A Test Drive of Voting Methods

By William Poundstone

William Poundstone is author of Gaming the Vote: Why Elections Aren't Fair (and What We Can Do About It) (Hill and Wang, 2008). E-mail: william.poundstone@sbcglobal.net

Mathematics is not normally a subject that inspires heated controversies. A notable exception is the mathematics of voting. "What is the fairest way of voting?" is a simple question, and you might hope that it would have a simple answer. It doesn't. This issue is the biggest, longest-running dispute in voting theory.

There is little controversy over what voting method is *worst*. That would be the plurality vote — the system used in almost all American elections. Any voting theorist will tell you that the plurality vote is especially vulnerable to vote splitting. When two candidates' base of support overlaps, both candidates are penalized. This is most easily seen in a case like the 1912 Presidential election, where two Republicans (William Howard Taft and Teddy Roosevelt) ran against each other, splitting the Republican vote. The result was the election of a Democrat (Woodrow Wilson) who probably wouldn't have won otherwise. A more common example of vote splitting is the familiar political phenomenon of a spoiler. In 2000, Ralph Nader almost certainly took enough of the liberal vote away from Al Gore in the crucial state of Florida to hand the election to George W. Bush.

Spoilers aren't uncommon. Of the U.S. presidential elections since 1828, at least five were probably decided by spoilers. In effect, we've elected the "wrong" president eleven percent of the time.

So what's the solution? Everyone agrees that we need to collect a little more information from voters. The ballot should be designed so that voters can express how they feel about the candidates who weren't their first-place choices. It is also necessary to find a way to make use of that additional information in tallying the ballots.

There are two obvious approaches. One is to ask the voters to rank the candidates in order of preference. Dozens of ranked-ballot methods have been devised. The two most historically important are named for Jean-Charles Borda and the Marquis de Condorcet, rival members of the French Academy in the eighteenth century. Borda and Condorcet sparred over the merits of their respective systems, initiating a feud whose skirmishes continue to the present day. Though Borda's and Condorcet's schemes can use an identical ballot, their methods of tallying ranked votes are completely different, and sometimes, so are the winners.

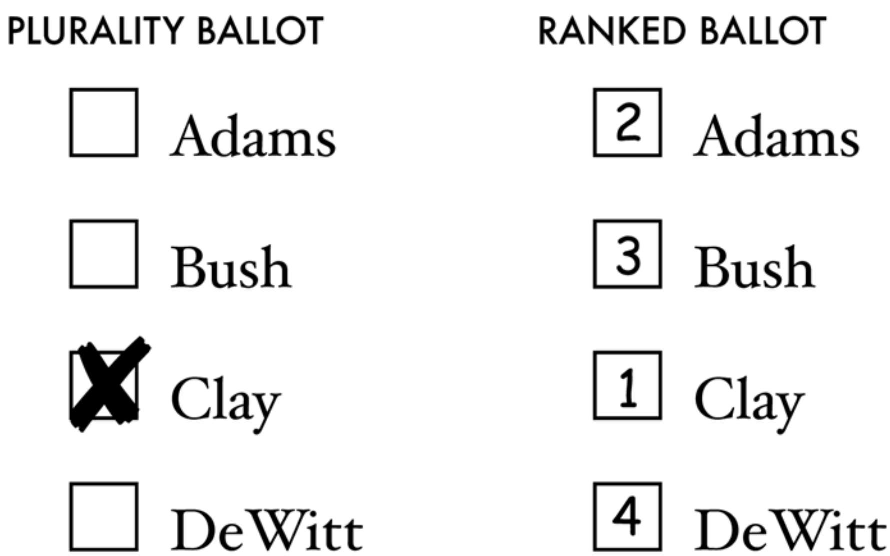

In Borda's system, the rankings are converted into points. In an election with four candidates, every first-place ranking is worth 3 points, second-place is 2 points, third-place is 1 point, and fourth (last) place is worth nothing. The points are added, and the candidate with the most points wins.

Sounds reasonable? Borda thought so. Condorcet objected that the winner of an election should be able to beat every other candidate in two-way votes — something not guaranteed by Borda's system. Today, computers can readily determine the Condorcet winner by using the ballot rankings.

Another important ranked system is instant runoff voting (IRV), a nineteenth-century invention that is now used in Australia, Ireland, and other nations. IRV simulates a series of runoffs in which the least popular candidate of each round is eliminated. Each time a candidate is eliminated, ballots ranking that candidate highest are transferred to the highest-ranked of the remaining candidates. That way, votes for minor candidates are not "wasted" and ultimately count toward the voter's preferred front-runner.

Though these capsule descriptions of IRV, Condorcet, and Borda may all sound perfectly reasonable, it is possible to have an election in which the three systems produce three different winners from the same set of ranked ballots.

A fundamentally different approach to voting is to ask voters to rate or score the candidates. In approval voting, the ballot is essentially a report card with pass-or-fail grades. The candidate who gets the most passing grades ("approval votes") wins. A generalization of approval voting called range voting has voters rate the candidates on a numerical scale — say, 1 to 10. The candidate with the highest average score wins. Approval and range voting are much more closely related to each other than the ranked systems are but, yes, it's possible for range and approval voting to have different winners.

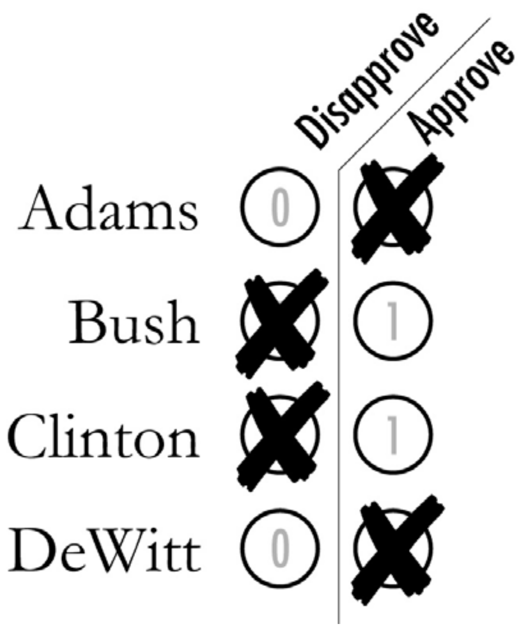

Which system is best? There is no easy answer because we are dealing with the paradoxical mathematics of voting, the human behavior of voters and candidates, and philosophical questions about what it means for a voting method to be fair. Over the past twenty years, the ensuing debates have resulted in frayed nerves, strained personal relationships, and very little consensus. Donald G. Saari, a Borda proponent, and Steven Brams, an approval voting advocate, have found it necessary to avoid discussing their differences in order to remain friends. No such pact of silence restrains the raging war of words between range voting theorist Warren D. Smith and

Rob Richie, a political activist whose organization promotes IRV. The ongoing lack of agreement has surely hampered efforts to replace the plurality vote with something better.

I'd like to showcase a novel way of visualizing a few of these abstract issues. It was devised by Ka-Ping Yee, a recent graduate of Berkeley's computer science Ph.D. program. Yee became interested in voting in 2004, when Berkeley adopted IRV for its city elections and the Presidential race raised concerns about electronic voting. He read up on Kenneth Arrow's impossibility theorem and tried to convince the members of his Berkeley residence, Kingman Hall, to adopt approval voting. The motion failed, but the residence later adopted Condorcet voting, another method Yee had discussed.

Much of the voting literature focuses on what can go <u>wrong</u> with electoral methods — on mathematical paradoxes (often rare) and aberrant voter behavior (often conjectural). These are unquestionably important matters to discuss. But Yee took the opposite tack. Suppose that every voting method works exactly the way its proponents want it to work. Which method would be fairest then?

To answer this, Yee used computer simulations to generate colorful "maps" of multicandidate elections. In these simulations, the public's views on political issues form a normal distribution. That is the familiar bell-shaped curve that describes a wide range of natural variations. Don't get too hung up on this normal distribution business. Yee isn't saying that political beliefs follow a normal distribution, only that this provides a good test drive for voting systems. It eliminates the more contrived paradoxes and restricts attention to a simple case where it is possible to have a good intuitive sense of what an election's outcome ought to be.

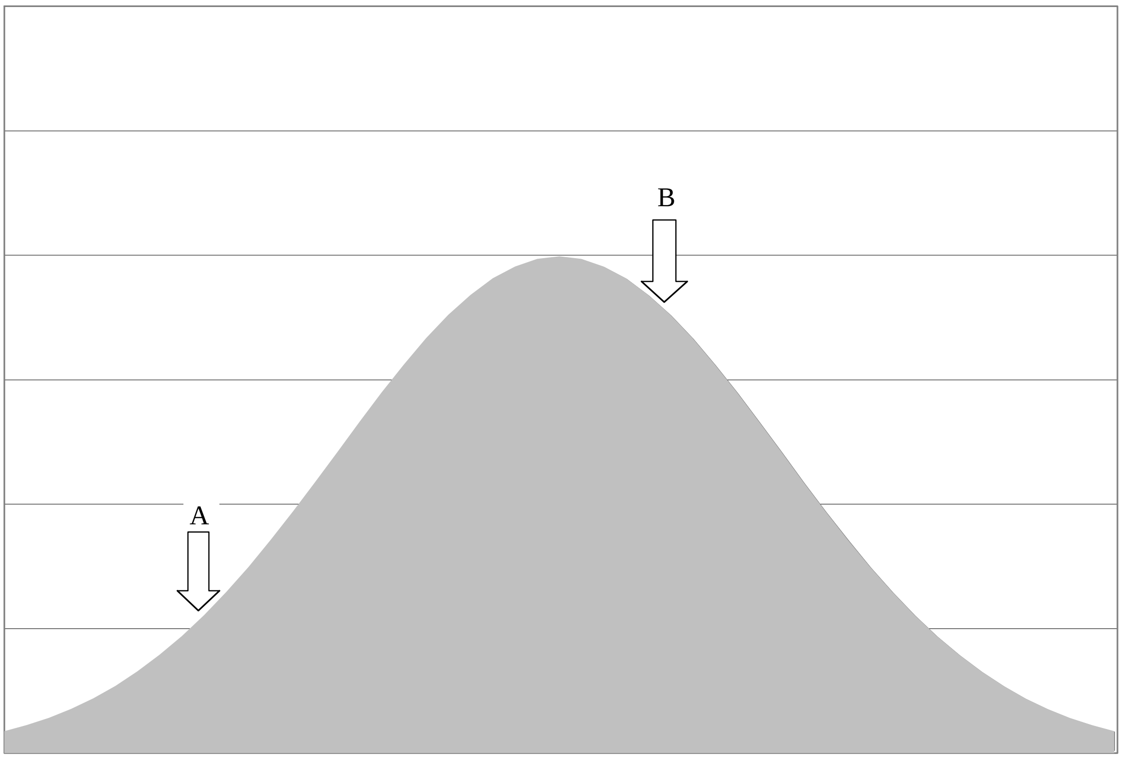

This diagram's horizontal scale represents the range of possible convictions on a particular issue. The height of the curve at any point tells how many voters favor that particular position. The curve is highest somewhere in the middle. This represents the center of opinion. As you go further to the left or right, the height of the curve diminishes.

There are two candidates, A and B. B is a little to the right of center, and A is well to the left of center. Should the candidates otherwise be equally qualified and appealing, and should the issue represented by the chart be the only one that matters, you would expect that B would beat A under any reasonable voting system. B is closer to the political mainstream.

In the diagrams below, Yee postulates two independent political issues. After all, most campaigns have many issues at stake. You might think of the two here as small v. large government and social conservative v. social liberal. A candidate who takes a given position on one issue is free to take any position at all on the other issue. You therefore need a two-dimensional square to chart the range of possible positions.

With the additional dimension, the normal curve becomes a three-dimensional bell or a rounded hillock in a flat plane. The top of the bell or hill is the center of voter opinion.

Yee's simulations use 100,000 virtual voters, a mid-sized city's worth, scattered in that kind of symmetrical bell curve. Each voter marks a ballot based on distance between the voter's and the candidates' views in the two-dimensional issue space. Every voter is assumed to vote honestly, without any strategic calculation. In Yee's diagrams, the color of the square at any particular point indicates which candidate would win, should the center of voter opinion coincide with that point. If that sounds confusing, don't worry. The best way to get the hang of it is to look at some examples. (You can see more at zesty.ca/voting/sim.)

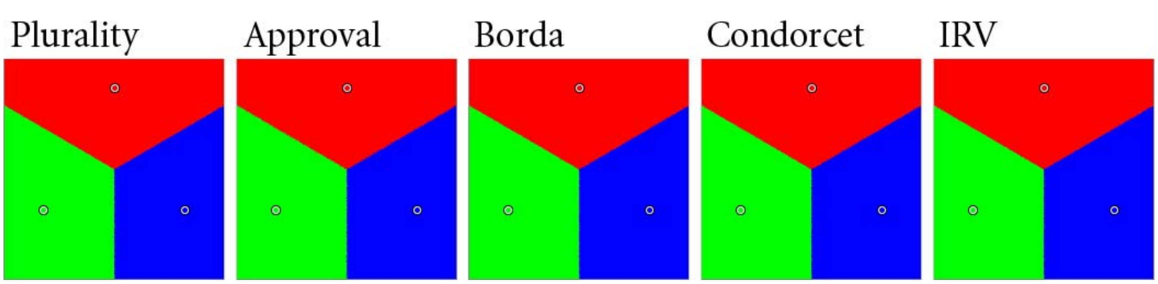

Each square represents the same election as it would be tallied under five voting systems. The candidates are shown as little circles. In this particular case there are three candidates forming a triangle. This could represent a three-way race between a left-wing candidate, a right-wing candidate, and independent candidate running on a single issue like abortion or immigration.

In this three-way election, all the squares are near-identically partitioned into three sectors. Each candidate is located in the center of his or her own sector. That means that if the average voter's views are close to those of a particular candidate, that candidate would win. That's exactly the way we'd want it to be. The fact that there are three candidates and three sectors means that all of the candidates have a fair shot at winning.

Because all the squares look the same, the outcome of this election does not depend on the choice of voting system. It depends only on the expressed will of the voters. This is how democracy is supposed to work. The bad news is, this is a very special case.

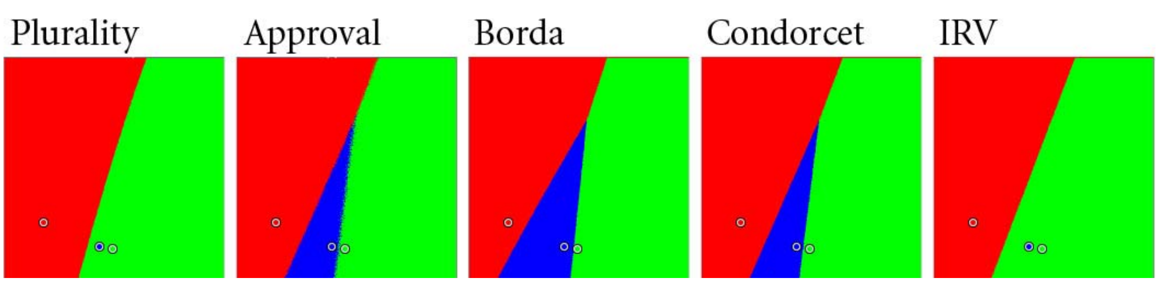

These five diagrams show an election with two "clones." Clones are candidates so similar that they appeal to the same voters. An example might be two similar Democrats running against one Republican. This could result in the clones splitting the Democratic vote and ensuring the election of the

Republican. Here the clones are the two close-spaced circles at lower center. Any voter who likes one will like the other nearly as much. A third candidate is well off to the left.

In the plurality vote chart, this leads to vote splitting. There are only two sectors and thus only two possible winners. The clone closer to the center (shown as blue) is squeezed out entirely. He cannot win, not even when public sentiment is dead-centered on his platform. The boundary between the red and green territories is close to the clone candidates. This reflects vote splitting. The red candidate wins unless the voters overwhelmingly favor the clones.

Approval, Borda, and Condorcet do better by the blue candidate. In these diagrams, there is a slice for each candidate, and each candidate is within his own slice. Should the political center of gravity coincide with any candidate's views, that candidate wins.

The IRV diagram looks much like the plurality diagram. But notice that the boundary between the red and green candidates is now about halfway between them. This is more equitable than with the plurality vote. It reflects the fact that the green clone will be the second-place choice of the shut-out (blue) clone's voters. IRV will redistribute the blue votes to the green candidate. Of course, the blue candidate may not find this fair — he can't win, no matter what.

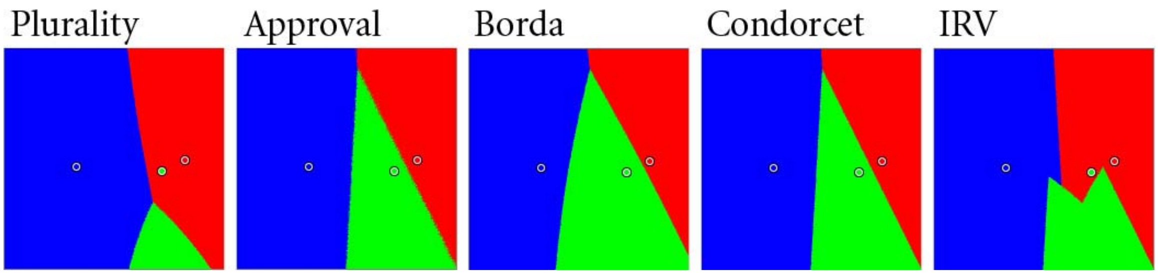

This is another case of vote splitting. The weird thing here is the M-shaped win region of one of the candidates under IRV. The chart reflects the "winner-turns-loser paradox."

This was described in a 1977 article with the pointed title, "Single transferable vote: An example of a perverse social choice function." ("Single transferable vote" is the multiple-winner form of IRV.) The authors, Gideon Doron and Richard Kronick, showed that it is possible for a voter to make an IRV candidate lose by ranking him *higher*. Huh?

Here's an example, supplied by Yee. Imagine an alternate universe in which Ralph Nader is the most popular candidate of 2000, and IRV is the voting system. Thirty-nine percent of the voters are for Nader, 31 percent are for Gore, and 30 percent are for Bush. The Bush voters rank Gore second and Nader last. The Gore voters are split about evenly between Nader and Bush as their second choice.

With IRV, when no candidate has a majority of first-place votes, the lowest-ranking candidate will be eliminated. That's Bush. All of Bush's votes would be transferred to the Bush voters' second-place choice, Gore, giving Gore an easy 61 to 39 percent victory over Nader.

Okay. Now say that Bush gives a really impressive speech in a Nader stronghold. A few Nader voters (2 percent of the total electorate) are so swept away that they switch their votes to Bush. Instead of ranking Bush last, they rank him first.

The resulting numbers are now 37 percent of first-place votes for Nader, 31 percent for Gore, and 32 percent for Bush. See what happens? Gore is now in third place. It’s Gore who is eliminated. About half his vote goes to Bush, and half to Nader. Nader beats Bush by about 52.5 percent to 47.5 percent.

Voting for Bush instead of Nader caused Nader to win. This is crazier than the spoiler effect. In principle, you could have a candidate taking out ads telling people not to vote for him… because polls show that, if he gets any more votes, he’ll lose!

The weird M shape in the diagram is the visual expression of this paradox. Imagine that public opinion is centered within the left corner of the M-shape. That area is colored green, meaning that the corresponding candidate is winning. But the little dot representing the green candidate is outside of the M shape. As public opinion drifts closer toward the green candidate’s position, it will move out of the M shape and cause one of his opponents to win.

When Yee’s territories are simple slices, changes in the public’s views translate logically into changes in the candidates’ fortunes. When a candidate is winning, and something happens to make him even more popular, the candidate will still be winning. When territories are more complex, changes can have the *opposite* effect. Increases in popularity can cause a winner to lose. Decreases can cause a loser to win.

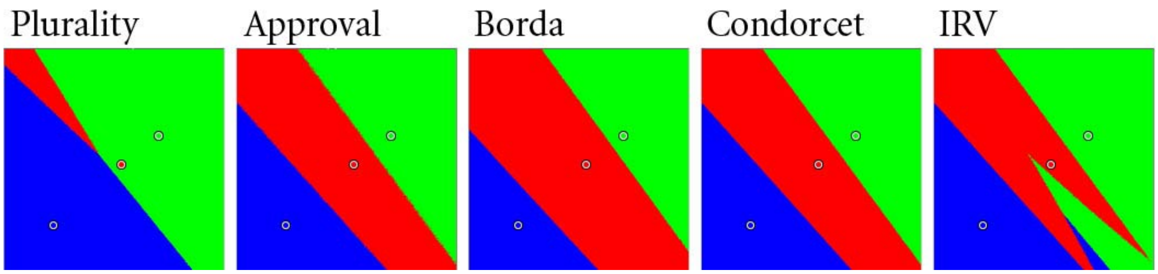

Here's another example of an oddly shaped IRV territory. It's all the easier to see how the center of voter opinion could shift in a straight line, causing the IRV winner to flip-flop a way that doesn't make any sense.

IRV is especially prone to paradox with four or more candidates. Look at the IRV win region here:

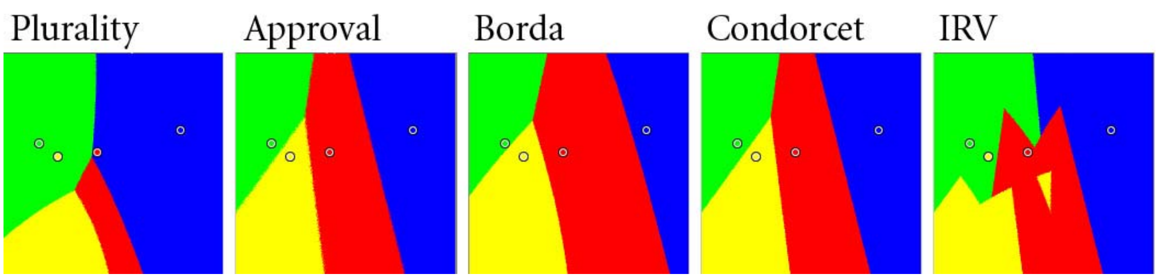

And here:

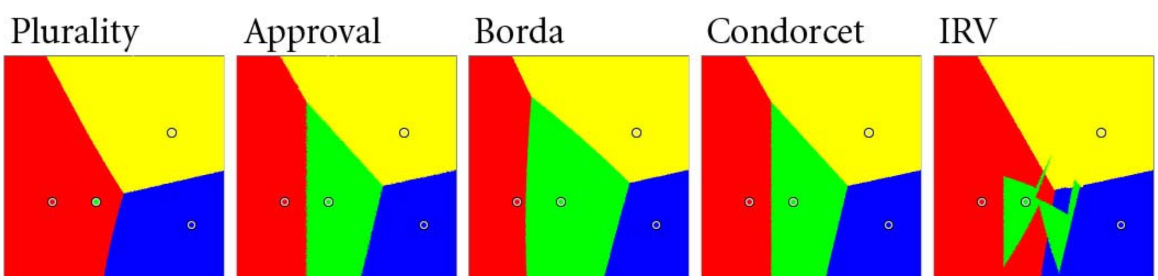

In this last race, one of the four candidates, shown as green, is within the triangle formed by the other three. This would be a compromise candidate. You might expect that a reasonable voting system would elect this candidate whenever the center of public opinion is sufficiently close to his views. With IRV this candidate has a crushed-beer-can-shape win region. But that's probably better than the plurality vote, where the compromise candidate has no win region at all.

Yee's Borda charts don't look all that different from the Condorcet and approval charts (which look nearly identical — and in case you're wondering, range voting would look nearly the same, too). This may come as a surprise in view of the two-century feud over Borda and Condorcet, and the more recent debates over approval and range voting.

What you see isn't necessarily what you get. For these simulations, Yee is assuming that everyone votes sincerely. Many believe that Borda is especially prone to strategic voting, for instance, so this diagram in particular might be taken with a grain of salt.

The approval and Condorcet diagrams aren't quite identical. Look closely, and you will notice that the boundaries between regions are grainy with approval voting and sharp with Condorcet (or the other systems).

This reflects the fact that there is more than one way of casting a sincere approval ballot. The voter is free to make a judgment call on which candidates are worthy of approval. To model this, Yee randomly assigned each voter a threshold distance, using a log-normal distribution. The voter approves only those candidates who fall within his threshold. You might say that each of Yee's virtual voters has a unique sense of political entitlement. Some voters are very picky and judge no candidate good enough to approve unless he is very close to their views. Others are more generous and tend to approve lots of candidates.

Because Yee assigns a random element to approval voters and not to the other systems, the boundaries of the approval voting territories blur. When there is a near-tie between two candidates, the race can be tipped by a few voters being generous or stingy with approvals. The diagram thus shows some “wrong color” points near the boundaries. (A chart for sincere range voting would eliminate the grainy boundaries.)

Not so expected is the virtual congruence of approval and Condorcet voting. Approval voting’s proponents have sometimes claimed that the system tends to elect the Condorcet winner. This argument has often been founded on the assumption that approval voters will be well-informed and strategic. In Yee’s simulation, approval voting elects the Condorcet winner even when no one acts strategically.

Critics have charged that approval voting would elect the least-common denominator, tofu not beef, the haircut who mouths platitudes. The conception of Condorcet voting is almost the opposite. Donald Saari has puckishly compared its premise to the ethos of a Western gunslinger: Condorcet voting is a shoot-out to be the last man standing. No one seems to worry that the Condorcet winner would be a mealy-mouthed politician.

These reputations may deserve a rethinking. Under some circumstances, the two systems are closer to being functionally equivalent than has been imagined.

Look at the difference between Borda and approval/Condorcet in the diagrams. Candidates in the middle, hemmed in by two or more rivals, command a bigger slice of territory under Borda than they do under Condorcet/approval.

This illustrates an anti-spoiler effect of the Borda system. With the plurality vote, it's *bad* for a candidate to have a clone or a spoiler. With Borda, it's *good*. Clones and spoilers increase a candidate's Borda total.

Here's the simplest example. There are two candidates, Kennedy and Nixon, and two voters. The first voter prefers Kennedy to Nixon. The second prefers Nixon to Kennedy. It's a tie.

Then a third candidate enters the race. Call him Dixon. Dixon is a clone of Nixon, parroting his every opinion and position paper. The only difference is that Dixon has a bad comb-over and his family is kind of obnoxious.

The first voter now ranks the candidates Kennedy > Nixon > Dixon. The second voter ranks them Nixon > Dixon > Kennedy.

In this three-candidate Borda count, we award 2 points for first place, 1 point for second place, and no points for third and last place. Kennedy gets 2+0=2 points. Nixon gets 1+2=3 points. And Dixon gets 0+1=1 point. Nixon, who was tied with Kennedy before, is now the winner. Instead of being hurt by a clone, he's been helped.

A spoiler effect is unfair, and so is an anti-spoiler effect. The Condorcet and approval voting diagrams carve up the square more equitably. Draw a line midway between candidates A and B. That determines the boundary between the region where A wins and where B wins. Draw another line equidistant from B and C. That divides their territory. The winner is whoever is closest to the center of voter opinion.

(I don't mean to leave the impression that Condorcet and approval — or range voting — are interchangeable. Use a more complex distribution of voter opinion and throw in the variables of human behavior, and they can produce different results. But that would take us beyond the scope of this quick test drive.)

IRV has something that all other alternative voting methods conspicuously lack: political momentum. In recent years, a nonprofit organization known as FairVote: The Center for Voting and Democracy, cofounded by Rob Richie in 1992, has spearheaded successful referenda campaigns to adopt IRV for local elections in Berkeley, San Francisco, Oakland, Minneapolis, and the state of North Carolina. Along the way, IRV has been endorsed by some heavyweight politicians — among them, Barack Obama and John McCain.

Like many voting theorists, Yee believes that FairVote has latched onto the "wrong" system. The diagrams make it easy to see why he feels this way. Even a quick scan shows that IRV looks different from the other systems. These differences demonstrate that IRV has a problem delivering a fair or reasonable outcome with three or more candidates. It's true of course that Yee's model of politics is vastly simpler than the human reality. But Yee's assumptions are generous ones, giving IRV and the other systems the benefit of the doubt. This is how IRV performs when everyone ranks the candidates honestly, just as they're supposed to do.

The "winner turns loser" paradox is common to all procedures using candidate eliminations and runoffs. You would have similar glitches with two sequential plurality votes, a primary followed by a runoff. This pattern describes much of our political process, of course. The American communities that have adopted IRV have mostly used it to replace two plurality elections with a single IRV election. The sales pitch has been that IRV saves money and improves voter turnout. These are accurate claims and worthy goals. The bigger picture is that runoffs — instant or otherwise — are a relatively poor fix for the problems of a single plurality vote.

IRV's political traction may have much to do with an issue that is all-too-often overlooked. What would any of these "improved" voting systems do to party politics?

In his farewell address, George Washington pleaded with the young nation to avoid having political parties. He thought that voters didn't need parties to tell them who to vote for. Washington's advice was ignored; he was arguing against the inevitable math of plurality voting. Without a two-party system to limit the choices to two, the plurality vote is prone to chaotic vote splitting. In the U.S. today, it's almost suicidal for a candidate to break ranks and run without a major party's nomination. Such a candidate ends up being a spoiler and hurting the very people who voted for him or her.

All of the improved voting systems promise to remove this spoiler penalty. Should we adopt <u>any</u> such system, politics is going to be different. Neither theory nor experiment can predict the future evolution of party politics under a voting method.

That can be a scary prospect. Democracy is such an important thing that we are properly risk-averse. We don't want to make a change unless we are sure it's a positive improvement with no downside whatsoever. It is tough to get that kind of assurance.

A few conjectures may be relatively secure. Without the spoiler penalty, there'd be a little more reason for party members who didn't get the nomination to run anyway — like Teddy Roosevelt or Joe Lieberman did. There'd be more reason for independents to run. This wouldn't mean the end of major parties. Like it or not, the realities of fund-raising would still rule politics. But a spoiler-proof method would mean that the major parties would have a little less power to be the gatekeepers of politics. We'd probably see a few more people on the ballot in important races. They wouldn't be spoilers, and some of them could have a shot at winning.

Of all the popularly advocated ranking and rating systems, IRV would probably change American politics the *least*. Notice that, in Yee's diagrams, the IRV chart often resembles the plurality chart. Both IRV and plurality often shut "moderates" out of a race. (I put "moderate" in quotes because it means here only a candidate who is ideologically between two others.) One example of this might occur in an election with a Republican, a Democrat, and an independent who expects to appeal to members of both parties — as an independent would have to do, to stand much chance of winning. Under IRV or plurality, the independent can fall victim to vote splitting from two sides — a so-called "center squeeze." That can take the independent out of the race, even when the other methods find the independent to have the strongest support of all.

IRV’s proponents say that more people will feel free to vote for third-party candidates, knowing that their votes will not be wasted. This is surely true. However, strong independents would still be penalized and find it unfairly difficult to win. Their supporters would have to be content with knowing that their votes were transferred to one of the major parties. This at least isn’t speculation. Australia and Ireland, which have used IRV for some time, have two-party systems much like we do. It’s unusual to see anyone without a major party nomination win.

There is then a fundamental philosophical difference IRV on the one hand, and systems like approval, range, and Condorcet on the other. This difference has not always been articulated and has often been obscured by the squabbling (important squabbling!) on how well these systems deliver on their promises.

IRV would fix the most glaring fault of the plurality vote system — namely, third-party spoilers unfairly tipping the race from one major party to the other. It would do this with minimal change to the two-party system. That’s because IRV can penalize popular independents, much as the plurality vote does.

The two-party system is, you might say, an unintended consequence of the plurality vote. The issue is not the ideologies of the parties, nor the fine people and hard work they help inspire. Rather, it’s the center squeeze, a mathematical quirk that none of the nation’s founders likely recognized. Having two popular candidates, one a little to left of center and the other a little to right of center, can prevent a yet-more-popular candidate from winning. Do we want to keep this feature of our politics?

I imagine George Washington would have said no. But you don't have to be a major-party hater to feel this way. You just have to believe that the voters, rather than a bias of the voting method, should decide how much power our two major parties deserve to have. Systems like approval, range, and Condorcet aspire to a level playing field for every candidate, no matter how many candidates there are. I think you can make a case that this is the way people always intended democracy to work — we just didn't have the right voting system to make it possible.

One thing is certain: neither the math nor the experts can tell us what kind of democracy we should want. That's something we'll all have to decide for ourselves.

Sources

Doron, Gideon, and Kronick, Richard (1977). "Single transferable vote: An example of a perverse social choice function." American Journal of Political Science 21, 303-311.

Yee, Ka-Ping (2005). "Voting Simulation Visualizations." zesty.ca/voting/sim