On Designing Very Good Voting Methods

Roy Minet – 2021/03/12

Introduction and Definitions – Designing very good voting methods has turned out to be much trickier than most would have expected. In spite of near-universal agreement that Plurality is awful, much effort, spanning nearly 250 years, has failed to produce a clear and strong consensus as to its best replacement. With the pernicious effects of Plurality becoming ever more corrosive, a solution is urgent. It may well be time to rethink the approaches to solving the problem.

As is always the case, *before* designing anything it is critical to clearly and unambiguously state the objective as well as what criteria the design should satisfy. We begin with:

The primary design objective for an election mechanism must be for it to most consistently render the best possible decisions (with the caveat that decision-making power be kept reasonably dispersed).

The best possible decision is that result which maximizes voter satisfaction, net of dissatisfaction, when summed over all voters who voted.

Since we are going to depend on voters to make these important decisions, we have no choice but to lightheartedly assume that, collectively, they possess the knowledge and wisdom to make them. The definition as stated assures that election outcomes are the best match to and manifestation of voters' collective wisdom. This is *the* overriding consideration.

However, in passing, it should also be noted that, although "fairness" (however defined) is *not* a consideration, the chosen definition happens to also be a great definition of the fairest outcome.

Note that adoption of this definition requires abandonment of the majoritarian principle.

A Perfect Voting Method – Suppose that a machine, call it a "satometer," is available which is trusted by voters and which can accurately measure each voter's sincere satisfaction regarding the issue at hand on some absolute scale of "sats" (dissatisfaction would be negative "sats"). If the net total sats for all voters voting on a referendum is positive, the result will be "yes" and if total sats is negative, the result will be "no" (a zero-sat tie would be resolved by falling back to majority rule, subtracting the number of voters who registered negative sats from those who registered positive sats, or in the last resort, to a stochastic process). This hypothetical referendum illustrates the ideal best possible decision that we aspire to most closely approach in real-world voting. It is based upon complete, accurate, comparable and sincere cardinal information measuring the satisfaction of each voter on the issue or on each candidate. Such accurate and complete information enables correct decisions to be made every time.

In the general case where voters are to select n of m options, we set up our voting booth satometer to grab a quick reading on each of the m options from each voter. In essence, we conduct a separate referendum on each option. This is indeed a complete and very large data set. The winners will be the n options with the highest net sat totals (or in extremely sad cases, the positive-most or least negative totals, assuming we choose to allow such "winners").

Characteristics of Real-World Voter Data – Lacking a satometer, obtaining such information-rich and *sincere* cardinal data from voters is extremely problematic. (At least one method has been proposed (True Weight Voting) which hopes to obtain and utilize such data.)

Voters do not hesitate to vote strategically if they think it will "weaponize" their ballot to exert increased influence in a direction they favor. Enough voters have to be able to figure out a strategy in order for it to have an impact. Various voter strategies may offset each other to some extent. We know strategic voting can affect any voting method to some degree. It could be an acceptably small degree or a large level which greatly degrades the worth of the method (e.g., voting for the "lesser evil" with Plurality). There have even been instances where some percentage of strategic voters improves the ability of a certain method to identify the correct winner. Ideally, we prefer a voting method which works well with sincere voters and minimizes strategic opportunities/motivations.

A voter's satisfaction for a given candidate may be strongly positive, strongly negative or anything in between. It can be zero either because they have no strong opinion regarding a candidate, or it can be zero because they simply know nothing about that candidate. Write-in candidates must be handled. There are a surprisingly large number of zero opinions for voter/candidate pairs in real elections. There is no way to obtain a meaningful ranking of all candidates from all voters.

It is an unjustifiable gross simplification to say that voters' satisfaction with various candidates is a function of the location of the candidate relative to the voter along a one-dimensional left-right line. The two-dimensional Nolan chart is much, much better, but even it is quite an oversimplification. In fact, there exists a multi-dimensional "issue space" with voters' coefficients of importance for various issues varying dramatically. Any worthwhile voting method must not crumble in the face of some pesky Condorcet cycle.

Axiomatic Approach – Perhaps the predominant approach over the centuries has been to write down a set of axioms that a voting method should satisfy. The axioms exclude various members of the set of all voting methods and it is hoped that at least one remains standing. The chosen axioms then *are* the definition of "best possible decision," not the simple definition as stated above. This could be useful, but seems like a rather indirect approach. Of course, the selection of axioms is crucial.

Much importance has been given to a set of “fairness” axioms. As previously stated, fairness isn’t a criterion, but the adopted definition does happen to provide a good and sufficient definition of it. In the axiomatic case, the definition of fairness is instead provided by the chosen axioms.

The inclusion of an ordinality axiom cannot be justified on the basis of eliminating strategic voting (a hand-waving argument at best), and it would summarily exclude substantially all of the best voting methods.

An axiom which *should* be included in the collection is that no candidate for whom total voter satisfaction is negative should ever win an election. A weaker version might also be acceptable: No candidate that a majority of voters dislike should ever win an election.

Assuming that more than one method remains after eliminations, the axiomatic approach does not provide much help choosing the best by quantitatively comparing their ability to consistently choose the correct winner. There has also been the “Condorcet winner worship” phenomenon. A Condorcet winner surely should command considerable respect, but it is easy to construct a hypothetical election which has a Condorcet winner that no reasonable person would say should win the election; and the underlying reason why Condorcet winners are not always the correct winner is understandable.

Data-driven Approach – Physicists have an annoying penchant for insisting upon, you know, some actual data to confirm or refute a hypothesis, or to guide further inquiry. However, acquiring reliable and useful data from actual elections is extremely difficult. Furthermore, gathering the large volume of data required to have confidence in conclusions for such a huge statistical problem seems prohibitively daunting and hazardous.

Fortunately, modern digital computers provide a powerful tool to model entire complex elections and to quickly gather detailed statistics on any desired aspect to any reasonable statistical significance. It is entirely practical to test many voting methods and compare them quantitatively. Some forms of strategic voting can easily and accurately be simulated; other kinds can be simulated with more effort and risk. For simulated elections, we actually do have a “satometer.” It is as though we have an X-ray or CT machine and can see inside elections.

Surprisingly, it appears there has been little serious effort in this area. There was a project circa 2000, but it unfortunately had some known problems and shortcomings. How elections are simulated does matter. More recently, there was a project in 2019 and one in 2020. Those produced a plethora of intriguing results which certainly pass the common sense test.

In 100,000 four-candidate elections with 10,000 sincere voters voting in each, there are typically 51,007 Condorcet winners, but 2,167 of them are *not* the correct winner. There are typically 21,385 majority winners, but 500 of those are *not* correct winners. The most important single datum is each voter's choice of the best candidate. The second most important single datum is each voter's choice of the worst candidate. After the top two, few voters have any additional useful information to contribute and that information is not of much additional help in picking the correct winner. IRV is only a small improvement over Plurality; it still exhibits 75% of the error Plurality has. Two new methods were designed which allow voters the option of picking a best and a worst candidate; these two methods have only 21% of Plurality's error.

The first new method is Approve/Approve/Disapprove Voting (AADV). Voters have the option of approving up to two candidates and disapproving one. It is directly scored and very simple in all respects. It can be viewed as the corrected version of Approval Voting.

The second new method is Best/Alternate/Worst Voting (BAWV). Like IRV, it is an iterative eliminations method, still simple for voters, but more complex to tally. BAWV is somewhat more immune to strategic voting and can be viewed as the corrected version of IRV. Voters have the option of picking up to one Best candidate and one Worst candidate. Those who choose a Best candidate (substantially all voters) have the option of also indicating an Alternate candidate. Alternate choices have absolutely no impact on anything unless and until the voter's Best choice is eliminated, at which time the Alternate is promoted and counted exactly as though it had been the original Best choice.

Open-source software is already available which supports the use of BAWV, AADV or Plurality for each race in real elections.

Replacing Plurality with either AADV or BAWV would have a profoundly beneficial effect on elections. With Plurality, political parties' best strategy to win is to nominate candidates who appeal to and motivate their base voters to get out and vote, thus exacerbating the extreme polarization we are experiencing. With AADV or BAWV, polarizing candidates with "high negatives" would have a much harder time winning. Parties would quickly realize that in order to win they need to nominate candidates with few negatives and much broader appeal.

Any voting method which does not give voters a way to express their dissatisfaction with any candidate is prone to committing the unforgivable blunder of electing a candidate that most voters dislike! Plurality actually does make this blunder, much to everyone's detriment.

Recent election simulation results certainly appear to have a lot of validity. They are the fruits of perhaps two person-years of work. (This is miniscule compared to what is more like 10,000 person-years of axiomatic efforts.) However, *election simulation results can be relied upon only to the extent that real-world elections actually were meaningfully simulated*. Other simulation projects need to be conducted to confirm, refute or extend the results to-date. In addition to other measures, the relative performance of voting methods has been cast in terms of a pure number: the percentage of the error of a “standard voting method” which simply selects the winner randomly. It is hoped that this measure may be reasonably independent of the particular way in which simulations are done and thus facilitate making comparisons.