

Proof. First, let us suppose that for some $1 \leq i \leq k$, $P_f(V, a_i) \neq P_f(V \cup \{v'\}, a_i)$. Without loss of generality, suppose that for any $1 \leq j < i$, $P_f(V, a_j) = P_f(V \cup \{v'\}, a_j)$. Consider the utility vector $\boldsymbol{u} = (1-\epsilon, 1-2\epsilon, \ldots, 1-i\epsilon, (m-i)\epsilon, (m-i-1)\epsilon, \ldots, \epsilon)$. First, let us suppose that $P_f(V, a_i) < P_f(V \cup \{v'\}, a_i)$. Then, if the true preferences are given by V, the voter casting v has utility vector $\boldsymbol{u}$, and ϵ is sufficiently small, then the voter casting v has an incentive to cast v' as well. This is because (as $\epsilon \to 0$) she effectively seeks to maximize the probability of one of $a_1, \ldots, a_i$ winning, and casting v' as well does not affect the probabilities of $a_1, \ldots, a_{i-1}$ winning and increases that of a_i. On the other hand, suppose that $P_f(V, a_i) > P_f(V \cup \{v'\}, a_i)$. Then, if the true preferences are given by $V \cup \{v'\}$, the voter casting v' has utility vector $\boldsymbol{u}$, and ϵ is sufficiently small, then the voter casting v' has an incentive to not participate. This is because (as $\epsilon \to 0$) she effectively seeks to maximize the probability of one of $a_1, \ldots, a_i$ winning, and not participating does not affect the probabilities of $a_1, \ldots, a_{i-1}$ winning and increases that of a_i. Hence, for any $1 \leq i \leq k$, $P_f(V, a_i) = P_f(V \cup \{v'\}, a_i)$.

The case where $P_f(V, a_i) \neq P_f(V \cup \{v'\}, a_i)$ for some $k+1 \leq i \leq l$ can be shown to contradict either false-name-proofness or participation by a symmetric argument (where, supposing without loss of generality that $P_f(V, a_j) = P_f(V \cup \{v'\}, a_j)$

for all $i < j \leq l$, the voter casting v or v' effectively tries to *minimize* the probability of one of the *last* $l - i + 1$ alternatives winning).

We obtain the following corollary, which states that it does not matter if the same vote is cast more than once. (This corollary is usually not powerful enough to use instead of the more general Lemma 1, but it provides some insight. In particular, for any fixed number of alternatives, this leaves only a finite number of multisets of votes to consider.)

Corollary 1 *For an anonymity-proof rule f, given that a vote is cast at least once, it does not matter how often it is cast.*

Proof. This follows from setting $B = \emptyset$ in Lemma 1.

The next few lemmas (2, 3, and 4) demonstrate how to transform any multiset of votes into a multiset of only two votes, without affecting one given alternative a's probability of winning (assuming that the rule is anonymity-proof).

Lemma 1 allows us to prove the following lemma, which states that reordering the alternatives after a given alternative a in a vote, as well as reordering those before a, does not affect a's probability of winning, unless we move alternatives past a.

Lemma 2 *Consider a multiset of votes V, and suppose that for some $v \in V$, v is of the form $B \succ a \succ C$. Let v' (not necessarily in V) be another vote of the form $B \succ a \succ C$ (that is, it is identical to v except for the internal ordering of B and C). Then, if f is anonymity-proof, $P_f(V, a) = P_f((V - \{v\}) \cup \{v'\}, a)$. That is, we can permute the alternatives on either side of a in a vote without affecting a's probability of winning.*

Proof. Suppose first that we permute only C, that is, that B is ordered the same way in both v and v'. Then, we can apply Lemma 1 (letting a correspond to a_k in that lemma) to obtain $P_f(V, a) = P_f(V \cup \{v'\}, a)$, and similarly $P_f((V - \{v\}) \cup \{v'\}, a) = P_f(V \cup \{v'\}, a)$, hence $P_f(V, a) = P_f((V - \{v\}) \cup \{v'\}, a)$. The case where we permute only B can be proven symmetrically. But then, in the general case where both B and C are permuted, we can transform v into v' in two steps, as follows. Let v'' be the vote of the form $B \succ a \succ C$ that agrees with v on B but with v' on C. By the above, we have $P_f(V, a) = P_f((V - \{v\}) \cup \{v''\}, a) = P_f((V - \{v\}) \cup \{v'\}, a)$.

The next lemma shows that we *can* move an alternative b past a given alternative a in a vote, without affecting a's probability of winning, *if* the other votes disagree on the relative ranking of a and b.

Lemma 3 *Consider a multiset of votes V, and suppose that for some $v \in V$, a is ranked before b. Additionally, suppose there is another vote $v' \in V$ that ranks a before b, and a third vote $v'' \in V$ that ranks b before a. Let v''' be a vote (not necessarily in V) that is obtained from v by improving b's position, placing it somewhere ahead of a (while not changing the order in any other way). Then, if f is anonymity-proof, $P_f(V, a) = P_f((V - \{v\}) \cup \{v'''\}, a)$. That is, we can move b to the other side of a in a vote without affecting a's probability of winning, if there are other votes that rank a before b and b before a.*

Proof. Let us first assume that a and b are adjacent in v and v'''. That is, a is ranked *directly* before b in v, and v''' is obtained from v simply by swapping a and b. By Lemma 1 (letting $\{a, b\}$ correspond to B in that lemma), for any alternative $c \notin \{a, b\}$, $P_f(V, c) = P_f(V \cup \{v'''\}, c)$, and also $P_f((V - \{v\}) \cup \{v'''\}, c) = P_f(V \cup \{v'''\}, c)$. Now, if we suppose that $P_f(V, a) < P_f(V \cup \{v'''\}, a)$, then, if the true preferences are given by V, the voter corresponding to v' would be better off casting v''' as well (since it will only affect the probabilities of a and b being elected, and v' prefers a). Conversely, if $P_f(V, a) > P_f(V \cup \{v'''\}, a)$, then the voter corresponding to v'' would be better off casting v''' as well. Hence, since f is false-name-proof, $P_f(V, a) = P_f(V \cup \{v'''\}, a)$. It similarly follows that $P_f((V - \{v\}) \cup \{v'''\}, a) = P_f(V \cup \{v'''\}, a)$ (since v' and v'' are still present in $(V - \{v\}) \cup \{v'''\}$). Hence, $P_f(V, a) = P_f((V - \{v\}) \cup \{v'''\}, a)$.

Now let us return to the general case where a and b are not necessarily adjacent in v and v'''. Let v'''' be the result of improving b's position in v to just after a, and let v''''' be the result of swapping a and b in v''''. Using Lemma 2, $P_f(V, a) = P_f((V - \{v\}) \cup \{v''''\}, a)$; using the above argument, $P_f((V - \{v\}) \cup \{v''''\}, a) = P_f((V - \{v\}) \cup \{v'''''\}, a)$; and using Lemma 2 again, $P_f((V - \{v\}) \cup \{v'''''\}, a) = P_f((V - \{v\}) \cup \{v'''\}, a)$.

In the next lemma, we use the previous lemmas to reduce a set of votes to a particular pair of votes, without affecting a's probability of winning. (The proofs of the remaining lemmas and corollaries are omitted due to space constraint.)

Lemma 4 *Given a nonempty multiset of votes V and a distinguished alternative a, let B be the set of alternatives that are ranked before a by every vote in V, let C be the set of alternatives that are ranked before a by some votes in V and after a by others, and let D be the set of alternatives that are ranked after a by every vote in V. Let v (not necessarily in V) be a vote of the form $B \succ a \succ C \cup D$, and let v' (not necessarily in V) be a vote of the form $B \cup C \succ a \succ D$. Then, if f is anonymity-proof, $P_f(V, a) = P_f(\{v, v'\}, a)$.*

It should be noted that Lemma 4 does not cover the case where $V = \emptyset$; in this case, neutrality demands that an alternative be chosen uniformly at random. The next lemma characterizes the behavior of an anonymity-proof voting rule when only a single vote is cast.

Lemma 5 *Let $v = a_1 \succ \ldots \succ a_m$. Let f be anonymity-proof and neutral, and let $p_f^i = P_f(\{v\}, a_i)$. Then, for some constant $0 \leq k_f \leq 1$, $p_f^i = k_f/m + (1 - k_f)(m - i) \cdot 2/(m(m-1))$. That is, with probability k_f the rule chooses an alternative at random, and with probability $1 - k_f$ it draws a pair of alternatives at random and chooses the preferred one.*

The final lemma characterizes the probability of a winning in the special pair of votes from Lemma 4, using Lemma 5.

Lemma 6 *Let v be a vote of the form $B \succ a \succ C \cup D$, and let v' be a vote of the form $B \cup C \succ a \succ D$. Then, if f is anonymity-proof and neutral, $P_f(\{v, v'\}, a) = k_f/m + (1 - k_f)(2|D| + |C|)/(m(m - 1))$, where k_f is defined as in Lemma 5. That* *is, the probability that a wins is the same as under the following rule for selecting the winner: with probability k_f the rule chooses an alternative at random; with probability $1 - k_f$ it draws a pair of alternatives at random, and if every vote prefers the same alternative between the two, it chooses that alternative, otherwise it flips a fair coin to decide between the two alternatives.*

Using the last three lemmas, the main result is now easy to prove. It states that any anonymity-proof neutral rule is either the rule that chooses an alternative at random, or the rule that draws two alternatives at random and runs the unanimity rule on these two alternatives, or a convex combination of these two rules.

Theorem 1 *The class of voting rules f that are anonymity-proof and neutral consists exactly of the following rules.*

- *With some probability $k_f \in [0, 1]$, the rule chooses an alternative uniformly at random.*
- *With probability $1 - k_f$ it draws a pair of alternatives uniformly at random;*
  - *If every vote prefers the same alternative between the two (and there is at least one vote), then it chooses that alternative.*
  - *Otherwise, it flips a fair coin to decide between the two alternatives.*

(All these rules are also false-name-proof in a stronger sense where the voter need not cast any vote with her true preferences, and this also implies that they are all strategy-proof.)

Proof. Let us first show that these rules indeed have the desired properties. They are clearly neutral. Conditional on a single random alternative being chosen, voters have no incentive to use false names or to not participate. Conditional on a random pair a, b of alternatives being drawn, there are four possibilities for a voter (who, without loss of generality, prefers a):

1. There are no other votes. In this case, the voter has a strict incentive to participate so that a is chosen, and no incentive to use false names.
2. All other votes prefer a. In this case, the voter has no incentive to use false names or not participate, since a will be chosen in any case.
3. All other votes prefer b. In this case, the voter has a strict incentive to participate so that at least a coin is flipped, and no incentive to use false names.
4. There are other votes that prefer a and other votes that prefer b. In this case, the voter has no incentive to use false names or not participate, since a coin will be flipped in any case.

We now show that there are no other rules with the desired properties. Let f be anonymity-proof and neutral. Lemma 5 defines k_f for this rule. Now, for an arbitrary multiset of votes V and an arbitrary alternative a, Lemma 4 shows how to convert V to a particular set of two votes $\{v, v'\}$, in a way that preserves a's probability of winning, and also preserves a's relationship to any other alternative b in the following sense:

- If all votes prefer a to b in V, the same is true in $\{v, v'\}$.

– If all votes prefer b to a in V, the same is true in $\{v, v'\}$.
– If some but not all votes prefer a to b in V, the same is true in $\{v, v'\}$.

Finally, Lemma 6 shows that for this set of two votes $\{v, v'\}$, alternative a's probability of winning is as in the claim of this theorem. Because of the preservation properties of the conversion, this must also be true for the original set of votes V.

4 Discussion

In this section, we study some corollaries of the main result, and make some comparisons to rules that are only strategy-proof.

The characterization makes it clear that the optimal anonymity-proof rule (in any reasonable sense of the word "optimal") is the one corresponding to $k_f = 0$, since this rule maximizes the probability that we can at least choose the better of two alternatives (if all votes agree). Even this rule is limited in the extent to which it can respond to the votes:

Corollary 2 *Under an anonymity-proof rule, the probability of any given alternative a winning is at most $2/m$ (for any multiset of votes). This probability is attained if and only if $k_f = 0$ and all votes rank a first.*

This is in sharp contrast to the class of strategy-proof rules. For example, it is strategy-proof to draw one of the votes at random and choose its most-preferred alternative (often referred to as the "random-dictator" rule). Under this rule, if an alternative ranks first in all votes, it will be chosen with probability 1. Also, within the class of strategy-proof rules, there is no rule that is clearly optimal. For example, it is also strategy-proof to draw a pair of alternatives at random, and choose the one that is preferred by more voters. Unlike the random-dictator rule, if there is an alternative that ranks first in all votes, this rule does not necessarily choose it; on the other hand, unlike the random-dictator rule, this rule does not run the risk of choosing an alternative that is ranked last by almost every vote (but first by a few).

Another sharp contrast between strategy-proof rules such as the above two and any anonymity-proof rule is the following. For the winning alternative not to be chosen uniformly at random, anonymity-proof rules require complete agreement on at least one pair of alternatives:

Corollary 3 *If V and a are such that for any $b \neq a$, there is a vote in V that prefers a to b, as well as one that prefers b to a, then for any anonymity-proof voting rule, $P_f(V, a) = 1/m$.*

5 Extension: group strategy-proofness

A stronger notion than strategy-proofness is *group strategy-proofness*. A mechanism is group strategy-proof if there is never a coalition of agents that can jointly misreport their preferences so that they are all better off. An analogous result to Gibbard's characterization of strategy-proof voting rules has been given for group strategy-proofness [Barbera, 1979b].

Neither of group strategy-proofness and anonymity-proofness implies the other. For example, with two alternatives, the majority rule is group strategy-proof. On the other hand, as it turns out, not all of the rules in Theorem 1 are group strategy-proof. The following theorem shows how the characterization in this paper changes if group strategy-proofness is added as a requirement.

Theorem 2 *The class of voting rules f that are anonymity-proof, group strategy-proof, and neutral consists exactly of the following rules.*

- *For two alternatives, the rules that satisfy the conditions are the same as in Theorem 1.*
- *For three or more alternatives, only the rule that chooses an alternative uniformly at random satisfies the conditions.*

Proof. For two alternatives, under any of the rules from Theorem 1, to increase the probability of one alternative winning, it is necessary to get some of the voters that prefer the other alternative to change their votes—but of course they have no incentive to do so. Hence, these rules are group strategy-proof.

For three or more alternatives, all we need to show is that if $k_f < 1$, then the rule is not group strategy-proof. (The $k_f = 1$ rule is group strategy-proof because it completely ignores the votes.) For three alternatives, consider the following profile of preferences: voter one prefers $a \succ b \succ c$, with utilities $3, 1, 0$, respectively; voter two prefers $c \succ b \succ a$, also with utilities $3, 1, 0$, respectively. If both voters vote truthfully, then there is no agreement on any pair of alternatives, so that the winner will be chosen uniformly at random, and each voter obtains an expected utility of $4/3$. However, if the voters cast the votes $a \succ c \succ b$ and $c \succ a \succ b$ instead, then the probability that b wins is $k_f/3$, whereas the probability for each of a and c is $k_f/3 + (1 - k_f)/2$. This results in an expected utility of $1(k_f/3) + 3(k_f/3 + (1 - k_f)/2) = 3/2 - k_f/6$ for each voter, which is strictly more than $4/3$ when $k_f < 1$. Hence the rule is not group strategy-proof. This example is easily extended to more than three alternatives (for example, by placing the additional alternatives at the bottom of each voter's preferences).

6 Future research

Although Theorem 1 completely characterizes anonymity-proof neutral voting rules, much remains to be done in future research. The most natural next direction to take is to consider settings where the space of possible preferences is restricted. It is well-known that such restrictions can introduce very satisfactory *strategy-proof* rules. For example, in many settings there is a natural order on the alternatives (*e.g.*, in political elections, we can order candidates by how far to the left of the political spectrum they are). In such a setting, a voter's preferences are said to be *single-peaked* if she always prefers alternatives that are closer to her most-preferred alternative to alternatives that are further away (when these alternatives are on the same side of the most-preferred alternative) [Black, 1948]. It is well-known that when preferences are single-peaked, choosing the most preferred alternative of the *median* voter (the voter that, if we sort the voters by their most preferred alternatives, ends up in the middle) is strategy-proof, and (if the number of voters is odd) this alternative will be preferred to any other alternative by more than half of the voters (*i.e.*, it is the *Condorcet winner*). Single-peakedness can only be of limited help for anonymity-proofness: for example, when there are only two alternatives, single-peakedness does not restrict preferences at all, so we cannot do anything more than in the general case. Specific application settings can also allow for more positive results, as has already been shown to be the case for combinatorial auctions. In a sense, such settings correspond to a very special way of restricting preferences. Other directions for future research include dropping the requirement of neutrality, and extending the result to allow voters to express indifferences.

Finally, if no good anonymity-proof mechanisms turn out to exist for a setting that we are interested in, then we need to consider other options. One natural solution is to verify agents' identities, that is, to check whether multiple preference reports came from the same agent. It is generally not necessary to verify the identities of all agents; rather, it suffices to verify those of a select few based on the submitted preference reports [Conitzer, 2007]. Another option is to suppose that each additional identifier used comes at a small cost to the manipulating agent. Much more positive results can be obtained in that setting [Wagman and Conitzer, 2008].[5] In either case, the results in this paper provide a natural starting point for analysis. A final approach is to try to stop the problem at the source and make it impossible or impractical for an agent to sign up for more than one account. It seems difficult to do so without compromising the anonymity of the Internet, though it is not inconceivable: see Conitzer [2008] for one possible approach to achieving this using memory tests (which is, for now, far from practical).

Acknowledgements

I would like to thank Hervé Moulin for a helpful discussion; audiences at two Dagstuhl workshops, the INFORMS Annual Meeting, the World Congress of the Game Theory Society, Microsoft Research, Duke University, Hebrew University, Bar-Ilan University, and the California Institute of Technology for helpful feedback; and the Alfred P. Sloan Foundation and the National Science Foundation for support (through a Research Fellowship and award number IIS-0812113, respectively).

[5] Alternatively, it has been suggested to consider the setting where each voter can cast at most some constant k votes (at no cost). This setting seems much less problematic because we can simply *expect* every voter to cast k votes, so that all the false-name manipulations cancel out. (Admittedly, things are not quite this simple, because it may be optimal for a voter to make some of her k votes different from her other votes—but it is easy to see that there is no reason to do so in, say, a majority election between two alternatives.)

Bibliography

Salvador Barbera. Majority and positional voting in a probabilistic framework. *The Review of Economic Studies*, 46(2):379–389, 1979.

Salvador Barbera. A note on group strategy-proof decision schemes. *Econometrica*, 47:637–640, 1979.

Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.

Ed H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

Vincent Conitzer. Limited verification of identities to induce false-name-proofness. *TARK*, pages 102–111, 2007.

Vincent Conitzer. Using a memory test to limit a user to one account. *AMEC*, 2008.

Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:587–602, 1973.

Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45:665–681, 1977.

Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

Roger Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, 41(1), 1979.

Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

Baharak Rastegari, Anne Condon, and Kevin Leyton-Brown. Revenue monotonicity in combinatorial auctions. *AAAI*, pages 122–127, 2007.

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

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

Liad Wagman and Vincent Conitzer. Optimal false-name-proof voting rules with costly voting. *AAAI*, pages 190–195, 2008.

Makoto Yokoo, Yuko Sakurai, and Shigeo Matsubara. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence*, 130(2):167–181, 2001.

Makoto Yokoo, Yuko Sakurai, and Shigeo Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.

Makoto Yokoo, Toshihiro Matsutani, and Atsushi Iwasaki. False-name-proof combinatorial auction protocol: Groves mechanism with submodular approximation. *AAMAS*, pages 1135–1142, 2006.

Makoto Yokoo. The characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. *IJCAI*, pages 733–742, 2003.