Many-Candidate Nash Equilibria for Elections Involving Random Selection

by Jeffrey S. Rosenthal, University of Toronto

(Version of July 24, 2017)

Abstract. We consider various voter game theory models which involve some form of random selection, including random tie-breaking, and single elimination, and runoff of the top two candidates. Under certain rules for resolving ties, we prove that with any number of candidates, each such model has a Nash equilibrium in which all candidates attempt to contest the election at the median policy. For models which do not permit ties, we prove that each such model has a Nash equilibrium in which the number of candidates contesting the election is essentially equal to the ratio of the positive payoff for winning divided by the negative of the payoff for losing. All of these model variations thus predict lots of candidates. This result contrasts with Duverger's Law, which asserts that only two (major) candidates will contest the election at all, and which has been confirmed in some other voter game theory models. However, it is consistent with recent primary and leadership and runoff elections where the number of major candidates reached two figures. We close with a simulation study showing that, through repeated elections and averaging and tweaking, candidates' actions will sometimes converge to their predicted equilibrium behaviour.

Keywords: Voter model; Game theory; Nash equilibrium; Multiple candidates; Runoff election; Random selection.

Mathematics Subject Classification (2010): Primary 91A06; Secondary 60E15, 91F10.

1 Introduction

Duverger's Law (Duverger, 1951; Riker, 1982; Schlesinger and Schlesinger, 2006) claims that certain vote systems tend towards outcomes in which only two parties contest elections while all other parties stay out. Empirically, this law is roughly consistent with elections in the United States with just two major parties (Democratic and Republican), though somewhat less so for elections in Canada and the United Kingdom and other countries which have more than two major parties.

By contrast, many real-world party primaries and leadership conventions and run-off style elections have many more candidates. For example, the 2016 U.S. Republican Party presidential primary featured a total of 17 mostly credible candidates (much more than two!). Similarly, an impressive 14 candidates entered the 2017 Conservative Party of Canada leadership contest. And, 11 candidates entered the first round of the 2017 French presidential election.

Mathematically, vote systems are often modelled using game theory, with candidate actions governed by Nash equilibria. It is known (see Section 2 below) that some standard voter models lead to conclusions consistent with Duverger's Law, i.e. with only two significant candidates. By contrast, this paper considers Nash equilibria for different game theory voter models, focusing on the question of which models lead to many different candidates contesting elections. We shall show that for some slightly modified models, many more than two candidates will enter.

Our model modifications below correspond to different ways of determining the election winner from among many candidates, including: multiple candidates at the same exact position are first widowed down to one, either through random selection or through forming alliances; the lowest candidate is eliminated sequentially (as in many party primaries); or the winner is chosen by a runoff of the top two candidates (as for e.g. the French presidency). We consider these model modifications either with (Section 4) or without (Section 7) the common but counter-intuitive (see Section 6) game theory rule that equal candidates can "tie" for winning the election, and each receive a fraction of the (positive) payoff.

With that "ties" rule, we prove (Theorem 1) that for each of our modified models, there is a Nash equilibrium in which *all* of the candidates all contest the election, all coming in at the same median position. Without that rule, we prove (Theorem 2) that for each of our modified models, there is a Nash equilibrium in which the number of candidates contesting the election is essentially equal to the ratio of the positive payoff for winning an election, divided by the negative of the payoff for losing. (More precisely, it is equal to the floor of this ratio plus one, assuming that many potential candidates are available.) Finally, we present (Section 9) a simulation study showing that, through repeated elections and averaging and tweaking, candidates' actions will sometimes converge to their predicted equilibrium behaviour.

In short, our results predict that in the models we study, there will be *many* candidates who contest the election, not just two. These large numbers of candidates are in sharp contrast with Duverger's Law, however they are consistent with the real-world examples discussed above, and they do follow naturally from the models considered herein.

2 Background

Nash equilibria (Nash, 1951) are often used to model strategies of political candidates in elections. In one standard model (see e.g. Hotelling, 1929; Downs, 1957; Osborne, 2003), political positions are considered to be real numbers, and voters' opinions are distributed according to some fixed probability density v on the real line, and each of n different political candidates either stakes out some position $x_i \in \mathbf{R}$, or chooses not to contest the election at all by selecting "OUT". Then, each voter votes for the candidate whose position is closest to their own (dividing their vote equally between multiple equally-close candidates). The payoff for each political candidate is 0 if they do not contest the election, or $+1$ if they receive the most votes, or -1 if they contest the election and receive fewer votes than some other candidate, or $1/k$ if they *tie* for most votes with a total of k different candidates. (For precise definitions, see Section 3 below.)

With just $n = 2$ candidates, the Nash equilibrium for this model is straightforward. Namely, both candidates will contest the election with political position equal to the median value m of v, i.e. $x_1 = x_2 = m$ where m is chosen such that $\int_{-\infty}^{m} v(t)\,dt = 1/2$. In this way, they each tie and thus receive equal expected payoff (in any reasonable model). On the other hand, if either candidate deviates to another position $x_i \neq m$, then (assuming v is positive in a neighbourhood of m) they will receive *fewer* votes than the other candidate, and thus receive payoff of -1 which is less than $1/2$. Or, if either candidate deviates to $x_i = \text{OUT}$, then they will receive payoff of 0 which is again less than $1/2$. So, neither candidate can gain by deviating from the position $x_i = m$. It follows that the choice $x_1 = x_2 = m$ is a Nash equilibrium (in fact *strict*, since any deviation leads to a *lower* payoff).

However, if $n \geq 3$, then the situation is very different. For example, if the first two candidates each choose the median position $x_1 = x_2 = m$, then the third candidate can deviate to a position x_3 which is slightly more than m. Candidate 3 will thus win nearly half of the votes, while Candidates 1 and 2 will divide the remaining votes equally and thus each win just over a quarter of the votes. So, Candidate 3 will receive payoff of $+1$, which is more than the $1/3$ payoff they would receive by coming in at $x_3 = m$. Indeed, it is known (see e.g. Eaton and Lipsey, 1975; Shaked, 1982; Osborne, 1993) that there is no pure (i.e., non-random) Nash equilibrium in the standard model with $n = 3$. By contrast, Hug (1995) has shown the existence of Nash equilibria with three candidates for a different model in which candidates attempt to maximise their votes (without any negative payoff for losing), but the policy they will later enact is randomly distributed about their intended choice, and voters attempt to minimise a quadratic loss function based on this random distribution. And, Rosenthal (2016b) presents a modification in which candidates' positions are perceived with uncertainty, which admits a three-candidate Nash equilibrium with $x_1 = x_2 = x_3 = m$.

Another modification (Osborne, 1996) has the candidates choose their positions *sequentially*, so that candidate i can base their position on the already-chosen positions of candidates $1, 2, \ldots, i-1$. For this model, it is conjectured that just the first and last candidates will contest the election, each with position equal to the median m, while the other $n-2$ candidates will all stay out. Osborne's conjecture has been proven for $n = 2$ and $n = 3$ and $n = 4$, and verified by simulation for n as large as 7 (Rosenthal, 2016a). However, it has been disputed when $n = 12$ (de Vries, 2015; de Vries et al., 2016), though even there it still appears that exactly two candidates – the second and twelfth in this case – will enter the election at position m while the other $n-2$ candidates choose OUT, still consistent with Duverger's Law.

These various results indicate that Duverger's Law holds in the standard model with $n = 2$ as discussed above, and also in the context of Osborne's sequential conjecture (as far as it has been verified), though it does not hold in some other models. In addition, in the context of runoff-style elections, Bouton and Gratton (2015) consider a different model in the case of only $n = 3$ candidates, and argue that for their model there exist "Duverger's law equilibria" in which only two candidates obtain nonzero expected vote share, again essentially consistent with Duverger's Law.

In the present paper, we consider a number of simple modifications of the standard model, each of which admits many candidates contesting an election. This is in strong contrast to Duverger's Law, but is consistent with empirical evidence from primaries and runoff elections as discussed above, to it might be more useful for modelling such contests.

3 The Standard Voting Model

We next define a precise standard game theory model for such voting. We assume throughout that there is some voter probability density v on $\mathbf{R}$, with median value m, such that v positive in some neighbourhood of m. (For example, perhaps v is the density of the Uniform[0,1] density, with median $m = 1/2$.) We further assume there is some positive number $n \in \mathbf{N}$ of candidates. Each candidate simultaneously chooses a policy position $x_i \in \mathbf{R}$ on which to contest the election, or else chooses $x_i = \text{OUT}$.

Each candidate who chooses OUT does not contest the election at all, and receives a payoff of 0. Each candidate who does not choose OUT contests the election, with candidate i receiving the vote share w_i given by

$$w_i \;=\; \frac{\int_{t \in R_i} v(t)\, dt}{\#\{j : x_j = x_i\}}\,. \tag{1}$$

Here $R_i = \{t \in \mathbf{R} : |t - x_i| \le |t - x_j| \text{ for all } 1 \le j \le n\}$ is the vote region won (or tied for winning) by Candidate i, i.e. the set of policy positions which are as close to x_i as to any other candidate position x_j. And, $\#\{j : x_j = x_i\}$ is the number of candidates who chose precisely the same position as Candidate i.

Then, if one candidate wins the election outright (i.e., obtains the strictly largest vote share), then they receive a payoff of $+1$. If $k \ge 2$ candidates all *tie* for the largest vote share, then they each receive a payoff of $1/k$. If a candidate contests and loses the election, then they receive a payoff of -1.

In this paper, we consider some modifications of this model. We divide our models into two classes. The first, like the above standard model, allows for the possibility that multiple candidates will *tie* for winning the election and share the positive payoff.

4 Ties Case: Models

We next define several different but related voter models, all of which (like the standard model) allow for the possibility of ties. In all cases, we assume there are a total of n candidates, that each candidate can choose a position $x_i \in \mathbf{R}$ or OUT, and that there is a voter density function v which is positive in a neighbourhood of its median value m.

4.1 Model #1: Random Selection

In our first model, if multiple candidates all choose the exact same position x_i, then one of those candidates is selected uniformly at random to contest the election; all other candidates having the same value of x_i are removed from the election and receive a payoff of 0.

Once those tying candidates have been removed, then the remaining candidates have their win regions R_i and the vote shares w_i computed as usual as in equation (1). Then, as usual, if k candidates tie for the highest vote share then they each receive a payoff of $1/k$, while any candidates who remain in the election but obtain less than the highest vote share receive a payoff of -1. (And once again, any candidate who chooses OUT receives a payoff of 0.)

In summary, this game proceeds exactly like the usual model, except that when multiple candidates choose the same position x_i then duplicate entries are removed from the election at random. This model thus breaks ties from equal positions of multiple candidates by an arbitrary random selection among the multiple candidates.

4.2 Model #2: Combined Alliances

Our second model is similar to our first, but with one difference. This time, if $r \geq 2$ candidates all choose the exact same position x_i, then those r candidates are all *combined* into a single alliance. That is, they are all replaced by a single candidate. The election then proceeds as usual with this single candidate, according to the usual formula (1). Once the election is concluded, each of the r original candidates receive $1/r$ of whatever payoff their combined single candidate obtains.

4.3 Model #3: Two-Round Runoff

Our third model is as follows. As usual, each candidate who chooses OUT receives a payoff of 0. Then, the vote share w_i of each candidate who is not OUT is computed as usual as in (1). If all of the w_i are equal, then all k of the remaining candidates tie for winning the election, and each receive payoff of $1/k$. Otherwise, the two candidates with the highest vote shares (selected uniformly randomly if necessary) are advanced to the runoff election, while all the other candidates who are not OUT are eliminated and receive a payoff of -1. The two remaining candidates' vote shares are re-computed again as in (1) (but now disregarding all but the two remaining candidates), and the candidate with the lowest vote share is again eliminated (and also receives a payoff of -1). The one remaining candidate is declared the winner (and thus receives payoff of $+1$).

In practice, this two-round runoff system is used in many places, for example to elect the presidents of France and of Argentina.

4.4 Model #4: Sequential Elimination

This model proceeds by repeatedly removing the candidate with the fewest votes. That is, each candidate who chooses OUT receives a payoff of 0 as usual. Then, at the first step, the vote share w_i of each candidate who is not OUT is computed as usual as in (1). If all of the remaining positions x_i are all equal, then all k of the remaining candidates tie for winning the election, and each receive payoff of $1/k$. Otherwise, whichever candidate has the lowest value of w_i is eliminated from the election, and thus loses and receives a payoff of -1. (Note: if some number $r \geq 2$ of candidates all *tie* for having the lowest vote share, then one of those r tying candidates is selected *uniformly at random* for elimination; the other $r-1$ tying candidates are *not* eliminated, and proceed to the next step of the elimination.)

The remaining candidates' vote shares are then re-computed again as in (1) (but now disregarding all candidates who have already been eliminated), and the candidate with the lowest vote share is again eliminated. This process is repeated until either only one candidate remains so that candidate is declared the winner (and thus receives payoff of $+1$), or some number k of candidates remain and all have the same exact position x_i, in which case they all tie for winning (and each receive payoff of $1/k$).

In practice, this sequential elimination system can be accomplished either through a single preferential ballot (sometimes called "instant-runoff"), or through multiple rounds of voting. It is in fact used in many cases around the world, e.g. for Australian parliamentary elections and for some party leadership contests in Canada and in the U.K. It also "essentially" applies to U.S. presidential primary contests, since candidates there who perform poorly in early state contests tend to drop out sequentially as the long primary process continues.

5 Ties Case: Main Result

We now show the existence of a Nash equilibrium for each of the above models, for any number n of candidates.

Theorem 1 *For each of the above Models #1 and #2 and #3 and #4, for any number $n \in \mathbf{N}$ of candidates, the set of positions $x_1 = x_2 = \ldots = x_n = m$ is a Nash equilibrium, i.e.*

with this choice no one candidate can increase their expected payoff by changing their action while the other $n-1$ *actions remain fixed.*

Proof. If $n = 1$ the statement is trivial, and if $n = 2$ then it follows as in the standard model since neither candidate can improve their expected payoff to more than 1/2 while the other candidate still comes in at m. Thus, we assume that $n \geq 3$.

Suppose $x_1 = x_2 = \ldots = x_n = m$. Our models have all been constructed so that in this case, each candidate receives expected payoff $1/n$. We need to show that no one candidate can increase their expected payoff above $1/n$. Now, if a candidate switches to OUT then their expected payoff becomes 0, which is less than $1/n$, so they do *not* benefit by switching. It remains to consider the case where one candidate, say candidate 1, changes to some other position $x_1 \neq m$, while the remaining candidates' positions remain at m. We do this separately for each of the models.

We begin with Model #1. Since all other $n-1$ candidates still have $x_i = m$, therefore $n-2$ of them (chosen at random) will be eliminated from the election. We will be left with one single candidate i at position $x_i = m$, plus candidate 1 at position $x_1 \neq m$. In that case, since m is the median value of v, it follows (using the positivity of v near m) that Candidate i will obtain a strictly greater vote share than Candidate 1. Thus, Candidate 1 will receive a payoff of -1, which is less than $1/n$, so that Candidate 1 does *not* benefit by switching.

The argument for Model #2 is similar. Indeed, there, if $x_1 \neq m$ while $x_2 = x_3 = \ldots = x_n = m$, then Candidates 2 through n will be combined into a single alliance, which will then defeat Candidate 1 as above. So, Candidate 1 will again receive a payoff of -1, which is again less than $1/n$.

For Model #3, several cases must be considered. If Candidate 1 moves far enough away, then they will have fewer votes than the other $n-1$ candidates. In this case, they will not make the runoff at all. If they move exactly the right distance away, then they will *tie* with the other $n-1$ candidates, and thus have probability $2/n$ of making the runoff. Or, if they move less far away, then they will definitely make the runoff. But in any case, if they do not make the runoff then they lose and receive payoff -1, while if they do make the runoff, then as above they will receive fewer vote shares than the other runoff candidate (who will be at position m) and thus will still lose and still receive payoff -1. In short, if Candidate 1 moves to any position $x_1 \neq m$, then they will eventually lose the election (either immediately or in the runoff) and thus receive expected payoff of -1, so in any case they will not benefit by switching.

For Model #4, depending on how far away Candidate 1 moves, they might be eliminated earlier or later in the sequential elimination. However, even if they survive the first $n-2$ elimination steps, at the final step they will be competing against a single candidate at position m, and at that point they will lose as above. So, in any case, they will end up losing the election, and will thus again receive expected payoff of -1 which does not benefit them. (Furthermore, until Candidate 1 is eliminated, there will always be at least one candidate at position $x_1 \neq m$, and at least one candidate at position m, so the rule about "if all of the remaining positions x_i are all equal" will *not* be invoked. For more about this issue, see Section 6 below.)

In summary, for each of the four models, if $x_1 = x_2 = \ldots = x_n = m$, then no candidate can improve their expected payoff by switching to another action, i.e. by going OUT or by choosing a different value x_i. Thus, for any n, the set of actions $x_1 = x_2 = \ldots = x_n = m$ is a Nash equilibrium for this model, as claimed. ∎

6 Ties: Discussion

So far, we have only considered models which, like the standard model, allow for the possibility of ties. However, this may be unrealistic, since most real elections require that a unique winner be chosen, and indeed have some prescribed mechanism (such as a coin toss, or a new election) for selecting among tied candidates.

Furthermore, the rules for resolving ties have to be specified in just the right way for Theorem 1 to hold. This is best illustrated in Model #4 (sequential elimination). There, it was specified that if all k remaining candidates all have the exact same position x_i then they each receive payoff $1/k$. It might be more natural to say that whenever all k remaining candidates all simply have the same vote share w_i then they each receive payoff $1/k$, but this would actually lead to a different result!

Specifically, consider Model #4 with this slight modification, with $n \geq 4$, and with $x_1 = x_2 = \ldots = x_n$. Then we claim that Candidate 1 can increase their payoff from $1/n$

to 1/3 by switching from $x_1 = m$ to $x_1 = z$, where $z < m$ is chosen precisely so that $\int_{-\infty}^{(z+m)/2} v(t)\, dt = 1/3$. Indeed, in that case, $n-3$ other candidates at position $x_i = m$ will first be eliminated, at which point there will be 3 candidates left, with positions (z, m, m). At this point, according to the usual formula (1), each of the three candidates will receive vote share $w_1 = w_2 = w_3 = 1/3$. So, with the slight modification, at that point a three-way tie would be declared, and each candidate would thus receive payoff 1/3, which is better than the original $1/n \leq 1/4$.

Similarly, it was necessary to specify in Model #1 (random selection) that additional candidates at the same position are "removed" from the election and receive payoff 0. It might be more realistic to assume that such candidates actually *lose* the election and instead receive payoff -1. But again, this would change the results, since if e.g. $n = 3$ with $x_1 = x_2 = x_3$, then under this modification each candidate would receive expected payoff $(1/3)(+1) + (2/3)(-1) = -1/3 < 0$ and would thus prefer to change to OUT.

Such considerations led us to re-consider the "ties" rules entirely. We eventually decided that the most realistic models should produce a *single* winner, with no ties. To achieve this, any apparent ties should be broken randomly and uniformly. In practice this might correspond to an actual coin toss or other random resolution. Alternatively, it might correspond to the realisation that even if candidates are exactly tied in terms of their electoral support, certain minor randomness – such as a few citizens getting flat tires and failing to vote – would usually be sufficient to break the tie somehow.

When considering models without ties, a major difference arises even in the case where $x_1 = x_2 = \ldots = x_n = m$. In the standard model, and in each of the previous modifications, this set of actions gives a payoff of $1/n$ to each candidate. However, without ties, each candidate would then have probability $1/n$ of winning the election and receiving payoff of $+1$, and probability $1 - (1/n)$ of losing the election and receiving payoff of -1. This gives each candidate an overall expected payoff of $(1/n) - (1 - (1/n)) = (2/n) - 1$. In particular, for $n > 2$, this expected payoff is *negative*, so it is now better for each candidate to stay OUT rather than to contest the election even when every candidate chooses the same exact position.

To address this issue, we next consider some models in which the candidate who wins the election receives a payoff of α, for some constant $\alpha > 0$, not necessarily $+1$. The payoff for losing remains -1, so α corresponds to the *ratio* of how good it is to win, divided by how bad it is to lose. It seems reasonable that α could be quite large, since winning an election is much more significant and impactful than merely contesting and losing it. In the standard model, the precise value of α is completely irrelevant as long as it is positive, since randomness is never used to determine winners and losers. In our models with ties, the value of α is again mostly irrelevant, since only in certain limited cases does a candidate have positive probabilities of both winning and losing. However, in models *without* ties, the value of α is crucial, as we shall see. Roughly, in the above scenario, each candidate now has expected payoff $(1/n)(\alpha) - (1 - (1/n)) = (\alpha + 1 - n)/n$, which is non-negative whenever $n \leq \alpha + 1$. This suggests that the number of candidates who contest the election should be governed by the value of α, which is indeed the case (as we shall see in Theorem 2 below).

7 No-Ties Case: Models

We next define several different voter models, similar to the previous models, but now not allowing for the possibility of ties. We again assume there are a total of n candidates, that each candidate can choose a position $x_i \in \mathbf{R}$ or OUT, and that there is a voter density function v which is positive in a neighbourhood of its median value m. We also assume a fixed constant $\alpha > 0$ corresponding to the payoff for winning the election.

7.1 Model #1*: Random Selection Without Ties

In our first model, if multiple candidates all choose the exact same position x_i, then one of those candidates is selected uniformly at random to contest the election. All other candidates having the same position x_i lose the election and receive a payoff of -1.

Once those other candidates have been removed, then the remaining candidates have their win regions R_i and the vote shares w_i computed as in equation (1). Then, if k candidates tie for the highest vote share then one of them is selected uniformly at random and wins the election and receives payoff α, while the other $k - 1$ tying candidates, as well as any other candidates who remain in the election but obtain less than the highest vote share, lose and receive a payoff of -1. (As usual, any candidate who chooses OUT receives a payoff of 0.)

7.2 Model #2*: Combined Alliances Without Ties

In this model, if $r \geq 2$ candidates all choose the exact same position x_i, then those r candidates are all *combined* into a single alliance candidate. The election then proceeds with the alliance candidates and all candidates with unique positions. Vote shares for those candidates are determined according to (1). As above, if k candidates tie for the highest vote share then one of them is selected uniformly at random and wins the election and receives payoff α, while the other $k-1$ tying candidates, as well as any other candidates who remain in the election but obtain less than the highest vote share, lose and receive a payoff of -1. (And, as usual, any candidate who chooses OUT receives a payoff of 0.)

Once the election is concluded, each of the r original candidates who were combined into a single alliance each receive $1/r$ of whatever payoff their combined single alliance candidate obtained.

7.3 Model #3*: Two-Round Runoff Without Ties

This model proceeds as follows. Each candidate who chooses OUT receives a payoff of 0. The vote share w_i of each candidate who is not OUT is computed as in (1). Among the candidates who have or are tied for the highest or second-highest (counted with repetition) vote share, two are selected uniformly at random to advance to the runoff election, while all the other candidates who are not OUT are eliminated and receive a payoff of -1. The two remaining candidates' vote shares are re-computed again as in (1) (but now disregarding all but the two remaining candidates), and the candidate with the lowest vote share is again eliminated (and also receives a payoff of -1). The one remaining candidate is the winner and receives payoff α.

7.4 Model #4*: Sequential Elimination Without Ties

In this model, once again, each candidate who chooses OUT receives a payoff of 0. Then, at the first step, the vote share w_i of each candidate who is not OUT is computed as in (1). Whichever candidate has the lowest value of w_i is eliminated from the election (with ties broken by a uniform random selection), and thus loses and receives a payoff of -1. The remaining candidates' vote shares are re-computed again as in (1) (but now disregarding all candidates who have already been eliminated), and the candidate with the lowest vote share is again eliminated and receives payoff -1. This process is repeated until only one candidate remains, at which point that candidate is declared the winner and receives payoff α.

8 No-Ties Case: Main Result

We now show the existence of a Nash equilibrium for each of the above models, for any value of n. (Below, $\lfloor \alpha \rfloor$ is the *floor* of α, i.e. the greatest integer not exceeding α.)

Theorem 2 *For each of the above Models #1* and #2* and #3* and #4*, for any number $n \in \mathbf{N}$ of candidates, and for any subset $S \subseteq \{1, 2, \ldots, n\}$ with $|S| = \min(n, \lfloor \alpha \rfloor + 1)$, the set of actions in which all candidates i for $i \in S$ contest the election at position $x_i = m$, and all candidates i for $i \notin S$ stay OUT, is a Nash equilibrium, i.e. with this choice no one candidate can increase their expected payoff by changing their action while the other $n-1$ actions remain fixed. Furthermore, if $n \geq 2$ and $\alpha > 1$ and α is not an integer, then this set of actions is a* strict *Nash equilibrium, i.e. any one candidate's change of action strictly* decreases *their expected payoff.*

To prove Theorem 2, we begin with a lemma.

Lemma 1 *For each of the above Models #1* and #2* and #3* and #4*, for any number $n \geq 2$ of candidates, if there are r candidates at position m where $1 \leq r \leq n-1$, and one candidate j at a different position $x_j \neq m$, and the remaining $n-r-1 \geq 0$ candidates are all OUT, then Candidate j will lose the election.*

Proof. We begin with Model #1. Since r candidates are at position m, therefore $r-1$ of them (chosen at random) will be eliminated from the election. We will be left with one single candidate i at position $x_i = m$, plus candidate j at position $x_1 \neq m$. In that case, it follows (as in the proof of Theorem 1) that Candidate i will obtain a strictly greater vote share than Candidate j, so Candidate j will lose the election.

Model #2 is similar. In this case, all the candidates at position m will be combined into a single alliance at position m, which will then defeat Candidate j as above, so Candidate j will again lose.

For Model #3, as before, if Candidate 1 moves far enough away, then they will have fewer votes than the other $n-1$ candidates and lose immediately, or if they move exactly the right distance away then they will *tie* with certain other candidates and thus have a certain probability of making the runoff, or if they move less far away then they will have the largest vote share and thus definitely make the runoff. But even if they do make the runoff, then in the runoff they will compete with a single other candidate at position m, so as above they will receive fewer vote shares and thus still lose and still receive payoff -1.

For Model #4, the argument is very similar to the corresponding proof of Theorem 1. Again, depending on how far x_j is from m, Candidate j might be eliminated earlier or later in the sequential elimination. But even if they survive until only two candidates remain, nevertheless at the final step they will be competing against a single candidate at position m, and at that point they will lose the election and receive -1 just as above. ■

Proof of Theorem 2. If $n = 1$, then for any $\alpha > 0$, $\min(n, \lfloor \alpha \rfloor + 1) = 1$ so we must have $S = \{1\}$, i.e. Candidate 1 comes in at m and receives payoff $\alpha > 0$. This is clearly a Nash equilibrium (though not strict), since if they deviate to OUT then they reduce their payoff to 0, while if they deviate to some other position $x_1 \in \mathbf{R}$ then they still win the election and still receive the same payoff α.

If $n \geq 2$ and $0 < \alpha < 1$, then again $\min(n, \lfloor \alpha \rfloor + 1) = 1$, so just one of the candidates comes in at m. Suppose for definiteness that $x_1 = m$ while $x_2 = \ldots = x_n =$ OUT. Then Candidate 1 wins the election and receives payoff of $\alpha > 0$, while each Candidate i for $i \geq 2$ stays out and receives payoff of 0. If Candidate 1 deviates to some other value $x_1 \in \mathbf{R}$, then they still win the election and still receive the same payoff α, while if they deviate to OUT then they receive the smaller payoff 0. If Candidate i for some $i \geq 2$ deviates to $x_2 = m$, then they have probability 1/2 of winning the election, so their expected payoff is $(1/2)(\alpha) + (1/2)(-1) = (\alpha - 1)/2 < 0$ which is strictly worse than staying out. Or, if Candidate i for some $i \geq 2$ deviates to some other value $x_i \in \mathbf{R} \setminus \{m\}$, then they will lose the election to Candidate 1 by Lemma 1, so their payoff will be $-1 < 0$. So, this is indeed a Nash equilibrium (though not strict).

Finally, we consider the case where $n \geq 2$ and $\alpha \geq 1$. In this case, the claimed equilibrium has some number $\min(n, \lfloor \alpha \rfloor + 1) \geq 2$ of candidates at position m, while the remaining $n - \min(n, \lfloor \alpha \rfloor + 1)$ candidates stay OUT. We have to show that any one candidate's change of action does not increase their expected payout, and strictly decreases it if α is not an integer.

Consider first a candidate at position m. They have probability $1/\min(n, \lfloor\alpha\rfloor + 1)$ of winning the election, or $1 - (1/\min(n, \lfloor\alpha\rfloor + 1))$ of losing, and thus initial expected payoff

$$\left(\frac{1}{\min(n, \lfloor\alpha\rfloor + 1)}\right)(\alpha) + \left(1 - \frac{1}{\min(n, \lfloor\alpha\rfloor + 1)}\right)(-1) \ = \ \frac{\alpha + 1 - \min(n, \lfloor\alpha\rfloor + 1)}{\min(n, \lfloor\alpha\rfloor + 1)}$$
$$\geq \ \frac{\alpha + 1 - (\lfloor\alpha\rfloor + 1)}{\min(n, \lfloor\alpha\rfloor + 1)} \ = \ \frac{\alpha - \lfloor\alpha\rfloor}{\min(n, \lfloor\alpha\rfloor + 1)},$$

which is ≥ 0, and strictly > 0 if α is not an integer. If they deviate to OUT, then their payoff becomes 0, which is not more than their initial expected payoff (and strictly less if α is not an integer). If instead they deviate to some other position $\neq m$, then they will lose the election by Lemma 1, and thus receive payoff -1, which is strictly less than their initial expected payoff.

If $\min(n, \lfloor\alpha\rfloor + 1) < n$, i.e. $\lfloor\alpha\rfloor + 1 \leq n - 1$, then there are also some candidates who are initially OUT, with payoff 0. Then, if one such candidate deviates to coming in at a position $\neq m$, then by the Lemma 1 they will lose the election and receive payoff -1. If instead one such a candidate deviates to coming in at position m, then they will be one of $\min(n, \lfloor\alpha\rfloor + 1) + 1 = \lfloor\alpha\rfloor + 2$ candidates at m. They will then have probability $1/(\lfloor\alpha\rfloor + 2)$ of winning the election, and $(\lfloor\alpha\rfloor + 1)/(\lfloor\alpha\rfloor + 2)$ of losing, so their expected payoff will be

$$\left(\frac{1}{\lfloor\alpha\rfloor + 2}\right)(\alpha) + \left(\frac{\lfloor\alpha\rfloor + 1}{\lfloor\alpha\rfloor + 2}\right)(-1) \ = \ \frac{\alpha - \lfloor\alpha\rfloor - 1}{\lfloor\alpha\rfloor + 2}$$

which is < 0 since we always have $\alpha - \lfloor\alpha\rfloor < 1$. Hence, in either case, their expected payoff will strictly decrease from 0 to a negative value.

In summary, for each of the four models, if $x_i = m$ for $i \in S$ while $x_i = \text{OUT}$ for $i \notin S$ where $|S| = \min(n, \lfloor\alpha\rfloor + 1)$, then if any one candidate switches to another action then their expected payoff will not increase (and will strictly decrease under the specified conditions). This proves the result. ∎

9 Finding Equilibria Through Repeated Tweaks

The previous results all concern equilibrium behaviour, and state that *if* candidates are behaving a certain way, then they have no incentive to change. However, this does not address the question of how the candidates might *find* such equilibrium behaviour in the first place.

Specifically, we imagine that elections will be repeated many times, and the candidates see the results each time. Occasionally a candidate will try *changing* their action, and see if that seems to improve their expected payoff. If it does, then they will stick with the new action; if not, then they will revert to their previous action. If this updating is repeated many times, will the candidates converge to equilibrium behaviour as in the above theorems?

9.1 Computer Program

To test this, we wrote a computer program in C (available at probability.ca/runoff.c) to simulate this repeated candidate updating algorithm.

For simplicity, we stick to our Model 4* above (sequential elimination without ties), which we consider to be our most interesting model, though similar simulations could be performed for our other models too. For definiteness, we let the voter distribution v be the Uniform[0,1] density throughout. To create a discrete model for easier simulation, we restrict all of our candidate actions to either OUT or one of the eleven fixed positions 0.0, 0.1, 0.2, ..., 1.0.

The program then proceeds by repeatedly selecting a candidate at random, having them try a random new action for some number av of elections, computing (random) payoffs each time according to the rules for Model 4*, and then averaging all of the resulting payoffs together. (Thus, the candidate does not know their true expected payoff value, but merely a random unbiased estimate of it. This is somewhat similar to "pseudo-marginal MCMC"; see Andrieu and Roberts, 2009.) The selected candidate then checks to see if their average payoff appears to have increased as a result of their action tweak. If it does, then that candidate switches to the new action, otherwise they revert to their previous action. This updating is repeated a total of $reps$ times. The corresponding algorithm is summarised as Algorithm 1:

Algorithm 1 Algorithm to Search for Equilibria (given n, α, $reps$, and av)

```
for i = 1, 2, ..., n do
    x_i ← random action
end for
curpayoffvec ← AVRPAYOFFVEC(x_1, x_2, ..., x_n).
sumvec ← (0, 0, ..., 0)
numcounted ← 0
for r = 1, 2, ..., reps do
    Choose γ ∈ {1, 2, ..., n} uniformly at random.
    x'_γ ← random action
    newpayoffvec ← AVRPAYOFFVEC(x_1, x_2, ..., x_{γ-1}, x'_γ, x_{γ+1}, ..., x_n).
    if newpayoffvec[γ] > curpayoffvec[γ] then
        x_γ ← x'_γ.
        curpayoffvec ← newpayoffvec
    end if
    if r > burnin then
        sumvec = sumvec + curpayoffvec
        numcounted = numcounted + 1
    end if
end for
return (sumvec / numcounted)

procedure AVRPAYOFFVEC(z_1, z_2, ..., z_n)
    for a = 1, 2, ..., av do
        Assume the candidate actions are given by {z_i}.
        while number of candidates competing in the election > 1 do
            Compute all candidate payoffs as in (1).
            Select one candidate from among those with the lowest vote share.
            Eliminate that one candidate from the election.
        end while
        Assign payoff α to the remaining candidate, −1 to all eliminated, and 0 to all OUT.
    end for
    return The average of the av different payoff vectors for the n candidates.
end procedure
```

9.2 Results

We next compare the average results from our algorithm simulations to the theoretical equilibria values, for a variety of parameter settings. The algorithm generally does a fairly good job of finding the equilibrium behaviour. In particular, it usually (though not always) relegates all candidate actions to either 0.5 (i.e. the median value of v) or to OUT. In some cases, it finds approximately the correct number of candidates to come in at 0.5, while in other cases it tends to have too many candidates stay OUT.

Below, we graph the comparison between the theoretical average counts (blue, left) and the simulation average counts (red, right) for each possible action, for different parameter choices. (We set the value of *burnin* to be *reps*/2 throughout.)

We begin with the case of $n = 5$ candidates, and win payoff $2 < \alpha < 3$. In that case, Theorem 2 predicts that $\min(n, \lfloor\alpha\rfloor + 1) = 3$ of the 5 candidates will enter at 0.5, with the other 2 candidates staying OUT. With $reps = 1{,}000$, and $\alpha = 2.5$, and with the number of averages av equal to just 10, the algorithm already does *fairly* well at finding this equilibrium, though it slightly underestimates the average number of candidates at 0.5 and correspondingly overestimates the number who stay OUT. By contrast, increasing α to 2.9 (which makes no difference to Theorem 2 since $\lfloor\alpha\rfloor$ is unchanged, but makes it easier for the algorithm to detect when coming in at 0.5 is slightly better on average than staying OUT), and increasing av to 100, the algorithm performs almost identically to the theory:

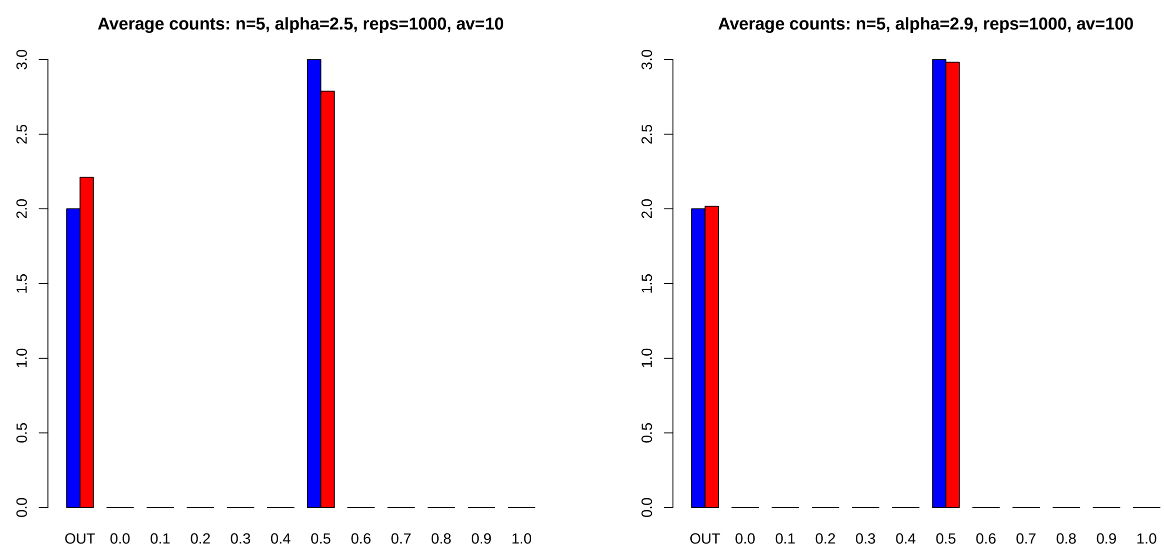

We next consider the case where there are $n = 15$ candidates, with win payoff $\alpha = 7.9$. In this case, Theorem 2 predicts that $\min(n, \lfloor\alpha\rfloor + 1) = 8$ of the 15 candidates will enter at 0.5, with the other 7 candidates staying OUT. Now, if the *reps* and *av* settings are too small, then the results do not match the theory too well: too many candidates staying OUT, and (in some cases) some candidates come it at other positions like 0.4 and 0.6 instead of 0.5:

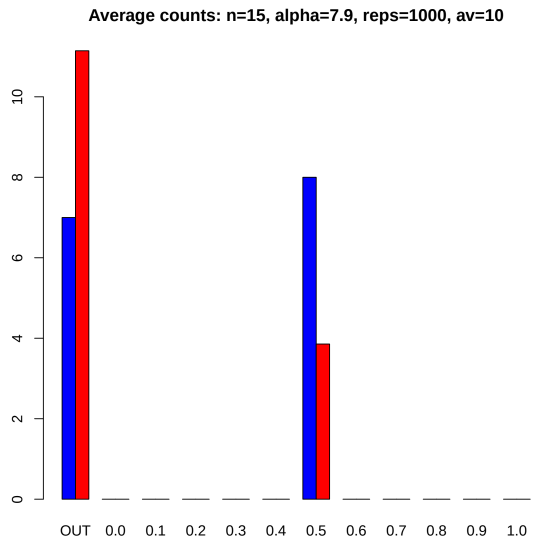

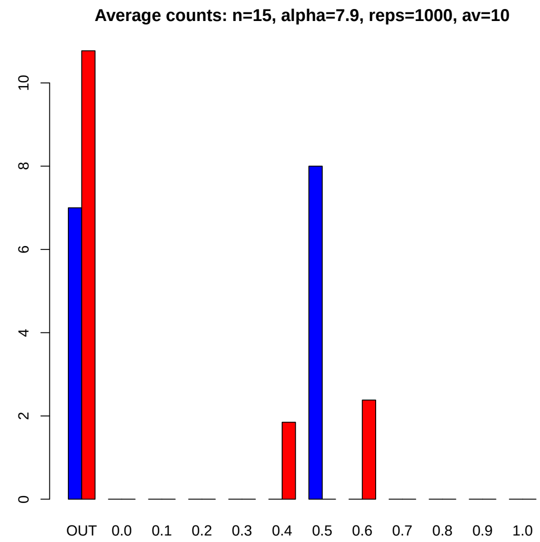

However, if the value of av is sufficiently large, like 1,000 or even 10,000, then the simulation results tend to be much closer to the theoretical predictions:

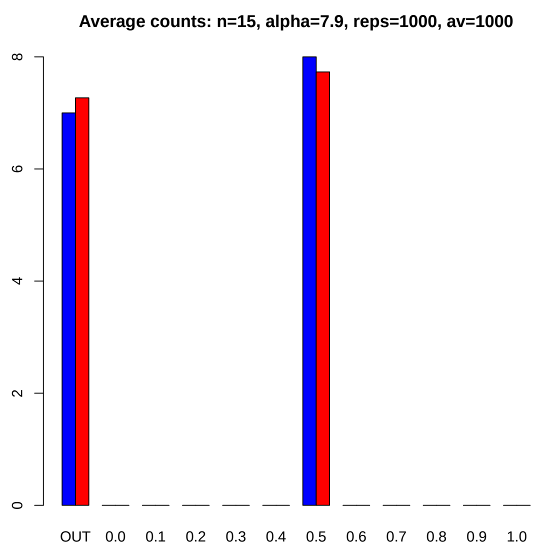

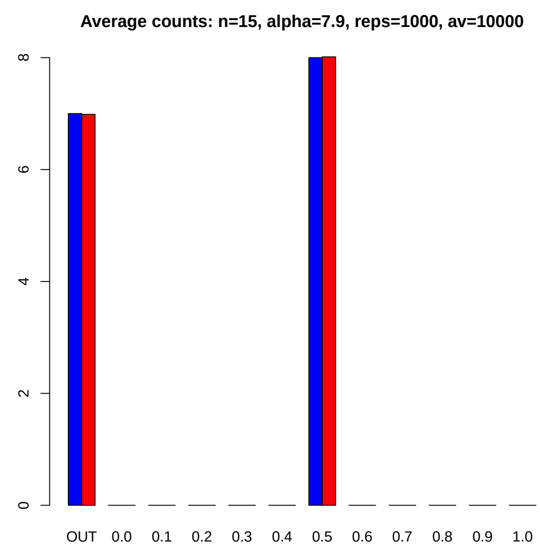

We conclude from this that, with enough election repetitions and averaging, it is indeed possible for the candidates to randomly tweak their actions and eventually converge to equilibrium behaviour as predicted by Theorem 2. Further scenarios can be tested as desired using the same C program.

Acknowledgements. I thank Martin J. Osborne for introducing me to this topic and for many helpful discussions. This research was partially supported by NSERC of Canada.

References

C. Andrieu and G.O. Roberts (2009), The pseudo-marginal approach for efficient Monte Carlo computations. Ann. Stat. **37(2)**, 697–725.

L. Bouton and G. Gratton (2015), Majority runoff elections: Strategic voting and Duverger's hypothesis. *Theoretical Economics* **10**, 283–314.

J.-P. de Vries (2015), Duverger's (f)law: Counterproof to the Osborne Conjecture. MSc thesis. Available at: https://thesis.eur.nl/pub/17645/

J.-P. de Vries, J.J.A. Kamphorst, M.J. Osborne, and J.S. Rosenthal (2016), On a conjecture about the sequential positioning of political candidates. Work in progress.

A. Downs (1957), An Economic Theory of Democracy. New York: Harper & Row.

M. Duverger (1951), Les Partis Politiques. Paris: Armand Colin.

B.C. Eaton and R.G. Lipsey (1975), The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition. *Review of Ecnomic Studies* **42**, 27–49.

H. Hotelling (1929), Stability in competition. *Economics Journal* **39**, 41–57.

S. Hug (1995), Third parties in equilibrium. *Public Choice* **82**, 159–180.

J. Nash (1951), Non-cooperative games. *The Annals of Mathematics* **52(2)**, 286–295.

M.J. Osborne (1993), Candidate positioning and entry in a political competition. *Games and Economic Behaviour* **5**, 133–151.

M.J. Osborne (1996), A conjecture about the subgame perfect equilibria of a model of sequential location. Available at: https://www.economics.utoronto.ca/osborne/research/conjecture.html

M.J. Osborne (2003), An Introduction to Game Theory. Oxford University Press.

W. Riker (1982), The two-party system and Duverger's law: an essay on the history of political science. *American Political Science Review* **76(4)**, 753-766.

J.S. Rosenthal (2016a), Stochastic Simulation of Sequential Game-Theory Voting Models. *Communications in Statistics: Simulation and Computation*, to appear.

J.S. Rosenthal (2016b), Nash Equilibria for Voter Models with Randomly Perceived Positions. Submitted for publication.

J.A. Schlesinger and M.S. Schesinger (2006), Maurice Duverger and the study of political parties. *French Politics* **4**, 58–68.

A. Shaked (1982), Existence and computation of mixed strategy Nash equilibrium for 3-firms location problem. *Journal of Industrial Economics* **31**, 93–96.