# Optimal False-Name-Proof Voting Rules with Costly Voting*

**Liad Wagman**
Duke University
Dept. of Economics
Durham, NC, USA
liad.wagman@duke.edu

**Vincent Conitzer**
Duke Unversity
Depts. of Computer Science and Economics
Durham, NC, USA
conitzer@cs.duke.edu

## Abstract

One way for agents to reach a joint decision is to vote over the alternatives. In open, anonymous settings such as the Internet, an agent can vote more than once without being detected. A voting rule is false-name-proof if no agent ever benefits from casting additional votes. Previous work has shown that all false-name-proof voting rules are unresponsive to agents' preferences. However, that work implicitly assumes that casting additional votes is costless. In this paper, we consider what happens if there is a cost to casting additional votes. We characterize the optimal (most responsive) false-name-proof-with-costs voting rule for 2 alternatives. In sharp contrast to the costless setting, we prove that as the voting population grows larger, the probability that this rule selects the majority winner converges to 1. We also characterize the optimal group false-name-proof rule for 2 alternatives, which is robust to coalitions of agents sharing the costs of additional votes. Unfortunately, the probability that this rule chooses the majority winner as the voting population grows larger is relatively low. We derive an analogous rule in a setting with 3 alternatives, and provide bounding results and computational approaches for settings with 4 or more alternatives.

## Introduction

In multiagent systems, a general approach for the agents to make a joint decision is for all of them to report their preferences over the alternatives; then, based on these reported preferences, an outcome is chosen according to a *mechanism* (or *rule*). One serious complication is that if this is to their benefit, self-interested agents will lie about their preferences. *Mechanism design* is the study of how to design the mechanism so that good outcomes will be chosen in spite of such self-interested behavior. By a result known as the *revelation principle* (Gibbard 1973; Green & Laffont 1977; Myerson 1979; 1981), we can, in a sense, without loss of generality focus our attention on mechanisms that choose outcomes in such a way that no agent has an incentive to lie. There are multiple ways of making this idea precise; the strongest is to require that the mechanism is *strategy-proof*, that is, each agent is always best off reporting its true preferences, no matter what those preferences are and what the other agents do.

In a general *social choice* (or *voting*) setting, there is a (usually finite) set of alternatives, and each agent reports ordinal preferences over these alternatives. For this very general setting, negative results are known: for example, the Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975) states that if there are at least three alternatives and preferences are unrestricted, then there exists no deterministic voting rule (mechanism) that is nondictatorial (more than one agent's preferences are taken into account), onto (for every alternative, there are some votes that make that alternative win), and strategy-proof. Still, there are also some positive results. For example, if there are only two alternatives, then the majority rule (choose whichever alternative receives more votes) is strategy-proof. Also, if preferences are *single-peaked* (Black 1948) (roughly, alternatives are ordered on a line and agents always prefer alternatives closer to their most-preferred alternative), then choosing the most-preferred alternative of the median voter is strategy-proof. If we assume additional structure—namely, each agent can make or receive payments and its preferences are *quasilinear*—then there are many strategy-proof mechanisms, for example, VCG mechanisms (Vickrey 1961; Clarke 1971; Groves 1973). But we will not consider mechanisms that use payments in this paper.

Unfortunately, in open, anonymous environments such as the Internet, new manipulations are possible: namely, an agent can vote multiple times without being detected. If the mechanism is such that there is never an incentive for an agent to do so, the mechanism is said to be *false-name-proof* (Yokoo, Sakurai, & Matsubara 2004).[1] False-name-proofness is a much more restrictive requirement than strategy-proofness: for example, not even the majority rule is false-name-proof. Recently, an extremely negative result has been obtained for false-name-proofness in voting

*Conitzer is supported by an Alfred Sloan Research Fellowship.


[1]As is the case for strategy-proofness, false-name-proofness is a dominant-strategies criterion, that is, using only one identifier should be optimal regardless of one's preferences and regardless of what the other agents do. A (weaker) Bayesian definition can also be given, but the dominant-strategies version is the one that has been studied so far, and the related impossibility results for misreporting preferences concern strategy-proofness (that is, a dominant-strategies criterion), so we focus on false-name-proofness in the dominant-strategies sense in this paper. This also makes our positive results stronger.

settings (Conitzer 2007a): this result implies that, unless the agents unanimously agree that one alternative is preferred to another, then the winning alternative must be chosen uniformly at random. Given this result, it would seem reasonable to give up on false-name-proofness in general social choice settings. (For combinatorial auction settings, some reasonable false-name-proof mechanisms do exist (Yokoo, Sakurai, & Matsubara 2001; Yokoo 2003; Yokoo, Matsutani, & Iwasaki 2006), but here, too, impossibility results are known (Yokoo, Sakurai, & Matsubara 2004; Rastegari, Condon, & Leyton-Brown 2007).) Indeed, it has been proposed to perform some limited verification of agents' identities to remove the incentive to use false names (Conitzer 2007b).

In this paper, we consider another way around the impossibility result for false-name-proofness in voting settings. The impossibility result implicitly assumes that there is no cost to voting multiple times; however, it is reasonable to assume that casting an additional vote comes at a small cost. For example, to cast an additional vote, an agent may need to sign up for another (say, free e-mail) account, and this requires some effort (*e.g.* most free e-mail providers require a user to solve a CAPTCHA (von Ahn *et al.* 2003; von Ahn, Blum, & Langford 2004) before obtaining an account). In this context, a rule is false-name-proof (with costs) if the cost of casting additional votes always outweighs the benefit to an agent. We characterize the optimal rule with two alternatives in this setting, and show that as the number of agents grows, the probability that this rule chooses the alternative that the majority rule *would have* chosen, *if* there had been no false-name manipulation, converges to 1—a very positive result, in sharp contrast to the costless setting. However, we also consider the possibility that multiple agents collude and share the cost of the additional votes. We characterize the resulting optimal *group false-name-proof* mechanism and show that unfortunately, it does not have the convergence property. We also obtain some results for a 3-alternative setting, and propose a linear programming technique for settings with more alternatives, in the spirit of automated mechanism design (Conitzer & Sandholm 2002).

## Definitions

In this section, we consider a setting with 2 alternatives, $A$ and $B$. We assume that each agent strictly prefers one of the alternatives; which alternative is preferred is private information to the agent. We normalize each agent's utility function so that the agent receives utility $1$ if its favorite alternative is selected, and $0$ otherwise. Votes can be cast either for $A$ or for $B$. We assume that each agent can cast one vote for free; for each additional vote, the agent incurs a cost of $c$. This model can be generalized in several ways without changing the results; we discuss this in a later section. For simplicity of presentation we focus on the simpler model first. This model is reasonable, for example, when every agent already has an account (for other reasons), but signing up for an additional account is costly (either because there is a monetary cost, or because some effort needs to be exerted, e.g. a CAPTCHA needs to be solved to obtain the account). We will only consider voting rules that are *anonymous*, that is, each vote is treated in the same way. Hence, we need only keep track of the number of votes for each alternative.

**Definition 1 (State)** *A* state *consists of a pair $(x_A, x_B)$, where $x_j \geq 0$ is the number of votes for $j \in \{A, B\}$.*

Voting rules are defined over states. In this paper, we allow for randomized voting rules.

**Definition 2 (Voting rule)** *A* voting rule *is a mapping from the set of states to the set of probability distributions over outcomes. The probability that alternative $j \in \{A, B\}$ is selected in state $(x_A, x_B)$ is denoted by $P_j(x_A, x_B)$; we must have $P_A(x_A, x_B) + P_B(x_A, x_B) = 1$.*

A voting rule is *neutral* if alternatives are treated in the same way. We will only consider neutral rules in this paper.

**Definition 3 (Neutrality)** *A voting rule is neutral if $P_A(x, y) = P_B(y, x)$ for all $(x, y)$.*

Let $(x_A, x_B)$ be the state that results from all the votes with the exception of agent $i$'s vote(s). Let $t^i_A$ and $t^i_B$ denote the number of times that agent $i$ votes for alternatives $A$ and $B$, respectively. Then, if agent $i$ prefers alternative $j$, $i$'s expected utility is $u_i(x_A, x_B, t^i_A, t^i_B) = P_j(x_A + t^i_A, x_B + t^i_B) - (t^i_A + t^i_B - 1)c$. We now define some standard concepts from mechanism design. In these definitions we only consider alternative $A$, but the analogous statement for $B$ follows from neutrality.

**Definition 4 (Voluntary participation)** *A (neutral) voting rule satisfies* voluntary participation *if for an agent $i$ who prefers $A$, for all $(x_A, x_B)$, $u_i(x_A, x_B, 1, 0) \geq u_i(x_A, x_B, 0, 0)$.*

**Definition 5 (Strategy-proofness)** *A (neutral) voting rule is* strategy-proof *if for an agent $i$ who prefers $A$, for all $(x_A, x_B)$, $u_i(x_A, x_B, 1, 0) \geq u_i(x_A, x_B, 0, 1)$.*

**Definition 6 (False-name-proofness)** *A (neutral) voting rule is* false-name-proof (with costs) *if for an agent $i$ who prefers $A$, for all $(x_A, x_B)$, for all $t^i_A \geq 1$ and $t^i_B$, $u_i(x_A, x_B, 1, 0) \geq u_i(x_A, x_B, t^i_A, t^i_B)$.*

In other words, a rule is false-name-proof if, given that an agent casts a vote for its most preferred alternative, it does not benefit from casting additional votes.

**Lemma 1** *A (neutral) rule satisfies voluntary participation and false-name proofness if and only if for all $x_A, x_B \geq 0$,*

1. $P_A(x_A + 1, x_B) - P_A(x_A, x_B) \geq 0$,
2. $P_A(x_A, x_B) - P_A(x_A, x_B + 1) \geq 0$,
3. $P_A(x_A + 2, x_B) - P_A(x_A + 1, x_B) \leq c$, *and*
4. $P_A(x_A, x_B + 1) - P_A(x_A, x_B + 2) \leq c$.

*Such rule is also strategy-proof.*

We omit the proof of Lemma 1 due to space constraint. If $c = 0$, then the optimal (in a sense to be made precise later) neutral false-name-proof rule that satisfies voluntary participation is the following "unanimity" rule:

- If one alternative gets all the votes, select it.
- Otherwise, select an alternative uniformly at random.

The disadvantage of this rule is clear: even if one alternative receives 100 votes and the other, 1 vote, then a coin is flipped to determine the winner. That is, the rule is not very *responsive* to votes. In some sense, the "most" responsive rule is the majority rule, which chooses the alternative that receives more votes (and flips a coin if there is a tie), thereby maximizing the sum of the utilities.[2] However, the majority rule is not false-name-proof. As we will see shortly, when $c > 0$, there are false-name-proof rules that are more responsive (more like majority) than the unanimity rule above. Our objective is to maximize responsiveness under the constraint of false-name-proofness.

One may wonder how we should compare two rules if one is more responsive for some states, and the other is more responsive for other states. However, this turns out not to matter, because we will find a rule that is *strongly optimal*, that is, most responsive for all states.

**Definition 7 (Strong optimality)** *A neutral false-name-proof voting rule $P$ that satisfies voluntary participation is* strongly optimal *if for any other neutral false-name-proof voting rule $\widetilde{P}$ that satisfies voluntary participation, for any state $(x_A, x_B)$ where $x_A \geq x_B$, we have $P_A(x_A, x_B) \geq \widetilde{P}_A(x_A, x_B)$.*

It should be noted that by neutrality, it follows that for such a rule we also have that if $x_A \leq x_B$, then $P_A(x_A, x_B) \leq \widetilde{P}_A(x_A, x_B)$. Also, there cannot exist two different strongly optimal rules.

## The optimal false-name-proof rule

We are now ready to present the (strongly) optimal false-name-proof rule with 2 alternatives, $FNP2$.

**Definition 8 ($FNP2$)** *Suppose without loss of generality that $x_A \geq x_B$. Rule $FNP2$ sets $P_A(x_A, x_B) = 1$ if $x_A > x_B = 0$, and $P_A(x_A, x_B) = \min\{1, \frac{1}{2} + c(x_A - x_B)\}$ if $x_A \geq x_B > 0$ or $x_A = x_B = 0$.*

For example, let us consider $FNP2$ over the states $(x_A, x_B)$, $x_A, x_B \leq 5$, when $c = 0.15$:

| | | | | | | |
|---|---|---|---|---|---|---|
| 5 | 0 | 0 | .05 | .2 | .35 | .5 |
| 4 | 0 | .05 | .2 | .35 | .5 | .65 |
| 3 | 0 | .2 | .35 | .5 | .65 | .8 |
| 2 | 0 | .35 | .5 | .65 | .8 | .95 |
| 1 | 0 | .5 | .65 | .8 | .95 | 1 |
| 0 | .5 | 1 | 1 | 1 | 1 | 1 |
| $x_B/x_A$ | 0 | 1 | 2 | 3 | 4 | 5 |

From Definition 8, it is easy to see that rule $FNP2$ satisfies strategy proofness: an agent can never increase the probability of its preferred alternative being selected by voting for its less preferred alternative. Neutrality also follows directly, since the labelling of alternatives is irrelevant to $FNP2$, and since $P_A(x, x) = 1/2$ for any $x \geq 0$. As the following theorem proves, not only does $FNP2$ satisfy the properties of neutrality, strategy proofness, and false-name proofness, but it is also uniquely strongly optimal.

**Theorem 1** *Rule $FNP2$ is the unique strongly optimal neutral false-name-proof voting rule with 2 alternatives that satisfies voluntary participation.*

**Proof**: The rule is neutral by definition. We first prove that it satisfies the conditions of Lemma 1. We begin with Condition 1. If $x_A \geq x_B > 0$, then $P_A(x_A + 1, x_B) - P_A(x_A, x_B) = \min\{1, \frac{1}{2} + c(x_A + 1 - x_B)\} - \min\{1, \frac{1}{2} + c(x_A - x_B)\} \geq 0$. If $x_B > x_A > 0$, then $P_A(x_A + 1, x_B) - P_A(x_A, x_B) = (1 - P_B(x_A + 1, x_B)) - (1 - P_B(x_A, x_B)) = P_B(x_A, x_B) - P_B(x_A + 1, x_B) = \min\{1, \frac{1}{2} + c(x_B - x_A)\} - \min\{1, \frac{1}{2} + c(x_B - x_A - 1)\} \geq 0$. Finally, if $x_A = 0$ or $x_B = 0$, it is easy to check that $P_A(x_A + 1, x_B) - P_A(x_A, x_B) \geq 0$. This proves Condition 1; Condition 2 follows by symmetry.

Next, we prove Condition 3. If $x_A + 1 \geq x_B > 0$, then $P_A(x_A + 2, x_B) - P_A(x_A + 1, x_B) = \min\{1, \frac{1}{2} + c(x_A + 2 - x_B)\} - \min\{1, \frac{1}{2} + c(x_A + 1 - x_B)\} \leq c$. If $x_B \geq x_A + 2$, then $P_A(x_A + 2, x_B) - P_A(x_A + 1, x_B) = (1 - P_B(x_A + 2, x_B)) - (1 - P_B(x_A + 1, x_B)) = P_B(x_A + 1, x_B) - P_B(x_A + 2, x_B) = \min\{1, \frac{1}{2} + c(x_B - x_A - 1)\} - \min\{1, \frac{1}{2} + c(x_B - x_A - 2)\} \leq c$. Finally, if $x_B = 0$, then $P_A(x_A + 2, x_B) - P_A(x_A + 1, x_B) = 1 - 1 = 0$. This proves Condition 3; Condition 4 follows by symmetry. Hence, $FNP2$ satisfies the conditions of Lemma 1, so it is in fact a neutral false-name-proof voting rule that satisfies voluntary participation.

All that remains to show is strong optimality, that is, for any other such rule $\widetilde{P}$, for any state $(x_A, x_B)$ with $x_A \geq x_B$, $P_A(x_A, x_B) \geq \widetilde{P}_A(x_A, x_B)$, where $P$ is $FNP2$. (We recall that the analogous statement when $x_B \geq x_A$ follows by symmetry.) Neutrality requires that for any $x \geq 0$, $\widetilde{P}_A(x, x) = 1/2$. Next, for any $x > 0$, false-name proofness requires that $\widetilde{P}_A(x + 1, x) - \widetilde{P}_A(x, x) \leq c$, so that $\widetilde{P}_A(x + 1, x) \leq 1/2 + c$. Similarly, $\widetilde{P}_A(x + 2, x) - \widetilde{P}_A(x + 1, x) \leq c$, so that $\widetilde{P}_A(x + 2, x) \leq \widetilde{P}_A(x + 1, x) + c \leq 1/2 + 2c$. Continuing in the same manner, for any $t > 0$, $\widetilde{P}_A(x + t, x) \leq 1/2 + tc$ must hold. Also, naturally, $\widetilde{P}_A(x + t, x) \leq 1$. So $\widetilde{P}_A(x + t, x) \leq \min\{1, 1/2 + tc\}$. But $P_A(x + t, x) = \min\{1, 1/2 + tc\}$. Hence $FNP2$ is strongly optimal. ∎

By Lemma 1, $FNP2$ is also strategy-proof.

### Responsiveness in the limit

In this section, we investigate the following question: as $n$, the number of agents, goes to infinity, is it more likely that $FNP2$ will choose the majority winner? The "majority winner" here refers to the winner that the majority rule *would have* produced if every agent voted exactly once. Of course, if we actually used the majority rule, agents would likely use false names. Ideally, $FNP2$ would usually choose the majority winner, in which case we get the best of both worlds: false-name-proofness and high responsiveness.

To make this concrete, suppose agents' preferences are drawn i.i.d.: with probability $p$ an agent prefers $A$, with $1-p$,

[2] One can argue about the precise definition of responsiveness. For example, the rule that chooses $A$ if the total number of votes is odd and $B$ otherwise is more "responsive" in the sense that each additional vote changes the outcome. However, such rules violate neutrality and voluntary participation.

$B$. If $c = 0$, we need to use the aforementioned unanimity rule; now, as $n$ grows, the probability of all agents agreeing goes to $0$, and the probability of flipping a coin goes to $1$. Hence, the rule becomes completely unresponsive as $n \to \infty$. But what if $c > 0$? Here, we get the opposite result:

**Theorem 2** *As $n \to \infty$, the probability that $FNP2$ chooses the majority winner converges to* $1$.

We omit the remaining proofs due to space constraint, but the intuition for this result is simple. $FNP2$ randomizes only if the election is close to tied, that is, $|x_A - x_B| < 1/(2c)$. However, as $n$ increases, the binomial distribution over agent preferences converges to a flatter and flatter normal distribution. Hence, the probability that the election is close to tied goes to $0$. When $FNP2$ does not randomize, it chooses the majority winner.

Of course, Theorem 2 only shows what happens in the limit. Figure 1 illustrates how often $FNP2$ fails to choose the majority winner as a function of $n$.[3] The probability that $FNP2$ fails to choose the majority winner also depends on $c$ (the higher $c$, the less often it will do so), as illustrated in Figure 2. We note that with $c \geq 1/2$, the majority winner is always chosen. Finally, figure 3 illustrates how often $FNP2$ fails to choose the majority winner as a function of $p$.

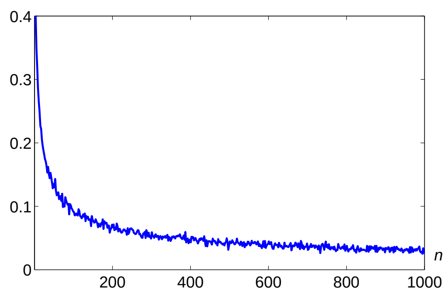


Figure 1: The average probability that $FNP2$ and majority disagree as a function of $n$ (for fixed $c = .1$, $p = 1/2$, and $n$ odd).

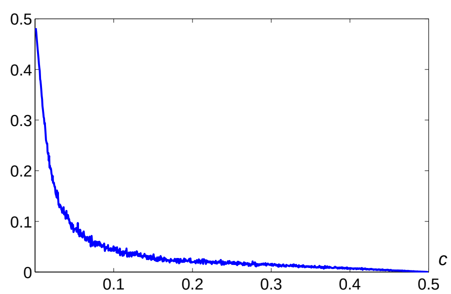


Figure 2: The average probability that $FNP2$ and majority disagree as a function of $c$ (for fixed $n = 501$ and $p = 1/2$).

## Group false-name proofness

In this section, we consider a stronger notion of false-name proofness. So far, we have only considered the possibility of a single agent casting multiple votes; in that case, the burden (cost) of these votes must be assumed by that single agent. However, other agents may benefit from such a manipulation. Hence, it may be possible for a coalition of such agents

[3]In simulating Figures 1-3, we use $1500$ trials per parameter triplet $(c, p, n)$. When $n$ is even, $FNP2$ performs better.

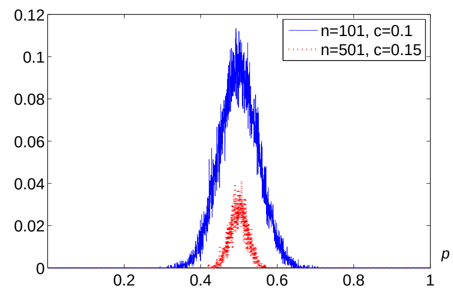


Figure 3: The average probability that $FNP2$ and majority disagree as a function of $p$ (for fixed $c$ and $n$).

to share the cost of this false-name manipulation—either implicitly (each agent in the coalition casts some of the additional votes) or explicitly (say, one agent casts all the additional votes, but the other agents in the coalition compensate that agent). This requires that the manipulating agents are able to commit to such a plan.

We call a rule in which coalitions never benefit from such a manipulation *group false-name-proof* (analogously to group strategy-proof). It is easy to see that $FNP2$ is not group false-name-proof: in the example with $c = .15$ given above, if $x_A = x_B = 2$ (for the true preferences), then the two agents that prefer $A$ have an incentive to make a contract to each cast an additional vote for $A$, resulting in a probability of $.8$ that $A$ wins, so that each of these agents is $.3 - .15 = .15$ better off. Alternatively, one of them can cast a single additional vote, and the other can compensate the first one $.075$; then the first agent is $.15 - .15 + .075 = .075$ better off, and the second one is $.15 - .075 = .075$ better off as well. (We assume quasilinear preferences.)

In this section, we only consider the variant where the manipulating agents are able to commit to transfers among themselves, so that they can share the cost equally. This is equivalent to the variant where they cannot make transfers, but commit to stochastic manipulations (*e.g.* the manipulating agents flip a coin over who casts the additional vote).

**Definition 9 (Group false-name-proofness)** *A (neutral) rule is* group false-name-proof (with costs and transfers) *if for all $k \geq 1$, for all $(x_A, x_B)$, for all $t_A \geq k$ and $t_B$, $P_A(x_A{+}k, x_B) \geq P_A(x_A{+}t_A, x_B{+}t_B) - c(t_A{+}t_B{-}k)/k$.*

The following lemma is analogous to Lemma 1.

**Lemma 2** *A (neutral) rule satisfies voluntary participation and group false-name proofness if and only if for all $x_A, x_B \geq 0$,*

1. $P_A(x_A + 1, x_B) - P_A(x_A, x_B) \geq 0$,
2. $P_A(x_A, x_B) - P_A(x_A, x_B + 1) \geq 0$,
3. $P_A(x_A + 2, x_B) - P_A(x_A + 1, x_B) \leq c/(x_A + 1)$, *and*
4. $P_A(x_A, x_B + 1) - P_A(x_A, x_B + 2) \leq c/(x_B + 1)$.

*Such rule is also group strategy-proof and false-name-proof.*

We now present the strongly optimal group false-name-proof rule (strong optimality is defined similarly as before).

**Definition 10** *[$GFNP2$] Suppose without loss of generality that $x_A \geq x_B$. Rule $GFNP2$ sets $P_A(x_A, x_B) = 1$ if $x_A > x_B = 0$, $P_A(x_A, x_B) = 1/2$ if $x_A = x_B = 0$, and $P_A(x_A, x_B) = \min\{1, 1/2 + \sum_{k=x_B}^{x_A-1}(\frac{c}{k})\}$ if $x_A \geq x_B > 0$.*

For example, let us consider $GFNP2$ over the states $(x_A, x_B)$, $x_A, x_B \leq 5$, when $c = 0.15$:

| | | | | | | |
|---|---|---|---|---|---|---|
| 5 | 0 | .8125 | .6625 | .5875 | .5375 | .5 |
| 4 | 0 | .775 | .625 | .55 | .5 | .5375 |
| 3 | 0 | .725 | .575 | .5 | .55 | .5875 |
| 2 | 0 | .65 | .5 | .575 | .625 | .6625 |
| 1 | 0 | .5 | .65 | .725 | .775 | .8125 |
| 0 | .5 | 1 | 1 | 1 | 1 | 1 |
| $x_B/x_A$ | 0 | 1 | 2 | 3 | 4 | 5 |

**Theorem 3** *Rule $GFNP2$ is the strongly optimal neutral group false-name-proof voting rule with 2 alternatives that satisfies voluntary participation.*

### Responsiveness in the limit

If $c$ is sufficiently large (at least $1/2$), then $FNP2$ coincides with the majority rule. However, there is no (finite) $c$ such that $GFNP2$ coincides with the majority rule, because for any $c$ there there exists some sufficiently large $k$ such that a coalition of size $k$ would like to group false-name-manipulate the majority rule in (for example) tied states.

We recall that if preferences are drawn i.i.d., then as $n \to \infty$, we will almost always choose the majority winner under $FNP2$. However, for $GFNP2$, this turns out not to (always) be the case. We note that the highest probability with which a neutral rule that satisfies voluntary participation fails to select the majority winner is $1/2$ (*e.g.* the unanimity rule). Recall that $p$ is the probability that an agent prefers alternative $A$.

**Theorem 4** *Let $w \in (0, 1/2)$. If $p \in [\frac{2c}{4c-2w+1}, 1 - \frac{2c}{4c-2w+1}]$, then as $n \to \infty$, $GFNP2$ fails to select the majority winner with probability at least $w$.*

For example, if $c = .1$ and $p \in [1/3, 2/3]$, Theorem 4 states that as $n \to \infty$, $GFNP2$ yields the opposite outcome from the majority rule at least $40\%$ of the time. If $c = .01$ and $p \in [1/3, 2/3]$, this becomes $49\%$. If $c = .01$ and $p \in [10/21, 11/21]$, this becomes $49.9\%$. Finally, if $p = 1/2$, Theorem 4 states that as $n \to \infty$, $GFNP2$ yields the opposite outcome from the majority rule $50\%$ of the time, implying that it arbitrarily selects an alternative. We omit figures due to space constraint.

## 3 alternatives

We now move on the case of 3 alternatives. Here, we assume that each agent has a utility of 1 for their most preferred alternative, and 0 for the other alternatives. Hence, each agent simply votes for their most preferred alternative. While this is without loss of generality in the 2-alternative case, it is not so here. With this assumption, reasonable strategy-proof rules are possible, for example, the plurality rule (choose the alternative with the most votes). Without this assumption, no reasonable strategy-proof rules exist (Gibbard 1973; 1977; Satterthwaite 1975). We only study false-name-proofness, not group false-name-proofness.

For technical reasons, we make one more assumption in the remainder of this paper: every alternative receives at least one vote. For example, if the alternatives are political candidates, presumably they would vote for themselves.

We now generalize strong optimality.

**Definition 11 (Strong optimality)** *A neutral, strategy-proof, and false-name-proof voting rule $P$ that satisfies voluntary participation is* strongly optimal *if for any other neutral, strategy-proof, and false-name-proof voting rule $\widetilde{P}$ that satisfies voluntary participation, for any state $(x_A, x_B, x_C)$ where $x_A \geq x_B \geq x_C \geq 1$, either $P_A(x_A, x_B, x_C) > \widetilde{P}_A(x_A, x_B, x_C)$; or $P_A(x_A, x_B, x_C) = \widetilde{P}_A(x_A, x_B, x_C)$ and $P_B(x_A, x_B, x_C) \geq \widetilde{P}_B(x_A, x_B, x_C)$.*

(We emphasize that we are restricting attention to the case where every alternative receives at least one vote.) It is not hard to see that if $c \geq \frac{2}{3}$, the plurality rule is the strongly optimal voting rule.

**Definition 12 ($FNP3$)** *Suppose without loss of generality that $x_A \geq x_B \geq x_C \geq 1$. Rule $FNP3$ is defined as follows. $P_A(x_A, x_B, x_C) = \min\{1, \frac{1}{2} + c(x_A - x_B) - \frac{1}{2}\max\{0, \frac{1}{3} - c(x_B - x_C)\}\}$, $P_C(x_A, x_B, x_C) = \max\{0, \frac{1}{3} - c(\frac{x_A+x_B}{2} - x_C)\}$, $P_B(x_A, x_B, x_C) = 1 - P_A(x_A, x_B, x_C) - P_C(x_A, x_B, x_C)$.*

**Theorem 5** *$FNP3$ is the strongly optimal neutral strategy-proof false-name-proof voting rule with 3 alternatives that satisfies voluntary participation.*

The following lemma provides some intuition about $FNP3$.

**Lemma 3** *Under rule $FNP3$, $P_j(x_A, x_B, x_C) = P_j(x_A - x_C + 1, x_B - x_C + 1, 1)$ for all $j \in \{A, B, C\}$.*

Lemma 3 allows us to represent rule $FNP3$ on a two-dimensional grid, because we only need to consider $P_A(x_A, x_B, 1)$. For example, following is $P_A(x_A, x_B, 1)$ under rule $FNP3$ when $c = 0.2$ and $1 \leq x_A, x_B \leq 6$:

| | | | | | | |
|---|---|---|---|---|---|---|
| 6 | 0 | 0 | 0 | .10 | .30 | .50 |
| 5 | 0 | 0 | .10 | .30 | .50 | .70 |
| 4 | .03 | .17 | .30 | .50 | .70 | .90 |
| 3 | .13 | .33 | .50 | .70 | .90 | 1 |
| 2 | .23 | .43 | .63 | .83 | 1 | 1 |
| 1 | .33 | .53 | .73 | .93 | 1 | 1 |
| $x_B/x_A$ | 1 | 2 | 3 | 4 | 5 | 6 |

Interestingly, under $FNP3$, sometimes a vote for one alternative increases the winning probability of another alternative (but not enough to violate strategy-proofness)—for example, $P_B(4, 2, 2) > P_B(4, 2, 1)$ when $c = .3$.

## 4+ alternatives

Unfortunately, we were unable to generalize the strongly optimal rule to $k \geq 4$ alternatives. We can, however, obtain an upper bound on the probability of choosing the plurality winner that must hold for any false-name-proof rule. We continue to assume that agents strictly prefer one of the alternatives and equally dislike all other alternatives.

**Procedure 1 (Upper bound)** *Let $(x_1, ..., x_m)$, where $m \geq 2$, denote the state, such that $x_1 \geq x_2 \geq ... \geq x_m \geq 1$. An upper bound $B_1(x_1, ..., x_m)$ on $P_1(x_1, ..., x_m)$ can be derived using the following recursion.*

1. *Base condition:*[4]

[4] $(x_k, ..., x_k, x_{k+1}, ..., x_m)$, where $k \in \{1, ..., m-1\}$, is the state where the first $k$ alternatives receive $x_k$ votes, and alternatives $k+1, ..., m$ receive $x_{k+1}, ..., x_m$ votes, respectively.

$B_m(x_{m-1}, ..., x_{m-1}, x_m) = \max\{0, \frac{1}{m} - c(x_{m-1} - x_m)\}$

2. *For* $k \in \{1, ..., m-1\}$, $B_k(x_k, ..., x_k, x_{k+1}, ..., x_m) = \frac{1}{k}\big(1 - \max\{0, B_{k+1}(x_{k+1}, ..., x_{k+1}, x_{k+2}, ..., x_m) - c(x_k - x_{k+1})\}\big)$.

Procedure 1 is a recursive application of the following observations. By neutrality, the first $k+1$ alternatives are selected with the same probability at any state $(x_{k+1}, ..., x_{k+1}, x_{k+2}, ..., x_m)$. It follows from false-name proofness that at a state $(x_k, ..., x_k, x_{k+1}, x_{k+2}, ..., x_m)$, where $x_k \geq x_{k+1}$, $P_{k+1}(x_k, ..., x_k, x_{k+1}, x_{k+2}, ..., x_m) \geq P_{k+1}(x_{k+1}, ..., x_{k+1}, x_{k+2}, ..., x_m) - c(x_k - x_{k+1})$.

## A general linear programming approach

While we were unable to give a general characterization of the optimal rule for 4+ alternatives, in this section, we do propose a linear programming approach for finding an optimal false-name-proof voting rule given a specific value of $c$, an upper bound $Z$ on the number of votes for each alternative, and a prior distribution $\pi$ over states. We continue to assume that each alternative receives at least one vote. We show how to find the rule that minimizes the expected distance to plurality (or any other rule $P'$). (Standard tricks can be used to linearize the absolute value operator.)

**Procedure 2 (Linear program)**

**Minimize** $\sum_{x_1=1}^{Z} \cdots \sum_{x_m=1}^{Z} \pi(x_1, \ldots, x_m)\big[ \sum_{k=1}^{m} |P_k(x_1, \ldots, x_m) - P'_k(x_1, \ldots, x_m)|\big]$ **subject to**

1. *Participation:* $\forall x_1, \ldots, x_m \in \{1, \ldots, Z\}$, $\forall k \in \{1, \ldots, m\}$, $P_k(x_1, \ldots, x_{k-1}, x_k+1, x_{k+1}, \ldots, x_m) - P_k(x_1, \ldots, x_m) \geq 0$.
2. *Neutrality:* $\forall x_1, \ldots, x_m \in \{1, ..., Z\}$, *for any permutation* $(w_1, \ldots, w_m)$ *of* $(x_1, \ldots, x_m)$, $\forall j, k \in \{1, \ldots, m\}$, *if* $w_j = x_k$ *then* $P_j(w_1, \ldots, w_m) = P_k(x_1, \ldots, x_m)$.
3. *Strategy proofness:* $\forall x_1, \ldots, x_m \in \{1, \ldots, Z\}$, $\forall j, k \in \{1, \ldots, m\}$, $P_k(x_1, \ldots, x_{k-1}, x_k+1, x_{k+1}, \ldots, x_m) \geq P_k(x_1, \ldots, x_{j-1}, x_j+1, x_{j+1}, \ldots, x_m)$.
4. *False-name proofness:* $\forall x_1, \ldots, x_m \in \{1, \ldots, Z\}$, $\forall k \in \{1, \ldots, m\}$, $P_k(x_1, \ldots, x_{k-1}, x_k+1, x_{k+1}, \ldots, x_m) - P_k(x_1, \ldots, x_m) \leq c$.

We note that because we require strategy-proofness, it never makes sense to cast a false-name vote for another alternative, which simplifies constraint (4).

## Extensions and future work

In the above, we have assumed that all agents have the same cost $c$ for casting an additional vote. In fact, all that is needed for all of the above results to go through is that we know that each agent has a cost of *at least* $c$ for each additional vote (and this is the greatest lower bound that we know). In fact, a single agent can have different costs for, say, her first and second additional votes, as long as we know the cost for each to be at least $c$. Additionally, we have assumed the first vote is free (for example, everyone already has one e-mail account). The analysis can be extended to settings where an agent's first vote is costly. However, in such settings, voluntary participation cannot be satisfied.

Future work can take on a number of directions. An immediate direction is to more robustly extend our results to settings with 3 or more alternatives. This can also be done under different assumptions. For instance, one could derive the optimal false-name-proof rule that does not necessarily satisfy voluntary participation or strategy-proofness. One could also consider *dichotomous* preferences (Inada 1964), for which responsive strategy-proof rules exist. (Under such preferences, each voter (equally) approves of a set of alternatives and (equally) disapproves of the remaining alternatives.) Another direction is to extend the group false-name-proofness results to the setting where agents cannot use transfers and where only deterministic contracts are allowed. Yet another direction is to consider weaker (e.g. Bayes-Nash equilibrium) notions of false-name-proofness.

## References


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

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

Conitzer, V., and Sandholm, T. Complexity of mechanism design. *UAI*, 103–110, 2002.

Conitzer, V. Anonymity-proof voting rules. In *Computational Social Systems and the Internet #07271, Dagstuhl Workshop*, 2007.

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

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

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

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

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

Inada, K. A note on the simple majority decision rule. *Econometrica* 32(4):525–531, 1964.

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

Myerson, R. Optimal auction design. *Math of OR* 6:58–73, 1981.

Rastegari, B.; Condon, A.; and Leyton-Brown, K. Revenue monotonicity in combinatorial auctions. *AAAI*, 2007.

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

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

von Ahn, L.; Blum, M.; Hopper, N.; and Langford, J. CAPTCHA: Using hard AI problems for security. *EUROCRYPT*, 2003.

von Ahn, L.; Blum, M.; and Langford, J. Telling humans and computers apart automatically: How lazy cryptographers do AI. *Comm. of the ACM* 47(2):56–60, 2004.

Yokoo, M.; Matsutani, T.; and Iwasaki, A. False-name-proof combinatorial auction protocol: Groves mechanism with submodular approximation. *AAMAS*, 2006.

Yokoo, M.; Sakurai, Y.; and Matsubara, S. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence* 130(2):167–181, 2001.

Yokoo, M.; Sakurai, Y.; and Matsubara, S. The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *GEB* 46(1):174–188, 2004.

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