

REFLECTIONS ON ARROW'S THEOREM AND VOTING RULES

Nicholas R. Miller
Department of Political Science
University of Maryland Baltimore County (UMBC)
Baltimore MD 21250 USA
nmiller@umbc.edu

September 2017
Revised January 2018

Abstract

These reflections, written in honor of Kenneth Arrow, sketch out how one political scientist thinks about Arrow's theorem and its implications for voting rules. The basic claim is that Arrow's theorem means that all real-world voting rules are problematic in two quite specific ways — namely, they can be neither 'strategyproof' nor 'spoilerproof'. However, Condorcet's pairwise version of majority rule, while not a fully specified voting rule because of the cyclical majorities problem, is itself both strategyproof and spoilerproof. Moreover, the cycling problem seems to occur only rarely in practice.

For the special issue of *Public Choice* in honor of Kenneth Arrow
edited by Elizabeth Maggie Penn

REFLECTIONS ON ARROW'S THEOREM AND VOTING RULES

Though Kenneth Arrow's 'impossibility theorem' has occupied my mind off and on for some 50 years, I was never introduced to it as an undergraduate government major at Harvard from 1960 to 1963. (But then I was never introduced to 'political behavior' either.) However, I subsequently discovered that Arrow (along with Duncan Black and Anthony Downs) appeared under the topic of 'Voting as a Means of Optimizing' on the syllabus of a course on Political Economizing taught by Edward Banfield that I collected during the first week of the Fall 1961 semester.¹ Once I began my graduate studies in political science at Berkeley in Fall 1963, I soon became aware of Arrow's theorem (and certainly of political behavior). But, as I recall, I did not actually acquire my copy of Arrow's book (now very well worn) until the latter part of my graduate career when, having passed preliminary exams in three traditional political science fields, I began to undertake my largely self-taught education in the newly developing field of public choice/social choice/positive political theory.² Since then my teaching and scholarly work has been regularly influenced by Arrow. Regrettably, I never actually met Arrow, though I heard him speak at the 2001 Public Choice Society meeting in San Antonio on the occasion of the fiftieth anniversary of the publication of the first edition of *Social Choice and Individual Values* (1951).

This essay sketches out how I, as a more or less conventional political scientist with an interest in social choice theory, think that Arrow's theorem, and its implications for voting rules, may best be conveyed to undergraduates and beginning graduate students, as well as to political science colleagues unfamiliar with the literature on social choice and voting. The exposition may be of some interest to the specialist readers of this journal as well. In any case, I focus entirely on the implications of Arrow's theorem for voting rules and do not address its implications for welfare economics. Therefore, I think of Arrow's 'alternatives' as candidates or possible legislative outcomes (not 'complete descriptions of social states') and his 'individuals' as voters (in an electorate or a legislative body); his 'preference orderings' become 'ballot rankings', and his 'social welfare functions' underlie voting rules for choosing candidates or policies. Of course, this is a fairly standard way of thinking about Arrow's theorem (cf. Samuelson 1967).

Various common claims concerning the implications of Arrow's theorem for voting and democracy seem to me considerably overdrawn, imprecise, or otherwise unsatisfactory. The most sweeping claim is that *Arrow's theorem means that democracy is impossible*. Democracy entails much more than voting rules, and clearly many political systems commonly characterized as democratic actually exist and often seem to function satisfactorily. While the 'late' Riker (1982) — in contrast to the 'early' Riker (1953) — would say that these are (merely) 'liberal democracies', not 'populist democracies', and that Arrow's theorem implies that populist democracy, in which the

¹ Harvard students did not actually register for classes until the end of second week of the semester, so it was customary to shop around for courses during the first week or so.

² I was not entirely self-taught; in 1969 I took a seminar on Formal Models in Politics taught by the newly arrived and very junior assistant professors Michael Leiserson and Robert Axelrod — probably one of the first such political science courses taught anywhere outside of Rochester.

'popular will' (or at least 'majority will') is faithfully translated into public policy, is impossible, I would attribute failure to realize 'populist' democracy much more to the kinds of evidence set out recently by Achen and Bartels (2016) than to Arrow's theorem. A less sweeping claim is that *Arrow's theorem means that there is no perfect voting rule*. However, this claim is hardly surprising, while Arrow's theorem is generally said to be surprising. Moreover, a voting rule that satisfied democratic conditions no stronger than 'non-dictatorship' and 'non-imposition' along with the other Arrow conditions would hardly qualify as 'perfect'. A still less sweeping claim (implied by the subtitle *Why Elections Aren't Fair* of Poundstone, 2008) is that *Arrow's theorem means that every voting rule is unfair*, but this begs the question of the meaning of fairness. My preferred one-sentence claim is that *Arrow's theorem means that all real-world voting rules are problematic in two quite specific ways — namely, they are neither 'strategyproof' nor 'spoilerproof'*. But at the same time, many voting rules are problematic in other ways that cannot be blamed on Arrow's theorem.³

We consider a finite set of *voters* and a finite (and typically small) set of *alternatives* (e.g., candidates or proposals). Each voter has a *preference ordering* over the alternatives, and each voter submits a *ballot ranking* of the alternatives, which may or may not correspond to the voter's preference ordering. A collection of ballot rankings, one for each voter, is a *ballot profile*. A *preference aggregation rule* maps every possible ballot profile into a *social preference relation* between every pair of alternatives. A *social choice rule* selects 'best' or 'winning' alternatives for every ballot profile and subset of alternatives. A *voting rule* selects a unique 'best' or 'winning' alternative for every ballot profile and subset of alternatives.⁴

I will use the following now standard notation to indicate social preference between pairs of alternatives: $x \succ y$ means that 'x is better than y' (or 'x is strictly preferred to y'), $x \succeq y$ means that 'x is at least as good as y' (or 'x is preferred or indifferent to y'), and $x \sim y$ means that 'x and y are equally good' (or 'x and y are indifferent'). If $x \succeq y$ and $y \succeq z$ imply $x \succeq z$, the social preference relation is *transitive* and produces a social ordering of alternatives (similar to an individual preference ordering or ballot ranking) and assures that 'best' or 'most socially preferred' alternatives exist.

We shall give particular consideration to the two longtime rival preference aggregation rules proposed by Borda (1784) and Condorcet (1785). Let $n(xy)$ be the number of voters who rank x above y on their ballots. Under (*relative*) *majority rule* (as advocated by Condorcet), $x \succ y$ if $n(xy) > n(yx)$, $y \succ x$ if $n(yx) > n(xy)$, and $x \sim y$ if $n(xy) = n(yx)$. Under the *Borda rule*, an alternative x earns one *Borda point* for each alternative that it is ranked above on each ballot and a *Borda score* $B(x)$ equal to its points summed over all ballots, and $x \succ y$ if $B(x) > B(y)$, $y \succ x$ if $B(y) > B(x)$, and $x \sim y$ if $B(x) = B(y)$.

³ For a catalog of problematic (or 'paradoxical') features of many voting rules, see Felsenthal (2012).

⁴ Thus I define a voting rule as incorporating tie-breaking and similar mechanisms, though I will avoid the problem of ties in these reflections. I also acknowledge that this definition excludes approval and range voting as voting rules, since they do not operate on ballot profiles as defined in the Arrovian setup.

To give a lay audience a sense of Arrow's theorem, I find it useful first to present May's (1952) theorem, even though May follows the earliest version of Arrow's work by several years. May's theorem deals with the special case of social choice between just two alternatives — for example, a two-candidate election or a proposal that may be either adopted or rejected. May identified three conditions that an aggregation rule may meet and which are intuitively understandable and appealing in many (though not all) circumstances.

First, a rule is *anonymous* if it does not discriminate among voters — that is, if it depends only on the number, not the identity, of voters submitting each ballot ranking. Second, a rule is *neutral* if it does not discriminate among alternatives — that is, if two alternative switch places in the ballot rankings of every voter, social preference between them is likewise switched. Third, a rule is *positively responsive* if, whenever $x \succeq y$ and a voter then raises x relative to y in his ballot ranking (while all other rankings remain the same), the result is that $x \succ y$; thus social indifference is a knife-edge condition that is broken positively when any voter changes the way he ranks the two indifferent alternatives.

May's theorem says this: *given two alternatives, a preference aggregation rule is anonymous, neutral, and positively responsive if and only if it is (relative) majority rule.*⁵ Variants of majority rule include *absolute majority rule*, under which $x \succ y$ only if $n(xy) > n/2$, and various *supramajority rules* (up to and including *unanimity rule*), under which $x \succ y$ only if $n(xy) \geq q \times n$ for some q such that $n/2 < q \leq 1$ (e.g., $q = 2/3$). Such rules are neutral but not positively responsive if $x \sim y$ in the event that $n(xy) < q \times n$ and $n(yx) < q \times n$; they are positively responsive but not neutral in the event that (say) $y \succ x$ unless $n(xy) > q \times n$, e.g., where y is some favored (constitutional, legislative, or other) status quo. *Weighted relative majority rule* (as used in stockholders' meetings, international organizations, etc.) is neutral and positively responsive but (by design) not anonymous.

Arrow (1951) considered the multi-alternative case and therefore introduced two additional conditions that become relevant only in this case. One is *independence of irrelevant alternatives* (IIA), which requires that social preferences over any subset of alternatives depend only on the ballot rankings over the same subset. Since 'any subset' includes any pair, IIA in effect requires that social choice be based on pairwise individual rankings, and it thus rules out the Borda rule and many (arguably all) other real-world voting rules. For example, if voters 1 and 2 both rank x over y over z while voter 3 ranks y over x over z , x (with a Borda score of 5) is socially preferred to y (with a Borda score of 4) but, if 3 reverses his ranking of x and z , x and y (both with Borda scores of 4) are now socially indifferent, even though no voter's ranking of x and y has changed.

IIA is a subtle condition that does not have the immediate intuitive appeal of May's conditions. Moreover, it has been the source of considerable confusion and controversies of interpretation (e.g., Ray 1973, Bordes and Tideman 1991, McLean 1995, 2003). Even though Arrow (1951: 23) himself identified 'the rank-order method of voting frequently used in clubs' (i.e., the

⁵ Of course, many distinct voting rules — for example, plurality rule, Borda rule, Instant Runoff Voting (IRV), etc. — are equivalent to majority rule in the special case of two alternatives.

Borda rule) as a method that does not satisfy IIA, he later (1963:110) said that 'every known electoral system satisfies this condition [IIA]'. In contrast, Penn (2015) observed in her recent survey of Arrow's theorem and its descendants that 'virtually every real-world voting system violates it [IAA]'.

It is sometimes said that IIA rules out consideration of preference 'intensity'. In some sense this may be true, but more fundamentally the Arrowian framework itself rules out most intensity considerations, because a preference aggregation rule is defined as mapping profiles of ballot *rankings* (or preference *orderings*) into social preferences. Certain comparisons of preference intensity may be possible in an ordinal framework. For example, it seems that we can say that a voter who ranks x over y over z expresses a stronger preference for x over z than for either x over y or y over z (which the Borda rule takes account of but which majority rule ignores). But it does not follow that such a voter expresses a stronger preference for x over z than a voter who ranks z over x over y expresses for z over x (which the Borda rule in effect presumes but majority rule does not).

Given its pairwise nature, majority rule (along with its variants) clearly satisfies IIA. So we arrive at this multi-alternative variant of May's theorem: *a preference aggregation rule satisfies anonymity, neutrality, positive responsiveness, and IIA if and only if it is majority rule (applied to each pair of alternatives)*.⁶

The other condition that Arrow introduces and is relevant only when there are more than two alternatives is that the social preference relation generated by a preference aggregation rule is *transitive* (like individual preference orderings or ballot rankings) and therefore produces a social ordering of the alternatives — or, in Arrow's (perhaps unfortunate) language, that the aggregation rule is a *social welfare function*.⁷ Clearly the Borda rule meets this condition but, as is (now) well known, majority rule has a problem in this respect. Some 165 years before Arrow wrote, Condorcet (1785) discovered that majority rule can produce 'contradictory' social preferences, as exemplified by the following preference profile illustrating what Arrow (1951: 2) referred to as the 'well-known "paradox of voting"': voter 1 ranks x over y over z ; voter 2 ranks z over x over y ; and voter 3 ranks y over z over x . Majority rule maps this profile into an intransitive social preference relation such that $x \succ y$ and $y \succ z$ but $z \succ x$. Such 'cyclical' social preferences were independently rediscovered (and so named) by Dodgson (1876). A few years later, Nanson (1882) indicated that he was familiar with Condorcet's work (though not Dodgson's) and with cyclical majorities. At essentially the same time as Arrow was formulating the problem of social choice, Black (1948) independently rediscovered the

⁶ The Borda rule satisfies all of these conditions except IIA. Another definition of neutrality (e.g., Sen 1970, p. 72) is stronger in the multi-alternative case and itself implies IIA. However, the Borda rule is neutral in the commonsensical way defined above, even though it violates IIA.

⁷ Since I state Arrow's theorem in terms of preference aggregation rules, I (like Penn 2015) treat transitivity of social preference as an independent condition. On the other hand, I take the domain of a preference aggregation rule to be all logically possible ballot profiles, so I do not treat Arrow's 'universal domain' condition as an independent condition.

problem of cyclical majorities (and, as reported in Black 1958, subsequently rediscovered Condorcet's and Dodgson's work). But there is little or nothing else to indicate that the paradox of voting was actually 'well-known' at the time of publication of Arrow's book. Shortly thereafter, Dahl and Lindblom (1953; 422), citing Arrow, characterized the paradox of voting as 'a minor difficulty in voting that people with a mathematical turn of mind enjoy toying with'. The practical implications of Arrow's theorem (and of social choice theory generally) for voting rules — as well as one's assessment of the dispute between Riker (1982) and Mackie (2003) concerning the possibility of 'populist democracy' — turn in large part on whether Dahl and Lindblom's causal dismissal of the paradox is justified.

The transitivity condition rules out majority rule (and other pairwise rules), while IIA rules out non-pairwise rules. Taken together, these considerations give us what we may call 'May's impossibility theorem': *given two or more voters and three or more alternatives, no preference aggregation rule simultaneously satisfies the conditions of anonymity, neutrality, positive responsiveness, independence of irrelevant alternatives, and transitivity.*

Finally, we turn to Arrow's theorem itself. What Arrow (1951) demonstrated is that, even if May's conditions are radically weakened and even if we look beyond majority rule to all possible aggregation rules, the impossibility remains. First, Arrow weakened the condition of anonymity to *non-dictatorship*, which rules out only the most radical kind of inequality among voters — namely, the existence of a voter (a 'dictator') whose ballot ranking wholly determines social preference (regardless of the rankings of other voters). Second, Arrow weakened the condition of neutrality to *non-imposition*, which rules out only the most radical kind of inequality among alternatives — namely, the existence of a pair of alternatives x and y such that x is never socially preferred to y (regardless of the rankings of voters). Third, Arrow weakened positive responsiveness to *non-negative responsiveness*, which requires that, whenever $x \succ y$ and a voter raises x relative to y in his ballot ranking (while all other rankings remain the same), it remains true that $x \succ y$; and that, whenever $x \sim y$ and a voter raises x relative to y in his ranking (while all other rankings remain the same), the result is that $x \succeq y$.

Arrow's theorem, as presented in Arrow (1951), then states the following: *given two or more voters and three or more alternatives, no preference aggregation rule simultaneously satisfies the conditions of non-dictatorship, non-imposition, non-negative responsiveness, independence of irrelevant alternatives, and transitivity.* In the second edition of his book, Arrow (1963) replaced the combination of non-imposition and non-negative responsiveness with the implied but weaker condition of *unanimity* (or *weak Pareto*), which requires that $x \succ y$ whenever individual ballots unanimously rank x over y .⁸

There are now many proofs of Arrow's theorem, including recent ones based on Barbera's (1980) notion of 'pivotal voters' (e.g., Geanakoplos 2005, Fey 2014). But I like best Sen's (2014, pp. 35-37) proof, which as he says 'follows Arrow's own line of reasoning, but through some

⁸ Arrow's conditions can be further weakened in various ways; for an accessible survey, see Penn (2015).

emendations that make it agreeably short and rather easy to follow'. Like most proofs, it assumes that all of Arrow's conditions hold other than non-dictatorship and shows that this assumption implies that the aggregation rule is dictatorial. It proceeds in four steps (and I follow Sen's exposition closely).

- (1) *Definition of decisiveness.* Given a preference aggregation rule, a set D of voters is *decisive for a pair of alternatives* $\{x,y\}$ if, whenever all voters in D rank x over y (or y over x), it follows that $x \succ y$ (or $y \succ x$), regardless of the ballots of voters not in D . Note that this is merely a definition and does not itself claim that decisive sets exist under any particular aggregation rule.
- (2) *Decisiveness is 'contagious'.* If a set D of voters is decisive for $\{x,y\}$, it follows from Arrow's conditions that D is *decisive for every pair of alternatives*. Consider the following partially specified ballot profile with respect to four alternatives $x, y, a,$ and b : all voters in D rank a over x over y over b , while all voters not in D rank a over x and y over b . By Arrow's unanimity condition (or the combination of his non-imposition and non-negative responsiveness conditions), $a \succ x$ and $y \succ b$; by the decisiveness of D for $\{x,y\}$, $x \succ y$; and by Arrow's transitivity condition, $a \succ y$, $x \succ b$, and in particular $a \succ b$. Given that it must be that $a \succ b$ while we have specified ballot rankings of a and b only for voters in D , by Arrow's IIA condition D is decisive also for $\{a,b\}$. Since the same argument can be made for any other pair of alternatives other than $\{x,y\}$, D is decisive for every pair of alternatives.
- (3) *A decisive set contains a decisive subset.* Given a decisive set D of two or more voters, it follows from Arrow's conditions that D contains a proper subset that is also decisive. Partition D into two non-empty subsets D' and D'' and consider the following partially specified ballot profile with respect to three alternatives $x, y,$ and z : all voters in D' rank x over y and x over z , and all voters in D'' rank x over y and z over y . (Voters not in D may have any rankings.) By the decisiveness of D , $x \succ y$. We separately consider these two possibilities: (i) $z \succeq x$ or (ii) $x \succ z$. Given (i), by Arrow's transitivity condition, $z \succ y$. Given that $z \succ y$ and we have specified ranking of z and y only for voters in D'' , by Arrow's IIA condition D'' is decisive for $\{z,y\}$ and therefore by step (2) is decisive for every pair of alternatives. Given (ii) and the fact that we have specified rankings over x and z only for voters in D' , by Arrow's IIA condition D' is decisive for $\{x,z\}$ and therefore is decisive for every pair. Thus one or other of the two proper subsets of D is itself decisive.
- (4) *There is a one-voter decisive set.* By Arrow's unanimity condition (or the combination of his non-imposition and non-negative responsiveness conditions), the set of all voters is decisive. Given that the number of voters is finite, repeated application of step (3) leads to a one-voter decisive set, i.e., a dictator.

The trouble with this and other proofs of Arrow's theorem is that they do not convey an intuitive sense as to how Arrow's conditions combine to rule out all preference aggregation rules, though they do suggest that transitivity and IIA do much of the heavy lifting. Less formal argument provides more insight.

Arrow's (and May's) impossibility theorems demonstrate the incompatibility of five conditions on preference aggregation rules; three are intuitive and appealing (minimal or strict) 'democratic' conditions (i.e., non-dictatorship or anonymity, non-imposition or neutrality, and non-negative or positive responsiveness), while two are less intuitive and appealing 'technical' conditions (i.e., IIA and transitivity). But for all practical purposes, the 'Arrow problem' lies in the near incompatibility of the two 'technical' conditions, for they are in effect at war with each other.⁹ While IIA demands that social preference be based on pairwise individual rankings, rankings aggregated in a pairwise manner violate transitivity at some ballot profiles. The 'democratic' conditions by themselves have little to do with Arrowian impossibility; they simply rule out resolving the conflict between transitivity and IIA in quite unsatisfactory ways — by means of a dictatorial rule or a rule that imposes a fixed (as by 'sacred tradition') social ordering. While dictatorial and imposed rules satisfy both 'technical' conditions, it is a stretch to call either a 'preference aggregation rule'.

If even minimally democratic preference aggregation rules violate one or other of the two 'technical' conditions, the next question is what consequences these violations have for real-world voting rules. In fact, these consequences can be specifically identified and they are quite bothersome. In undergraduate classes, I have regularly made this point by referring to a letter sent to the *Washington Post* in 2004 by an annoyed local organizer for the Green Party, which asked the following question.

[Electoral engineering] isn't rocket science. Why is it that we can put a man on the moon but can't come up with a way to elect our president that allows voters to vote for their favorite candidate, allows multiple candidates to run and present their issues and . . . [makes] the 'spoiler' problem . . . go away?

The letter writer was asking for a voting rule that (1) allows for more than two candidates (or other options), (2) never gives voters an incentive to submit ballot rankings that do not reflect their 'honest' preferences, and (3) avoids the 'spoiler' problem in its several variants.¹⁰ The answer to the letter writer's question is that electoral engineering meeting these stipulations is in fact more difficult than rocket science, since Arrow's theorem implies that no minimally democratic voting rule can meet them. Moreover, it is precisely the combination of the two 'technical' conditions that is the source of the problem.

⁹ This point is forcefully illustrated in a working paper by Dougherty and Heckelman (2017). They generate large samples of simulated three-alternative ballot profiles with varying numbers of voters derived from an 'impartial culture', an 'impartial anonymous culture', and ANES thermometer scores for candidates in several US presidential elections with significant third candidates. They then apply seven common preference aggregation rules (plurality, anti-plurality, Hare/IRV, Nanson, Borda, Copeland, and pairwise majority rule) to these profiles, and determine how frequently each rule violates one or more of Arrow's (1963) conditions. With very few exceptions, at most one condition is violated at each profile, and it is either IIA or transitivity. More specifically, majority rule occasionally violates transitivity, while the other rules frequently violate IIA.

¹⁰ The letter writer in fact advocated the use of IRV, which indeed comes closer to meeting requirements (2) and (3) than ordinary plurality rule does but does not fully meet them and, moreover, fails to meet one other requirement (non-negative responsiveness) that plurality rule does meet.

Requirement (2) is commonly called *strategyproofness*, and the Gibbard-Satterthwaite (Gibbard 1973; Satterthwaite 1975) theorem famously states the following: *given three or more alternatives, every strategyproof voting rule is dictatorial*. Gibbard's version of the theorem explicitly invokes Arrow's theorem in its proof. More recently, Reny (2001) presented side-by-side proofs of the Arrow and Gibbard-Satterthwaite theorems to highlight their close connection. Elsewhere Muller and Satterthwaite (1977) demonstrated the equivalence strategyproofness and a condition they call *strong positive association*. Given two ballot profiles P and P' such that every voter who ranks x over any other alternative y in P also ranks x over y in P' (though rankings of other alternatives may differ), this condition requires that, if x is the social choice given P , x is also the social choice given P' . Strong positive association in effect combines the two Arrow conditions of non-negative responsiveness and IIA. A voting rule based on a preference aggregation function that violates non-negative responsiveness encourages 'dishonest' voting in an obvious way. That a voting rule based on a preference aggregation function that violates of IIA may induce voters to misrepresent their preferences is evident from the three-voter Borda rule example given earlier to demonstrate violation of IIA — if voter 3 truly prefers y to x to z , he can benefit by ranking y over z over x , so that x is no longer a clear winner.

On the other hand, majority rule complies with both IIA and positive responsiveness and, as a preference aggregation rule, is clearly strategyproof. Suppose $x \succ y$. No voters with the opposite preference can misrepresent their preferences in anyway to change this social preference. Changing their ballot rankings between x and y cannot help because majority rule is positively responsive; changing any other rankings cannot help either because majority rule satisfies IIA.

Of course, the rub here is that majority rule is merely a preference aggregation rule; since it may produce cyclical social preferences at some profiles, it is not a fully specified voting rule. A voting rule based on majority rule must specify a winning alternative in the event there is no (Condorcet) winner but rather a 'top cycle'. The rule for breaking the cycle — whether Black's (1958, p. 66) suggested use of Borda's method as a fallback, Dodgson's or Kemeny's methods, incomplete pairwise voting of the parliamentary type, or anything else — must in one way or other itself violate IIA and thus opens the rule up to strategic misrepresentation of preferences.¹¹

Requirement (3) may be called *spoilerproofness*, and it pertains to the consistency of social choice as the set of alternatives ('the agenda') actually voted on expands or contracts while ballot rankings remain unchanged. When the agenda expands, a consistent voting rule either selects the same winning alternative or selects one of the newly available alternatives (but not a different previously available alternative); likewise, when the agenda contracts, a consistent rule selects the same winning alternative provided that it remains available. An inconsistent voting rule is distinctly bothersome and encourages various kinds of mischief, as the outcome of voting depends unreasonably on the nature of the agenda and on the agents or forces that shape the agenda. It is well

¹¹ Indeed, even if honest ballot rankings imply a clear majority rule winner, it may be open to one or more voters to 'contrive a (top) cycle' by misrepresenting their preferences over other alternatives, which may then be resolved in a way favorable to their honest preferences.

known that many voting rules, particularly including the familiar (in the US, UK, and Canada) plurality rule, frequently violate this condition.¹²

The connection between Arrow's theorem and spoilerproofness is less clearly established than the connection between the theorem and strategyproofness — indeed, it has been a matter of confusion going back to Arrow's original discussion of IIA (1963: 26):

Suppose that an election is held, with a certain number of candidates in the field, each individual filing his list of preferences [i.e., ballot rankings], and then one of the candidates dies. Surely the social choice should be made by taking each of the individual's preference lists, blotting out completely the dead candidate's name, and considering only the orderings of the remaining names in going through the procedure of determining the winner.

The Borda rule, for example, does not follow this 'sure' expectation. However, Mackie (2003: 130), among others, has suggested that it might be more reasonable to 'blot out' the dead candidate's name from the social preference ordering generated when the Borda rule is applied to the original set of candidates. While this 'global' Borda rule still violates IIA, it is consistent in the face of agenda contraction. One can argue whether it is more reasonable to 'blot out' a dead candidate's name before or after the application of the aggregation rule (and whether this has anything to do with IIA), but what we really want is a rule under which it does not matter at which stage the dead candidate's name is 'blotted out'. Furthermore, Mackie's suggestion that we might use the 'global' Borda rule to avoid consistency problems does not help when the agenda expands rather than contracts. More generally, any practical notion of a 'global' Borda rule is ill-defined because the 'global' set of alternatives is difficult or impossible to specify in the context of real-world voting.

In any event, majority rule (and its variants) has this desired property. Blotting out a dead candidate's name from every ballot ranking and then applying majority rule produces the same social preference relation as simply blotting out the candidate's name from the original social preference relation. Moreover, inserting a new candidate's name into ballot rankings and then applying majority rule does not affect the original social preference relationships. Again, the rub is that majority rule is not a fully specified voting rule; once some mechanism for breaking cycles is specified, inconsistencies may reappear.

The connection between consistency and IIA is contested. Ray (1973) claims they are logically independent. Dasgupta and Maskin (2008: footnotes 6 and 7) and Sen (2014: endnote 6) both suggest they are virtually interchangeable. Bordes and Tideman (1992: 183) claim that Arrow's theorem means that 'all reasonable, real-world voting rules violate C [consistency]' but that this follows from Arrow's requirement that social choice derive from a transitive social welfare function, not the IIA condition itself. I leave it to expert social choice theorists to sort this out but remain

¹² The susceptibility of US elections to 'spoilers' (such as Nader in 2000) is the dominant and recurring complaint in William Poundstone's book on *Gaming the Vote: Why Elections Aren't Fair* (2008). In more formal social choice theory, 'strategic candidacy' (e.g., Dutta et al. 2001) and 'independence of clones' (Tideman 1987) pertain to related problems.

confident that Arrow's theorem implies that lack of spoilerproofness joins lack of strategyproofness as the two unavoidable problematic features of all minimally democratic voting rules.¹³

As a preference aggregation rule, majority rule avoids both problems but it is not a fully specified voting rule. But while the failure of non-pairwise voting rules to comply with IIA seems to be pretty generic, the failure of majority rule to comply with transitivity seems to be considerably less generic. Indeed, Dasgupta and Maskin (2008) show that majority rule is 'robust' in that it satisfies all the conditions in May's (and therefore Arrow's) impossibility theorem over a more extensive domain of profiles than any other practicable aggregation rule. However, this result leaves open the question of how extensive that most extensive domain is.

Much analysis of the frequency of the cyclical majorities phenomenon has focused on the proportion of all logically possible profiles (or anonymous profiles) that generate majority rule cycles. Given three alternatives and many voters, about 8.7% of all profiles (and 6.3% of anonymous profiles) produce cycles (Gehrlein and Lepelley 2011: 21) but these frequencies increase as the number of alternatives increases. However, assuming that every logically possible profile is equally likely to arise (the so-called 'impartial culture' assumption) is in effect to assume the absence of any 'culture' at all — that is, the absence of any characteristic structuring of voter preferences. As such, this assumption is not only unrealistic but, because it entails a virtual majority tie between almost every pair of alternatives, it is also in effect the worst-case scenario with respect to the frequency of cycles (Tsetlin et al. 2003). For example, Feld and Grofman (1986 and 1988) show that even a slight tendency toward single-peakedness greatly reduces the frequency of cycles. Moreover, even if majority rule over some very large set of 'potential' alternatives is highly cyclic, it remains true that it is rather unlikely that majority rule over a small agenda of alternatives drawn out of this very large set (as actual candidates in an election or actual proposals before a legislature) produces a cycle.¹⁴

Apart from such theoretical studies, only a few (perhaps a couple of dozen; see Gehrlein and Lepelley 2011: 13) empirical examples of cyclical majorities have been identified, and Mackie (2003) has disputed many of these claims. In part, the failure to find more empirical examples results from the fact that appropriate data — that is, ballot profiles of three or more alternatives — is rarely available. Election results and surveys typically produce only information about the distribution of first preferences, and parliamentary voting procedures do not take enough pairwise votes to reveal cycles. In the few cases in which appropriate data are available — for example, individual ballots under a voting rule that requires voters to rank candidates — cyclical majorities are occasionally found but appear to be quite exceptional and rarely if ever appear as top cycles (e.g., Chamberlin et al. 1984, Feld and Grofman 1992; Tideman 2006: 99-106). Perhaps Dahl and Lindbloom's early claim that the cyclical majority phenomenon is only a minor difficulty for majority rule is justified. It

¹³ However, these two problematic features may tend to counteract one another; in any case, this is true with respect to plurality rule (Dowding and Van Hees 2008).

¹⁴ In particular, the notorious McKelvey (1976, 1979) 'global cycling theorem' pertaining to majority rule over a space of two or more dimensions does not imply that a cycle exists (or is even likely to exist) over a small set of arbitrarily selected points in the space.

is noteworthy that two eminent social choice theorists (Maskin and Sen 2017) recently advocated majority rule for US presidential primaries and elections and barely bother to take note of the cyclical majorities problem.

In summary, Arrow's theorem implies that, given three or more alternatives, no minimally democratic voting rules can be either strategyproof or spoilerproof. As a preference aggregation rule, majority rule avoids both problems, but it is not a fully specified voting rule because of the cyclical majorities problem. Yet this problem may in practice be very modest in scope.

To reinforce this conclusion, I end with an anecdote. About twenty years ago, a colleague in another social science department at UMBC contacted me with a query (which here I paraphrase and no doubt embellish somewhat): 'Nick, I understand you are an expert on election methods, so maybe you can give me some advice. I'm in charge of our department's balloting to rank and choose candidates for faculty positions. Of course, we use the usual rank-order [i.e., Borda] method, but I've noticed it has two problems. First, it often encourages faculty members to submit ranking of candidates that clearly do not reflect their actual preferences. Second, how we end up ranking candidates sometimes depends on exactly which candidates we include on the ballot. Is there some other voting method that we could use to avoid these problems?' I replied that in principle the answer to his question was 'No' (Arrow's theorem and all that). But I suggested that he announce to his colleagues that henceforth faculty candidates would be ranked on the basis of majority rule. I noted that it was logically possible that this method might fail to rank the candidates but I thought it was unlikely his department would ever confront this problem in practice, and I asked him to let me know if the problem ever arose. I never heard back (which is just as well, since I do not know what I would have recommended in that event).

Acknowledgements: For helpful comments, I thank Jac Heckelman, Dan Felsenthal, and Michel Le Breton.

References

- Achen, Christopher H., and Larry M. Bartels (2016). *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton: Princeton University Press.
- Arrow, Kenneth A. (1951). *Social Choice and Individual Values*. New York: Wiley.
- Arrow, Kenneth A. (1963). *Social Choice and Individual Values*, 2nd ed. New York: Wiley.
- Barberá Salvador (1980). Pivotal voters: A new proof of arrow's theorem. *Economics Letters*, 6/1, 13-16.
- Black, Duncan (1948). On the rationale of group decision-making. *Journal of Political Economy*, 56/1, 23-34.
- Black, Duncan (1958). *The Theory of Committees and Elections*. Cambridge: Cambridge University Press.
- Borda, Jean-Charles de (1784). *Mémoire sur les Élections au Scrutin*. Translation in Iain McLean and Arnold B. Urken (eds.), *Classics of Social Choice*. Ann Arbor: University of Michigan Press, 1995, pp. 83-89.
- Bordes, Georges, and Nicolaus Tideman (1991). Independence of irrelevant alternatives in the theory of voting. *Theory and Decision*. 30/2, 163–186.
- Chamberlin, John R., Jerry L. Cohen, and Clyde H. Coombs (1984). Social choice observed: Five presidential elections of the American Psychological Association. *Journal of Politics*, 46/2, 479-502.
- Condorcet, Marquis de (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Translation in Iain McLean and Arnold B. Urken (eds.), *Classics of Social Choice*. Ann Arbor: University of Michigan Press, 1995, pp. 91-112.
- Dahl, Robert A., and Charles E. Lindbloom (1953). *Politics, Economic, and Welfare*. New York: Harper & Row.
- Dasgupta, Partha, and Eric Maskin (2008). On the robustness of majority rule. *Journal of the European Economic Association*, 6/5, 949-973.
- Dodgson, Charles (1976). *A Method of Taking Votes on More Than Two Issues*. Reprinted in Iain McLean and Arnold B. Urken (eds.), *Classics of Social Choice*. Ann Arbor: University of Michigan Press, 1995, pp. 288-297./
- Dougherty, Keith L., and Jac C. Heckelman (2017). *The Probability of Violating Arrow's Conditions*. Working paper, Department of Political Science, University of Georgia.
- Dowding, Keith, and Martin Van Hees (2008). In praise of manipulation. *British Journal of Political Science*, 38/1, 1-15.

- Dutta, Bhaskar, Matthew O. Jackson, and Michel Le Breton (2001). Strategic candidacy and voting procedures. *Econometrica*, 69/4, 1013-1037.
- Feld, Scott L., and Bernard Grofman (1986). Partial single-peakedness: An extension and clarification. *Public Choice*, 51/1, 71-80.
- Feld, Scott L., and Bernard Grofman (1988). Ideological consistency as a collective phenomenon. *American Political Science Review*, 82/3, 773-788.
- Feld, Scott L., and Bernard Grofman (1992). Who's afraid of the big bad cycle? Evidence from 36 elections. *Journal of Theoretical Politics*, 4/2, 231-237.
- Felsenthal, Dan S. (2012). Review of paradoxes afflicting procedures for electing a single candidate. In Dan S. Felsenthal and Moshé Machover (eds.), *Electoral Systems: Paradoxes, Assumptions, and Procedures*. Heidelberg: Springer.
- Fey, Mark (2014). A straightforward proof of Arrow's theorem. *Economics Bulletin*, 34/3, 1792-1797.
- Geanakoplos, John (2005). Three brief proofs of Arrow's impossibility theorem. *Economic Theory*, 26/1, 211-215.
- Gehrlein, William V. (2006). *Condorcet's Paradox*. Berlin: Springer.
- Gehrlein, William V., and Dominique Lepelley (2011). *Voting Paradoxes and Group Coherence*. Berlin: Springer.
- Gibbard, Allan (1973). Manipulation of voting schemes: A general result. *Econometrica*, 41/4, 587-601.
- Mackie, Gerry (2003). *Democracy Defended*. Cambridge: Cambridge University Press.
- Maskin, Eric, and Amartya Sen (2014). *The Arrow Impossibility Theorem*. New York: Columbia University Press.
- Maskin, Eric, and Amartya Sen (2017). A new electoral system? *The New York Review of Books*, January 19, 2017, pp. 8-10.
- May, Kenneth O. (1952). A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20/4, 680-684.
- McKelvey, Richard D. (1976). Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory*, 12/3, 472-482.
- McKelvey, Richard D. (1979). General conditions for global intransitivities in formal voting models. *Econometrica*, 47/5, 1085-1112.
- McLean, Iain (1995). Independence of irrelevant alternatives before Arrow. *Mathematical Social Sciences*. 30/2, 107-126.
- McLean, Iain (2003). *The Reasonableness of Independence: A Conversation from Condorcet and*

- Borda to the Present Day*. Nuffield College Politics Working Paper 2003-W6, University of Oxford.
- Muller, Eitan, and Mark A. Satterthwaite (1977). The equivalence of strong positive association and strategy-proofness. *Journal of Economic Theory*, 14/2, 412–418.
- Nanson, E. J. (1882). *Methods of Election*. Reprinted in Iain McLean and Arnold B. Urken (eds.), *Classics of Social Choice*. Ann Arbor: University of Michigan Press, 1995, pp. 321-359.
- Penn, Elizabeth Maggie (2015). Arrow's theorem and its descendants. In Jac C. Heckelman and Nicholas R. Miller (eds.), *Handbook of Social Choice and Voting*. Cheltenham UK and Northampton MA: Edward Elgar, pp. 237-262.
- Poundstone, William (2008). *Gaming the Vote: Why Elections Aren't Fair*. New York: Hill and Wang.
- Ray, Paramesh (1973). Independence of irrelevant alternatives. *Econometrica*, 41/5, 987-991.
- Reny, Philip J. (2010). Arrow's theorem and the Gibbard-Satterthwaite theorem: A unified approach. *Economics Letters*, 70/1, 99–105.
- Riker, William H. (1953). *Democracy in the United States*. New York: Macmillan
- Riker, William H. (1982). *Liberalism Against Populism: A Confrontation Between the Theory of Democracy and the Theory of Social Choice*. San Francisco: W. H. Freeman and Company.
- Samuelson, Paul (1967). Arrow's mathematical politics. In Sidney Hook (ed.), *Human Values and Economic Policy*. New York: New York University Press.
- Satterthwaite, Mark Allen (1975). Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10/2, 187-217.
- Sen, Amartya K. (1970). *Collective Choice and Social Welfare*. San Francisco: Holden-Day, Inc.
- Sen, Amartya (2014). Arrow and the impossibility theorem. In Eric Maskin and Amartya Sen, *The Arrow Impossibility Theorem*. New York: Columbia University Press, pp. 29-55.
- Tideman, T. Nicolaus (1987). Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4/3 (1987), 185-206.
- Tideman, Nicolaus (2006). *Collective Decisions and Voting: The Potential for Public Choice*. Aldershot, UK: Ashgate.
- Tsetlin, Ilia, Michel Regenwetter, and Bernard Grofman (2003). The impartial culture maximizes the probability of majority cycles. *Social Choice and Welfare*. 21/3, 387-398.