Random tie-breaking in STV

Jonathan Lundell
jlundell@pobox.com

1 Introduction

The resolution of ties in STV elections is not a settled question. On the contrary, it remains a topic of lively discussion, with several papers published on the subject in these pages; see Earl Kitchener's note, "A new way to break STV ties in a special case" [1] for a summary.

Ties can arise in any STV election during exclusion. With some methods ties can arise at other stages as well; Jeffrey O'Neill [2] lists the cases.

O'Neill also lists four tie-breaking methods. Two methods use the first or last difference in prior rounds to break a tie, and two methods use later preferences—Borda scores or most (fewest) last-place preferences. Brian Wichmann [3] proposes to examine all possible outcomes.

None of these tie-breaking methods is guaranteed to break a tie, since they can themselves result in a tie, or in the case of [3] become so computationally expensive as to be impractical. These cases (strong ties) are typically broken randomly. Some election methods, eg, the Algorithm 123 version of Meek's method [4], rely exclusively on random tie-breaking.

Objections to random tie-breaking fall into two categories. One is a worry that voters and candidates will object to election decisions being made by chance instead of by voter preference. Thus Wichmann [3]: "When a candidate has been subject to a random exclusion in an election, he/she could naturally feel aggrieved." Other objections adduce examples in which it appears intuitively preferable to break a tie based on some measure of voter preference.

All STV election methods rely on random tie-breaking (or at least tie-breaking based on some consideration other than voter preference) to break strong ties. (Ties in first-past-the-post elections are often broken randomly as well, by coin toss, drawing straws, or drawing a high card.)

2 Prior-round tie-breaking

The rationale for forwards tie-breaking (using O'Neill's terminology) appears to be that it gives greatest weight to first preferences. O'Neill [2] argues for backwards tie-breaking:

> A more important problem, is that forwards tie-breaking does not use the most relevant information to break the tie. The most relevant information to break a tie is the previous stage and not all the way back to the very first stage. By immediately looking to the first stage to break the tie, the ERS97 rules allow the tie-breaking to be influenced by candidates eliminated very early in the process and also by surpluses yet to be transferred. Instead, if we look to the previous stage to break a tie, candidates eliminated early on in the process will have no influence in breaking the tie. In addition, it allows for surpluses to be transferred which gives a more accurate picture of candidate strength.

Carrying O'Neill's argument to its logical conclusion, however, the "most relevant information" is not in any prior round, but rather in the current round—and the current round declares a tie.

Prior-round tie-breaking encourages insincere voting. Consider this election fragment, with two candidates to exclude:

5	A
4	B
1	CB

Excluding C, we have:

5 A
5 B

and must now break the tie. Prior-round tie-breaking requires that we exclude B, since A led B 5-4 in the previous round. So voter CB, believing that the first choice (C) is likely to be excluded, is encouraged to insincerely vote B (or BC) so as not to jeopardize B's chances in the event of an A–B tie.

Prior-round tie-breaking is especially troublesome in the context of Meek rules, since it violates Meek's Principle 1: If a candidate is eliminated, all ballots are treated as if that candidate had never stood. But if C had never stood, A and B would have been tied.

3 Later-preference tie-breaking

Kitchener [5] points out a problem case for random tie-breaking:

> An extreme case can arise where there is one seat and the electors are the same as the candidates; for example, if a partnership is electing a senior partner. Each candidate may put himself first, and all, except candidate A, put A second. Under most present rules, one candidate then has to be excluded at random, and it may be A. There is no way of getting over this unreasonable result without looking at later preferences. . . .

The smallest such election:

1 A
1 B A
1 C A

Prior-round tie-breaking methods are of no help in the first round, and a random choice excludes A, the consensus choice, one third of the time. Kitchener proposes to use Borda scores to break the tie; we must still randomly break a strong B-C tie, but A survives and is elected.

This case is related to a problem with STV in general, pointed out by Meek [6]. "A related point, and probably the strongest decision-theoretic argument against STV, is the fact that a candidate may be everyone's second choice but not be elected."
. . . and also related to the general problem of premature exclusion.

Kitchener concedes that there is a problem with Borda tie-breaking, as there is with any tie-breaking method that relies on later preferences.

> It is a fundamental principle of STV that later preferences should not affect the fate of earlier ones; this encourages sincere voting, but means that some arbitrary or random choice must be made to break ties, which can give unreasonable results.

Responding to the Borda tie-breaking suggestion, David Hill [7] objects: "What matters is that tactical considerations have been allowed in, where STV (in its AV version in this case) is supposed to be free of them."

This point is crucial. In any election system, the rules, including the method of breaking ties, must of course be specified in advance. When we look at the partnership election example above, we interpret the ballots as the sincere expression of the voters, and so read the ballots as favoring A. But as both Hill and Kitchener observe, once later-preference tie-breaking is introduced, we must expect insincere voting. In the face of later-preference tie-breaking, B and C, to maximize their chances of winning (after all, each is their own first choice) must resort to bullet voting (American English—one might say characteristically AmE—for plumping). The ballots would then read,

1 A
1 B
1 C

. . . and we're forced to resort to a random choice. This seems a shame, since it does appear from the presumably sincere ballots in the initial profile that both B and C prefer A to the other. The partners might be well advised to adopt a special rule forbidding each to vote for herself. In that case, we would have:

1 abstain
2 A

. . . and A wins outright.

4 Random tie-breaking

An advantage claimed by Meek [6] for STV is that "There is no incentive for a voter to vote in any way other than according to his actual preference." One of Meek's motivations for proposing a new STV method is to come closer to that ideal. Likewise Warren [8], "It is one of the precepts of preferential voting systems that a later preference should neither help nor harm an earlier preference."

Any election method relies for its properties on the implicit assumption that voters will vote sincerely, that is, that their ballots will reflect, within the limitations of the specific method, their true preferences. Without sincere votes, any election method fails to reflect the will of the electorate, on the principle of garbage in, garbage out. It is perverse to use tie-breaking methods that reintroduce incentives for voters to vote insincerely. Hill and Gazeley [9]:

> In considering this, we need to take into account, among other things, that the true aim of an election should not be solely to match seats as well as possible to votes, but to match seats to the voters' wishes. Since we do not know the wishes we must use the votes as a substitute, but that makes it essential that the votes should match the wishes as far as possible. That, in turn, makes it desirable that the voters should not be tempted to vote tactically.

5 Voter psychology

One might counter that, except in small elections, the chances of a tie are sufficiently small that a voter ought to ignore the possibility of a tie altogether and vote sincerely. This argument is problematic on two fronts. First, our methods should work with small elections as well as large ones (and the line between small and large elections is not well defined). Second, especially in a high-stakes election, the voter's estimation of the risk associated with voting sincerely is likely to be wrong.

Computer security authority Bruce Schneier, interviewed in *CSO Magazine* [10], comments:

> Why are people so lousy at estimating, evaluating and accepting risk?... Evaluating risk is one of the most basic functions of a brain and something hard-wired into every species possessing one. Our own notions of risk are based on experience, but also on emotion and intuition. The problem is that the risk analysis ability that has served our species so well over the millennia is being overtaxed by modern society. Modern science and technology create things that cannot be explained to the average person; hence, the average person cannot evaluate the risks associated with them. Modern mass communication perturbs the natural experiential process, magnifying spectacular but rare risks and minimizing common but uninteresting risks. This kind of thing isn't new—government agencies like the [US] FDA were established precisely because the average person cannot intelligently evaluate the risks of food additives and drugs—but it does have profound effects on people's security decisions. They make bad ones.

For our purposes, read *tactical voting decisions* for *security decisions*. Rational insincere voting is bad enough; insincere voting based on faulty information or poor tactics is even worse.

6 A note on weighting votes in later-preference tie-breaking

Consider this election profile (BC rules, two to be elected, quota 10):

12	AB
7	BC
9	C
2	D

A is elected, and D is excluded, leaving B and C tied with nine votes each in the third round. If we break the tie with Borda scores:

A	36	(elected)
B	24+21 =	45
C	14+27 =	41
D	6	(excluded)

C is excluded, and B is elected as the last candidate standing for the second seat.

Notice in particular that while B receives only the two transferable votes from the AB voters (a quota of 10 being retained by A, who is elected), B gets full credit for all 12 AB votes in the Borda tiebreaker.

I suggest that the AB voters, having elected A, must carry only the transferable weight of their votes in calculating the tie-breaking Borda score. Otherwise these voters *double dip*, not only electing A, but also participating disproportionately and decisively in the tie-breaking elimination of C and subsequent election of B.

If we calculate the Borda scores using the weight of transferable votes (that is, votes currently allocated to hopeful candidates), we have:

A		(elected)
B	4+21 =	23
C	14+27 =	41
D		(excluded)

Calculated with the vote weights that give rise to the tie itself, the Borda score now breaks the tie to eliminate B, and C is elected.

The same argument applies to any method that breaks ties with later preferences. Votes committed to already-elected candidates should not be counted again in breaking subsequent ties.

7 A better later-preference tie-breaking method

The chief problem with STV tie-breaking with Borda scores is that it violates the principle of later-no-harm, and it does so in an especially egregious way. Suppose that six candidates are in the running, that I have voted ABC, and that B and C are tied for elimination. The Borda scores for B and C pick up four and three points, respectively, from my ballot. If the three points that my ballot contributes to C's Borda score is the margin for C's victory over B in the Borda tiebreaker, then my later mention of C has led directly to the defeat of B, even though I prefer B to C.

Consider an alternative later-preference tiebreaker. For the sake of simplicity, I will describe it for two-way ties, and then extend it to *n*-way ties. To break a tie, compare the ballots that prefer B to C to the number of ballots that prefer C to B, weighted as described in the note above. Exclude the less-preferred candidate. Break strong ties randomly.

This method, like all later-preference methods, violates later-no-harm, but it preserves a property that I will call *later-no-direct-harm*. My ranking of ABC will not harm B's chances in a BC tie. In the case of a BC tie, my ballot will either have no effect (the margin of B over C or vice versa without my ballot is sufficient that my ballot makes no difference), or it will cause the BC tie to be broken in favor of B, my preferred candidate in the tie (B and C are strongly tied without my ballot), or my ballot will convert a one-vote C advantage (without my ballot) to a strong tie (with my ballot), giving B an even chance in a random tiebreak.

That is, my ABC ballot either has no effect on breaking a BC tie, or it benefits B.

By *later-no-direct-harm*, I mean that the fact that I have ranked the later preferences BC will not harm my favorite in the potential tie between B and C. Later-no-harm is not avoided; my ABC preference could break a tie in favor of B, and B could subsequently defeat my first preference, A, whereas A might have prevailed had C won the BC tiebreaker. Any harm to A, however, will come indirectly, in a later round—and it would be rude for me to complain that the BC tie was broken on the basis of my preference for B over C.

Generalizing to breaking an *n*-way tie for exclusion:

1. Find the first mention of any member of the tied set of candidates on each ballot, and calculate the total such mentions for each of the candidates, using the transferable weight of each ballot. Ignore ballots that do not mention at least one tied candidate.

2. If all *n* candidates are still tied, exclude one tied candidate at random; *finis*.

3. Otherwise, remove from consideration for exclusion the candidate (or a random choice from the tied set of candidates) with the highest score from step 1.

4. If only one candidate remains, exclude that candidate; *finis*.

5. Otherwise, *n* is now the remaining number of tied candidates (that is, less the reprieved candidates from step 3); continue at step 1.

If the tie is for a winner rather than an exclusion, then remove from consideration the candidate with the lowest rather than the highest score. This is simply single-winner STV (AV or IRV) with weighted ballots, and suggests an alternative to the proposed algorithm for breaking a tie for exclusion: break an n-way tie for exclusion by counting an STV election (again with weighted ballots) with n candidates and $n-1$ winners; exclude the single loser.

It's worth noting that a similar procedure based on lowest preferences (along the lines of Coombs tie-breaking) does not satisfy the principle of later-no-direct-harm. For example, if candidates X, Y and Z are tied for exclusion and I have ranked those candidates XYZ, it's possible that my preference for Y over Z is decisive in favor of Y, and that Y but not Z beats X in a head-to-head tiebreaker; thus my preference for Y over Z decides the tiebreak in favor of Y over X, contrary to my preferences.

Likewise, Condorcet ranking is equivalent to the proposed method for two-way ties, but violates later-no-direct-harm in the general *n*-way-tie case.

The proposed tie-breaking method—let's call it *weighted first preference*—differs from prior-round tie-breaking methods in that it considers the preferences of all voters (suitably weighted), and not only voters who have ranked the tied candidates first (after elections and exclusions) in a prior round.

Hill and Gazeley [9] observe, in the context of Sequential STV:

> With this new version, should it be recommended for practical use? That depends upon whether the user is willing to abandon the principle that it should be impossible for a voter to upset earlier preferences by using later preferences. Many people regard that principle as very important, but reducing the frequency of premature exclusions is important too. We know that it is impossible to devise a perfect scheme, and it is all a question of which faults are the most important to avoid.
>
> In considering this, we need to take into account, among other things, that the true aim of an election should not be solely to match seats as well as possible to votes, but to match seats to the voters' wishes. Since we do not know the wishes we must use the votes as a substitute, but that makes it essential that the votes should match the wishes as far as possible. That, in turn, makes it desirable that the voters should not be tempted to vote tactically.
>
> They would not be so tempted if they felt confident that later preferences were as likely to help earlier ones as to harm them, and if they could not predict the effect one way or the other. At present, we see no reason to doubt that these requirements are met.

The proposed method for breaking ties satisfies the same criteria: later preferences are as likely to help earlier ones as to harm them, and voters cannot predict the effect one way or the other. This is not the case for other preference-based tie-breaking methods discussed in these pages.

Whether this slight opening of the door to a violation of later-no-harm is justified by the benefit of breaking ties non-randomly (in most cases) is, in David Hill's words [7], a matter of judgment.

8 Summary

Arguments for various nonrandom tie-breaking implicitly assume sincere voters. But the introduction of those very methods undermines that crucial precondition, and without sincere voters the arguments fail.

When O'Neill argues [2] that "forwards tie-breaking does not use the most relevant information to break the tie," and that later rounds reflect better information, the logical conclusion of his argument is that the most relevant information is not in a prior round at all, but rather in the current round that gives rise to the tie. That information is, simply, that the candidates have equal support, by the means we've chosen to measure that support.

Meek [6] drives this point further home with his Principle 1: "If a candidate is eliminated, all ballots are treated *as if that candidate had never stood*." Prior-round tie-breaking typically, though not exclusively, depends on preferences for candidates who have been excluded in the tie-breaking round. To consider those preferences violates Meek's Principle 1.

Later-preference tie-breaking (eg. Borda or Coombs) encourages insincere voting by violating the later-no-harm principle.

The encouragement of insincere voting is too high a price to pay for partially excluding chance from STV election methods. We should prefer random tie-breaking in all cases.

If preferences must be considered in breaking ties, then ties should be broken on the basis of overall earliest preferences, using transferable ballot weights.

9 References

[1] Earl Kitchener. A new way to break STV ties in a special case. *Voting matters*. Issue 20, pp 9–11.

[2] J C O'Neill. Tie-Breaking with the Single Transferable Vote. *Voting matters*. Issue 18, pp 14–17.

[3] Brian Wichmann. Tie Breaking in STV. *Voting matters*. Issue 19, pp 1–5.

[4] I. D. Hill, B. A. Wichmann and D. R. Woodall. Algorithm 123—Single Transferable Vote by Meek's method. *Computer Journal*. Vol 30, pp 277–281, 1987.

[5] Earl Kitchener. Tie-breaking in STV. *Voting matters*. Issue 11, pp 5–6.

[6] B L Meek. A New Approach to the Single Transferable Vote. *Voting matters*. Issue 1, pp 1–7.

[7] I D Hill. Tie-breaking in STV. *Voting matters*. Issue 12, pp 5–6.

[8] C H E Warren. An example showing that Condorcet infringes a precept of preferential voting systems. *Voting matters*. Issue 5, pp 5–6.

[9] I D Hill and Simon Gazeley. Sequential STV - a new version. *Voting matters*. Issue 15, pp 13–16.

[10] The Evolution of a Cryptographer, *CSO Magazine,* September 2003. Available from the McDougall web site.