A new method for approximating proportional representation in multi-seat elections based on range or approval voting

Thomas Bewley and Paolo Luchini

October 2020

This paper develops a new method, dubbed Transferable Range Voting (TRV), for approximating proportional representation in multi-seat elections. The TRV method combines range voting, in which voters score all candidates independently (on an absolute scale), with the key idea of transferable votes, as implemented by the Single Transferable Vote (STV) approach, in which excess votes (beyond those required to get elected) are transferred to other candidates. The scores that voters assign to candidates in this range voting approach may be either real numbers or integers over a predefined range $[0, S]$, and thus includes the important special case of approval voting, with admissible scores of 0 (disapprove) or 1 (approve) only; the present method is referred to as Transferable Approval Voting (TAV) when applied in this simplified setting. As it is implemented today, STV is based on preferential voting, in which voters rank order their preferred candidates; this paper thus generalizes the key (and, powerful) notion of transferable votes, as used by STV, to the range voting and approval voting settings.

1 Introduction

The selection of voting systems is one of the most consequential foundational underpinnings of any representative-based democratic institution, which must carefully balance majority opinions with minority interests. The very purpose of voting in such institutions is that, when there are different like-minded groups within the electorate who tend to vote for candidates who represent their own dominant interests, then these groups should be represented among the winners of the election, with numbers approximately proportional to their strength within the electorate. This important property is referred to as *proportional representation* [1]. In this regard, it is broadly recognized that many currently-implemented strategies, such as single-seat elections with first-past-the-post voting methods in highly gerrymandered legislative districts, as implemented in the US, UK, Canada, India, Pakistan, and elsewhere, leave significant room for improvement [2, 3, 4].

1.1 Single-seat election methods and their properties

As a starting point, consider the setting of single-seat elections. There are several well-defined, mathematically precise properties that are useful in characterizing the many voting methods available for such elections. These properties are studied extensively in the electoral studies literature; to set the stage for the discussion that follows, we mention briefly only two of them here (see, e.g., [5]):

(A) the *Condorcet* property: any candidate who would win in every two-candidate election against each remaining candidate (in those situations for which this happens to be the case) would also defeat these candidates in the corresponding multi-candidate election.

(B) the *Independent of Irrelevant Alternatives* (IIA) property: assume in an election given the choices $\{A, B\}$, candidate A is preferred over candidate B; introducing a third candidate C and holding a new election over the choices $\{A, B, C\}$, all other factors being the same as before, candidate B can not now be preferred over candidate A.

The Condorcet and IIR properties are two precise classifications of voting methods that guarantee specific insensitivity to what is more generally known as the *spoiler effect*. To illustrate, imagine that candidate A is preferred to candidate B in a $\{A, B\}$ matchup, by 50.3% to 49.7%, in a certain election. Now introduce a minor candidate C, with support from only 1.6% of the population in a three-person $\{A, B, C\}$ matchup, and further assume that C draws 70% of his support from voters that would otherwise vote for A in this matchup, 30% of his support from voters that would otherwise vote for B, and all other conditions are the same as before. In this hypothetical three-person $\{A, B, C\}$ election:

$49.7 - 0.30 * 1.6 = 49.22\%$ would list B as their first choice.

$50.3 - 0.70 * 1.6 = 49.18\%$ would list A as their first choice, and

1.6% would list C as their first choice.

Thus, if only the first choice preference of each voter mattered (as in *First Past the Post* voting, which is neither Condorcet nor IIR), B would win, despite the fact that A is actually preferred to B by 0.6% of the electorate, and only 1.6% of the electorate voted for the minor candidate C. If, on the other hand, the voting method used is Condorcet, A would win, as A is preferred to both B and C in individual two-candidate matchups. Similarly, if the voting method used is IIR, then again A would defeat B, because A is preferred to B in an $\{A, B\}$ matchup, so the introduction of C can not change this ordering in an $\{A, B, C\}$ election.

The above tangible example is mentioned here simply to emphasize that the identification of Condorcet and IIR voting methods, which can be something of a subtle exercise, can be quite significant in the outcome of practical voting schemes aimed at reflecting the will of the electorate, especially in close elections (which are actually quite common in well-functioning democracies); they are not simply academic exercises. Further, the very willingness of an informed voter to "risk" voting for a "minor" candidate like C, with only a (perceived) outside chance of winning, but who (possibly) much better reflects that voter's own dominant interests in the election, is inherently tied to the voting method implemented. If the voting method is Condorcet and/or IIR, then informed voters can vote their conscience (dubbed "sincere" voting) without fear of "wasting" their vote on a candidate with a perceived reduced chance of winning. At the very least, this assurance has been seen to significantly increase overall voter participation, by up to 12% [6, 7], in actual elections. Without this assurance, overall voter interest and participation in elections is reduced, and the electorate inevitably devolves into a two-party system—and the incomplete representation of the various minority interests of the electorate which a two-party system implies.

Preferential Voting is an election method in which each voter rank orders the candidates. There are dozens of techniques available to tally the votes from such an election in order to determine the winner(s). Some of these methods, such as *Tideman's Ranked Pairs* method [8] and the *Schulze* method [9], are well suited for both the election of a single winner, as well as for developing a rank ordering of the candidates via majoritarian considerations. Both the Ranked Pairs method and the Schulze method satisfy the Condorcet property; these two schemes (and the many other available Condorcet schemes for tallying preferential votes) differ in terms of the procedures by which they reconcile the peculiar but very real possibility that, after all votes are tallied, the electorate prefers A over B, B over C, and C over A, a situation dubbed *Condorcet's voting paradox* [10]. Unfortunately, by *Arrow's impossibility theorem* [11], no preferential voting scheme in which all ballots are accounted for equally can satisfy the IIA property.

Range Voting, on the other hand, is an election method in which voters are asked to (and, in the analysis, assumed to) score the candidates independently, on an absolute scale. The scores may be real numbers (between 0 and 1, between 0 and 10, between -10 and 10, etc; all such scales are actually equivalent in the end) or quantized (e.g., to integers between 0 and 10, or in the case of *Approval Voting*, to the two integers 0 or 1). The scores for each candidate are summed over all of the ballots, and the candidate with the highest tally wins[1]. This simple approach satisfies both the Condorcet and IIA properties.

1.2 Multi-seat election methods

Small electorates (shareholders, homeowners associations, etc) generally need to elect a handful of seats to a board in a manner which ensures proportional representation, thus fairly balancing competing minority interests. Large electorates in democratic societies, on the other hand, generally need to balance geographically local representation with the proportional representation of minority interests in the election of a representative body (parliament, etc.). To achieve this balance, a strong consensus appears to have formed [1, 2, 3, 4, 12] suggesting that approaches based on First Past the Post voting methods in single-seat districts, the lines of which are gerrymandered [13] by those in power, must be forgone in favor of elections in larger multi-seat districts, with voting schemes that ensure proportional representation of minority interests within each district. In either case, the development and broad implementation of voting systems capable of achieving proportional representation of minority interests in multi-seat elections is of significant interest[2].

The two major categories of election systems most broadly implemented today to achieve proportional representation [4, 12] in multi-seat elections[3] are *List PR* and the *Single Transferable Vote* (STV). With List PR methods, which are fundamentally tied to political parties, each political party publishes an ordered list of candidates on the ballot. Voters vote for a party, and parties receive seats in proportion to their overall share of votes, with winning candidates taken from these lists, in order of their positions within them. There are many variations of this general approach; in most of them, voters forfeit a significant component of the decision-making process to the political parties responsible for compiling the lists. Alternatively, in the List PR setting, *primary* elections may be held amongst those in the electorate that self-identify with each political party, in order to determine (using, in turn, some mult-seat election method) the ordered list of candidates for each political party to be used in the subsequent *general* election. [See also §1.3 for related discussion.]

The Single Transferable Vote (STV), on the other hand, is a remarkably clever class of closely-

[1]If seeking to seat multiple candidates using a simple range voting or approval voting method, one might simply assign the candidate with the second-highest tally to the second seat, etc. Note, however, that this is a simple majoritarian approach; that is, in a strongly polarized election with 60% of the population self-identifying with Party A and 40% of the population self-identifying with Party B, 100% of the elected candidates would likely be from Party A following this approach. A proportional representation approach like that developed in this paper, in contrast, is designed to achieve, in such an example, about 60% of the elected candidates from Party A and 40% from Party B. Numerical confirmation that the methods developed in this paper actually achieve such proportional representation is reported in §3.2 and §3.3.

[2]A counter-argument to this position is provided in [14], which suggests that majoritarian (rather than proportional) elections, and the two-party systems which they engender, significantly empower consumers relative to producers.

[3]A handful of other election systems are sometimes used in multi-seat elections, including the *Single Non-Transferable Vote* (SNTV), in which each ballot votes for just one candidate, and the m candidates with the most votes win, and *cumulative voting*, in which each ballot is assigned m voting "tokens", which may be distributed by the voter to m candidates, or piled up on one or more candidates. Such schemes are generally known to be inferior to STV in terms of the proportionality of the representation ultimately achieved, and are characterized by the important flaw of *strategic voting* mechanisms, in which groups of voters can collude to achieve their desired outcomes disproportionately; see, e.g., [1] for details. Voting schemes without such collusion mechanisms present, which instead motivate informed voters to simply vote their conscience, and which then apply carefully-constructed algorithms which balance the election results in a manner that inherently achieves proportional representation of minority interests to the maximum extent possible, are strongly preferred in representative-based democratic institutions.

related preferential voting methods [4, 12] in which proportional representation in multi-seat elections may be well approximated, without ceding a significant component of the decision making process to political parties. STV is thus one of the most popular class of proportional representation voting methods available today. Methods in this class assign a single vote to each ballot, with each of these votes initially allocated to its corresponding ballot's most preferred candidate; as the tally proceeds and candidates are either elected (by achieving a sufficient threshold of votes) or eliminated (by not receiving enough votes to advance further), the votes are proportionally transferred, as appropriate, to other candidates further down each ballot. The details of how these proportional transfers of votes are performed vary somewhat among different methods in this class; one common variant is summarized in §1.2.1.

In the election of a single candidate, STV (referred to as *Instant Runoff Voting* in this setting) is neither Condorcet nor IIA; on the other hand, simple range and approval voting methods (upon which the TRV and TAV methods are based), under the sincere voting assumption, are both Condorcet and IIA. At each step of the vote tallying process, the TRV and TAV methods developed in §2 modify the traditional range voting and approval voting methods using the idea of transferable votes (as inspired by STV) in order to discount at that step the individual ballots that rate highly candidates that have already won, thereby giving greater consideration to individual ballots that are not yet well represented among the current list of winners. This approach thus combines the favorable properties of the range voting and approval voting methods with the proportional representation goals achieved by STV.

1.2.1 Description of the Gregory variant of the STV algorithm

The STV approach is a preferential voting tallying process that may be thought of as being coordinated by the assignment and distribution of fictitious voting "currency". To be explicit, assume there are n voters, p candidates, and m seats to be assigned in an election; the Gregory (a.k.a. Senatorial Rules) variant of the STV method [12], may then be described in this framework as follows:

0. Each ballot i is initially assigned $c_i = 1$ voting "currency".

1. An "offer" of "voting currency"[4] is made from each ballot i to the highest-ranked candidate j on that ballot which is still left under consideration; this offer is given simply by the total amount of voting currency which ballot i has left at this iteration.

2a. The offers to each candidate j from all the ballots are then summed, and the highest total offer r compared with a threshold quota required to be named a winner, given by the Droop quota [15] $q = 1 + n/(m + 1)$; if $r < q$, the candidate with the lowest total offer at this round is removed from further consideration, and the process repeated from step 1.

2b. Otherwise, the candidate with the highest total offer r "cashes in" as the winner of this round. At this point, the amount by which r exceeds q is returned to the respective ballots that elected this candidate in a proportional sense. In other words, if the candidate winning a given round has been offered twice the threshold quota required to be named a winner, the amount of voting currency deducted from each ballot during this round is reduced to half of the amount that it initially offered to the winning candidate.

[4] It is important to point out that this "offer" of "voting currency" is, of course, solely an internal bookkeeping mechanism used by the STV algorithm itself, which in the end is fairly simple, to tally the votes. As STV is a preferential voting algorithm, all the voter does in an election which uses STV to tally is to rank order the candidates, from their first choice to their last acceptable choice.

3. The process then proceeds again from step 1, with the winning candidate at the previous iteration removed from further consideration, until m winning candidates are identified.

It is seen that, in an STV contest, winning candidates need to have both *strong* support, to not be eliminated in step 2a in the early rounds of tallying, and *broad* support, eventually garnering enough voting currency to exceed the threshold and cash in. In the $m = 1$ setting of single-seat elections, STV (known as Instant Runoff Voting in this setting) completes as soon as the above algorithm reaches step 2b the first time. In the $m > 1$ setting of multi-seat elections, on the other hand, the otherwise somewhat elusive goal of proportional representation is well approximated, as those ballots at each later round that are not yet well represented amongst the current winning candidates are counted more heavily in the later rounds of the tallying.

1.3 Primary elections, pivoting, and radicalization

Note that primary + general election approaches, as mentioned in §1.2 in the context of List PR methods for multi-seat elections, are used extensively in the setting of single-seat elections (as for the US president), where the "list" determined by each party at the primary stage is just a single candidate. This approach is far from being either Condorcet or IIA, though, as very different metrics are needed to win the primary and general election contests, especially in the case of single-seat elections when the voters within the electorate are not necessarily loyal to their specific self-identified political parties in the general election, which is common. Typically, candidates vye to "rally their base" in the run-up to the primary election, in order to place well amongst their party's constituents in the primary, then traditionally (at least, to an extent) attempt to "pivot" to more broadly electable postures in the run-up to the general election. At best, this pivoting on the issues of the day between these two distinct contests is duplicitous. At worst, little-to-no effort is actually even made to accomplish such a pivot to somewhat more moderate stances after the primary, and the general election (and, the manner of governance thereafter) becomes a polarized "battle" between non-compromising radicalized postures that were solidified by the candidates during the primary contests.

The deviation of the primary + general election approach from the Condorcet and IIA properties is particularly pronounced when voters are generally lethargic and disenfranchised about voting, often not showing up to vote at all unless a particularly radical candidate (on one end of the political spectrum or the other) is in the running.

We thus identify some important problems that compel broad populations (not just academics writing papers) to advocate replacing various existing voting methods with new methods that satisfy the Condorcet and IIA properties:

1. the *spoiler effect*, which is effectively eliminated by Condorcet and IIA methods, as exemplified by the A v. B v. C example highlighted in §1.1,

2. *insincere/strategic voting* [that is, rather than simply voting for those who you think would do the best job from your own personal perspectives, instead either voting for someone else who you think has a better possibility of winning or, worse, colluding with others to achieve your desired outcomes disproportionately], which is essentially eliminated by implementing tallying algorithms which inherently and robustly achieve proportional representation themselves,

3. *voter disenfranchisement* [that is, reduction in voter turnout] due to the perception that sincere voting would run a high probability of "wasting" your vote on candidates who might otherwise be viewed as "minor", and the concomitant *inviability of third parties*, both of which are alleviated by the implementation of Condorcet and IIA methods, which assure informed voters that their votes are not "wasted" when voting sincerely, and

4. a *tendency towards radicalization*, which the primary + general election process appears now to increasingly engender, with the process of "rallying the base" in primary elections tending to solidify strong and uncompromising positions on either end of the political spectrum.

The effect highlighted in point 4 seems to have intensified in recent years, perhaps as a result of politically-leaning 24-hour news stations repeatedly playing clips of postures made during the primary process, which seem to make it much less viable for candidates to pivot from "rallying the base" early in an electoral contest to more centrist personas as the general election (and, the time for actual governance) approaches.

The manner in which Condorcet and IIA methods address problems 1, 2, and 3 above highlights their inherent "fairness", and the manner in which Condorcet and IIA methods that bypass separate "primary" constests address problem 4 (see in particular the last paragraph of §1.2.1) highlights what might be identified as their "centrist" tendencies. It is also worthy to note that voting methods that effectively address problems 2 and 3 above do not benefit those already in political power, and thus the shift to Condorcet / IIA voting methods, with the "fair" and "centrist" tendencies described above, is expected to be met with significant resistance (via any of a number of dubious us-versus-them arguments) by the leadership of both political parties in existing two-party systems. To be successful, such a shift would thus likely need to be advocated for by grass-roots voter-based efforts, in situations where direct referendums by voters are viable avenues for change in democratic institutions.

2 Description of the new TRV/TAV algorithm

Assume again that there are n voters, p candidates, and m seats to be assigned in an election. The Transferable Range Voting (TRV) approach proposed here is a range voting method in which each of the n voters is assumed to score each of the p candidates independently (on some absolute scale), with each voter's ballot subsequently tallied in a transferable vote manner, akin to the STV method summarized in §1.2.1, such that proportional representation is ultimately approximated. The score $s_{i,j}$ given by each voter i to each candidate j may be taken in the discussion that follows as a real number in the range $[0, S]$; it is often convenient to take $S = 1$. Note, however, that using scores that are on a different scale (with different values of S), and/or using scores that are quantized (e.g., as integers in the range $[0, 10]$, which is perhaps most natural), may be handled by precisely the same tallying algorithm. In the extreme case in which each score used is either a 0 or a 1, for disapprove or approve, the method proposed is referred to as Transferable Approval Voting (TAV).

Analogous to the description of the Gregory variant of STV in §1.2.1, TRV/TAV is now introduced as a similar process coordinated via transferable voting currency as follows:

0. Each ballot i is assigned $c_i = C$ in voting "currency" (nominally, we take $C = S$, but see §4 for discussion of other possibilities).

1. An "offer" of "voting currency"[5] is made from each ballot i to each candidate j, given by the minimum of the amount of voting currency which ballot i has left at this iteration, and the score $s_{i,j}$ which ballot i assigns to candidate j (cf. step 1 of the method described in §1.2.1).

[5]As in Footnote 4, it is important to point out that this "offer" of "voting currency" is, again, solely an internal bookkeeping mechanism used by the TRV/TAV algorithm itself, which in the end is again fairly simple, to tally the votes. As TRV is a range voting algorithm, and TAV is an approval voting algorithm, all the voter does in an election which uses TRV or TAV to tally is to score the candidates over the range specified; that is, in the range $[0, S]$ for the TRV method, or with a 0 (unacceptable) or 1 (acceptable) in the TAV method.

2a. The offers to each candidate j from all the ballots are then summed, and the candidate with the highest total offer r "cashes in" and wins this round.

2b. At this point, r is compared with a threshold quota required to be named a TRV winner, initially taken as $q = C \cdot n/m$ (cf. the Droop quota used in step 2 of the method described in §1.2.1). The amount by which r exceeds q is then returned to the respective ballots that elected this candidate in a proportional sense. In other words, if the candidate winning a given round has been offered twice the threshold quota required to be named a winner, the amount of voting currency deducted from each ballot during this round is reduced to half of the amount that it initially offered to the winning candidate.

3. The process then proceeds again from step 1, with the winning candidate at the previous iteration removed from further consideration, until m winning candidates are identified.

Note that if all n voters ultimately contribute all of their money to winning candidates, the threshold quota required to be named a TRV/TAV winner works out to be precisely $q = C \cdot n/m$, as suggested in step 2b. However, this usually does not work out to be the case, as some ballots only wind up indicating significant support to candidates who ultimately lose. The process described above is thus iterated, with q reduced slightly at each iteration in a simple convergent manner until the "cash in" costs ultimately made by each of the winning candidates work out to be precisely equal. The resulting TRV/TAV method, which incorporates this q adjustment iteration, is illustrated via a simple working Matlab code in Algorithm 1. [In practical tests, this overall q adjustment iteration is seen to play a rather minor role in the operation of the TRV/TAV algorithm, as discussed in §3.1.]

It is noted here that TAV is not the first attempt at proportional representation leveraging approval voting. A few computationally-complex competing algorithms have been suggested, as discussed in [16], based on minisum, minimax, and weighted sum comparisons between ballots. A primary reason that TAV is attractive in comparison to these schemes is its relative simplicity and transparency, which is achieved by incorporating the powerful notion of transferable votes, as implemented in STV.

3 Numerical tests

All executable Matlab codes used to produce the results discussed below are available at
https://github.com/tbewley/TRV

3.1 TRV with random votes

The approach outlined in §2, with the q adjustment iteration implemented as described (that is, iterating until the "cash in" cost of all winning candidates is approximately equal), is codified in Algorithm 1. Note that Algorithm 1 may be tested quite simply in Matlab or Octave with random data, as shown in Algorithm 2. Adding a few print statements to Algorithm 1, most tests of this sort, for the values $\{n = 20, p = 10, m = 5\}$ tested by Algorithm 2 as provided, indicate that the iterative refinement of q actually changes the value of q only slightly (typically by less than 5%), and that this refinement of q usually (at least 98% of the time, when running Algorithm 2 a large number of times) has zero effect on the list of winners returned by TRV. The interested reader is encouraged to run several random tests of Algorithm 2 for the values of $\{n, p, m\}$ of his/her particular interest in order to develop similar statistics. A representative result generated by such random tests, in the (fairly rare) case that the iterative refinement on q actually affects the final list of winners, is summarized as follows:

Algorithm 1: Transferable Range (or, Approval) Voting in executable Matlab syntax.

```
function [winner] = TRV(n,p,m,s)
% Inputs: n = # of voters, p = # of candidates, m = # of winners,
% s(i,j) \in [0,1] = Score that ballot i assigns to candidate j (ties ok),
%                    normalized to lie on the range [0,1]; that is, S=1.
% Output: winner(k) = list of m winners
C = 1;               % initial voting currency assigned to each ballot
total_cost = C*n;    % initialize total_cost as the total currency
epsilon = 0.00001;   % tolerance used to refine threshold quota q
while 1
   q=total_cost/m;   % (re-)initialize threshold quota required to be named a winner
   c(1:n) = C;       % (re-)initialize voting currency assigned to each ballot
   for k=1:m         % for each winner...
      for j=1:p      % for each candidate...
         offer(:,j)=min(c(:),s(:,j));      % amount voters offer to candidate j
         total_offer(j)=sum(offer(:,j)); % total candidate j is offered
      end
      % remove offers to winning candidates, so they don't win again
      for kbar=1:k-1, total_offer(winner(kbar))=0.0; end
      [r(k),winner(k)] = max(total_offer);        % identify winner(k)
      cost(k)=min(q,r(k)); % cost(k) is minimum of quota q and winning offer r(k)
      c(:)=c(:)-offer(:,winner(k))*cost(k)/r(k); % cash in winner(k)
   end
   previous_total_cost=total_cost; total_cost=sum(cost); % adjust q if necessary
   if abs(total_cost-previous_total_cost) < epsilon*total_cost, break, end
end                                                   % repeat until converged
% end function TRV
```

Algorithm 2: Test script for TRV for a population voting randomly.

```
% Matlab script to test TRV with random votes.
n=20, p=10, m=5,  s=rand(n,p); s=s/max(max(s)); [winner]=TRV(n,p,m,s)
```

A. total offers (in order) to each of the 10 candidates at the first iteration:

   {7.9247, 8.4641, 10.1731, 8.3701, 8.9339, 8.5880, 7.9209, 12.0758, 9.7347, 12.2787}

B. initial ranking of the 10 candidates based on these offers (i.e., the "majoritarian" result, equivalent to the standard range voting order):

   $\{10, 8, 3, 9, 5, 6, 2, 4, 1, 7\}$

C. winners of both initial pass and first q refinement of TRV (taking $q_0 = Cn/m = 4$, and $q_1 = 3.86$):

   $\{10, 8, 3, 6, 4\}$

D. winners after each of the second through tenth q refinement iterations of TRV (with the q_i converging quickly towards $q_\infty = 3.8055$):

   $\{10, 8, 3, 6, 9\}$

E. cost associated with each of the winners after convergence of q:

   {3.8055, 3.8055, 3.8055, 3.8055, 3.8055}

Note in the above representative example that:

1. The standard range voting ordering obtained simply via majoritarian considerations, given in B, differs from the proportional representation result determined by TRV, given in D. However, the first winner in both is always the same.

2. The iterative refinement on q changed the list of winners in after the first q refinement in this example (compare C and D). Note that all changes to the list of winners due to the refinement of q in TRV that we have witnessed thus far have happened in the first few iterations on q; the need to iterate q all the way to convergence, by selecting ϵ very small in Algorithm 1, is thus not seen to be a significant concern.

3. Once the refinement iterations on q converge, the "costs" associated with each of the winning candidates (that is, the minimum of the quota q and the winning offer $r(k)$ at each iteration), as shown in E, are essentially equal, as expected.

3.2 TRV with votes representative of strong party affiliations

Algorithm 3 suggests a new benchmark test problem for a population with party affiliations, to test the veracity of proportional representation voting schemes. This test problem, as provided, assumes that a given electorate is 40% Party A, 30% Party B, 20% Party C, and 10% Party D, and that each of these four political parties has put up 10 candidates to run in an election for a total 10 open seats.

In short, it is found in this example that, if each voter scores a $w_2 = 0.8$ or more for candidates from his/her own political party, and scores a $w_1 = 0.2$ or less for candidates from some other party, then TRV almost always returns in this example 4 winners that are from Party A, 3 that are from Party B, 2 that are from Party C, and 1 that is from Party D, thus providing proportional representation along party lines. Recall that TRV combines the information from all of the votes on all of the ballots; if w_2 is significantly reduced and/or w_1 is significant increased in this example, reflecting a less polarized electorate in which voters support candidates from other parties in addition to their own, then these distributions begin to change, as well they should.

3.3 TAV with votes representative of strong party affiliations

As illustrated in Algorithm 4, modifying the benchmark proposed in the previous section to use approval voting (with all votes restricted to be 0 or 1) is entirely straightforward. Analogous to the results reported in the TRV case, it is again found that if each voter scores a 1 for candidates from his/her own political party a sufficiently large (w_2) fraction of the time, and scores a 1 for candidates from a different political party a sufficiently small (w_1) fraction of the time, TAV nearly always results in representation with proportions which reflect the party affiliations of the electorate itself.

4 Discussion

The Transferable Range Voting (TRV) method codified in Algorithm 1 combines the process of range voting with the idea of transferable votes, as implemented in the Gregory variant of STV (a preferential voting method), with the specific goal of achieving proportional representation in multiple candidate elections. The first candidate elected in this approach (see steps 0, 1, and 2a of the scheme as described in §2) is chosen via simple range voting, and thus this choice is both Condorcet and IIA. Additional winners are sought in a manner which is specifically not Condorcet nor IIA; additional winners are instead chosen via a range voting approach which is modified in a

Algorithm 3: Test script for TRV for votes modeling a population with party affiliations.

```
% Matlab script to test TRV for an example representing party affiliations.
% We assume 4 political parties, each with 10 candidates running.

n=1000; p=40; m=10; n1=0.4*n; n2=0.3*n; n3=0.2*n; n4=0.1*n;  % populations
w1=0.2; % maximum support given to someone outside your own party
w2=0.8; % minimum support given to someone within your own party

% Note that proportional representation is given (within quantization) by:
% n1/n Party A candidates in the range (1 :10)
% n2/n Party B candidates in the range (11:20)
% n3/n Party C candidates in the range (21:30)
% n4/n Party D candidates in the range (31:40)
% If w1 is almost 0 and w2 is almost 1, TRV well approximates these ratios,
% as generally shown in the numerical results below, when run several times.

% initialize random votes for candidates outside one's party...
s=w1*rand(n,p);
% ... then adjust votes higher when voting inside one's party
s(1          :n1,          1 :10)=w2+(1-w2)*rand(n1,10); % Party A
s(n1+1       :n1+n2,       11:20)=w2+(1-w2)*rand(n2,10); % Party B
s(n1+n2+1    :n1+n2+n3,    21:30)=w2+(1-w2)*rand(n3,10); % Party C
s(n1+n2+n3+1:n1+n2+n3+n4,31:40)=w2+(1-w2)*rand(n4,10); % Party D

[winner]=TRV(n,p,m,s)   % Finally, tally votes.
```

Algorithm 4: Modifications required to Algorithm 3 to test TAV scenarios.

```
...
w1=0.02; % fraction of time you approve of a candidate outside your own party
w2=0.8;  % fraction of time you approve of a candidate within your own party
...
s=0.5+0.5*sign(rand(n,p)-(1-w1));
...
s(1:n1,1:10)=0.5+0.5*sign(rand(n1,10)-(1-w2)); % Party A
% etc.
```

manner which discounts individual ballots that rate highly candidates that have already won, thereby giving greater consideration to individual ballots that are not yet well represented among the current list of winners, thereby achieving the goal of proportional representation, as demonstrated in §3.2. The scores used by voters in TRV may be real numbers or integers over a specified range $[0, S]$. Approval voting is a special case, in which, all approved candidates on a ballot are scored with a 1, and all unapproved candidates are scored with a 0; when applied in this setting, this algorithm is referred to Transferable Approval Voting (TAV).

As seen in the numerical results of §3.1, TRV can produce a substantially different list of winners than standard range voting, though the first winner (chosen in a Condorcet, IIA manner) will always be the same. Taking $S = C = 1$, the threshold quota q required to be named a TRV winner is close to, but usually slightly less than, n/m, reflecting the fact that most, but not quite all, of the voting "currency" initially assigned to the individual ballots is "cashed in" during the TRV tallying process. The precise value of q used in TRV is found iteratively; upon convergence, the "costs" associated with "cashing in" by all of the individual winners are essentially equal. As seen in the numerical benchmark problem proposed in §3.2, if the electorate is highly politically polarized (that is, if voters generally score quite high candidates from their own parties, and score quite low candidates from other parties), then the TRV approach accurately achieves proportional representation along party lines (subject to the inaccuracies due to quantization). If the electorate is less polarized (generally, a desirable situation), the votes of people from different political parties are more tightly integrated during the TRV tallying process, and proportional representation along strict party lines becomes less accurate. In this case, the vote then becomes more about the individual candidates themselves, rather than party lines, but the concept of proportional representation when selecting from among these candidates, as distinct from majoritarian representation as achieved by simple range voting, is still quite applicable. Analogously, it is seen in the numerical benchmark problem proposed in §3.3 that, if each voter approves of candidates from his/her own political party a sufficiently large fraction of the time, and disapproves of candidates other different political partis a sufficiently small fraction of the time, TAV also results in proportional representation along party lines.

The value of total voting currency C assigned to each ballot in TRV/TAV is nominally taken as S, the maximum score allowable to any given candidate. [The value of S in TRV, in turn, is simply a scaling factor, and is taken w.l.o.g. as $S = 1$ in this paper.] If we instead take $C \geq S \cdot m$, then no ballot will every be short on voting currency during the tallying process, and the offer made to each candidate will be simply the score which that ballot assigned to that candidate, thereby reducing TRV to simple range voting, and TAV to simple approval voting (see Footnote 1). Selecting C to be somewhere between S and $S \cdot m$, then, provides a mechanism to compromise between proportional representation and simple majoritarian outcomes.

Interestingly, the TRV/TAV method has a built-in quantification of under-represented minority interests. If $n \gg p > m$, which is often the case, the first-choice candidates of all voters can not necessarily wind up as TRV winners. If upon convergence of Algorithm 1, however, the converged threshold quota q, works out to be equal or close to $q_0 = C \cdot n/m$, then most of the voters at least support the pool of winning candidates enough to spend almost all of their voting "currency". The converged value of $R = 1 - q/q_0$ represents the extent to which some voters do not support the pool of candidates that ultimately won the election. A value of R near zero means that almost all voters have identified candidates that they support strongly enough to spend all of their voting currency, whereas a value of R near, e.g., 0.1 means that, averaged over all of the ballots, 10% of the voting currency associated with each ballot has gone unspent on the pool of winning candidates. A more direct way of quantifying under-represented minority interests is to take statistics on the unspent voting currency c_i assigned to each ballot i after the voting is complete. The mean value $(1/n)\sum_i (c_i/C)$ is simply the metric R mentioned above. The rms value $\sqrt{(1/n)\sum_i (c_i/C)^2}$ provides an alternative measure of unspent voting currency that might also be of interest. We are unaware of any other voting schemes for proportional representation that have associated metrics for quantifying under-represented minority interests.

It is important to emphasize that the STV, TRV, and TAV methods are all fairly simple algorithmically, and may be implemented (by modifying the executable Matlab codes included in this article) in a straightforward cellphone app, webpage, Excel spreadsheet, SQL database, etc. The actual mechanisms for voting (rank ordering the candidates in the case of STV, scoring the candidates on a scale of $[0, S]$ in the case of TRV, and approving/disapproving the candidates in the case of TAV) are all quite natural to the voter. For those actually interested in the inner workings of the tallying algorithm, the notions upon which all three are based (the assignment to, and transference of, "voting currency," an equal amount of which is initially allocated to each ballot) is both transparent and objectively fair.

Finally, we comment that the TAV approach deserves particular attention, as the ballot used in approval voting (in which a voter may mark *all* of his/her approved candidates with an X) constitutes a particularly straightforward extension of the ballot used in simple plurality voting (in which a voter may mark only a single candidate with an X). Thus, for the important purpose of introducing multi-seat districts with proportional representation of various minority interests to an electorate accustomed only to plurality voting, like the long-standing tradition in most of the USA, TAV (based on approval voting) might be somewhat easier to introduce than both TRV (based on range voting) and STV (based on preferential voting). On the other hand, more sophisticated electorates might well prefer the enhanced degree of voter expression that may be accomodated using TRV (based on range voting) as compared with that possible using STV (based on preferential voting) and TAV (based on approval voting).

References

[1] Felsenthal, D. (1990) *Topics in Social Choice: Sophisticated voting, efficacy, and proportional representation*. Praeger.

[2] *Fairvote.org* http://www.fairvote.org

[3] Gray, C.P.G. http://www.cgpgrey.com/politics-in-the-animal-kingdom/

[4] *ACE Encyclopaedia: Electoral Systems* http://aceproject.org/ace-en/topics/es/

[5] Taylor, A. (2005) *Social choice and the mathematics of manipulation*. Cambridge.

[6] Banducci, S.A., & Karp, J.A. (2009) 'Electoral systems, efficacy, and voter turnout'. In Klingemann, H.-D. (ed) *The comparative study of electoral systems*. Oxford, 109-134.

[7] Franklin, M.N. (1996) 'Electoral participation'. In LeDuc, L., Niemi, R., & Pippa Norris (eds.) *Comparing Democracies: Elections and Voting in Global Perspective*. Thousand Oaks, CA: Sage, 216-35.

[8] Tideman, T.N. (1987) Independence of clones as a criterion for voting rules *Social Choice and Welfare* **4**, 185-206.

[9] Schulze, M. (2011) A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method *Social Choice and Welfare* **36**, 267–303.

[10] Nurmi, H. (2002) *Voting paradoxes and how to deal with them*. Springer.

[11] Arrow, K. (1962) *Social Choice and Individual Values*. Wiley.

[12] Lakeman, E. (1974) *How Democracies Vote: Study of Electoral System*. Faber and Faber.

[13] Katz, R. (1997) *Democracy and Elections* Oxford.

[14] Chang, E., Kayser, M., Lionzer, D., & Rogowski, R. (2011) *Electoral systems and the balance of consumer-producer power*. Cambridge.

[15] Benoit, K. (2000) Which Electoral Formula Is the Most Proportional? A New Look with New Evidence *Political Analysis*, **8**, 381-388.

[16] Kilgour, D.M., Brams, S.J., & Sanver, M.R. (2006) 'How to elect a representative committee using approval balloting.' In Simeone, B, & Pukelsheim, F. (eds) *Mathmatics and democracy; recent advances in voting systems and collective choice*. Springer, 83-95.