

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

CLOSENESS MATTERS: MONOTONICITY FAILURE IN IRV ELECTIONS WITH THREE CANDIDATES

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

February 2012
Revised August 2013
Second revision January 2016
Third revision September 2016
Fourth revision February 2017

Forthcoming in *Public Choice*

Abstract

A striking attribute of Instant Runoff Voting (IRV) is that it is subject to *monotonicity failure* — that is, getting more (first preference) votes can cause a candidate to lose an election and getting fewer votes can cause a candidate to win. Proponents of IRV have argued that monotonicity failure, while a mathematical possibility, is highly unlikely to occur in practice. This paper specifies the precise conditions under which this phenomenon arises in three-candidate elections, and then applies them to a number of large data sets in order to get a sense of the likelihood of IRV's monotonicity problem in varying circumstances. The basic finding is that the monotonicity problem is significant in many circumstances and very substantial whenever IRV elections are closely contested by three candidates.

An earlier version of this paper was presented at the Second World Congress of the Public Choice Societies, Miami, March 8-11, 2012. This revision uses more extensive simulations, replaces tables with graphs, and corrects several errors, but it entails no substantial changes in results. The conference version of the paper has been cited in several subsequent works, which are in turn cited here. I thank Robert Z. Norman, Joseph Ornstein, and Dan Felsenthal for helpful comments.

CLOSENESS MATTERS: MONOTONICITY FAILURE IN IRV ELECTIONS WITH THREE CANDIDATES

1. Introduction

It is generally recognized that ordinary plurality voting (or first-past-the-post) is problematic in elections with three or more candidates. Instant Runoff Voting (IRV) — also known as the Alternative Vote, Hare Rule, and Ranked Choice Voting — has been proposed as a less troublesome alternative. Under IRV, voters rank the candidates in order of preference. A candidate supported by a majority of first preferences is elected. Otherwise, the candidate supported by the fewest first preferences is eliminated and his or her ballots are transferred to other candidates on the basis of second preferences. This process is repeated until one candidate is supported by a majority of ballots. Here we consider only three-candidate contests, so IRV is limited to a single ‘instant runoff’ and is formally equivalent to plurality voting with a runoff between the top two candidates in the event the leading candidate is not supported by a majority of votes.

A striking feature of IRV is that getting more (first-preference) votes may result in defeat for a candidate who would otherwise win and getting fewer votes may result in victory for a candidate who would otherwise lose. Most voting rules, including plurality voting, never exhibit this anomaly and accordingly are *monotonic*. But many years ago, Smith (1973) showed that rules that incorporate runoffs are subject to *monotonicity failure*. Several years later, Doron and Kronick (1977; also see Straffin 1980) observed that Smith’s class of runoff systems includes the Single Transferable Vote (STV), the single-winner variant of which is IRV. This finding attracted some attention among political scientists (Riker 1982, pp. 49-50; Brams and Fishburn 1983; Fishburn and Brams 1983). Recently, Felsenthal and Tideman (2013, 2014) have confirmed by examples that other runoff systems are subject to monotonicity failure — indeed, to what in Section 3.3 of this paper is dubbed ‘double monotonicity failure’.

Many proponents of IRV (and others) have argued that monotonicity failure under IRV, while a mathematical possibility, is highly unlikely to occur in practice. Thus, Amy (2002, p. 55) says: ‘While it is clear that nonmonotonicity can theoretically occur in an IRV election, most experts believe that the conditions needed for this paradox to occur are so special that it would be an extremely rare occurrence. One statistical study [Allard 1996] found that if IRV-like elections were to be held throughout the United Kingdom, a nonmonotonic result would occur less than once a century’. In a ‘hands-on assessment of STV’ based on his years of experience as the Chief Electoral Officer for Northern Ireland since the introduction of STV in 1973, Bradley (1995) reported that ‘the experience of the use of STV in Northern Ireland over the past 22 years, involving a range of election types and sizes, reveals no evidence to support *in practice* the lack of monotonicity’. A standard text on electoral systems (Farrell 2001, p. 150) cites both Allard and Bradley to support the claim that ‘there is no evidence that it [i.e., monotonicity failure under IRV] is a common occurrence’. Fair Vote (2009; also see Poundstone 2008, pp. 267-268), a US electoral reform group that advocates IRV (but calls it Ranked Choice Voting), claims that, ‘in terms of the frequency of non-monotonicity in real-world elections, ‘there is no evidence that this has ever played a role in any IRV election. . . .’

Since monotonicity failure is a striking and counterintuitive phenomenon, it may be helpful to provide a (more or less) real-world example — namely, a simplified version of the 2009 IRV election for mayor of Burlington, Vermont.¹ The Republican candidate was supported by 39% of the first preferences, the Democratic candidate by 27%, and the Progressive (left-of-Democrat) candidate by 34%. The Democrat therefore was eliminated, with his ballots transferring to one or other surviving candidate on the basis of second preferences, which were 37% for the Republican and 63% for the Progressive, representing 10% and 17% of the total electorate. Thus, in the instant runoff the Republican got $39\% + 10\% = 49\%$ and the Progressive won the election with $34\% + 17\% = 51\%$. Now consider a wholly make-believe sequel. A third of the Republicans (13% of all voters) are so traumatized by the prospect of a Progressive mayor that they flee Burlington for more politically hospitable climes and are replaced by a like number of newcomers attracted by the prospect of a Progressive mayor. At the next election, all votes are cast exactly as before, except for the 13% of the electorate once made up of Republicans now replaced by Progressives. The Progressive candidate won a squeaker before, so with this augmented support he will surely win more comfortably this time. But, in fact, he does not win at all. The Republican candidate now has 26% of the vote, the Democrat 27%, and the Progressive 47%, so the instant runoff is between the Democrat and Progressive, and the Democrat wins handily by gaining the second preferences of the (remaining) Republican voters (who find the prospect of a Democratic mayor at least marginally more tolerable than a Progressive one). The consequence of the Progressive candidate's first-preference support being augmented by 13% of the electorate is that he loses where before he won.

Section 2 introduces basic notation and terminology and formally defines the phenomenon of monotonicity failure under IRV. Section 3 specifies the precise conditions under which variants of monotonicity failure arise in three-candidate elections. These theoretical results tell us that this problem can occur only in relatively competitive three-candidate elections, but they do not directly indicate its frequency in such elections. Thus, Section 4 applies the conditions to various sets of simulated election results and shows that the problem arises quite frequently in competitive elections and the frequency increases as elections become still more closely contested. Section 5 offers some concluding observations.

2. Preliminaries

A three-candidate IRV *ballot profile* is a set of n rankings of candidates X , Y , and Z , one for each of n voters. We assume that all voters rank all three candidates and that the number of voters is large (i.e., electorate-sized) and fixed.² Given a particular ballot profile B , the candidate with the

¹ The description here is simplified in that there were other minor candidates and some voters cast 'truncated' ballots (that did not rank all candidates). Complete vote tallies may be found at <http://rangevoting.org/Burlington.html>.

² If the number is small, the problem of ties becomes prominent, but there is no standard way to break ties under IRV (or either plurality voting). If the number of voters can vary, IRV is subject to additional types of monotonicity problems; see, for example, Norman (2010), Felsenthal and Tideman (2013) and Felsenthal and Nurmi (2016).

most first preferences is the *plurality winner*, the candidate with the second most first preferences is the *plurality runner-up*, and the candidate with the fewest first preferences is the *plurality loser*. Let $n(PW)$, $n(P2)$ and $n(PL)$ be the number of ballots that rank the plurality winner, the plurality runner-up, and plurality loser first; always $n(PL) < n/3 < n(PW)$. If $n/2 < n(PW)$, the plurality winner is also a *majority winner*.

Given a three-candidate ballot profile B , let x , y and z be the number of ballots giving first-preference support to candidates X , Y and Z , respectively. Let x_y be the number of ballots showing a first preference for X and second preference for Y (and therefore a third preference for Z), and likewise for x_z and for other candidates. Given a different ballot profile B' , the candidates have x' , y' and z' first preferences, respectively, and likewise for the other notation.

Given a ballot profile B , if a majority of voters rank X over Y , i.e., if $x + z_x > y + z_y$, we say that X *beats* Y , and likewise for other pairs of candidates. A *Condorcet winner* beats both of the other candidates, and a *Condorcet loser* is beaten by both; otherwise there is a *Condorcet cycle* such that X beats Y beats Z beats X or the reverse.

The IRV winner is the majority winner if one exists and otherwise is either the plurality winner or the plurality runner-up, depending on which beats the other. A Condorcet winner fails to be the IRV winner if and only if it is also the plurality loser. If the plurality runner-up beats the plurality winner, IRV and plurality voting produce different winners and we call such a profile *critical*.

Profiles in which the plurality loser has at least one quarter of the first-preference support turn out to be especially significant. We call such profiles *competitive*; it can be checked that such profiles have no majority winner.

Our aim is to specify conditions under which an IRV ballot profile B is *vulnerable to monotonicity failure*, that is:

- (a) X wins under B , but there is some other profile B' (i) that differs from B only in that some voters rank X higher in B' than in B and (ii) under which X loses (*Upward Monotonicity Failure* or UMF), or
- (b) X loses under B , but there is some other profile B' (i) that differs from B only in that some voters rank X lower in B' than in B and (ii) under which X wins (*Downward Monotonicity Failure* or DMF).

In either event, every voter ranks Y and Z the same way under B and B' . Following Norman (2010), we call B and B' *companion* profiles.

Given that Z is the plurality loser under profile B , two conditions are necessary and jointly sufficient to make B vulnerable to (upward or downward) monotonicity failure in the event that X is moved up or down in some voter orderings.

Condition 1. This condition pertains to the *runoff pair* and requires that the ballot changes that produce B' from B must deprive Y of enough first preferences (for UMF), or give Z enough additional first preferences (for DMF), that the runoff that had been between X and Y is now between X and Z .

Condition 2. This condition pertains to the *runoff outcome* and requires that X , which won (for UMF) or lost (for DMF) the runoff against Y under B , must lose (for UMF) or win (for DMF) the runoff against Z under B' .

3. Conditions for monotonicity failure

As a step along the way to providing analytical expressions for calculating the proportion of three-candidate ballot profiles in the ‘impartial anonymous culture’ that are vulnerable to monotonicity failure under both plurality and anti-plurality runoff voting, Lepelley et al. (1996) specified — very concisely and using rather formidable notation — conditions for UMF and DMF failure under both voting rules. The conditions for IRV are presented here more transparently, in terms that highlight the operation of IRV and the plurality and Condorcet concepts reviewed in the previous section.³

3.1 Upward monotonicity failure

Let X be the IRV winner and let Z be the plurality loser under ballot profile B . It can be checked that ballot changes that move X upwards from third to second place, or to first place on ballots that had Z in first place, cannot change the IRV winner. Therefore, the essential difference between B and any companion profile B' that produces UMF is that X is ranked first on some ballots in B' on which Y was ranked first in B .

If profile B is vulnerable to UMF, Condition 1 requires that X can gain enough first-preference ballots at Y 's expense that two things are simultaneously true in the resulting ballot profile B' : (i) X is still not a majority winner, and (ii) Y becomes the plurality loser instead of Z . This can occur whenever $n/2 - x > y - z$, which simplifies to $z > n/4$. Condition 2 requires that Z beat X under B' and thus under B . Condition 1 can be simplified and, given Condition 1, Condition 2 can be restated in terms of the original profile B (see Appendix).

Proposition 1. A three-candidate ballot profile B in which X is the IRV winner and Z is the plurality loser is vulnerable to UMF if and only if:

Condition 1U: B is competitive, i.e., $z > n/4$; and

Condition 2U: the plurality loser beats the IRV winner, i.e., $z + y_z > x + y_x$.

UMF can therefore occur only in relatively closely contested elections such that all three candidates receive between 25% and 50% of the first preference votes. However, both critical and non-critical profiles may be vulnerable to UMF.

3.2 Downward monotonicity failure

Let Y be the IRV winner and Z be the plurality loser under ballot profile B . It can be checked that ballot changes that move X downwards from second to third place, or increase Y 's first preferences

³ While Lepelley et al. (1996) provide a proof of their version of Proposition 2 in their Appendix A, they do not provide a proof of their version of Proposition 1 (but rather cite Berg and Lepelley, 1993), so one is provided in the appendix of this paper.

by moving X downwards, cannot change the IRV winner. Therefore, the essential difference between the initial ballot profile B and any companion ballot profile B' that produces DMF is that X is ranked second and Z first on some ballots in B' on which X was ranked first in B .

If profile B is vulnerable to DMF, Condition 1 requires that X can lose enough first-preference ballots in favor of Z that Y becomes the plurality loser rather than Z . This can occur whenever $x - y > y - z$, which simplifies to $n/3 > y$. Furthermore, in order for Z to gain these first preferences rather than Y , it must be that $x_z > y - z$. Condition 2 requires that X beat Z under B' . This requires that X beat Z under B and that X must still beat Z after $(y - z)$ first-preference ballots shift from X to Z , which simplifies to $y_z < n/2 - y$.

Proposition 2. A three-candidate ballot profile B in which Y is the IRV winner and Z is the plurality loser is vulnerable to DMF if and only if:

- Condition 1D:* (a) the IRV winner has the first-preference support of fewer than one-third of the voters, i.e., $y < n/3$, and (b) $x_z > y - z$; and
Condition 2D: $y_z < n/2 - y$.

Since $y < n/3$, Y cannot be the plurality winner, so only critical profiles are vulnerable to DMF. Since X and Z together have more than two-thirds of the first-preference support and X cannot be a majority winner, $z < n/6$, suggesting that elections vulnerable to DMF may be less closely contested than those vulnerable to UMF.

The following corollary states more transparent necessary (but not sufficient) conditions for DMF.

Corollary 2.1. A three-candidate ballot profile B is vulnerable to DMF under IRV only if

- (1) the IRV winner has the first-preference support of fewer than one-third of the voters; and
- (2) the plurality winner beats the plurality loser.

Proposition 1 and Corollary 2.1 together imply that a three-candidate ballot profile B in which the plurality winner is also a Condorcet winner is not vulnerable to either UMF or DMF.

3.3 Double monotonicity failure

There is an obvious connection between UMF and DMF. If profile B is vulnerable to UMF with respect to profile B' , profile B' in turn is vulnerable to DMF with respect to profile B . In this sense, UMF and DMF are the same phenomenon.

A further question is whether a *single* profile B can *simultaneously* be vulnerable to both UMF and DMF. We dub this vulnerability to *Double Monotonicity Failure* (2MF). Proposition 1 and Corollary 2.1 together imply that a three-candidate ballot profile is vulnerable to 2MF under IRV only if it is competitive, critical, and cyclical.⁴ Such a profile must satisfy Conditions 1U and 1D and also

⁴ Cyclicity follows because UMF implies that the plurality loser beats the IRV winner but DMF implies that it cannot be a Condorcet winner. However, Felsenthal and Tideman (2014) show that, given five or more candidates, cyclicity is not required for 2MF.

Conditions 2U and 2D. That such a ballot profile can exist is shown by the following example with $n = 100$, in which Z is the plurality loser and X is the IRV winner:

<u>38</u>	<u>32</u>	<u>30</u>
Y	X	Z
Z	Y	X
X	Z	Y

The profile is vulnerable to UMF because, if 9 of the 38 YZX voters move X to the top of their ballots, Y becomes the plurality loser instead of Z , and X loses to Z in the runoff. At the same time, the profile is vulnerable to DMF because, if 3 of the 38 YZX voters drop Y to second or third preference, Y remains the plurality winner but X becomes the plurality loser, and Y then beats Z in the runoff.

3.4 Monotonicity failure with single-peaked profiles

Suppose that one of the three candidates is a ‘centrist’ with respect to ideology or policy relative to the other two candidates, who in turn are (relatively) ‘extreme’ but in opposite directions (e.g., one to the ‘left’ and the other to the ‘right’ of the centrist candidate). If the centrist candidate is universally viewed as a ‘compromise’ between the two more extreme ones, no voters have the centrist as their third preference. Such a profile, which includes only four of the six possible rankings of three candidates, is called *single-peaked*; since one candidate is never ranked lowest, it can also be called *bottom-restricted*. In such a profile, only the centrist candidate can be a plurality loser that beats the IRV winner, giving the following:

Proposition 3. If a three-candidate ballot profile B is single-peaked, B is vulnerable to UMF under IRV if and only if

- (1) B is competitive, and
- (2) the centrist candidate is the plurality loser.

However, given single-peakedness, the conditions for DMF cannot be simultaneously fulfilled.

Proposition 4. If a three-candidate ballot profile B is single-peaked, B is not vulnerable to DMF (or 2MF) under IRV.

Proposition 2 stipulates that in profile B the plurality loser is Z and the IRV winner is Y . Suppose that Y is extreme: if X is centrist, Y can beat X only if Y is a majority winner, so B cannot be vulnerable to DMF; if Z is centrist, $y_z = y$ and DMF implies that $y < n/2 - y$, which simplifies to $y < n/4$ implying that X is a majority winner and contradicting the stipulation that Y is the IRV winner. Therefore, Y must be the centrist candidate. But this cannot be true either, because single-peakedness then implies that $x_z = 0$, contradicting the requirement for DMF that $x_z > y - z$.

3.5 Monotonicity failure with clone candidates

One important way in which plurality voting is deemed to be problematic is that it is susceptible to *spoiler effects* — that is, even though candidate X would win a ‘straight fight’ against

candidate Y , Y may win if a third candidate Z enters the contest. IRV is sometimes advocated as a voting rule that precludes this problem. Indeed, IRV does preclude what is often dubbed the ‘Nader problem’, referring to the fact that Ralph Nader’s place on the ballot in Florida in 2000 presumably drew more votes from Gore than Bush and therefore cost Gore Florida’s electoral votes (and the presidency). However, under IRV Gore would have won, since the instant runoff would have been equivalent to the straight fight that Nader ‘spoiled’. Moreover, profiles exhibiting the ‘Nader problem’ cannot be vulnerable to monotonicity failure under IRV, since the third candidate receives few votes.

But now consider a three-candidate election in which two candidates C_1 and C_2 have rather similar policy positions or otherwise appeal to the same group of voters, while a third candidate D has a distinctive policy position or otherwise appeals to a different group of voters. This implies that voters are partitioned into two subsets with substantially opposed preferences: those who prefer both C_1 and C_2 to D and those who prefer D to both C_1 and C_2 . C_1 and C_2 , being adjacent in every voter’s ballot ranking, may be called *clones* (Tideman 1987), while D is the *distinctive* candidate whom no one ranks second. Therefore, as with single-peakedness, profiles with clone candidates include only four of the six possible rankings of the three candidates but these preferences are *middle-restricted*.

The case in which D supporters are a large minority, i.e., $n/3 < d < n/2$, is of special interest. If C supporters are sufficiently equally divided between the two clones with respect to their first preferences, D becomes the plurality winner — and would be elected under plurality voting — even though D is also the Condorcet loser. Put otherwise, the clones are ‘mutual spoilers’ in that either clone wins if the other is not a candidate but, if both are candidates, each spoils the other’s chance of election. IRV also resolves the mutual problem to the advantage of the majority of voters favoring the clone candidates, because at least one clone must get into the instant runoff and thereby becomes the IRV winner.⁵ But this advantage of IRV now comes at the cost of vulnerability to UMF.

Proposition 1 restricted to the special case of middle-restricted preferences gives the following:

Proposition 5. If a three-candidate profile B is middle-restricted, B is vulnerable to UMF under IRV if and only if

- (1) B is competitive, and
- (2) the plurality loser is the clone candidate that beats the other clone.

However, given middle-restricted preferences, the conditions for DMF cannot be simultaneously fulfilled.

Proposition 6. If a three-candidate ballot profile B is middle-restricted, B is not vulnerable to DMF (or 2MF) under IRV.

⁵ However, IRV presents another spoiler problem: which clone candidate wins may depend on whether D enters the election.

Proposition 2 stipulates that Z is the plurality loser and Y is the IRV winner. Given a middle-restricted profile without a majority winner, Y must be one of the two clone candidates. If X is the other clone and Z is the distinctive candidate, it follows that $x_z = 0$; but, given that Z is the plurality loser, $y - z > 0$, so Condition 1D(b) cannot hold. If Z is the other clone and X is distinctive, $y_z = y$, so Condition 2D implies $y < n/4$ and that X is a majority winner, contradicting the stipulation that Y is the IRV winner.

Finally, we may note that bottom-restricted and middle-restricted preferences are instances of *value-restricted* preferences (Sen 1966), of which there is a third category, namely *top-restricted* preferences. But if only two candidates are ranked first, one or the other is a majority winner, so no such profiles are vulnerable to monotonicity failure.

4. Monotonicity failure in simulated ballot profiles

As noted earlier, it has been claimed that monotonicity failure, while mathematically possible, is highly unlikely to occur in practice. The preceding theoretical results tell us that three-candidate IRV ballot profiles are vulnerable to monotonicity failure only once specific thresholds of election closeness are crossed. However, they do not directly suggest how frequently such vulnerability occurs in competitive elections. Proposition 1 suggests (but does not prove) that, once they become competitive, vulnerability to UMF increases as elections become more closely contested, though Proposition 2 suggests that the opposite may hold with respect to DMF.

As also noted earlier, Lepelley et al. (1996) provide analytical results for the frequency of monotonicity failure given an *impartial anonymous culture*, in which each anonymous profile (i.e., each specification of the *number* of voters with each possible ballot ranking) is equally likely to occur. They find that, with a large number of voters, 4.51% of such ballot profiles are vulnerable to UMF and 1.97% to DMF. In unpublished work, Smith (2010) reports identical rates based on simulation results. Given an *impartial culture* (in which each distinct profile is equally likely to occur), other simulations by Smith (2010) show that 12.16% of profiles are vulnerable to UMF and 4.95% to DMF. More recently, Plassmann and Tideman (2014) generated a large sample of profiles for three-candidate elections based on ‘a statistical model that simulates voting situations that follow the same distributions as voting situations in actual elections’, and they examined how frequently different voting rules, including IRV, encounter various problems including monotonicity failure. In their data, only about 1% of profiles were vulnerable to monotonicity failure, reflecting the fact that ‘voting situations in actual elections’ typically are not competitive.

More directly relevant are the results of Ornstein and Norman (2014), building on earlier work by Ornstein (2010). Analyzing simulated elections generated by various configurations of voter ideal points in a two-dimensional space with three candidates competing for first-preference support, they find that the resulting profiles are vulnerable to UMF ‘in anywhere from 0.7% to 51% of all cases, and between 15% and 51% of competitive elections’ (Ornstein and Norman 2014, p. 6) with the higher rates in the most competitive profiles. However, they examine UMF only and their analysis focuses primarily on how vulnerability to UMF varies as the simulated candidate competitions proceed.

4.1 Monotonicity failure in random profiles

To supplement the work of Ornstein and Norman and to extend the analysis to DMF (and 2MF), we examine three large and diverse samples, RAN1, RAN2 and RAN3, each with 256,000 randomly generated ballot profiles for about 30 million voters. In each sample, each candidate receives on average about one-third of the first-preference votes, but RAN1 has little dispersion in this respect so almost all profiles are competitive, RAN2 has more dispersion and about two-thirds are competitive, and RAN3 has the most dispersion and fewer than half are competitive.⁶

Table 1 reports, for each sample, the percentage of profiles that are vulnerable to UMF, DMF and 2MF as well as ‘total’ monotonicity failure (TMF = UMF + DMF – 2MF). It also shows similar percentages for competitive, critical, and cyclical profiles, as well as profiles that meet the conditions given in Corollary 2.1. Vulnerability to monotonicity failure varies among the three samples in the expected manner.

However, Figure 1 shows that when we plot vulnerability to monotonicity failure against election closeness measured by the percentage of first-preference support for the plurality loser, all three samples look about the same.⁷ As soon as the logical threshold of 25% is crossed, about 10% to 15% of profiles are vulnerable to UMF, increasing to about 45% in the closest contests. In contrast, vulnerability to DMF does not become evident as soon as the logical threshold of about 17% is crossed but begins to appear when the plurality loser wins about 20% of the vote and increases to about 15% in the closest contests. Once the 25% threshold is crossed, about half of the profiles vulnerable to DMF are also vulnerable to UMF (and thus also to 2MF). When elections are contested most closely, more than 50% of profiles are vulnerable to some kind of monotonicity failure.

In summary, in this diverse — and perhaps most representative — set of simulated profiles, vulnerability to monotonicity failure in three-candidate IRV elections is hardly a rare event, increases as elections become more competitive, and occurs more than half the time in the most closely contested elections.

4.2 Monotonicity failure in single-peaked profiles

We next examine three large samples of single-peaked ballot profiles. Sample SP1 was generated in the same manner as RAN2, except that Y was assigned as the second preference on all

⁶ Specifically, each sample was generated by drawing the number of first preferences for candidate X from a normal distribution with a mean of 10 million and standard deviation of 1.2 or 2.4 or 3.3 million, subject to the constraint that $x \geq 0$, and rounding to the nearest integer. Then the number of such ballots ranking Y second was drawn from a normal distribution with a mean of $x/2$ and a standard deviation of $x/6$, subject to the constraint that $0 \leq x_y \leq x$, with Z ranked second in the remaining $x_z = x - x_y$ ballots. The numbers for the other rankings were determined in like manner.

⁷ This measure of election closeness directly ties in with the definition of competitiveness used in Proposition 1 and elsewhere. The percent of first-preference support for the plurality winner minus that for the plurality loser is an alternative measure, which entails no specific threshold for UMF; in figures based on this measure, UMF and DMF both appear gradually (though DMF appears earlier) as closeness increases, but otherwise they are very similar to those presented here.

ballots with *X* or *Z* ranked first (making *Y* the ‘centrist’ candidate). In the other two samples, preferences are restricted in the same way, but one candidate is on average less popular than the two others, with an average of 6 million first preferences while the other two average 12 million, thereby making fewer profiles competitive. Sample SP2 has a weak centrist candidate *Y* and SP3 has a weak extreme candidate *Z*.

Table 2 shows the rate of vulnerability to UMF for each sample. Perhaps the most striking feature of the table is the very high incidence of vulnerability among competitive profiles with a weak center candidate, which results because the center candidate is very likely to be the plurality loser; but since many fewer profiles are competitive, the overall rate of vulnerability is less than in the symmetric case.

Figure 2 shows the rate of vulnerability to UMF by election closeness in each sample. In the symmetric case, each candidate is equally likely to be the plurality loser, so the centrist candidate has this status in one-third of the profiles regardless of closeness. A weak centrist candidate typically, but not always, is the plurality loser, so almost all competitive profiles are vulnerable to UMF; however, as support for the plurality loser increases, the centrist candidate is less likely to have this status, so vulnerability actually declines. In contrast, when an extreme candidate is weak, the more typical pattern holds.

4.3 Monotonicity failure with clone candidates

We next examine two large samples of ballot profiles with clone candidates. Sample CL1 was generated in the same manner as RAN2, except that no ballots with *X* or *Z* ranked first had *Y* assigned as the second preference (making *Y* the distinctive candidate), whereas in sample CL2 the distinctive candidate has on average the first-preference support of 12 million voters compared to 9 million for each of the clones, making the spoiler problem under plurality voting especially prominent.

Table 3 suggests that a strong distinctive candidate increases the frequency of vulnerability to UMF. Figure 3 refines this conclusion by showing that profiles with a strong distinctive candidate are considerably more likely to be vulnerable to UMF than equally competitive profiles from the symmetric sample. As competitiveness increases, the frequency of vulnerability to monotonicity failure increases in both samples, but more rapidly in the latter, so that the frequencies almost converge at about 30% in the most competitive elections.

4.4 Monotonicity failure in English general elections: 1992-2010

Finally, we examine data derived from constituency-level results from the five U.K. general elections from 1992 through 2010.⁸ Only results from English constituencies are used, because virtually all constituency elections in England during this period were essentially three-party (Labour, Liberal Democrat and Conservative) affairs, while those in Wales, Scotland, and Northern Ireland almost always included significant (and sometimes winning) candidates of ‘nationalist’ parties as well (and in 2015 many constituencies in England included significant UK Independence Party candidates). The

⁸ These data come from Pippa Norris’s Shared Datasets website (<http://www.hks.harvard.edu/fs/pnorris/Data/Data.htm>).

very few English constituencies that did not closely fit a three-candidate pattern were excluded. The five general elections provide a sample of 2642 three-candidate elections, with an average of about 45,000 voters each.

An obvious problem is that these elections were conducted under plurality voting and therefore the results provide us only with (what we take to be) the first preferences of voters, while IRV requires that voters rank the three candidates in order of preference. This problem is addressed by allocating second preferences in each district and each year in proportion to second preferences nationwide, as determined by surveys that provide individual level data about second preferences for each of the five elections.⁹ These surveys indicate that English voter preferences were ‘partially single-peaked’ — that is, most but not all Labour (‘left-of-center’) voters had the Liberal Democrats (the ‘centrist’ party) as their second preference as did most but not all Conservative (‘right-of-center’) voters, while Liberal Democrat voters had more evenly divided second preferences. In addition to this sample (designated ENG1), two other samples with different second (and third) preferences were generated: in ENG2 preferences are strictly single-peaked and in ENG3 second preferences are assigned randomly in the same manner as the RAN samples.

Table 4 is set up in the same manner as Table 1. The most obvious and striking feature of the English data is that considerably fewer profiles (on the order of 1% to 2%) are vulnerable to monotonicity failure than in the wholly simulated data sets. This might suggest that frequencies based on the latter are largely irrelevant and misleading — once we look at (more or less) ‘real’ electoral data, the problem of monotonicity failure under IRV almost disappears, perhaps not to the vanishingly low level claimed by Allard (1996), but to a very low level indeed.

However, very few of these English elections — like the actual elections used to calibrate the statistical model of Plassmann and Tideman (2014) — were competitive three-candidate contests (in part because they were conducted under plurality voting, not IRV). In fact, 60% have a majority winner and 95% are non-competitive. But Table 5 shows that, among the small number of competitive three-candidate English elections, vulnerability to monotonicity failure was quite common.

The underlying similarity between the English and simulated data is evident in Table 5, which shows vulnerability to monotonicity failure in the English data by election closeness in the same manner as the preceding figures. (A table format with collapsed categories of closeness is used because of the small number of cases.) For a given level of closeness, vulnerability to monotonicity failure appears to be approximately as common in the English data as in the simulated data and, in particular, it approaches or exceeds 50% in the most closely contested elections.

5. Concluding remarks

This paper has set out the precise conditions under which vulnerability to monotonicity failure arises in three-candidate IRV elections and has applied them to several large and varied sets of

⁹ The survey estimates for 1992 through 2005 come from Curtice (2009) and for 2010 from Ritchie and Gardini (2012). Respondents who did not express a second preference or expressed a fourth-party second preference were excluded from these calculations.

simulated ballot profiles in order to get a sense of the severity of IRV's monotonicity problems in varying circumstances. Confirming and extending the work of Ornstein and Norman (2014), it demonstrates that vulnerability to monotonicity failure should not be dismissed as a phenomenon that is logically possible but occurs very rarely. In particular, a substantial proportion of competitive IRV ballot profiles — and upwards of 50% of the most closely contested profiles — are vulnerable to monotonicity failure.

A common argument in favor of IRV is that it mitigates the 'wasted vote' psychology that handicaps third (and additional) candidates under ordinary plurality voting, both by discouraging them from entering and discouraging voters from supporting them if they do enter — that is to say, IRV is intended to produce, and probably would produce, fewer essentially two-candidate elections and more relatively closely contested three-candidate (or multi-candidate) elections, exactly the sort most vulnerable to monotonicity failure.¹⁰

However, to demonstrate that vulnerability to monotonicity failure is a relatively common phenomenon does not establish that it is also an especially significant or worrisome phenomenon.¹¹ Having avoided issues pertaining to its significance thus far, I take note of some such considerations in these concluding remarks.

First, the phenomenon itself is often misstated or misunderstood. Advocates of IRV often say that there is little or no evidence that IRV produces 'non-monotonic election results'.¹² But non-monotonicity applies not to any particular 'election result' but to the IRV voting rule itself. Here I have been careful to say that an IRV ballot profile (effectively, an IRV 'election result') may be 'vulnerable to monotonicity failure'. This emphasizes that the problem entails an implicit comparison between companion profiles that are related in a non-monotonic and paradoxical fashion.

Second, I noted above that IRV is advocated in part to avoid the 'wasted vote' psychology under plurality voting that encourages supporters of third-party candidates to vote 'strategically' for the 'lesser of two evils' between the two major-party candidates. However, as discussed by Felsenthal and Tideman (2014), vulnerability to monotonicity failure under IRV provides an incentive for a different kind of strategic voting. This incentive is most evident and always exists given a profile vulnerable to DMF: candidate *X* loses when everyone votes sincerely but some supporters of *X* can cause *X* to win by strategically lowering *X* in their ballot rankings. A similar incentive can exist in a more indirect way given a profile vulnerable to UMF: *X* wins when everyone votes sincerely but *X* loses when some voters who have *Y* as their first preference and *Z* as their second strategically move *X* to the top of their rankings. Of course, all voting rules are vulnerable to strategic voting of one sort

¹⁰ While this analysis has been restricted to the case of three-candidate elections, it also applies to the penultimate stage of a multi-candidate IRV election in which three candidates survive. This suggests that vulnerability to monotonicity failure is no less frequent, and likely more frequent, in multi-candidate elections.

¹¹ However, Felsenthal (2012) characterizes non-monotonicity as an 'especially intolerable' property for a voting rule.

¹² For example, see the quotations in the introductory section.

or other, and the incentives for strategic voting under plurality voting are surely easier to discern than those under IRV.

Third, the non-monotonicity of IRV can also cause problems when some voters change their ‘true’ preferences — in particular, non-monotonicity means that successful persuasive efforts may backfire. Suppose that shortly before an election public opinion (as determined by a poll) corresponds to a profile that projects candidate *X* to be the IRV winner but is also vulnerable to UMF. Suppose further that candidate *X*’s campaign makes a final persuasive effort to secure his victory. This effort may be successful in that it raises candidate *X* in some voters’ rankings (while no other preference change), but at the same time it may be self-defeating in that this increased support leads to *X*’s defeat.

Finally, since monotonicity failure entails a (typically hypothetical) comparison with some companion profile, the phenomenon is for the most part hidden from view. It is thus not surprising that Chief Election Officer Bradley (1995) never saw ‘evidence’ of lack of monotonicity in his many years of experience in supervising elections. However, an election recount could bring a companion profile into explicit view. For example, suppose that the close part of the 2009 Burlington election had been its first-preference component, rather than the instant runoff, and suppose it had been close enough to trigger a recount. Suppose that it is first reported that the Republican candidate got 28% of the first-preference support, the Democrat 27%, and the Progressive 45%, but that the recount produces the following announcement: (1) the initial winner of the election, i.e., the Progressive candidate, was mistakenly denied, and the Republican candidate mistakenly credited with, 2% of the vote; and (2) therefore the Progressive actually lost the election. It would be fair to expect that this announcement would produce considerable confusion and consternation, perhaps coupled with demands for a change in the voting system.¹³

¹³ In fact, after the Burlington election local activists examined the IRV ballots with enough care to discover companion profiles that rendered the original profile vulnerable to monotonicity failure. This provoked enough controversy to lead to the enactment of a different voting system, namely ordinary plurality plus runoff — which has the same problem but hides it better.

Appendix: Proof of Proposition 1

Condition 1 requires that $n/2 - x > y - z$. Substituting $(n - y - z)$ into this expression in place of x , removing parentheses, and simplifying gives

$$n/2 - n + y + z > y - z,$$

which further simplifies to Condition 1U. Thus $z > n/4$ puts Z , rather than Y , into the runoff with X under B' .

Condition 2 requires that Z beat X in the runoff under B' . This implies that Z also beats X under B . Therefore, Condition 2U is necessary for UMF, but it needs to be shown that it is also sufficient.

In the event that $y_x \geq y - z$, all the first-preference ballots that X must gain at Y 's expense to make Y the Plurality Loser under B' can come from the y_x ballots that would in any case transfer to X in a runoff with Z under B , so Z beats X by the same margin under B' as under B . If $y_x < y - z$, it is evidently more difficult for Z to beat X under B' than under B because, to the extent that $y - z$ exceeds y_x , Z loses and X gains $[(y - z) - y_x]$ transferred ballots from Y in the runoff. Therefore, it must be that

$$z + y_z - [(y - z) - y_x] > x + y_x + [(y - z) - y_x].$$

Suppose to the contrary that

$$z + y_z - [(y - z) - y_x] \leq x + y_x + [(y - z) - y_x].$$

Removing parentheses and rearranging terms, we get

$$3z \leq x + y.$$

Substituting $(n - z)$ for $(x + y)$ and further simplifying, we get $z \leq n/4$, contradicting Condition 1U. Thus, given that Condition 1U holds, it follows that, if Z beats X under B , Z also beats X under B' and therefore implies UMF.

References

- Allard, C. (1996). Estimating the probability of monotonicity failure in a U.K. general election. *Voting Matters*, Issue 5, 1-7.
- Amy, D. (2000). *Behind the ballot box: A citizen's guide to voting systems*. Westport, CN: Praeger.
- Berg, S., and Lepelly D. (1963). Note sur le calcul de la probabilité des paradoxes du vote, *Mathématiques, Informatique et Sciences Humaines*. 121, 33-48.
- Bradley, P. (1995). STV and monotonicity: A hands-on assessment. *Representation*, 33/2, 46-47.
- Brams, S. J., and Fishburn, P. (1983). Some logical defects of the Single Transferable Vote. In A. Lijphart and B. Grofman, (Eds.), *Choosing an electoral system*. New York: Praeger, pp. 147-151.
- Curtice, J. (2009). Recent history of second preferences. (http://news.bbc.co.uk/nol/shared/spl/hi/uk_politics/10/alternative_vote/alternative_vote_june_09_notes.pdf).
- Doron, G., and Kronick, R. (1977). Single Transferable Vote: An example of a perverse social choice function. *American Journal of Political Science*, 21/2, 303-311.
- Fair Vote (2009), Monotonicity and IRV — Why the monotonicity criterion is of little import. (<http://archive.fairvote.org/monotonicity/>).
- Farrell, D. M. (2001). *Electoral systems: A comparative introduction*. Hampshire: Palgrave.
- Felsenthal, D. S. (2012). Review of paradoxes afflicting procedures for electing a single candidate. In D. S. Felsenthal and M. Machover (Eds.), *Electoral Systems: Paradoxes, Assumptions, and Procedures*, Berlin: Springer, 2012, pp. 19-91.
- Felsenthal, D. S., and Nurmi, H. (2016). Two types of participation failure under nine voting methods in variable electorates. *Public Choice*, 168/1-2, 115-135.
- Felsenthal, D. S., and Tideman, N. (2013). Varieties of failure of monotonicity and participation under five voting methods. *Theory and Decision*, 75/1, 59-77.
- Felsenthal, D. S., and Tideman, N. (2014). Interacting double monotonicity failure with strategic feasibility under five voting methods. *Mathematical Social Science*, 67, 57-66.
- Fishburn, P., and Brams, S. J. (1983). Paradoxes of preferential voting. *Mathematics Magazine*, 56: 207-214.
- Lepelley, D., Chantreuil, F., and Berg, S. (1996). The likelihood of monotonicity paradoxes in run-off elections. *Mathematical Social Sciences*, 31: 133-146.
- Norman, R. Z. (2010). 'The relationship between monotonicity failure and the no-show paradox. Paper presented at the 2010 Annual Meeting of the Public Choice Society, Monterey, CA, March 11-14, 2010.

- Ornstein, J. (2010). High prevalence of nonmonotonic behavior in simulated 3-candidate STV elections. Paper presented at the Annual Meeting of the Public Choice Society, Monterey, CA, March 11-14, 2010.
- Ornstein, J., and Norman, R. Z. (2014). Frequency of monotonicity failure under Instant Runoff Voting: Estimates based on a spatial model of elections. *Public Choice*, 161/1: 1-9.
- Plassmann, F., and Tideman, T. N. (2014). How frequently do different voting rules encounter voting paradoxes in three-candidate elections? *Social Choice and Welfare*, 42/1: 31-75.
- Poundstone, W. (2008). *Gaming the vote: Why elections aren't fair*. New York: Hill and Wang.
- Riker, W. H. (1982). *Liberalism against populism: A confrontation between the theory of democracy and the theory of social choice*. San Francisco: W.W. Freeman.
- Ritchie, K., and Gardini, A. (2012). Putting paradoxes into perspective — In defence of the Alternative Vote. In D. S. Felsenthal and M. Machover (Eds.), *Electoral Systems: Paradoxes, Assumptions, and Procedures*, Berlin: Springer, 2012, pp. 275-303.
- Sen, A. K. (1966). A possibility theorem on majority decisions. *Econometrica*, 34: 491-499.
- Smith, J. H. (1973). Aggregation of preferences with variable electorate. *Econometrica*, 41: 1027-1041.
- Smith, W. (2010). Three-candidate Instant Runoff Voting: Master list of paradoxes and their probabilities. Center For Range Voting (<http://rangevoting.org/IrvParadoxProbabilities.html>).
- Straffin, P. D., Jr. (1980). *Topics in the theory of voting*. Boston: Birkhauser.
- Tideman, T. N. (1987). Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4: 185-206.

<i>Random Ballot Profiles</i>		% of profiles	UMF	DMF	2MF	TMF
RAN1 [Lower Variability]	All	100.0%	31.2%	9.3%	4.2%	36.3%
	Competitive	98.2%	31.8%	9.4%	4.3%	36.9%
	Critical	37.0%	39.0%	25.1%	11.4%	52.7%
	Cyclical	19.9%	99.2%	21.2%	21.2%	99.2%
	Corollary 2.1	10.9%	44.7%	84.7%	38.4%	91.0%
RAN2 [Medium Variability]	All	100.0%	16.4%	5.5%	2.4%	19.5%
	Competitive	69.0%	23.8%	7.6%	3.4%	28.0%
	Critical	25.6%	28.9%	21.3%	9.2%	41.0%
	Cyclical	12.5%	88.3%	19.4%	18.9%	88.7%
	Corollary 2.1	6.8%	41.6%	81.0%	34.9%	87.7%
RAN3 [Greater Variability]	All	100.0%	10.2%	3.8%	1.5%	12.5%
	Competitive	48.3%	21.2%	7.3%	3.2%	25.4%
	Critical	19.3%	23.9%	19.5%	8.0%	35.4%
	Cyclical	8.6%	81.2%	18.7%	17.9%	82.0%
	Corollary 2.1	4.7%	39.1%	79.6%	32.8%	86.2%

Table 1 Vulnerability to monotonicity failure in random ballot profiles

<i>Single-Peaked Ballot Profiles</i>	SP1 [Symmetric]		SP2 [Weak Center]		SP3 [Weak Extreme]	
	% of profiles	UMF (= TMF)	% of profiles	UMF (= TMF)	% of profiles	UMF (= TMF)
All	100.0%	23.0%	100.0%	10.6%	100.0%	1.2%
Competitive	69.1%	33.3%	13.2%	80.5%	13.2%	9.5%
Critical	40.9%	16.4%	17.7%	13.5%	40.3%	0.5%

Table 2 Vulnerability to monotonicity failure in single-peaked ballot profiles

<i>Ballot Profiles with Clone Candidates</i>	CL1 [Symmetric]		CL2 [Strong Distinctive Candidate]	
	% of profiles	UMF (= TMF)	% of profiles	UMF (= TMF)
All	100.0%	13.4%	100.0%	16.9%
Competitive	69.1%	19.3%	57.6%	29.3%
Critical	57.6%	21.3%	65.5%	21.6%


Table 3 Vulnerability to monotonicity failure in ballot profiles with clone candidates

<i>English Ballot Profiles</i>		% of profiles	UMF	DMF	2MF	TMF
ENG1 [Survey Second Preferences]	All	100.0%	1.4%	0.3%	0.0%	1.7%
	Competitive	4.2%	33.0%	7.1%	0.0%	40.2%
	Critical	8.3%	5.5%	3.7%	0.0%	9.1%
	Cyclical	0.3%	0.0%	0.0%	0.0%	0.0%
	Corollary 2.1	1.2%	0.0%	25.8%	0.0%	25.8%
ENG2 [S-P Second Preferences]	All	100.0%	2.2%	0.0%	0.0%	2.2%
	Competitive	4.2%	50.9%	0.0%	0.0%	50.9%
	Critical	11.8%	3.2%	0.0%	0.0%	3.2%
ENG3 [Random Second Preferences]	All	100.0%	0.4%	0.3%	0.1%	0.6%
	Competitive	4.2%	9.8%	6.3%	2.7%	13.4%
	Critical	4.7%	4.0%	6.5%	2.4%	8.1%
	Cyclical	1.1%	32.1%	10.7%	10.7%	32.1%
	Corollary 2.1	0.3%	33.3%	88.9%	33.3%	88.9%

Table 4 Vulnerability to monotonicity failure in English ballot profiles

<i>English Ballot Profiles</i>		ENG1 [Survey Second Prefs.]				ENG2 [S-P]	ENG3 [Random Second Prefs.]			
<i>Support for PL</i>	<i>Freq.</i>	UMF	DMF	2MF	TMF	UMF	UMF	DMF	2MF	TMF
0-24%	2530	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
25%	35	25.7%	2.9%	0.0%	28.6%	51.4%	5.7%	2.9%	0.0%	8.6%
26%	32	21.9%	6.2%	0.0%	28.1%	53.1%	0.0%	3.1%	0.0%	3.1%
27%	19	42.1%	5.3%	0.0%	47.4%	47.4%	5.3%	10.5%	5.3%	10.5%
28%	12	41.7%	8.3%	0.0%	50.0%	41.7%	8.3%	0.0%	0.0%	8.3%
29-30%	11	63.6%	18.2%	0.0%	81.8%	63.6%	54.5%	9.1%	9.1%	54.5%
31-33%	3	33.3%	33.3%	0.0%	66.7%	33.3%	33.3%	67.7%	33.3%	66.7%

Table 5 Vulnerability to monotonicity failure of English ballot profiles by election closeness



CLOSENESS: PERCENT OF FIRST PREFERENCE SUPPORT FOR PLURALITY LOSER

Note: In all figures, data is aggregated over intervals one percentage point wide and the height of each MF line at X% indicates the rate of MF over the interval from X% to (X+1)%.

Fig. 1 Vulnerability to monotonicity failure of random profiles by election closeness


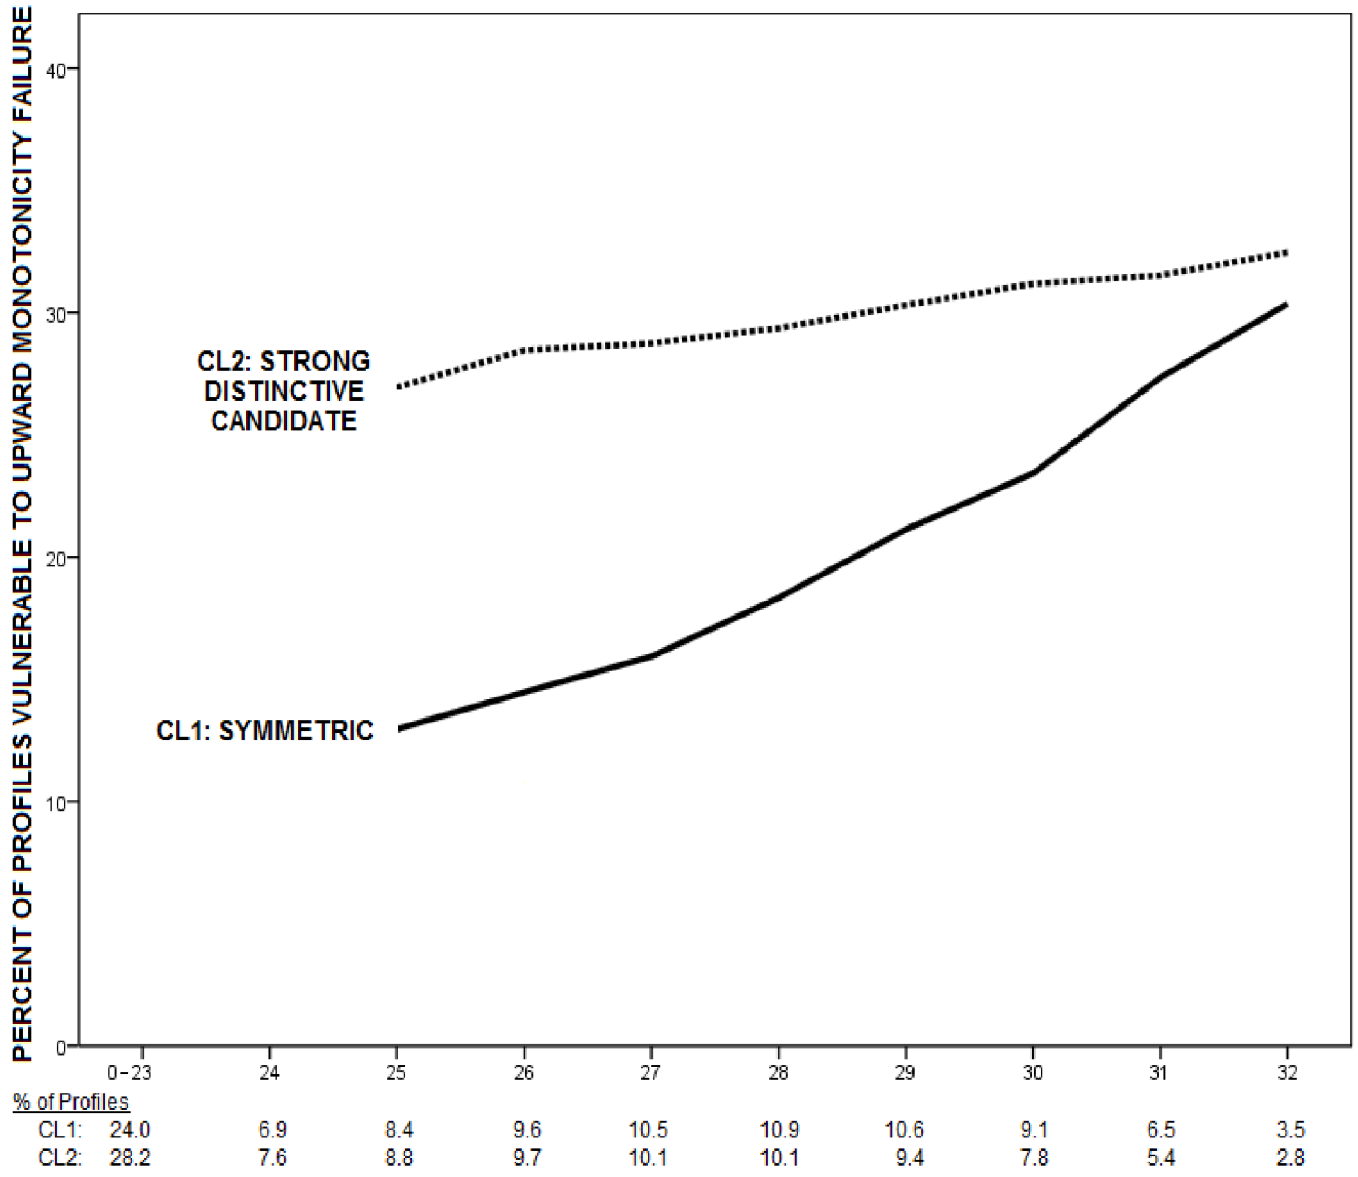


Fig. 2 Vulnerability to upward monotonicity failure of single-peaked profiles by election closenes



CLOSENESS: PERCENT OF FIRST PREFERENCE SUPPORT FOR PLURALITY LOSER

Fig. 3 Vulnerability to upward monotonicity of middle-restricted profiles by election closeness