Voter Sovereignty and Election Outcomes

Steven J. Brams
(Corresponding Author)
Department of Politics
New York University
New York, NY 10003
USA
steven.brams@nyu.edu

M. Remzi Sanver
Department of Economics
Istanbul Bilgi University
80310, Kustepe, Istanbul
TURKEY
sanver@bilgi.edu.tr

February 2003

Abstract

Voters are *sovereign* to the degree that they can express their approval for any set of candidates. While voter sovereignty is maximized under approval voting (AV), AV can lead to

- a plethora of outcomes, depending on where voters draw the line between acceptable and unacceptable candidates; and
- Condorcet losers and other lesser candidates, even in equilibrium.

But we argue that voters' judgments about candidate acceptability *should* take precedence over standard social-choice criteria, such as electing a Condorcet or Borda winner. Among other things, we show that

- sincere outcomes under all voting systems considered are AV outcomes, but not vice versa;
- a Condorcet winner's election under AV is always a strong Nash-equilibrium outcome but not under other systems; and
- outcomes that other systems cannot prevent can be prevented under AV.

Keywords: approval voting; elections; Condorcet winner/loser; Nash equilibria.

Voter Sovereignty and Election Outcomes[1]

1. Introduction

Our thesis in this paper is that *several* outcomes of single-winner elections may be acceptable. Perhaps the most dramatic recent example illustrating this proposition is the 2000 US presidential election, in which George W. Bush won the electoral vote—disputed though it was in Florida—and Al Gore won the popular vote. In the eyes of many analysts, the outcome in Florida, which was decisive for Bush's electoral-vote victory and which he won with less than 0.01 percent of the popular vote in that state, was essentially a tie.

In such a statistical dead heat, it seems fair to say that each candidate is as entitled as the other to win and thereby become the president. Although it is the electoral vote, not the popular vote, that counts in a presidential election, the divided verdict on the two different tallies, coupled with a divided 5-4 US Supreme Court decision that halted further recounts in Florida, lends credence to the notion that Al Gore had a strong claim on the presidency. Indeed, because there were 537,000 more Gore voters than Bush voters nationwide (0.5 percent of the total), one might contend that Gore was more acceptable than Bush even though Bush, constitutionally, was the winner

To be sure, the extreme closeness of this election was unusual. But many elections, especially those with three or more candidates, may have more than one acceptable outcome.

For example, even when there is a Condorcet winner, who can defeat every other candidate in pairwise contests, there may be a different Borda-count winner, who on the

[1] We thank Eyal Baharad and Shmuel Nitzan for valuable suggestions. Steven J. Brams acknowledges the

average is ranked higher than a Condorcet winner. If there is no Condorcet winner because of cyclical majorities, the Condorcet cycle may be broken at its weakest link to select the strongest candidate in the cycle, who need not be the Borda winner.

That different voting systems can give different outcomes is, of course, an old story. The observation that different outcomes may satisfy different social-choice criteria is also old hat (Nurmi, 1999, 2002, give many examples). What is new here is our claim that in an election with three or more candidates, other outcomes—not just the Condorcet winner, the Borda-count winner, or the strongest candidate in a cycle—may be *more* acceptable to the electorate. In fact, even a Condorcet loser, who would lose in pairwise contests to every other candidate, may turn out to be the most acceptable candidate.

To justify this last statement, we need to define some measure of "acceptability." If voters rank candidates from best to worst, where they draw the line in their rankings between acceptable and unacceptable candidates offers one such measure. It is precisely this information that is elicited under approval voting (AV), whereby voters can approve of as many candidates as they like or consider acceptable. This gives them the opportunity to be *sovereign* by expressing their approval for any *set* of candidates, which no other voting system permits.[2]

Call a candidate a *Pareto candidate* if there is no other candidate that all voters rank higher. We demonstrate that candidates selected under AV always include at least one Pareto candidate. In fact, AV *dominates* so-called scoring systems, including

---

support of the C.V. Starr Center for Applied Economics at New York University.

[2] Voter sovereignty should be distinguished from Arrow's (1963) condition of "citizen sovereignty," whereby for any two alternatives *a* and *b*, if all voters prefer *a* to *b*, *a* cannot be prohibited as the social choice. If voters are "sincere," AV satisfies citizen sovereignty, because all voters who approve of *b* will also approve of *a*. Note that voter sovereignty describes the behavior of individual voters whereas citizen sovereignty is a property of a voting system.

plurality voting (PV) and the Borda count (BC), with respect to the election of Pareto candidates: A Pareto candidate elected by a scoring system is always elected by AV for some sincere and admissible strategies, but not vice versa. This is also true for ranking systems that do not rely on scoring, including the Hare system of single transferable vote (STV) and the majoritarian compromise (MC), as we will show.

But if AV does a better job of finding Pareto candidates, doesn't it open the door to a plethora of possibilities? Isn't this a vice rather than a virtue, as some have argued (e.g., Saari and Van Newenhizen, 1988a; Saari, 1994, 2001)?[3]

This argument might have merit if the plethora of possibilities were haphazard choices that could easily be upset when voters are manipulative. But we show that AV often leads to Nash-equilibrium outcomes, from which voters with the same preferences will have no incentive to depart. Moreover, if voters with different preferences are able to coordinate their choices and none has an incentive to depart, AV *guarantees* the election of a unique Condorcet winner (if one exists).

The latter notion of stability is that of a *strong* Nash equilibrium, which yields outcomes that are invulnerable to departures by *any* set of voters. None of the other voting systems we assay guarantees that a unique Condorcet winner, and only a Condorcet winner, will be a strong Nash equilibrium outcome when voters are sincere. While AV offers this guarantee, however, it also allows for other Nash-equilibrium

[3] The critique of AV by Saari and Van Newenhizen (1988a) provoked an exchange between Brams, Fishburn, and Merrill (1988a, 1988b) and Saari and Van Newenhizen (1988b) over whether the plethora of AV outcomes more reflected AV's "indeterminacy" (Saari and Van Newenhizen) or its "responsiveness" (Brams, Merrill, and Fishburn); other critiques of AV are referenced in Brams and Fishburn (2003). Here we argue that *which* outcome is chosen should depend on voters' judgments about the acceptability of candidates rather than standard social-choice criteria, which—as we will show— may clash with these judgments.

outcomes, including even a Condorcet loser, who may be the most acceptable candidate, even in equilibrium.

In section 2, we define preferences and strategies under AV and give an example that illustrates the choice of sincere, admissible strategies. In section 3 we characterize AV outcomes, describing the "critical strategy profile" that produces them, and compare these outcomes with those given by other voting systems. Among other things, we show that no "fixed rule," in which voters vote for a predetermined number of candidates, always elects a unique Condorcet winner, suggesting the need for a more flexible system.

The stability of outcomes under the different voting systems is analyzed in section 4, where we show that Nash equilibria and strong Nash equilibria may vary from system to system. Also, Condorcet voting systems, which guarantee the election of Condorcet winners when voters are sincere, may not elect Condorcet candidates in equilibrium.

In section 5 we show that rational departures by voters from unstable outcomes under other voting systems may not induce AV outcomes, but rational departures under AV always do. Hence, outcomes under AV form a closed set.

Nonstrong Nash equilibria might be thought of as possessing a kind of local stability, whereas strong Nash equilibria possess a global stability. These different kinds of equilibria may coexist, which is to say that *which* stable outcome is chosen will depend on which candidates voters consider acceptable and whether they coordinate their choices. In large-scale public elections, coordination is typically done when voters draw inferences from polls, not by face-to-face communication, which is commonplace in smaller settings like committees.

That a Condorcet candidate is a globally stable choice under AV should not be surprising. What is more surprising is that such a candidate can be upset if (i) coordination is difficult and (ii) many voters consider another candidate more acceptable.

Speaking normatively, we believe that voters *should* be sovereign, able to express their approval of any set of candidates. Likewise, a voting system *should* allow for the possibility of multiple acceptable outcomes, especially in close elections. That AV more than other voting systems is responsive in this way we regard as a virtue.

That it singles out as strong Nash–equilibrium outcomes unique Condorcet winners may or may not be desirable. We discuss these and other questions related to the nature of acceptable outcomes in section 6, where we suggest that "acceptability" replace the usual social-choice criteria for assessing the satisfactoriness of election outcomes chosen by sovereign voters.

2. Preferences and Strategies under AV

Consider a set of voters choosing among a set of candidates. We denote individual candidates by small letters *a, b, c, ….* A voter's *strict preference relation* over candidates will be denoted by *P*, so *aPb* means that a voter strictly prefers *a* to *b*, which we will denote by the following left-to-right ranking (separated by a space): *a b*. We assume in the subsequent analysis that all voters have strict preferences, so they are not indifferent among two or more candidates.[4]

We assume that every voter has a *connected* preference: For any *a* and *b*, either *aPb* or *bPa* holds. Moreover, *P* is *transitive*, so *aPc* whenever *aPb* and *bPc*. The list of preferences of *all* voters is called a *preference profile* **P**.

An *AV strategy S* is a subset of candidates. Choosing a strategy under AV means voting for all candidates in the subset and no candidates outside it. The list of strategies of all voters is called a *strategy profile* **S**.

The number of votes that candidate *i* receives at **S** is the number of voters who include *i* in the strategy *S* that they select. For any **S**, there will be a set of candidates ("winners") who receive the greatest number of votes.

An AV strategy *S* of a focal voter is *admissible* if it is not dominated by any other strategy—that is, if there is no other strategy that gives outcomes at least as good as, and sometimes better than, *S* for all strategy profiles **S** of voters other than the focal voter. Brams and Fishburn (1978, 1983) show that admissible strategies under AV involve always voting for a most-preferred candidate and never voting for a least-preferred candidate.

An AV strategy is *sincere* if, given the lowest-ranked candidate that a voter approves of, he or she also approves of all candidates ranked higher. Thus, if *S* is sincere, there are no "holes" in a voter's approval set: Everybody ranked above the lowest-ranked candidate that a voter approves of is also approved; and everybody ranked below is not approved.[5] A strategy profile **S** is said to be admissible and sincere if and only if the strategy *S* that every voter chooses is admissible and sincere, based on each voter's preference *P*.

---

[4] This restriction simplifies the analysis; its relaxation to allow for voter indifference among candidates has no significant effect on our findings.

[5] Admissible strategies may be insincere if there are four or more candidates. For example, if there are exactly four candidates, it may be rational for a voter to approve of his or her first and third choices without also approving of a second choice (see Brams and Fishburn, 1983, pp. 25-26, for an example). However, the circumstances under which this happens are sufficiently rare and nonintuitive that we henceforth suppose that voters choose only sincere, admissible strategies under AV.

As an illustration of these concepts, assume that there are 7 voters who can be grouped into three different types, each having the same preference *P* over the set of four candidates {*a*, *b*, *c*, *d*}:

Example 1

1. 3 voters: *a b c d*
2. 2 voters: *b c a d*
3. 2 voters: *d b c a*

Thus, each of the 3 type (1) voters has preference *aPbPcPd.* The three types define the preference profile **P** of all 7 voters. We assume that all voters of each type choose the same strategy *S*.

Voters of type (1) have three sincere, admissible strategies: {*a*}, {*a*, *b*}, and {*a*, *b*, *c*}, which for convenience we write as *a*, *ab*, and *abc*. A typical sincere, admissible strategy profile of the 7 voters is **S** = (*a*, *a*, *a*, *bc*, *bc*, *dbc*, *dbc*), whereby the 3 voters of type (1) approve of only their top candidate, the 2 voters of type (2) approve of their top two candidates, and the 2 voters of type (3) approve of all candidates except their lowest-ranked. The number of votes of each candidate at **S** is 4 votes for *b*, 4 votes for *c*, 3 votes for *a*, and 2 votes for *d*. Hence, AV selects candidates {*b*, *c*} as the (tied) winners at **S**.

3. Election Outcomes under AV and Other Voting Systems

Given a preference profile **P**, we consider the set of all candidates that *can* be chosen by AV when voters use sincere, admissible strategies. We call this set *AV outcomes*. Clearly, a candidate ranked last by all voters cannot be in this set, because it is inadmissible for any voter to vote for this candidate.

Define an *AV critical strategy profile* for candidate *i* at preference profile **P** as follows: Every voter who ranks *i* as his or her worst candidate votes only for the candidate that he or she ranks top. The remaining voters vote for *i and* all candidates they prefer to *i*.

Let $C_i(\mathbf{P}, \mathbf{S})$ be the AV critical strategy profile of candidate *i*. In Example 1, the critical strategy profile for candidate *a* is $C_a(\mathbf{P}, \mathbf{S}) = (a, a, a, bca, bca, d, d)$, giving *a* 5 votes compared to 2 votes each for *b*, *c*, and *d*. It can easily be seen that $C_i(\mathbf{P}, \mathbf{S})$ is admissible and sincere. We next show that under AV candidate *i* cannot do better vis-à-vis the other candidates than at $C_i(\mathbf{P}, \mathbf{S})$.

Proposition 1. *Assume all voters choose sincere, admissible strategies. The AV critical strategy profile for candidate i, $C_i(\mathbf{P}, \mathbf{S})$, maximizes the difference between the number of votes that i receives and the number of votes that every other candidate j receives.*

Proof. Clearly, no other sincere, admissible strategy profile yields candidate *i* more votes than its AV critical strategy profile $C_i(\mathbf{P}, \mathbf{S})$. Now consider the number of votes received by any other candidate *j* at $C_i(\mathbf{P}, \mathbf{S})$. Candidate *j* will receive no fewer and sometimes more votes if there are the following departures from $C_i(\mathbf{P}, \mathbf{S})$:

(i) a voter who ranked candidate *i* last, and therefore did not vote for him or her, votes for one or more candidates ranked below his or her top-ranked choice (possibly including candidate *j*); or

(ii) a voter who did not rank candidate *i* last or next-to-last votes for one or more candidates ranked below *i* (possibly including candidate *j*).

In either case, candidate *j* never gets fewer, and may get more, votes when there are these departures from candidate *i*'s critical strategy profile $C_i(\mathbf{P}, \mathbf{S})$. Because (i) and (ii) exhaust the possible departures from $C_i(\mathbf{P}, \mathbf{S})$ that involve voting for some other candidate *j*, candidate *i* cannot do better vis-à-vis candidate *j* than at $C_i(\mathbf{P}, \mathbf{S})$. Q.E.D.

The next proposition provides a simple way to determine whether any candidate *i* is an AV outcome:

Proposition 2. *Candidate i is an AV outcome if and only if i is chosen at his or her critical strategy profile* $C_i(\boldsymbol{P}, \boldsymbol{S})$.

Proof. The "if" part is a direct consequence of the fact that $C_i(\mathbf{P}, \mathbf{S})$ is sincere and admissible. To show the "only if" part, suppose candidate *i* is not chosen by AV at $C_i(\mathbf{P}, \mathbf{S})$. By Proposition 1, $C_i(\mathbf{P}, \mathbf{S})$ maximizes the difference between the number of votes that *i* receives and the number of votes that any other candidate *j* receives, so there is no other sincere, admissible strategy profile at which *i* can be chosen by AV. Q.E.D.

Using Proposition 2, we next give a characterization of candidates that *cannot* be AV outcomes.

Proposition 3. *Given any preference profile* ***P*** *and any candidate i, i cannot be an AV outcome if and only if there exists some other candidate j such that the number of voters who consider j as their best choice and i as their worst choice exceeds the number of voters who prefer i to j*.

Proof. Given any preference profile **P** and any two candidates *i* and *j*, voters can be partitioned into three (disjoint) classes:

(i) those who prefer *i* to *j*;

(ii) those who consider *j* as the best choice and *i* as the worst choice; and

(iii) those who prefer *j* to *i* but do not fall into class (ii).

At critical strategy profile $C_i(\mathbf{P}, \mathbf{S})$, the voters in class (i) will vote for *i* but not *j*; those in class (ii) will vote for *j* but not *i*; and those in class (iii) will vote for both *i* and *j*. Setting aside class (iii), which gives each candidate the same number of votes, candidate *i* cannot be selected at $C_i(\mathbf{P}, \mathbf{S})$ if and only if the number of voters in class (ii) exceeds the number of voters in class (i). Hence, by Proposition 2 candidate *i* cannot be an AV outcome. Q.E.D.

In effect, Proposition 3 extends Proposition 2 by saying precisely when candidate *i* will be defeated by candidate *j* and cannot, therefore, be an AV outcome. On the other hand, if there is no candidate *j* such that the number of class (ii) voters exceeds the number of class (i) voters, candidate *i* is an AV outcome at its critical strategy profile $C_i(\mathbf{P}, \mathbf{S})$.

AV can generate a plethora of outcomes. Consider again Example 1, in which we showed earlier that AV selects candidate *a* at $C_a(\mathbf{P}, \mathbf{S})$. Similarly, AV selects candidates *b* and {*b*, *c*}, all with 7 votes, at critical strategy profiles $C_b(\mathbf{P}, \mathbf{S}) = \{ab, ab, ab, b, b, db, db\}$ and $C_c(\mathbf{P}, \mathbf{S}) = \{abc, abc, abc, bc, bc, dbc, dbc\}$. However, $C_d(\mathbf{P}, \mathbf{S}) = \{a, a, a, b, b, d, d\}$, so candidate *a* (3 votes) rather than candidate *d* (2 votes) is chosen at candidate *d*'s critical strategy profile.[6] In sum, the set of AV outcomes that are possible in Example 1 is {*a*, *b*, {*b*, *c*}}.

[6] That *d* cannot be chosen also follows from Proposition 3: More voters (3) consider *a* as their best choice and *d* as their worst choice than prefer *d* to *a* (2).

As noted earlier, a candidate is a *Pareto candidate* if there is no other candidate that all voters rank higher. Example 1 illustrates three things about the tie-in of Pareto candidates and AV outcomes:

- *a* and *b* are Pareto candidates and AV outcomes;
- *c* is not a Pareto candidate but is a component of an AV outcome (it ties with *b* at $C_c(\mathbf{P}, \mathbf{S})$); and
- *d* is a Pareto candidate but not an AV outcome.

These observations are generalized by the following proposition:

Proposition 4. *The following are true about the relationship of Pareto candidates and AV outcomes:*

(i) At every preference profile ***P****, there exists a Pareto candidate that is an AV outcome or a component of an AV outcome;*

(ii) Not every Pareto candidate is necessarily an AV outcome; and

(iii) A non-Pareto candidate may be a component of an AV outcome but never a unique AV outcome.

Proof. To show (i), take any preference profile **P**. Assume that every voter votes only for his or her top choice. Then the one or more candidates chosen by AV, because they are top-ranked by some voters, must be Pareto candidates. To show (ii), it suffices to check the critical strategy profile $C_d(\mathbf{P}, \mathbf{S})$ of Example 1, wherein candidate *d* is not an AV outcome but is a Pareto candidate because *d* is top-ranked by the 2 type (3) voters.

In Example 1, we showed that *c* is not a Pareto condidate but is a component of an AV outcome. To show that a non-Pareto candidate can never be a unique AV outcome

and prove (iii), consider any **P** at which there exists a non-Pareto candidate *i* that is a component of an AV outcome. Take any sincere, admissible strategy profile **S** where this outcome is selected. Because *i* is not a Pareto candidate, there exists some other candidate *j* that every voter prefers to *i*. Hence, every voter who voted for *i* at **S** must have voted for *j* as well, which implies that *i* and *j* tie for the most votes. Indeed, *all* candidates *j* that Pareto dominate *i* will be components of an AV outcome at **S**. Because at least one of the candidates *j* that Pareto-dominate *i* must be ranked higher by one or more voters than all other candidates *j*, AV picks a Pareto candidate that ties candidate *i*. Q.E.D.

In Example 1, candidate *b* is the *Condorcet winner*, who can defeat all other candidates in pairwise contests, and candidate *d* is the *Condorcet loser*, who is defeated by all other candidates in pairwise contests. Not surprisingly, *b* is an AV outcome whereas *d* is not. However, consider the following 7-voter, 3-candidate example:

Example 2

1. 3 voters: *a b c*
2. 2 voters: *b c a*
3. 2 voters: *c b a*

Notice that the 2 type (2) and the 2 type (3) voters prefer candidates *b* and *c* to candidate *a*, so *a* is the Condorcet loser. But because the critical strategy profile of candidate *a* is $C_a(\mathbf{P}, \mathbf{S}) = (a, a, a, b, b, c, c)$, *a* is an AV outcome—as are also candidates *b* and *c*, rendering all three candidates in this example AV outcomes.

We summarize the Condorcet properties of AV outcomes with our next proposition:

Proposition 5. *Condorcet winners are always AV outcomes, whereas Condorcet losers may or may not be AV outcomes.*

Proof. If candidate *i* is a Condorcet winner, a majority of voters prefer *i* to every other candidate *j*. This implies that fewer voters rank *j* as their best choice and *i* as their worst choice, which by Proposition 3 implies that candidate *i* is an AV outcome. That a Condorcet loser may not be an AV outcome is shown by candidate *d* in Example 1, whereas candidate *a* in Example 2 shows that a Condorcet loser may be an AV outcome. Q.E.D.

Define a *fixed rule* as a voting system in which voters vote for a predetermined number of candidates.

Proposition 6. *There is no fixed rule that always elects a unique Condorcet winner.*

Proof. Consider the following 5-voter, 4-candidate example:

Example 3

1. 2 voters: *a d b c*
2. 2 voters: *b d a c*
3. 1 voter: *c a b d*

Vote-for-1 elects {*a*, *b*}, vote-for-2 elects *d*, and vote-for-3 elects {*a*, *b*}. Thus, none of the fixed rules elects the unique Condorcet winner, candidate *a*. Q.E.D.

By contrast, several sincere, admissible strategies, including $C_a(\mathbf{P}, \mathbf{S}) = (a, a, bda, bda, ca)$, in which different voter types vote for different numbers of candidates, elect *a*. Clearly, the flexibility of AV may be needed to elect a unique Condorcet winner.

We next turn to scoring rules and analyze the relationship between the winner they select and AV outcomes. The best-known scoring rule is the Borda count (BC): Given that there are *n* candidates, BC awards $n - 1$ points to each voter's first choice, $n - 2$ points to each voter's second choice, …, and 0 points to his or her worst choice.

In Example 1, the BC winner is candidate *b*, who receives from the three types of voters a Borda score of 3(2) + 2(3) + 2(2) = 16 points. In Example 2, the BC winner is also candidate *b*, who receives from the three types of voters a Borda score of 3(1) + 2(2) + 2(1) = 9 points. In these examples, the BC winners coincide with the Condorcet winners, making them AV outcomes (Proposition 5), but this need not be the case, as we will illustrate shortly.

There are other *scoring rules* besides BC, so we begin with a definition. Given *m* candidates, fix a non-increasing vector $(s_1,..., s_m)$ of real numbers ("scores") such that $s_i \geq s_{i+1}$ for all $i \in \{1,..., m - 1\}$ and $s_1 > s_m$. Each voter's k^{th} best candidate receives score s_k. A candidate's score is the sum of the scores that he or she receives from all voters.

For a preference profile **P**, a scoring rule selects the candidate or candidates that receive the highest score. A scoring rule is said to be *strict* if it is defined by a decreasing vector of scores, $s_i > s_{i+1}$, for all $i \in \{1,..., m - 1\}$.

We next show that all scoring-rule winners, whether they are Condorcet winners or not, are AV outcomes, but candidates that are selected by no scoring rule may also be AV outcomes:

Proposition 7. *At all preference profiles* ***P****, a candidate chosen by any scoring rule is an AV outcome. There exist preference profiles* ***P*** *at which a candidate is not chosen by any scoring rule but is, nevertheless, an AV outcome.*

Proof. We begin by proving the first statement. Take any preference profile **P** and any candidate *i* chosen by a scoring rule at **P**. Let $(s_1,..., s_m)$ be the scoring-rule vector that results in the election of candidate *i* at **P**. By a normalization of the scores, we can without loss of generality assume that $s_1 = 1$ and $s_m = 0$.

Note that AV can be seen as a variant of a nonstrict scoring rule, whereby every voter gives a score of 1 to the candidates in his or her strategy set *S* (approved candidates) and a score of 0 to those not in this set. AV chooses the candidate or candidates with the highest score.[7]

Let $r_k(x)$ denote the number of voters who consider candidate *x* to be the k^{th} best candidate at **P**. Because candidate *i* is picked by the scoring rule $(s_1,..., s_m)$, it must be true that

$$s_1[r_1(i)] + s_2[r_2(i)] + ... + s_m[r_m(i)] \geq s_1[r_1(j)] + s_2[r_2(j)] + ... + s_m[r_m(j)] \qquad (1)$$

for every other candidate *j*.

To show that the scoring-rule winner, candidate *i*, is an AV outcome, consider *i*'s critical strategy profile $C_i(\mathbf{P}, \mathbf{S})$. There are two cases:

[7] Of course, AV is not a scoring rule in the classical sense whereby voters give scores to candidates according to the same predetermined vector. The restrictions on the vector that sincere, admissible strategies impose is that (i) the first component (score of the top candidate) be 1, (ii) the m^{th} component (score of the bottom candidate) be 0, (iii) all components representing candidates at or above the lowest candidate a voter approves of are 1, and (iv) all components below the component representing this candidate are 0.

Case (i): Voters rank candidate i last. Under a scoring rule, these voters give a score of 0 to candidate *i*, a score of 1 to their top choices, and scores between 0 and 1 to the remaining candidates. Under AV, these voters give a score of 0 to candidate *i*, a score of 1 to their top choices, and scores of 0 to the remaining candidates at $C_i(\mathbf{P}, \mathbf{S})$.

Thus, candidate *i* does the same under the scoring rule as under AV (left side of inequality (1)), whereas all other candidates *j* do at least as well under the scoring rule as under AV (right side of inequality (1)). This makes the sum on the right side for the scoring rule at least as large as, and generally larger than, the sum of votes under AV, whereas the left side remains the same as under AV. Consequently, if inequality (1) is satisfied under the scoring rule, it is satisfied under AV at $C_i(\mathbf{P}, \mathbf{S})$.

Case (ii): Voters do not rank candidate i last. Under a scoring rule, these voters give candidate *i* a score of s_k if they rank him or her k^{th} best. Under AV, these voters give a score of 1 to candidate *i* at $C_i(\mathbf{P}, \mathbf{S})$. Thus, every s_k on the left side of equation (1) is 1 for candidate *i* under AV, which makes the sum on the left side at least as large as, and generally larger than, the sum under a scoring rule. By comparison, the sum on the right side for all other candidates *j* under AV is less than or equal to the sum on the left side, with equality if and only if candidate *j* is preferred to candidate *i* by all voters. Consequently, if inequality (1) is satisfied under the scoring rule, it is satisfied under AV at $C_i(\mathbf{P}, \mathbf{S})$.

Thus, in both cases (i) and (ii), the satisfaction of inequality (1) under a scoring rule implies its satisfaction under AV at candidate *i*'s critical strategy profile, $C_i(\mathbf{P}, \mathbf{S})$. Hence, a candidate chosen under any scoring rule is also an AV outcome.

To prove the second statement, consider the following 7-voter, 3-candidate example (Fishburn and Brams, 1983, p. 211):

Example 4

1. 3 voters: *a b c*
2. 2 voters: *b c a*
3. 1 voter: *b a c*
4. 1 voter: *c a b*

Because candidate *b* receives at least as many first choices as *a* and *c*, and more first and second choices than either, every scoring rule will select *b* as the winner. But *a* is the Condorcet winner and, hence, an AV outcome by Proposition 5.[8] Q.E.D.

We next show the outcomes of two social choice rules that are not scoring rules, the Hare system of single transferable vote (STV) and the majoritarian compromise (MC), are always AV outcomes (at their critical strategy profiles), whereas the converse is not true—AV outcomes need not be STV or MC outcomes. Before proving this result, we illustrate STV and MC with a 9-voter, 3-candidate example:[9]

Example 5

1. 4 voters: *a c b*
2. 2 voters: *b c a*

[8] Example 4 provides an illustration in which BC, in particular, fails to elect the Condorcet winner.
[9] These two voting systems, among others, are discussed in Brams and Fishburn (2002). MC, which is less well known than STV, was proposed independently as a voting procedure (Hurwicz and Sertel, 1997; Sertel and Yilmaz, 1999; Sertel and Sanver, 1999) and as a bargaining procedure under the rubric of "fallback bargaining" (Brams and Kilgour, 2001). As a voting procedure, the threshold for winning is assumed to be simple majority, whereas as a bargaining procedure the threshold is assumed to be unanimity, but qualified majorities are also possible under either interpretation.

3. 3 voters: *c b a*

Under STV, candidates with the fewest first-choice—and successively lower-choice—votes are eliminated; their votes are transferred to second-choice and lower-choice candidates in their preference rankings until one candidate receives a majority of votes. To illustrate in Example 5, because candidate *b* receives the fewest first-choice votes (2)—compared with 3 first-choice votes for candidate *c* and 4 first-choice votes for candidate *a*—*b* is eliminated and his or her 2 votes go to the second choice of the 2 type (2) voters, candidate *c*. In the runoff between candidates *a* and *c*, candidate *c*, now with votes from the type (2) voters, defeats candidate *a* by 5 votes to 4, so *c* is the STV winner.

Under MC, first-choice, then second-choice, and then lower-choice votes are counted until at least one candidate receives a majority of votes; if more than one candidate receives a majority, the candidate with the most votes is elected. Because no candidate in Example 5 receives a majority of votes when only first choices are counted, second choices are next counted and added to the first choices. Candidate *c* now receives the support of all 9 voters, whereas *a* and *b* receive 4 and 5 votes, respectively, so *c* is the MC winner.

Proposition 8. *At all preference profiles **P**, a candidate chosen by STV or MC is an AV outcome. There exist preference profiles **P** at which a candidate chosen by AV is neither an STV nor an MC outcome.*

Proof. We start by showing that every STV outcome is an AV outcome. Suppose candidate *i* is not an AV outcome at preference profile **P**. By Proposition 3, there exists a candidate *j* such that the number of voters who rank *j* as their best candidate and *i* as their

worst candidate exceeds the number of voters who prefer *i* to *j*. A fortiori, the number of voters who consider *j* as their best candidate exceeds those who consider *i* as their best candidate.

This result holds for any subset of candidates that includes both *i* and *j*. Hence, STV will never eliminate *j* in the presence of *i*, showing that *i* cannot be an STV winner.

Neither can *i* be an MC winner, because *j* will receive more first-place votes than *i*. If this number is not a majority, the descent to second and still lower choices continues until at least one candidate receives a majority. Between *i* and *j*, the first candidate to receive a majority will be *j*, because *j* receives more votes from voters who rank him or her first than there are voters who prefer *i* to *j*. Thus, *j* will always stay ahead of *i* as the descent to lower and lower choices continues until *j* receives a majority.

To show that AV outcomes need not be STV or MC outcomes, consider Example 4, in which the Condorcet winner, candidate *c*, is chosen under both STV and MC. Besides *c*, AV may also choose candidate *a* or candidate *b*: *a* is an AV outcome at critical strategy profile $C_a(\mathbf{P}, \mathbf{S}) = (a, a, a, a, b, b, c, c, c)$; and *b* is an AV outcome at critical strategy profile $C_b(\mathbf{P}, \mathbf{S}) = (a, a, a, a, b, b, cb, cb, cb)$. Q.E.D.

So far we have shown that AV yields at least as many, and generally more, (Pareto) outcomes than any scoring rule and two nonscoring voting systems.[10] To be sure, one might question whether the three possible AV outcomes in Example 4 have an equal claim to being *the* social choice. Isn't candidate *c*, the Condorcet winner, BC winner, STV winner, and MC winner—and ranked last by no voters—the best overall

[10] While scoring rules, STV, and MC offer a good range of different systems to compare AV with, we will later describe two Condorcet procedures, one proposed by Duncan Black and the other by A. H. Copeland,

choice? By comparison, candidate *b* is only a middling choice; and candidate *a*, who is the plurality-vote (PV) winner, is the Condorcet loser.[11]

But just as AV allows for a multiplicity of outcomes, it also enables voters to prevent them.

Proposition 9. *There exist preference profiles **P** at which AV can prevent the election of every candidate, whereas scoring rules, STV, and MC cannot prevent the election of all of them.*

Proof. Consider the following 3-voter, 3-candidate example:

Example 6

1. 1 voter: *a b c*
2. 1 voter: *b a c*
3. 1 voter: *c b a*

Let "|" indicate each voter's dividing line between the candidate(s) he or she considers acceptable and those he or she considers unacceptable. If the three voters draw their lines as follows,

a | *b c* *b a* | *c* *c* | *b a*,

b and *c* will *not* be chosen (*a* will be). If the voters draw their lines as follows,

a | *b c* *b* | *a c* *c b* | *a*,

---

that guarantee the election of Condorcet candidates if voters are sincere. Recent works comparing different voting systems include Saari (1994, 2001), Nurmi (1999, 2002), and Brams and Fishburn (2002).

[11] Note that PV is a degenerate scoring rule, under which a voter's top candidate receives 1 point and all other candidates receive 0 points. By Proposition 7, sincere outcomes under PV are AV outcomes but not

a and *c* will *not* be chosen (*b* will be). Thus the voters can prevent the election of every one of the three candidates under AV.

By contrast, the Condorcet winner, *b*, wins under every scoring system, including BC, and also under MC. Under STV, either *a* or *b* may win, depending on which of the three candidates is eliminated first. Thus, only *c* is prevented from winning under these other systems, showing that AV is unique in being able to prevent the election of each of the three candidates. Q.E.D.

We have seen that AV allows for outcomes that BC, MC, and STV do not (e.g., *c* in Example 6 when there is a three-way tie). At the same time, it may preclude outcomes (e.g., *b* in Example 6) that other systems cannot prohibit. In effect, voters can fine-tune their preferences under AV, making outcomes responsive to information that transcends these preferences.

We next consider not only what outcomes can and cannot occur under AV but also what outcomes are likely to persist because of their stability. While we know that non-Pareto candidates cannot win a clear-cut victory under AV (Proposition 4), might it be possible for Condorcet losers to be AV outcomes *and* stable? To answer this question, we will distinguish two types of stability.

4. Stability of Election Outcomes

As earlier, we assume that voters choose sincere, admissible strategies under AV. Now, however, we suppose that they may not draw the line between acceptable and

---

vice-versa. As a case in point, candidate *a* is the sincere PV outcome in Example 5, whereas candidates *b* and *c* are also sincere AV outcomes.

unacceptable candidates as they would if they were truthful. Instead, they may vote strategically in order to try to obtain a preferred outcome.

To determine what is "preferred," we extend preference to sets. If a voter's preference is *a b*, he or she will prefer *a* to {*a*, *b*}, and {*a*, *b*} to *b*. If a voter's preference is *a b c*, he or she may prefer any of outcomes *b*, {*a*, *c*}, or {*a*, *b*, *c*} to any other. In assessing the stability of outcomes later, we will assume that these are all admissible preference over sets.

We define two kinds of stability, the first of which is the following: Given a preference profile **P**, an AV outcome is *stable* if there exists a strategy profile **S** such that no voters of a single type have an incentive to switch their strategy to another sincere, admissible strategy in order to induce a preferred outcome.[12] In analyzing the stability of AV outcomes, we need confine our attention only to those outcomes stable at $C_i(\mathbf{P}, \mathbf{S})$ because of the following proposition:

Proposition 10. *An AV outcome i is stable if and only if it is stable at its critical strategy profile,* $C_i(\mathbf{P}, \mathbf{S})$.

Proof. The "if" part follows from the existence of a strategy profile, $C_i(\mathbf{P}, \mathbf{S})$, at which outcome *i* is stable. To show the "only if" part, assume candidate *i* is unstable at $C_i(\mathbf{P}, \mathbf{S})$. At any other strategy profile $\mathbf{S}'$, candidate *i* receives no more approval votes and generally fewer than at $C_i(\mathbf{P}, \mathbf{S})$ by Proposition 1. Hence, those voters who switch to different sincere, admissible strategies to induce the election of a preferred candidate at **S** can also do so at $\mathbf{S}'$. Q.E.D.

[12] Treating voters of one type, all of whose members have the same preference, as single (weighted) voters provides the most stringent test of stability. This is because any outcome that can be destabilized by the

The strategies of voters associated with a stable AV outcome at $C_i(\mathbf{P}, \mathbf{S})$ define a Nash equilibrium of a voting game in which the voters have complete information about each others' preferences and make simultaneous choices.[13]

Neither candidate *a* nor candidate *b* is a stable AV outcome in Example 5. At critical strategy profile $C_a(\mathbf{P}, \mathbf{S}) = (a, a, a, a, b, b, c, c, c)$ that renders candidate *a* an AV outcome, if the 2 type (2) voters switch to strategy *bc*, candidate *c*, whom the type (2) voters prefer to candidate *a*, wins. At critical strategy profile $C_b(\mathbf{P}, \mathbf{S}) = (a, a, a, a, b, b, cb, cb, cb)$ that renders candidate *b* an AV outcome, the 4 type (1) voters have an incentive to switch to strategy *ac* to induce the selection of candidate *c*, whom they prefer to candidate *b*.

Although AV outcomes *a* and *b* in Example 5 are not stable at their critical strategy profiles, AV outcome *c* most definitely is stable at its critical strategy profile, $C_c(\mathbf{P}, \mathbf{S}) = (ac, ac, ac, ac, bc, bc, c, c, c)$: No switch on the part of the 4 type (1) voters to *a*, of the 2 type (2) voters to *b*, or of the 3 type (3) voters to *cb* can lead to a preferred outcome for any of these types—or, indeed, change the outcome at all (because candidate *c* is the unanimous choice of all voters at *c*'s critical strategy profile).

Not only can no single switch by any of the three types induce a preferred outcome for the switchers at $C_c(\mathbf{P}, \mathbf{S})$, but no *coordinated* switches by two or more types can induce a preferred outcome. Thus, for example, if the *ac*-voters switched from *ac* to *a*, and the *bc* voters switched from *bc* to *b*, they together could induce AV outcome *a*, which the 4 type (1) voters would clearly prefer to outcome *c*. But *a* is the worst choice of the 2

---

switch of individual voters (of one type) can be destabilized by the switch of all voters of that type, but the converse is not true.

[13] For an analysis of Nash equilibria in voting games under different rules and information conditions from those given here, see Myerson (2002) and references cited therein.

type (2) voters, so they would have no incentive to coordinate with the type (1) voters to induce this outcome.

That AV outcome *c* is, at the critical strategy profile of candidate *c*, invulnerable to coordinated switches leads to our second type of stability: Given a preference profile **P**, an outcome is *strongly stable* if there exists a strategy profile **S** such that no types of voters, coordinating their actions, can form a coalition *K*, all of whose members would have an incentive to switch their AV strategies to other sincere, admissible strategies in order to induce a preferred outcome.

We assume that the coordinating players in *K* are allowed to communicate to try to find a set of strategies to induce a preferred outcome *for all of them*. These strategies define a *strong* Nash equilibrium of a voting game in which voters have complete information about each others' preferences and make simultaneous choices.

Proposition 11. *An AV outcome i is strongly stable if and only if it is strongly stable at its critical strategy profile,* $C_i(\mathbf{P}, \mathbf{S})$.

Proof. Analogous to that of Proposition 10. Q.E.D.

What we have yet to show is that an AV stable outcome need not be strongly stable. To illustrate this weaker form of stability, consider AV outcome *a* in Example 1 and its critical strategy profile, $C_a(\mathbf{P}, \mathbf{S}) = (a, a, a, bca, bca, d, d)$. The 2 type (2) voters cannot upset this outcome by switching from *bca* to *bc* or *b*, nor can the 2 type (3) voters upset it by switching from *d* to *db* or *dbc*. However, if these two types of voters cooperate and form a coalition *K*, with the 2 type (2) voters choosing strategy *b* and the 2 type (3) voters choosing strategy *db*, they can induce the selection of Condorcet winner *b*, whom both types prefer to candidate *a*. At critical strategy profile $C_a(\mathbf{P}, \mathbf{S})$, therefore,

AV outcome *a* is stable but not strongly stable, whereas AV outcome *b* is strongly stable at its critical strategy profile, $C_b(\mathbf{P}, \mathbf{S}) = (ab, ab, ab, b, b, db, db)$.

If an AV outcome is neither strongly stable nor stable, it is *unstable*. Clearly, strongly stable outcomes are always stable, but not vice versa.

Proposition 12. *AV outcomes are strongly stable, stable, or unstable. All three kinds of AV outcomes may coexist.*

Proof. We have just shown that AV outcome *b* is strongly stable, and AV outcome *a* is stable, in Example 1. We now show that candidate *c* in this example is an unstable AV outcome. At the critical strategy profile of candidate *c*, $C_c(\mathbf{P}, \mathbf{S}) = (abc, abc, abc, bc, bc, dbc, dbc)$, candidates *b* and *c* tie with 7 votes each. Candidate *c* might therefore be selected under some tie-breaking rule. But $\{b, c\}$ is not a stable AV outcome: If either the 3 *a*-voters switch to *ab*, the 2 *bc*-voters switch to *b*, or the 2 *dbc*-voters switch to *db*, candidate *b* will be selected, whom all three types of voters prefer to $\{b, c\}$.[14] Q.E.D.

While Proposition 12 shows that strongly stable, stable, and unstable AV outcomes may coexist, it is important to know the conditions under which each kind of outcome can occur. For this purpose, we call a Condorcet winner *tied* if it comprises two or more candidates that tie against each other but defeat all other candidates in pairwise contests.

Proposition 13. *An AV outcome is strongly stable if and only if it is a (possibly tied) Condorcet winner.*

[14] As we will show in Proposition 14, unstable AV outcomes do not necessarily include, as here, non-Pareto candidates as components. Unique Pareto candidates can also be unstable AV outcomes.

Proof. To prove the "if" part, suppose candidate *i* is a unique Condorcet winner at **P**. Suppose there exists a coalition of voters *K*, comprising one or more types, that prefers some other candidate *j* to candidate *i* and coordinates to induce the selection of *j*. Because candidate *i* is a unique Condorcet winner, however, the cardinality of *K* is strictly less than the cardinality of coalition *L*, whose members prefer *i* to *j*. Hence, whatever sincere, admissible strategy switch the members of *K* consider at candidate *i*'s critical strategy profile to induce the election of candidate *j*, *j* will receive fewer votes than *i*, proving that *i* is a strongly stable AV outcome. If a Condorcet winner comprises two or more (tied) candidates, none can defeat any other, so all will be strongly stable.

Now it is possible that a majority of voters have preferences either of the form (1) *a b c*… or (2) *c b a*…. Moreover, *b* may be a Condorcet winner (because of the preferences of the other voters), whereas *a* and *c* are not. If both the type (1) and the type (2) voters prefer {*a*, *c*} to *b*, and coordinate their strategies by approving only of their first choices, they may be able to elect {*a*, *c*} initially—depending on the AV strategies of the other voters—even though *b* is a Condorcet winner. In the end, however, either *a* or *c* will be chosen, according to some tie-breaking rule. Because Condorcet winner *b* is able to defeat each of these candidates in separate pairwise contests, {*a*, *c*}—after the tie is broken—is not stable. This exhausts the situations in which a set, comprising two or more non-Condorcet candidates, potentially may destabilize a Condorcet winner under AV.

To prove the "only if" part, suppose that candidate *i* is not a Condorcet winner. Consequently, there exists a coalition of voters *K*, comprising one or more types, that prefers some other candidate *j* to *i* and coordinates to induce the election of *j*. Because *i*

is not a Condorcet winner, the cardinality of K is a strict majority. We will now show that i is not a strongly stable AV outcome at its critical strategy profile, $C_i(\mathbf{P}, \mathbf{S})$, which by Proposition 11 shows that i is not a strongly stable AV outcome. Suppose AV does not elect i at $C_i(\mathbf{P}, \mathbf{S})$. Then i is not an AV outcome and hence not a strongly stable one. Now suppose that AV elects i at $C_i(\mathbf{P}, \mathbf{S})$. Because the members of K can change their strategies to elect j, whom they prefer to i, i is not a strongly stable AV outcome. Q.E.D.

To illustrate the case of a tied Condorcet winner, assume that there are two candidates, i and j, and two voters, 1 and 2. If voter 1 ranks the candidates i j, and voter 2 ranks the candidates j i, then (i, j) is the critical strategy profile for both candidates. Obviously, neither one nor both voters can induce an outcome each prefers to the tied outcome $\{i, j\}$, so this outcome is strongly stable without there being a unique Condorcet winner.[15]

Our next proposition establishes that every AV outcome may be unstable.

Proposition 14. *There may be no strongly stable, or even stable, AV outcomes—that is, every AV outcome may be unstable.*

Proof. That there may be no stable AV outcomes is shown by the following 3-voter, 3-candidate example:[16]

Example 7

1. 1 voter: *a b c*
2. 1 voter: *b c a*

[15] Under a somewhat weaker definition of a strong Nash equilibrium, the equivalence of strong Nash-equilibrium outcomes and Condorcet winners is shown for a large class of voting rules in Sertel and Sanver (forthcoming).

3. 1 voter: *c a b*

Consider the critical strategy profile that selects candidate *a*, $C_a(\mathbf{P}, \mathbf{S}) = (a, b, ca)$. If voter (2) switches to *bc*, he or she can induce preferred outcome $\{a, c\}$. In a similar manner, it is possible to show that neither candidate *b* nor candidate *c* is a stable AV outcome. Q.E.D.

We next show that Condorcet losers as well as winners may be stable AV outcomes.

Proposition 15. *A unique Condorcet loser may be a stable AV outcome, even when there is a different outcome that is a unique Condorcet winner (and therefore strongly stable).*

Proof. Consider the following 7-voter, 5-candidate example:

Example 8

1. 3 voters: *a b c d e*
2. 1 voter: *b c d e a*
3. 1 voter: *c d e b a*
4. 1 voter: *d e b c a*
5. 1 voter: *e b c d a*

Candidate *a* is the Condorcet loser, ranked last by 4 of the 7 voters. But at its critical strategy profile, $C_a(\mathbf{P}, \mathbf{S}) = (a, a, a, b, c, d, e)$, candidate *a* is a stable AV outcome,

[16] Example 7 is the standard example of the Condorcet paradox, or cyclical majorities, in which there is no Condorcet winner.

because none of the four individual voters, by changing his or her strategy, can upset *a*, who will continue to receive 3 votes.

Consider the critical strategy profile of candidate *b*, $C_b(\mathbf{P}, \mathbf{S})$ = (*ab*, *ab*, *ab*, *b*, *cdeb*, *deb*, *eb*), who receives 7 votes, compared with 3 votes each for *a* and *e*, 2 votes for *d*, and 1 vote for *c*. Again, no single type of voter can upset this outcome, nor can any coalition, because candidate *b* is the unique Condorcet winner, making him or her strongly stable. Q.E.D.

We consider later whether a Condorcet loser, like candidate *a* in Example 8, "deserves" to be an AV outcome—and a stable one at that. But first we will compare the stability of outcomes under AV and under other systems.

For this purpose, we need stability definitions for these other voting systems. A *stable outcome* under each of these systems is one in which no single type of voter has an incentive to switch its ranking to another ranking in order to induce a preferred outcome in a voting game in which voters have complete information about each others' preferences and make simultaneous choices. A *strongly stable outcome* is one in which no types of voters, coordinating their actions, can form a coalition *K*, all of whose members have an incentive to switch their rankings in order to induce a preferred outcome.

Proposition 16. *Stable or strongly stable AV outcomes need not be stable or strongly stable scoring-system, STV, or MC outcomes. Conversely, stable or strongly stable scoring-system, MC, or STV outcomes—while always AV outcomes—need not be stable or strongly stable AV outcomes.*

Proof. In the proof of Proposition 7, we indicated that there is no scoring system that selects candidate *a* in Example 4, so obviously *a* cannot be a stable or strongly stable outcome under a scoring system. Because *a* is the unique Condorcet winner in this example, however, it is a strongly stable, and therefore a stable, AV outcome. Likewise, candidate *a* is a stable AV outcome in Example 1 (see proof of Proposition 11), but it is not the STV or MC outcome (candidate *b* is) so cannot be a stable outcome under these systems.

Next we show that both a scoring system and MC may select a stable outcome that is not a stable AV outcome. Consider the following 6-voter, 3-candidate example:

Example 9

1. 2 voters: *a c b*
2. 3 voters: *b c a*
3. 1 voter: *c a b*

Under BC, *c* gets 7 points, *b* gets 6 points, and *a* gets 5 points. Moreover, *c* is stable (and strongly stable): If either the 2 type (1) voters or the 3 type (2) voters interchange *c* and their last choice, the last choice wins, which is worse for them. Likewise under MC, *c* wins by getting 6 votes to 3 votes for *b* and 2 votes for *a* at level 2. Moreover, neither the type (1) nor the type (2) voters can obtain a preferred outcome by giving a different preference ranking.

By contrast, under AV, *c* gets 6 votes at its critical strategy profile, $C_c(\mathbf{P}, \mathbf{S})$ = (*ac, ac, bc, bc, bc, c*), whereas *b* gets 3 votes and *a* gets 2 votes. But by voting only for *b*, the 3 type (2) voters can induce {*b, c*}, which they prefer to *c* alone, so *c* is not a stable (or strongly stable) AV outcome.

Finally, we show that STV may elect a stable outcome that is not a stable AV outcome. Consider the following 6-voter, 4-candidate example:

Example 10

1. 3 voters: *a b d c*
2. 2 voters: *c b d a*
3. 1 voter: *d c b a*

Under STV, first *d* is eliminated, after which the type (3) voter's vote is transferred to *c*, which ties with *a* (3 votes each) to give {*a*, *c*} as the outcome. It is easy to see that neither the type (3) voter—whose second or fourth choice will be elected—nor the type (1) or type (2) voters can give (false) preference rankings that lead to a preferred outcome.

Under AV, there are two strategy profiles that elect {*a*, *c*}: (i) (*a*, *a*, *a*, *c*, *c*, *dc*) and (ii) (*a*, *a*, *a*, *c*, *c*, *dcb*). If the 3 type (1) voters prefer {*a*, *b*, *c*} to {*a*, *c*} in the case of (i), or *b* to {*a*, *c*} in the case of (ii), they will have an incentive to switch from strategy *a* to strategy *ab*. Hence, {*a*, *c*} is not stable under all admissible preferences over sets, proving that a stable STV outcome may not be a stable AV outcome. Q.E.D.

Thus, no system, including AV, dominates others with respect to producing stable outcomes: All the systems we have analyzed so far may produce stable or strongly stable outcomes, but this stability is not necessarily duplicated under other systems. We next ask whether voting systems that guarantee the election of Condorcet winners when voters are sincere do any better.

A *Condorcet voting system* is one that always elects a Condorcet winner, if one exists, when voters are sincere. This candidate, however, may not be a Nash-equilibrium outcome, much less a strong one (as under AV).

Proposition 17. *Condorcet winners may not be stable or strongly stable outcomes under Condorcet voting systems.*

Proof. Consider again Example 3, in which candidate *a* is the unique Condorcet winner and thus the strongly stable outcome under AV:

Example 3 (repeated)

1. 2 voters: *a d b c*
2. 2 voters: *b d a c*
3. 1 voter: *c a b d*

We consider two Condorcet systems and show that the 2 type (2) voters can induce an outcome they prefer to *a* by misrepresenting their preferences as (2):

1. 2 voters: *a d b c*
2 . 2 voters: *b d c a*
3. 1 voter: *c a b d*

Under *Black's system*, the winner is the Condorcet winner. If there is no Condorcet winner, the winner is the BC winner.

There is no Condorcet winner under the (2) misrepresentation; the BC winner is candidate *b* with 9 points (to 8 points for candidate *a*). Thus, the 2 type (2) voters can induce *b*, which they prefer to *a*.

Under *Copeland's system*, the winner(s) are the candidate(s) that defeat the most other candidates in pairwise contests. Because candidate *a* defeats two candidates (*b* and *d*) and candidate *b* also defeats two candidates (*c* and *d*) under the (2) misrepresentation, the outcome is {*a*, *b*}, which the 2 type (2) voters prefer to *a*. In sum, the sincere ranking of the type (2) voters is not in equilibrium under either Black's system or Copeland's system. Q.E.D.

Thus, the two Condorcet systems we have considered do not do what they purport to do *in equilibrium*, whereas AV always renders Condorcet winners strong Nash-equilibrium outcomes.

5. Rational Departures from Unstable Outcomes

The (nonstrong) stability of outcomes under the different systems that we have analyzed so far is very much a local property: It depends on the *inability* of voters of one type, by misrepresenting their preferences, to change the sincere outcome to their advantage.

AV probably allows for more of this kind of stability than any other system—in addition to rendering Condorcet winners strongly stable. But we have seen that there are preference profiles in which other systems give (strong) stability when AV does not (Proposition 16).

When AV outcomes are unstable, rational departures from them are *closed*—they always lead to other AV outcomes—but this is not necessarily the case with other systems, as we next show.

Proposition 18. *Consider the set of AV outcomes, and suppose that a voter type finds it rational to depart from its AV strategy to effect a preferred outcome. Then the preferred outcome will also be an AV outcome, making the set of AV outcomes closed with respect to rational departures by a voter type. But a rational departure under another system may not yield an AV outcome.*

Proof. If a voter type finds it rational to depart from an unstable outcome under AV, then its voters do so by approving of more or fewer candidates, using a different sincere strategy. But this strategy is also an AV strategy, so the preferred outcome that is induced is an AV outcome. Thus, the set of AV outcomes is closed with respect to rational departures by a voter type.

To show that a rational departure of a voter type under another system may not give an AV outcome, consider the following 2-voter, 5-candidate example:

Example 11

1. 1 voter: *a b c d e*
2. 1 voter: *d e c b a*

Under BC, *d* wins with 5 points. But if voter (1) falsely indicates his or her preferences to be (1) below,

1 . 1 voter: *b a c d e*

2. 1 voter: *d e c b a*

the outcome will be {*b*, *d*}, each of whose two members receives 5 points.

Although voter (1), based on his or her *true* preference, prefers this outcome to *d*, {*b*, *d*} is not an AV outcome. To show this, assume that the sincere strategies of voters

(1, 2) are (*ab, d*); then the outcome is {*a, b, d*}. If the sincere strategies of voters (1, 2) are (*abcd, decb*), then the outcome is {*b, c, d*}. While these 3-way ties that include *b* and *d* are AV outcomes, the 2-way tied outcome {*b, d*} is not—there are no sincere strategies that give this outcome. Hence, a rational departure under BC may not yield an AV outcome. Q.E.D.

That the set of rational departures is closed under AV does not say which AV outcomes are reachable from what others via rational departures. In future research, we plan to explore how, through manipulation, AV outcomes may be transformed into other AV outcomes. Of course, stable AV outcomes will be sinks, from which no rational departures will occur. But just as important as the local (and global) stability of outcomes are the dynamics of changes that might occur when outcomes are unstable.

6. Summary and Conclusions

We began by suggesting that more than one outcome of an election may be acceptable, which we illustrated with the 2000 US presidential election. But our focus was on elections with three or more candidates, wherein acceptability is determined by where voters draw the line between acceptable and unacceptable candidates.

This criterion of acceptability underlies the choices that voters make under AV. We summarize both our positive findings and our normative conclusions about how voter sovereignty, expressed through AV, affects election outcomes:

1. AV enables voters to indicate those points in their preference rankings above which candidates are acceptable, which we believe is information that *should* determine

their social choices. This information may either rule in or rule out candidates that would be chosen under other systems.

2. AV may select a multitude of candidates—even all candidates in a race—at their critical strategy profiles. This may include even Condorcet losers, whom most social choice theorists would condemn as egregious choices, especially if there is also a Condorcet winner in the race. Our view is different: The choice of even a Condorcet loser may, on occasion, be justified.

3. We do not disparage Condorcet winners, which no fixed rule may elect. Furthermore, there may be a candidate or candidates who stand in between a Condorcet winner and a Condorcet loser that are also AV outcomes, rendering several candidates in a race viable.

4. Grounding social choices on the notion of acceptability rather than on traditional social-choice criteria is a radical departure from the research program initiated by Borda and Condorcet in late 18th-century France (McLean and Urken, 1995). While we do not eschew these criteria, they should not be the be-all and end-all for judging whether outcomes are acceptable or not. Rather, the pragmatic judgments of sovereign voters about who is acceptable and who is not should be decisive.

5. Voters' judgments will be affected by the stability of different AV outcomes. Thus, a Condorcet loser may not be a stable AV outcome—much less a strongly stable one—so this candidate's viability is less than were he or she a stable outcome.

6. When there is a unique Condorcet winner, this candidate is always a strongly stable AV outcome. Thereby AV preserves the majority will, at least if there is some kind of strategic coordination among voters, through sincere voting.

7. In large-scale elections, this coordination is possible to a limited extent from information provided by polls (Brams and Fishburn, 1983, ch. 7). But strategizing by voters may not be perfect, allowing Condorcet winners, even under AV, to be defeated on occasion.

8. This failure may sometimes be salutary, especially when a BC winner differs from a Condorcet winner and "majority tyranny" is a concern (Baharad and Nitzan, 2002). In such a situation, the BC winner may be a more acceptable candidate, even if he or she is not, like the Condorcet winner, a strongly stable AV outcome.

9. While the stability and strong stability of outcomes facilitates their selection, even unstable AV outcomes should be considered acceptable, especially if there are no stable outcomes because of a Condorcet paradox (Miller, 1983).

10. Speaking normatively, AV provides a better way of finding consensus choices than do other voting systems because of the information that it both suppresses (preference rankings) and expands upon (who is acceptable and who is not in the rankings):

- It is simpler to use than the alternative voting systems, except possibly PV, because it does not require that the voters rank candidates (often an arduous task if there are more than about five candidates).

- It may actually make choices easier than PV for a voter who (i) is relatively indifferent among more two or more candidates or (ii) favors a candidate that is not competitive (e.g., Ralph Nader in 2000) and, hence, may also want to vote for a more viable second choice (e.g., Al Gore).

11. While different outcomes may be strongly stable, stable, or unstable under different systems, AV probably endows outcomes with more stability, on average, than do its competitors. Condorcet systems, in particular, do not always elect Condorcet candidates in equilibrium.

12. Rational departures from unstable AV outcomes are always AV outcomes, which is not true of rational departures from unstable outcomes under other systems.

13. The local stability of Nash equilibria and the global stability of strong Nash equilibria indicate that acceptability may be stabilized at different levels. No level is sacrosanct. Thus, we see no reason to insist that the strong stability of a Condorcet winner under AV should supersede the local stability of a different Borda winner, or even the possible instability of a Condorcet loser.

14. Ultimately, acceptability depends on the judgments of voters. AV provides a compelling means for them to exercise their sovereignty, both for and against candidates.

References

Arrow, Kenneth J. (1963). *Social Choice and Individual Values*, 2nd ed. New Haven, CT: Yale University Press.

Baharad, Eyal, and Shmuel Nitzan (2002). "Ameliorating Majority Decisiveness through Expression of Preference Intensity." *American Political Science Review* 96, no. 4 (December): 745-754.

Brams, Steven J., and Peter C. Fishburn (1978). "Approval Voting." *American Political Science Review* 72, no. 3 (September): 831-847.

Brams, Steven J., and Peter C. Fishburn (1983). *Approval Voting*. Boston: Birkhäuser.

Brams, Steven J., and Peter C. Fishburn (2003). "Going from Theory to Practice: The Mixed Success of Approval Voting." Preprint, Department of Political Science, New York University.

Brams, Steven J., and Peter C. Fishburn (2002). "Voting Procedures." In Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura (eds.), *Handbook of Social Choice and Welfare*, vol. 1. Amsterdam: Elsevier Science, pp. 175-236.

Brams, Steven J., Peter C. Fishburn, and Samuel Merrill III (1988a). "The Responsiveness of Approval Voting: Comments on Saari and Van Newenhizen." *Public Choice* 59 (November): 121-131.

Brams, Steven J., Peter C. Fishburn, and Samuel Merrill (1988b). "Rejoinder to Saari and Van Newenhizen." *Public Choice* 59 (September): 149.

Brams, Steven J., and D. Marc Kilgour (2001). "Fallback Bargaining." *Group Decision and Negotiation* 10, no. 4 (July): 287-316.

Fishburn, Peter C., and Steven J. Brams (1983). "Paradoxes of Preferential Voting."

Mathematics Magazine 56, no. 4 (September): 207-214.

Hurwicz, Leonid, and Murat R. Sertel (1997). "Designing Mechanisms, in Particular for Electoral Systems: The Majoritarian Compromise." Preprint, Department of Economics, Bogazici University, Istanbul, Turkey.

McLean, Ian, and Arnold B. Urken (eds.) (1995). *Classics of Social Choice*. Ann Arbor, MI: University of Michigan Press.

Miller, Nicholas R. (1983). "Pluralism and Social Choice." *American Political Science Review* 77, no. 3 (September): 734-747.

Myerson, Roger B. (2002). "Comparison of Scoring Rules in Poisson Voting Games." *Journal of Economic Theory* 103, no. 1 (March): 219-251.

Nurmi, Hannu (1999). *Voting Procedures and How to Deal with Them*. Berlin: Springer-Verlag.

Nurmi, Hannu (2002). *Voting Procedures under Uncertainty*. Berlin: Springer-Verlag.

Saari, Donald G. (1994). *Geometry of Voting.* New York: Springer-Verlag.

Saari, Donald G. (2001). *Chaotic Elections! A Mathematician Looks at Voting.* Providence, RI: American Mathematical Society.

Saari, Donald G., and Jill Van Newenhizen (1988a). "The Problem of Indeterminacy in Approval, Multiple, and Truncated Voting Systems." *Public Choice* 59 (November): 101-120.

Saari, Donald G., and Jill Van Newenhizen (1988b). "Is Approval Voting an 'Unmitigated Evil': A Response to Brams, Fishburn, and Merrill." *Public Choice* 59 (November): 133-147.

Sertel, Murat R., and M. Remzi Sanver (1999). "Designing Public Choice Mechanisms,"

in Imed Limam (ed.), *Institutional Reform and Development in the MENA Region*. Cairo, Egypt: Arab Planning Institute, pp, 129-148.

Sertel, Murat R., and M. Remzi Sanver (forthcoming). "Strong Equilibrium Outcomes of Voting Games are the Generalized Condorcet Winners. *Social Choice and Welfare*.

Sertel, Murat R., and Bilge Yilmaz (1999). "The Majoritarian Compromise Is Majoritarian-Optimal and Subgame-Perfect Implementable." *Social Choice and Welfare* 16, no. 4 (August): 615-627.