Utilitarian Approval Voting

by John C Lawrence

This is a PrePrint

j.c.lawrence@cox.net

April 26, 2025

Abstract

Approval voting (AV) is a system in which each voter gives one vote to each candidate that is approved and a zero to the others. Approval voting overcomes the spoiler effect that first-past-the-post or plurality voting creates. Utilitarian voting (UV) is a system in which each voter rates each candidate on some scale, and then the ratings are summed to determine the winner. Preference approval voting (PAV) combines approval voting with preference rankings to determine a winner. Utilitarian approval voting (UAV) involves placing a threshold in the individual utilitarian ratings and converting all those above threshold into approval votes. Where to place this threshold is left up to the voter in AV and PAV, but UAV combined with an optimal threshold mechanism (OTM) provides a method of placing the threshold in an optimal manner. Just as PAV combines approval and preference information to determine the winner, UAV combines approval and utility information to determine the utilitarian winner, the one who maximizes social utility.

Introduction

It is a well known phenomenon that third party candidates can cause the so-called spoiler effect in certain types of voting systems such as first-past-the-post or plurality voting which is the type used in almost all US elections. In this type of election the winner is the candidate with the most votes even if it is not a majority. The spoiler effect occurs when a third party candidate who has no possibility of winning the election draws votes away from the candidate with whom they are most closely politically aligned, and this results in the candidate with the least public support winning the election.

In the Presidential campaign of 2000, Ralph Nader ran as the nominee of the Green Party. The Democratic nominee was Al Gore, and the Republican nominee was George W. Bush. (Wikipedia) On election night, it was unclear who had won, with the electoral votes of the state of Florida still undecided. The returns showed that Bush won Florida by such a close margin that state law required a recount. A month-long series of legal battles led to the highly controversial 5–4 U.S. Supreme Court decision, Bush v. Gore, which ended the recount. Ultimately, Bush won Florida by 537 votes, a margin of 0.009%. Nader received 97,421 votes in Florida. Reform Party candidate Pat Buchanan and Libertarian party candidate Harry Browne received 17,484 and 16,415 votes respectively. Buchanan and Browne probably took votes away from George W. Bush since they were on the conservative side of the spectrum. However, Nader's campaign probably took many more votes from Gore than Buchanan and Browne took away from Bush. If Nader had not been in the race, Gore probably would have gotten the majority of Nader's votes in Florida, and since the Florida vote turned out to be critical in deciding the election, Nader voters did indeed tip the balance to George W. Bush. Ultimately, Bush won 271 electoral votes, one vote more than the 270 required to win, while Gore won the popular vote by 543,895 votes (a margin of 0.52% of all votes cast).

Approval Voting

Approval voting is a system which eliminates the spoiler affect. In approval voting, the voter casts one vote for every candidate they approve of. This in effect gives 1 vote to all approved candidates and 0 votes to all the rest. For example, let's say the candidates are Green, Red and Blue with Green and Blue being on the left side of the political spectrum and Red being on the right side. Therefore, Green is much closer to Blue, politically, than to Red. Liberal voters have 2 candidates to consider, and conservative voters have only one. If the votes of liberal voters are split between Green and Blue, with plurality voting, Red will win even if the total of Green and Blue votes is greater than the total of Red votes. If Green or Blue had dropped out of the race before the final vote, a liberal almost surely would have won. If the number of voters preferring Green to Blue were much less than the number preferring Blue to Green, then Green could be considered the spoiler.

Now consider the same race but with a different voting system: approval voting. Voters give one vote to every candidate they approve of so that liberal voters could vote for both Green and Blue with the result that Blue, having more votes than either Green or Red, would win. Approval voting eliminates the spoiler effect. So what is the drawback to approval voting? Surely, all a voter would have to do is to place all candidates on the left-right political spectrum, and either vote for all candidates on the left if one is a liberal or vote for all those on the right if one is a conservative. However, what if it's not clear where to place the candidates on the spectrum? For instance, a particular candidate might have conservative views on some issues and liberal views on others. What if a voter who mostly preferred liberal candidates actually liked some conservative candidates? Where would one draw the line between approved and unapproved candidates or would it even make sense to draw a line? Why not just approve all candidates that a voter likes whether they be liberal or conservative?

Utilitarian Voting

Utilitarian voting is a method of voting which allows a voter to give more information about each candidate. In utilitarian voting one rates all the candidates on some scale such as the scale of real numbers between zero and one. Then the ratings are added up over all voters to determine the outcome of the election. This would result in an ordered set over all candidates. If the election were for one winner, then just the candidate with the highest rating would win. However, there is also the possibility of an election to choose candidates for several seats such as a committee. In that case, if there are m seats, the top m candidates would be chosen for the winning set based on the final results of adding their utilitarian ratings over all voters. Claude Hillinger (2005: pp. 295-321) has made the case for utilitarian voting: "Utilitarian collective choice assumes that individual preferences are given as cardinal numbers; social preference is defined as the sum of these numbers."

Let's say there are two candidate, A and B, and the voters vote utilitarian style. We will assume that voter 1 has a utility of .8 for candidate A and .4 for candidate B. Obviously, this voter prefers that candidate A and not candidate B becomes the winner of the election. Rather than submit his sincere utilities to the voting system, this voter can vote strategically giving candidate A a utility of 1 and candidate B a utility of 0. Lehtinen (2008) has shown that strategic voting is actually desirable because it tends to produce the utilitarian winner, the winner who maximizes social utility. If there are several candidates in the race, a voter could list his sincere utilities for each candidate as numbers between 0 and 1. Then voting strategically, they would elevate some candidates to a vote of 1 and lower some of them to zero. Voting strategically with utilitarian voting devolves into approval voting. Some candidates are given a vote of "1" and some are given a vote of "0". The question is where to draw the line so as to maximize the power of this individual vote and such that their utility in the outcome of the election is maximized.

Preference Approval Voting (PAV)

The winner in a PAV election is determined by two rules:

1. If no candidate, or exactly one candidate, receives a majority of approval votes, then the PAV winner is the AV winner—that is, the candidate who receives the most approval votes.

2. If two or more candidates receive a majority of approval votes, then

(i) If one of these candidates is preferred by a majority to every other majority approved candidate, then he or she is the PAV winner—even if not the AV or Condorcet winner. (The Condorcet winner is the one who is preferred to every other candidate in binary comparisons.)

(ii) If there is not one majority-preferred candidate because of a cycle among the majority-approved candidates, then the AV winner among them is the PAV winner—

even if not the AV or Condorcet winner.

According to the Brams and Sandver paper, we are considering a multiple candidate, one winner election. First we count all the approval votes for each candidate. Secondly, we consider all candidates who received at least a majority of approval votes, two by two. For instance, if the candidates were A, B, C, D and E, and A, B and C got at least a majority of the approval votes, we would count the preferences of A compared to B, A compared to C. If A were preferred to B by more voters than those who preferred B to A, and A were preferred to C by more voters than those who preferred C to A,, then A would be the PAV winner by rule 2(i). If there is a cycle among majority preferred candidates (eg. A is preferred to B, B is preferred to C, and C is preferred to A), then the candidate with the most AV votes is the winner.

Utilitarian Approval Voting

With utilitarian approval voting (UAV) each voter rates the candidates on a scale, and then a threshold is placed such that all those above threshold get an approval vote of 1 and all those below get an approval vote of 0. Each voter expresses his utilities on a scale of his own choosing. By way of contrast, Preference Approval Voting (PAV) involves placing a threshold in an ordinal list of candidates, a ranked list as opposed to a rated list. In PAV a line is drawn between approved and unapproved candidates, but the placement of the line is strictly up to each voter based on their own intuition. If UAV is combined with an Optimal Threshold Mechanism (OTM) (Lawrence, 2023) the placement of the optimal threshold can be calculated rather than guessed at. We assume that there are multiple candidates and each individual voter has a utility for each candidate. We consider here that there is one winner of the election although this method can be extended to the case of multiple winners. The question is where the threshold is to be drawn such that those below threshold are given "0' votes each and those above are given "1" vote. The optimal threshold is the threshold which results in the maximization of the expected value of average utility of above threshold candidates for each voter. As the threshold increases, there are fewer candidates above threshold, the average utility rating of the set of candidates above threshold increases, and the probability of selection of any particular candidate in this set decreases. Conversely, as the threshold decreases, the number of candidates above threshold and the probability of selection of one of them increases while the average utility rating of the set of candidates above threshold decreases.

Lawrence (2023) has shown a method by which the voting system itself applies the strategy as part of the vote collection and amalgamation process so that sincere utilities are elicited from the voters. The voters then have no incentive to vote insincerely. This makes it possible to compute the actual social utility of the winner and also individual utilities of the outcome of the election for each individual voter.

In UAV with OTM all preferences are known to the system since it is assumed that all voter inputs are sincere utilities for all candidates. If A has more utility than B, then, obviously, A is preferred to B. The outcome of the election would be a ranked list of approval votes over all candidates. A list of all binary comparisons among the candidates could also be computed and rules 1 and 2 for PAV applied to find the PAV winner. However, there would not be any cycles among candidates since utilitarian voting would produce precise differences among all the candidates. There could be ties though, but that case will not be considered here. Therefore, Rule 2 of PAV would be unnecessary. It would only be necessary to compute binary comparisons among those candidates who received at least a majority of the votes. Let's say the candidates are A, B, C, D, and E. Also that candidate A is the AV winner having received the most AV votes. Also let's assume that A, B and C received at least a majority of votes. Since all the utility information from the individual voters is known to the system, we could compute the number of voters who preferred A to B, A to C and B to C. The maximum of these computations would be the PAV winner.

Similarly, the UAV winner is not necessarily the one with the largest number of approval votes. The approval winner is not necessarily the one that maximizes utility. Let's say that the top approval rated candidates are A, B, C, D and E. Let's assume that the number of approval votes for A is maximum with B, C, D and E the next highest approval vote getters in that order. Just because A has more approval votes than any other candidate does not necessarily mean that A is the utilitarian winner i.e. the one that maximizes utility. However, since we have exact utility information about the top ranked winners, we can calculate the utilities for all those who have gotten at least a majority of approval votes. Similarly, for PAV we considered the AV results for those candidates who got at least a majority of AV votes. Then we could choose that candidate who had the maximum utility to be the UAV winner even if they were not the one with the most approval votes.

Summary and Conclusions

We compare two voting systems: approval voting (AV), utilitarian voting (UV), and their variations, utilitarian approval hybrid (UAV) and preference approval (PAV) voting. AV solves the spoiler problem which affects plurality voting and which has determined the outcomes of several US Presidential elections, most recently the 2000 election in which the main candidates were George W Bush, Al Gore and Ralph Nader. Because of the spoiler effect, Ralph Nader drew votes away from Al Gore with the result that George W Bush won. Because AV voters vote for every candidate they approve, they would have been able to vote for both Gore and Nader with the result that Gore would have won.

AV leaves the problem of which candidates to approve and which to disprove up to the voters. They can expand or contract their approval sets strategically based on where they feel the line should be drawn to give them the best advantage.

In UV each voter assigns a utilitarian rating to each candidate, and then the ratings are summed over all voters and candidates to determine the winner. The problem with UV is that rather than submit sincere ratings to the system, it is advantageous to strategically alter the ratings in such a way as to maximize individual voter utility in the outcome of the election. In order to do this some candidates will be given a rating of "1" and some a rating of "0". Even if this is acceptable, the problem is where to place the threshold between the two sets of candidates in the optimal way.

PAV is a method which combines ranking and approval information to select a winner who may not be the most approved candidate. According to Brams and Sanver (2006), the AV winner may be the most popular candidate, but the PAV winner could have a more coherent point of view, and, therefore, would be the preferred winner. PAV and AV leave it up to the voter to strategize by expanding or contracting their approval sets.

UAV with the Optimal Threshold Mechanism (Lawrence, 2023) resolves the issue of where to place the threshold between approved and unapproved candidates. If the system itself implements this mechanism, it doesn't make sense for the individual voters to do anything but to vote sincerely. This makes it possible to compute accurate social and individual utilities based on the outcome of the election. Combining approval and utility information makes it possible to pick the candidate who maximizes utility as the winner rather than the one with the most approval votes. This candidate then would be the utilitarian winner.

References

1.Brams, Steven, Sanver, Remzi (2006) Voting Systems that Combine Approval and Preference. ANNALESDU LAMSADE N6 <hal-00119047>

2. Hillinger, Claude (2005) The Case for Utilitarian Voting. Homo Oeconomicus 22(3).

3. Lawrence, John (2024) Utilitarian Social Choice With a Maximin Provision. Preprint online at https://www.socialchoiceandbeyond2.com/uscmm.pdf

4. Lawrence, John C (2023) Proving Social Choice Possible, Preprint: https://www.socialchoiceandbeyond2.com/proving.pdf

5. Lehtinen, Aki (2008) The Welfare Consequences of Strategic Behaviour Under Approval and Plurality Voting. European Journal of Political Economy 24(3).

6. Wikipedia, https://en.wikipedia.org/wiki/Ralph_Nader_2000_presidential_campaign