N. R. Miller
05/01/97
5th rev. 8/22/06

VOTING TO ELECT A SINGLE CANDIDATE

This discussion focuses on *single-winner elections*, in which a single candidate is elected from a field of two or more candidates. Thus it pertains to legislative elections within *single-member districts* (one representative per district) and to the election of *executive officials* (President, Governor, etc.) under a separation-of-powers (as opposed to a parliamentary) system.

The simplest single-winner election occurs when there are precisely two candidates, producing what the British call a *straight fight*. In this case, voting by *Simple Majority Rule* (SMR) strikes most people as fair and reasonable. Each voter votes for one or other candidate (or perhaps abstains), and — apart from the possibility of a tie — one candidate must receive an (absolute) majority of votes cast and that candidate is elected. (In the event of a tie, maybe we flip a coin.)

A mathematician by the name of Kenneth May (1952) demonstrated that SMR, and only *SMR*, meets four conditions that we may want a voting rule to meet in a straight fight between two candidates (or other alternatives) *A* and *B*. Each voter can: (i) vote for A, (ii) vote for B, or (iii) abstain. The voting outcomes are: (i) A wins, (ii) B wins, and (iii) deadlock. Here are May's conditions (which I have slightly reformulated).

> *Anonymity* (of votes). We do not need to know who cast which vote to determine the winner. In other words, all votes (and voters) are treated the same way.
>
> *Neutrality* (between candidates). If every vote for *A* becomes a vote for *B* and vice versa, the winning and losing candidates are reversed (or remain deadlocked). In other words, the two candidates are treated the same way.

A condition that is appealing on practical grounds is the following.

> *Resoluteness*. Regardless of how votes are cast, deadlock is avoided and there is always a winner. In other words, we can't have a "hung electorate" (in the manner of a "hung jury" using unanimity rule).

However, it should be clear the three conditions we have identified are inconsistent. In the event that an even number of (non-abstaining) voters have equally divided preferences between two candidates, Anonymity and Neutrality together require the kind of symmetric deadlock (i.e., a tie) that resoluteness rules out. So we must weaken resoluteness as follows.[1]

> *Almost Resoluteness*. Regardless of how votes are cast, deadlock is almost always avoided and there is almost always a winner. More precisely, if deadlock does occur, it is removed if any voter changes his or her vote in any fashion (from one candidate

[1] On the other hand, Resoluteness can be preserved if we weaken Anonymity to allow one voter to have a tie-breaking "casting" vote or weaken Neutrality to give one candidate victory in the event of a tie.

to the other, or from abstention to either candidate, or from either candidate to abstention). In other words, deadlock is a "knife-edge" condition.

Non-Negative Responsiveness. In the event of a deadlock between candidates A and B, if a voter then switches his vote in A's favor (i.e., from *B* to *A*, or from abstain to *A*, or from *B* to abstain), *A* remains the winner. In other words, votes never count "negatively" if they count at all.

May demonstrated that SMR meets these four conditions and is the only voting rule that can do so. Moreover, when SMR is used in a straight fight, no voter ever has reason to consider voting other than for his or her more preferred candidate. That is, we can expect all voting to be *sincere* and to "honestly" represent voters' preferences. Put more formally, in a straight fight SMR is *strategyproof*; in that no voter can ever improve the outcome with respect to his or her true preferences by misreporting those preferences on the ballot.

But once the number of candidates expands to three or more, all sorts of problems arise. First of all, many different apparently fair and reasonable voting rules (including those discussed below, along with other more esoteric possibilities) are available (and many are in actual use). Each such procedure reduces to SMR in the two-candidate case, but different such procedures often select different winners in the multi-candidate case. Moreover, all such voting rules have evident flaws. Indeed, two important flaws are essentially unavoidable in elections involving three or more candidates.

First, as noted above, SMR (among other procedures) is strategyproof in the two-candidate case. But *no voting procedure whatsoever is strategyproof given three or more candidates*.

Second, all voting procedures are vulnerable to *spoiler effects* when the field of candidates expands or contracts — that, whether candidate *A* or *B* is elected may depend on whether some third candidate (the potential "spoiler") enters the field or not.[2] (This fact provides an argument in favor a two-party system that makes most elections straight fights.)

Note that, when we have three or more candidates, a voter's preferences are not specified simply by listing a most preferred (top-ranked) candidate; rather we must specify the voter's full *preference ordering* to all candidates in the field, i.e., a first preference, second preference, etc. (We will simplify the discussion by assuming that voters are never indifferent between any candidates.) A collection of preference orderings for all voters is called a *preference profile*.

Here is an example to focus on. We use British party labels to identify three candidates — Labour, Liberal, and Conservative — one of whom is to be elected. While there are six possible orderings of three candidates, we first consider a simple profile in which only three of these orderings are present and we indicate the popularity of each.

---

[2] For example, almost certainly Gore would have won Florida's electoral votes (and the Presidency) if Nader had not been on the ballot in Florida.

Preference Profile 1

# of voters	46	20	34
1st pref.	*Labour*	*Liberal*	*Conservative*
2nd pref.	*Liberal*	*Conservative*	*Liberal*
3rd pref.	*Conservative*	*Labour*	*Labour*

Under *Simple Plurality* voting (what the British call "first-past-the-post" or FTPT voting), such as is used in British parliamentary elections and most U.S. elections, each voter votes for exactly one candidate, and the candidate receiving the most votes wins. For the time being, let us assume that, under plurality voting, each voter votes for his or her most preferred candidate, i.e., votes "sincerely." Here is the *plurality ranking* for Profile 1.

Candidates	Votes Received (= First Preferences)
Labour	46 votes (*winner*)
Conservative	34 votes
Liberal	20 votes

A plurality election with sincere voting takes account of *first preferences* only — that is, the top line of the preference profile. The *plurality winner* is the candidate who has the most first preferences; in the example above, the Labour candidate is the plurality winner (and wins under sincere plurality voting).

A *majority winner* is a candidate who has an (absolute) majority of first preferences. Clearly a majority winner is also a plurality winner; equally clearly, the reverse is not always true. And if there are three or more candidates and first preferences are dispersed, no candidate will be the first preference of a majority of voters. In Profile 1, there is no majority winner.

In the event that simple plurality does not give one candidate an absolute majority of votes, *Plurality Runoff* voting prescribes a runoff vote between the top two candidates in the plurality ranking. Thus in Profile 1 there would be a runoff between Labour and Conservative, which Conservative wins because the voters who most prefer the eliminated Liberal candidates all prefer Conservative to Labour and they are sufficient in number to overcome the Labour margin over Conservative with respect to first preferences. (A second trip to the polls can be avoided if voters rank all the candidates on a single ballot. This is called *Instant Runoff Voting* or IRV.)

Under *Approval Voting* (Brams and Fishburn, 1983), voters can vote for any number of candidates, and the candidate with the most such "approval votes" wins. In the three candidate case, this means that a voter can vote for just one candidate (as under simple plurality) or for two. (It should be clear that voting for all three is effectively equivalent to abstaining.) While approval voting has some advantages, it can be highly indeterminate. For example, given Profile 1 sincere approval voting can select Labour (if each voter votes for his most preferred candidate only), Conservative (if only voters in the 20-voter bloc cast two approval votes), or Liberal (if only voters in the 34-voter bloc cast two approval votes or if all voters cast two approval votes).

Under *Borda Point Voting* (proposed by the French philosopher Jean-Charles de Borda), votes rank the candidates on the ballot, and (in a three-candidate contest) candidates are awarded three points for each ballot on which they are ranked first, two points for each ballot on which they are ranked second, and one point for each ballot on which they are ranked third. Here is the*Borda ranking* for Profile 1:

Candidates	Points Received
Liberal	220 points (*winner*)
Labour	192 points
Conservative	188 points

Finally, suppose we look at all possible *pairs* of candidates and see which candidate in each pair is supported by a majority of voters. (Apart from "knife-edge" ties, one or other candidate must have majority support.) In other words, let's examine all possible straight fights. For Profile 1, we see the following:

Liberal vs. *Conservative*:	*Liberal* wins by 66-34
Conservative vs. *Labour*:	*Conservative* wins by 54-46
Liberal vs. *Labour*:	*Liberal* wins by 54-46

Thus we can order candidates in terms of (pairwise) *majority preference* such that *A* is ranked over *B* if and only if a majority of voters prefers *A* to *B*. For the example above we get the following *majority* (or *Condorcet*) *ranking*:

	Majority Ranking
1st pref.	*Liberal* (Condorcet winner)
2nd pref.	*Conservative*
3rd pref.	*Labour* (Condorcet loser)

Notice that this "majority ranking" is precisely the *opposite* of the "plurality ranking" based on first preferences only and that it also differs from the "Borda ranking" based on full orderings.

More than two hundred years ago the Marquis de Condorcet, a French philosopher and mathematician, proposed examining pairwise majority preference in this fashion to produce the *Condorcet voting rule*, under which the candidate at the top of the majority ranking — called the *Condorcet winner* — is elected. More generally, a Condorcet winner is a candidate who can beat every other candidate is a straight fight.

You should be able to verify the following points, many of which are illustrated in Preference Profile 1. For any preference profile:

(1) a majority winner is always a Condorcet winner, but the reverse is not true;

(2) a plurality winner may not be a Condorcet winner; and

(3) a Condorcet winner may not be a plurality winner — indeed, a Condorcet winner may have the fewest first preferences (e.g., Liberal in Profile 1).

Although there may be a Condorcet winner in the absence of a majority winner, it is also true that a Condorcet winner does not always exist. It may seem puzzling how this can occur, since — apart from ties —every ranking must have a highest-ranked element. The explanation is that there may be no majority ranking at all. Consider Preference Profile 2.

Preference Profile 2

# of voters	46	20	34
1st pref.	*Labour*	*Liberal*	*Conservative*
2nd pref.	*Liberal*	*Conservative*	*Labour*
3rd pref.	*Conservative*	*Labour*	*Liberal*

Notice that in Profile 2 first preferences are unchanged from Profile 1, so the plurality winner is the same as before and (as before) there is no majority winner. Conservative remains the plurality runoff winner but Labour becomes the Borda point winner, while approval voting remains indeterminate. But another crucial difference is apparent when we look at the straight fights:

Liberal vs. *Conservative*:	*Liberal* wins by 66-34
Conservative vs. *Labour*:	*Conservative* wins by 54-46
Labour vs. *Liberal*:	*Labour* wins by 80-20

It is now impossible to construct a majority ranking. Instead we have *cyclical majority*.[3] Since there is no majority ranking of the three candidates, there is no Condorcet winner. Thus, we can add a fourth proposition concerning Condorcet winners:

(4) there may be no Condorcet winner.

There may be a Condorcet winner even in the presence of a majority cycle, provided the cycle does not encompass all candidates. This can occur if there are four or more candidates, as in this example.

Preference Profile 3

# of voters	35	33	32
1st pref.	*B*	*C*	*D*
2nd pref.	*A*	*A*	*A*
3rd pref.	*C*	*D*	*B*
4th pref.	*D*	*B*	*C*

Candidate *A* is the Condorcet winner, yet there is a cycle including *B*, *C*, and *D*. This example also shows that, with four of more candidates, a Condorcet winner may have *no* first preferences at all.

---

[3] This phenomenon has also been called "the Condorcet effect" (since Condorcet discovered this anomaly), "the paradox of voting,"and "the Arrow problem" (for Kenneth J. Arrow, who publicized it in his book on *Social Choice and Individual Values*).

A voting rule is *Condorcet consistent* if, given sincere voting, it always selects the Condorcet winner when one exists. While Condorcet voting is obviously Condorcet consistent, previous examples showed that Liberal may fail to win given Profile 1 under each of the other voting rules discussed, so none of them is Condorcet consistent. But since Condorcet voting does not always select a winner, it cannot be deemed a full-fledged voting rule comparable to the others discussed here. This is especially unfortunate because, in so far as Condorcet voting does select winners, it is (unlike the others) both strategyproof and not subject to spoiler effects.

Of course, to say that majority cycles may exist is not to say that they typically are present. Indeed, if preferences are structured in a simple way by ideology (or otherwise), cycles cannot occur. In British politics, the three major parties are generally perceived to be ideologically ranked from left the right in the following manner:

More leftwing:	*Labour*
Relatively centrist:	*Liberal*
More rightwing:	*Conservative*

If voters commonly perceive this ideological dimension and each ranks candidates according to how "close" they are to the voter's own (most preferred) position on this dimension, voter preference orderings are restricted to the following admissible ordering:

	Admissible Orderings				**Inadmissible Orderings**	
	leftwingers	centrists		rightwingers		
1st pref.	*Lab*	*Lib*	*Lib*	*Con*	*Con*	*Lab*
2nd pref.	*Lib*	*Lab*	*Con*	*Lib*	*Lab*	*Con*
3rd pref.	*Con*	*Con*	*Lab*	*Lab*	*Lib*	*Lib*

If preferences are restricted in this so-called "single-peaked" fashion, regardless of popularity each the admissible orderings, it is always possible to construct a majority ranking, so a Condorcet winner always exists (Duncan Black, 1948, 1958). You can check that Profile 1 draws orderings exclusively from the admissible types, while Profile 2 includes an inadmissible type.

Note the strength of the "centrist" (Liberal) candidate in the admissible orderings. While it may be that few voters most prefer the centrist, no one likes the centrist least. The consequence is that the centrist candidate must be the Condorcet winner unless an (absolute) majority of voters have the leftwing ordering or have the rightwing ordering. Put otherwise (in the three-candidate case), the centrist candidate fails to be the Condorcet winner only if one of the extreme candidates is a majority winner.[4]

---

[4] More generally, if all voters can be ranked from most leftwing to most rightwing with respect to their first preferences, no cyclical majority occurs, so some position on the ideological spectrum must be the Condorcet winner. This position corresponds to the first preference of the *median voter*, such that no more than half the voters are more leftwing and no more than half are more rightwing (Duncan Black, 1948 and 1958). The Hotelling-Downs theory of electoral competition (to be discussed later in the course) states that two competing vote-seeking parties or candidates achieve equilibrium only when both adopt the position that corresponds to the first preference of the median voter.

We have to this point assumed that voters vote sincerely. But any voting rule with three or more candidates may give voters incentives to vote otherwise than sincerely.

Consider Profile 1 again. As we saw, Labour wins under Plurality Voting if voters are sincere. But it is also true that a majority of 54 voters prefer both other candidates to Labour. If they all vote for the same other candidate (either all for Liberal or all for Conservative), that candidate wins — an outcome they all prefer to a Labour victory. But doing this requires some members of this majority of 54 to vote insincerely, i.e., for their second preferences. Thus simple plurality voting (as well as other voting systems) can encourage what the British call *tactical voting* and most political scientists call *strategic voting*, i.e., non-sincere voting.

Of course, the problem remains of *how* the 54 voter majority will coordinate their votes — that is, will they vote for Liberal or for Conservative? Notice that, while all 54 voters prefer to see Labour defeated, they disagree as to *how* to defeat him, i.e., by voting Conservative or by voting Liberal. It is generally believed that, in practice, tactical voting in Britain mostly leads Liberal supporters to shift their votes "tactically" to their second-preference (Labour or Conservative) candidate, because they typically observe pre-election polls showing Liberal trailing well behind both other candidates, and they therefore conclude that a Liberal vote is "wasted" and that they should vote for the one of the two leading (non-Liberal) candidates that they prefer.[5]

Under Plurality Runoff [IRV], the 46 voters who most prefer Labour would do better by ranking Liberal first, as this assures (in the absence of countermoves by other voters) a Liberal victory without a runoff, which outcome they prefer to the Conservative victory that otherwise results. Given Profile 1, no voters can change their Borda score ballots in a way that improves the outcome for them. Given Profile 2, if the bloc of 20 ranks Conservative first and the bloc of 34 ranks Labour third, then Conservative gets the most Borda points (208 vs. 200 for Liberal and 192 for Labour), an outcome all 54 such voters prefer to victory by the sincere Borda winner Labour. Given some other profiles, the opportunity for strategic manipulation under Borda point voting is far more glaring, as is illustrated by Profile 4.

Preference Profile 4

	46	54
1st pref.	*Labour*	*Conservative*
2nd pref.	*Liberal*	*Labour*
3rd pref.	*Conservative*	*Liberal*

[5] Notice that this can happen even though Liberal is the Condorcet winner, reflecting the fact that polls (almost always) ask only about first preferences and Liberal's great strength lies in second preferences. Notice also that, if Liberal supporters find Labour and Conservative to be equally objectionable, they have no incentive to vote tactically. Finally, if pre-election polls show something close to a tie for second place (or a three-way tie), tactical voting becomes far more conjectural.

Labour wins if voting is sincere (demonstrating that Borda Point Voting can deny victory to a majority [and Condorcet] winner, i.e., Conservative), but the 54 Conservative-preferring voters can elect Conservative if they shove Labour down to third place on their ballots. In turn, the 46 Labour-preferring voters can counteract this by moving Liberal to the top of their ballots (the resulting Liberal victory being preferable to the 46 voters to a Conservative victory). Note that if strategic manipulation stops at this point (though it need not), Liberal is elected even though *everyone* prefers Labor to Liberal. (An even more perverse example of such strategic manipulation under Borda voting is presented in the Appendix.)

We now examine spoiler effects. Consider an individual who, when given a choice between Conservative and Labour only, chooses Conservative. We would think this voter mighty peculiar if he changed his choice to Labour in the event Liberal is added as a third option. But a sincere electorate using Plurality Voting may do exactly this, as can be verified by checking Profiles 1 and 2 (or thinking about the Bush/Gore/Nader example referred to in footnote 2). So can a sincere electorate using Borda point voting, as can be verified by checking Profile 4. That is, these procedures are subject to spoiler effects.

Plurality Runoff (and especially Instant Runoff Voting) is sometimes advocated on the grounds that it precludes such spoiler effects. It is true that Plurality Runoff is a big improvement over Simple Plurality in this respect, in that a third candidate (such as Nader) with little first-preference support cannot act as a spoiler in what is essentially a straight fight between two major candidates, because the runoff will become precisely that straight fight. However, Plurality Runoff [IRV] does not eliminate the spoiler problem, as is illustrated by Profile 1. While Liberal would win a straight fight with Conservative, Liberal will not even make it into the runoff if Labour enters the field. So Labour is a spoiler to Liberal. This is not a distinctive flaw in Plurality Runoff voting, however; as previously noted, the problem is unavoidable with three or more candidates.

However, Plurality Runoff [IRV] does have another flaw that is distinctive (and avoidable). We wouldn't expect a "reasonable" voting rule to respond *negatively* when a candidate's position in a preference profile becomes more favorable — put otherwise, increased support in the electorate should never hurt a candidate. (This notion generalizes May's Non-Negative Responsiveness.) But Plurality Runoff [IRV] can fail on this score.

Suppose we have three candidates *A*, *B*, and *C*, among whom first preferences are fairly equally divided. Suppose that *A* and *B* go into the runoff, which is therefore decided by the second preferences of the voters who most prefer *C*. Suppose enough of these second preferences are for *A* that *A* wins the runoff. Now suppose the preference profile is revised in a way that makes "public opinion" even more favorable to *A* (without changing anyone's preferences between *B* and *C*). In particular, suppose that some voters who previously most preferred *B* now move *A* up to their first preference (but *A* still is not a majority winner). The result of this change may be that the number of first preferences for *B* falls below the number of first preferences for *C*, with the result that *A* and *C* are paired in the runoff, which is decided by the second preferences of the remaining voters who most preferred *B*. And it may be that enough of these second preferences are for *C* that C rather than *A* wins the runoff. Added support therefore costs *A* electoral victory. Here is a specific example.

Original Preference Profile 5

35	10	25	30
A	*B*	*B*	*C*
B	*A*	*C*	*A*
C	*C*	*A*	*B*

Revised Preference Profile 5

35	10	25	30
A	*A*	*B*	*C*
B	*B*	*C*	*A*
C	*C*	*A*	*B*

Here is a related peculiarity of Plurality Runoff [IRV] voting.

Preference Profile 6

5	6	4	[2]
B	*C*	*A*	[*A*]
C	*B*	*B*	[*B*]
A	*A*	*C*	[*C*]

The preference profile is as shown above, but the two individuals with the bracketed preference orderings fail to vote. Thus the election outcome is determined by the remaining 15 voters. Candidates *B* and *C* are paired in a runoff, which *B* wins. This is a somewhat disappointing outcome for the two individuals who failed to vote, in that their second preference won. They regret their failure to get to the polls, since they wonder whether their first preference *A* might have won if they had not failed to vote. But it can be checked that, if they had gotten to the polls and voted according to their preferences, the outcome would have been worse, not better, for them. (Candidates *A* and *C* would be paired for a runoff, which *C* would have won.)

Appendix: Voting Rules, "Clone" Candidates, and "Turkey Raising"

Consider the following preference profile, in which a Republican minority is united behind a single candidate R but the Democratic majority is split between the two "clone" candidates D1 and D2 (see Tideman, 1987).

Democrats		*Republicans*	
35%	25%	25%	15%
D1	D2	R	R
D2	D1	D1	D2
R	R	D2	D1

Simple Plurality voting is notorious for penalizing clone candidates. In this case, the Republican candidate would win due to the Democratic split, even though R is at the bottom of the majority ranking. (R is the *Condorcet loser*, beaten by both D1 and D2 in straight fights.) Of course, it is precisely the expectation of such outcomes under Simple Plurality voting that leads to party formation and party discipline, i.e., the Democrats have a huge incentive to hold a prior *nominating convention* or *primary* to choose between D1 and D2 and then send just one of the two clones forward against the Republican. Given the preference profile above, D1 would win the nomination and then the general election.

The question arises of whether there are other voting rules that can reduce, eliminate, or even reverse the self-defeating effect of running clone candidates.

First we may note that, given the profile above, Plurality Runoff (instant or otherwise) solves the clone problem. In effect, the first-round election functions as the (Democratic) primary and the runoff as the general election in which the Democratic majority gets its way. But if there are four or more candidates, Plurality Runoff does not treat clones so well and, as we have seen, it is subject to other problems in addition.

As noted previously, Steven Brams and Peter Fishburn advocate Approval Voting as a desirable voting rule that (among other things) does not punish clones. In the profile above, presumably (almost all) Democrats would vote for both D1 and D2, one of whom would be elected. Of course, by not penalizing clones, AV does not encourage party formation or party unity. For this reason, many political scientists are more inclined to support AV for primary elections and non-partisan elections than for partisan general elections.

A variation of one type of party-list PR (Proportional Representation) system presents another voting method that does not penalize clones who have the same party affiliation. Each voter votes for a single candidate, as under Simple Plurality, but this vote counts in two ways: first, as a *party vote* to determine which party wins the election and, second, as a *candidate vote* to determine which candidate of the winning party is elected. In the profile above, D1 would be elected.

Perhaps surprisingly, Borda Point Voting actually *rewards* the running of clones. Suppose that there are two candidates and Republicans are again in the minority.

60 voters	40 voters
D	R1
R1	D

With just two candidates, the Borda point rule is identical to Plurality Voting (and SMR), so the Republican candidate R1 loses. But, if Borda Voting is in use, the Republicans can reverse the outcome by nominating an *additional clone candidate* R2 whom everyone sees as identical to R1 with respect to issues and ideology but inferior with respect to (let's say) personal qualities.

60 voters	40 voters
D	R1
R1	R2
R2	D

Now R1 wins with 60×2 + 40×3 = 120 + 120 = 240 points, while D gets 60×3 + 40×1 = 180 + 40 = 220 points and R2 gets 60×1 + 40×2 = 60 + 80 = 140 points. Of course Democrats can counteract this by strategically ranking R2 above R2, thereby reducing R1 to 60×1 + 40×3 = 60 + 120 = 180 points and raising R2 to 60×2 + 40×2 = 120 + 80 = 200 points, allowing D to win with the unchanged 220 points. Alternatively, they can counteract the Republican stratagem by running their own clone. Though it has strong advocates, the Borda scoring system evidently is highly susceptible to strategic maneuvers of this sort (which, moreover, have the effect of expanding the candidate field rather than winnowing it down in the manner of Plurality Rule).

Here is a considerably worse thing that the Borda point rule can do.[6] Suppose are three candidates: a more or less reasonable Democrat *D*, a more or less reasonable Republican *R*, and a real "turkey" T. Everyone one ranks T last, except two deranged T supporters. The profile is:

50 voters	48 voters	1 voter	1 voter
D	R	T	T
R	D	D	R
T	T	R	D

Voting is by the Borda rule. It is easy to see right off that, if everyone votes sincerely, *D* wins (the same outcome as under Simple Plurality). Doing the arithmetic, the point totals would be D = 249, R = 247, and T = 104. Anticipating this defeat, Republican voters caucus and notice an interesting feature of Borda Point Voting — it can pay voters to engage in "turkey raising"(the term is originally due to Cox, 1997), i.e., to strategically raise the "turkey" in their ballot rankings, so as to push the rival "serious" candidate down in their rankings and increase the point spread between the two. Suppose the Republicans strategically modify all their ballots so as to produce the following ballot profile:

50 voters	48 voters	1 voter	1 voter
D	R	T	T
R	T	D	R
T	D	R	D

The point totals would now be D = 201, R = 247, and T = 152, producing a clear R victory. But suppose that before the actual balloting takes place, Democrats also notice this feature of Borda Voting and, concerned that Republicans may engage in turkey raising, they determine that they must engage in some turkey raising of their own in order to counteract the anticipated Republican stratagem. So the final ballot profile is:

50 voters	48 voters	1 voter	1 voter
D	R	T	T
T	T	D	R
R	D	R	D

And the final point scores are D = 201, R = 197, and T = 202. May the best turkey win!

[6] The following simple example is adapted from Monroe (2001).

References

Kenneth Arrow. *Social Choice and Individual Values*. Wiley, 1951.

Gary W. Cox, *Making Votes Counts: Strategic Coordination in the World's Electoral Systems* . Cambridge, 1997.

Duncan Black. "On the Rationale of Group Decision-Making," *Journal of Political Economy*, 1948

Duncan Black. *The Theory of Committees and Elections*. Cambridge, England: Cambridge University Press, 1958.

Jean-Charles de Borda, "On Elections by Ballot," 1784.

Steven Brams and Peter Fishburn. *Approval Voting*: Birkhauser, 1983.

Marqis de Condorcet. "An Essay on the Application of Analysis to the Probability of Decisions Rendered by a Plurality of Voting," 1785.

Kenneth May. "A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision," *Econometrica*, 1952.

Burt Monroe. "Raising Turkeys: An Extension and Devastating Application of Myerson-Weber Voting Equilibrium." Presentation to 2001 Annual Meeting of the American Political Science Association (http://accuratedemocracy.com/archive/condorcet/Monroe/004004MonroeBurt.pdf)

T. Nicolas Tideman, "Independence of Clones as a Criterion for Voting Rules," *Social Choice and Welfare*, 1987.