Comparative survey of multiwinner election methods

Warren D. Smith*
WDSmith@fastmail.fm

June 18, 2006

Abstract — **After ≈ 100 years of relative sterility, recent progress has both invented several new multiwinner voting systems, and allowed us to begin to judge which of them is the best. We define and compare them here and come to the tentative conclusion that W.D.Smith's unconventional "asset voting" system is the best multiwinner election system known, and can be argued, for fundamental reasons, to be unmatchable by any conventional voting system.**

1 Multiwinner versus single-winner voting systems

In my opinion, the situation with single-winner elections is comparatively simple because there is a clear yardstick, called "Bayesian regret," (which is a rigorously mathematically and statistically defined quantity) for deciding whether election system A is "better" than election system B. My computer measured Bayesian regrets in a vast number of simulated elections with various kinds of voters (honest, strategic, ignorant to different degrees, etc.) different numbers of voters and candidates, different "utility generators," and various kinds of election systems [13] with the conclusion that range voting was clearly the best system among about 30 competing single-winner systems in that study.

*Multi*winner election systems, however, are currently terra incognita. It seems very unclear how to measure and how to justify a claim system A is better than system B.

In single-winner systems we can imagine each candidate has some "election utility" from the point of view of each voter. The Bayesian regret is simply the expected gap between the achieved utility-sum and the maximum possible utility sum, under some probabilistic model of utilities and voter behavior. But in multiwinner elections, the societal utility need *not* be just the sum of the utilities of the elected candidates – because it matters how well the elected parliament will work together and to what extent they cancel one another's bad tendencies out. These notions can be hard to quantify, to say the least. It is usually seen as a good goal to elect a "representative" parliament containing a diversity of views, which is an entirely different goal than of just picking *one* winner with the *best* views.

Recently, though, enough light has been shed on the subject to allow us to begin to reach useful conclusions. Our goal is to explain that.

2 Statement of the problem

There are $V > 1$ voters and $C > 1$ candidates. The voters cast some sort of "votes" that somehow provide information about their preferences among the candidates. Those votes are somehow processed by the election system, with the result being the selection of W "winning candidates," $1 \le W < C$.

The question is what the votes and election system ought to be.

3 Properties for multiwinner election systems

As either an election system design tool, or as a method of discriminating better from worse election systems, one can try demanding the election system have certain *properties*.

In my opinion the property-based approach for discriminating among single-winner election systems, was largely a *wrong approach*, which led the whole political science community down the *wrong path* for decades. That is because it is not clear which properties are more important than others. So while understanding the properties of a voting system is good because it gives us more understanding, it is insufficient to make it clear that election system A is better than election system B. And indeed such results as "Arrow's impossibility theorem" (showing that no voting system with a certain combination of properties could exist) erroneously led political scientists to the wrong idea that there could be no "best" single-winner voting system. The *right path* is "Bayesian regret" measurement, which in a sense automatically considers all possible "properties" or "paradox" (property failure) scenarios – including properties that nobody has ever defined or named – and automatically weights them all with the correct weights that depend both on the likelihood of that kind of paradox and on its severity (utility-decreasing effect). Bayesian regret measurement can be automated. Several people (including me, although I came later than most) independently invented Bayesian regret and realized it could be used for automated election system comparison. But I did the most extensive such automated comparative study and mine was the only study that included "range voting." Therefore it was I who made

*Send non-electronic mail to: 21 Shore Oaks Drive, Stony Brook NY 11790.

the extremely important discovery of the *clear superiority* of range voting over all other voting systems in the study. This reversed decades of wrong views that no voting system was clearly "best."

However, in the *multi*winner case, nobody has proposed any concept like Bayesian regret, and hence we do not know a way to make a clear, automated comparison between different election systems. Therefore, we are forced back to the old-fashioned property-based approach.

We now list some important possible properties of multiwinner election systems.

MN: Monotonicity. If a voter (or subset of identical voters) "increases their vote for candidate X, while leaving his vote for all others the same (or decreasing)" then that cannot decrease X's chance of being in the winner set.

FA: Fairness – no political parties needed. Some democracies have adopted multiwinner proportional election systems based on "party lists." That is, voters vote not for candidates, but for political parties, and then party P gets a fraction of seats proportional to its number of received votes, and awards those seats to selected party members, often according to a prestated prioritized list. I consider that to be an abomination[1]: the voting system ought to give all candidates for office, a priori equal chances of being elected, regardless of what their political party affiliations (if any) are, and regardless of their status within those parties. And it is an abomination that the orderings of the lists are chosen by the parties[2] rather than by the voters. Of course, in real life in any election system, candidates usually will enjoy advantages or disadvantages caused by their party affiliation and status within that party. I am not denying that reality. I *am* denying that the election system *itself* should take those party affiliations and status into account as part of its method of determining the winners. I want any and all such advantages to be external to the election method, not internal to it.

Call an election system "fair" if it selects the winner set purely as a function of the votes (and totally ignores party affiliations), and in a manner symmetric under the $C!$ permutations of the order of the C candidates, i.e. treating all candidates equally even if their party affiliations are unequal.

PR: Proportionality. In a situation in which the voters and candidates all fall into a finite number of disjoint "camps" and there is straight party line voting – that is, voters in camp j always vote the maximum possible for all candidates of type j and the minimum possible for all candidates of types other than j – then (if enough candidates from each camp are available) the winner set should have the same demographics as the electorate itself (up to unavoidable deviations caused by rounding to integers).

Note that actually, many possible kinds of stronger or weaker proportionality statements could be made if we attempt to specify precisely what the "effects of roundoff to integers" are allowed to be. We here are intentionally leaving that issue vague.[3]

RE: Representativeness. Consider a situation in which each voter "approves" of some subset of the candidates. Depending on the network of approval-relationships, there might or might not exist a way to declare W among the C candidates "winners" in such a way that each voter approves of at least one winning candidate ("his personal representative"). An election system features "representativeness" if, whenever such a winner set exists, one is actually chosen.

HY: Hyperproportionality. (So-called because it is a stronger property than mere proportionality.) As usual in discussions of proportionality, we imagine disjoint camps of party-line voters. *But*, we now make our Tory voters *not* simply mechanically vote maximum for every Tory and minimum for every non-Tory. We make them a bit more interesting: imagine they vote maximum for a nontrivial *subset* of the Tories chosen by each voter independently, and vote minimum for everybody else. (The strategies for determining the favored subset could be systematically *different* for, e.g. Whig and Tory voters – for example Tories might systematically be "more choosy" than Whigs.)

Then in a "hyperproportional" voting system, the winners still will have the same party-proportions as the electorate (if enough candidates and voters are available, and *regardless* of the voter choices within their camps) up to effects of roundoff to integers.

Hyperproportionality also could be called "immunity to cloning" – a concept introduced by N.Tideman. That is, if "clones" of a candidate are introduced into the pool of candidates, Tideman (in the case of a single-winner election) wanted the winner to be unaffected except possibly for replacement by a clone. Hyperproportionality is an analogous demand for the winner *set* in multiwinner elections – except for the unavoidable proviso that hyperproportionality cannot be assured if, e.g. half the voters want a Whig candidate in camp A in a 16-winner election in which fewer than 8 Whigs are running. Clone-based problems such as "vote splitting" or "teaming" are a very common pathology in many election systems: in plurality elections cloning a frontrunner is often enough to split the vote and cause both clones to lose; in Borda elections a party that runs an enormous number of cloned candidates is assured a huge victory.

DC: Descriptive complexity. The system should be easy to describe as an algorithm. (We can roughly measure this numerically by counting the number of lines in the algorithm,

[1] Really, the party-list approach is not a multiwinner election system at all, but rather, a pathetic admission of defeat in trying to design such a system – and then, after admitting defeat, those designers "pass the buck" by just having the parties choose the winners instead of the election system itself.

[2] This oversimplifies. Procedures for choosing and ordering the list vary in different countries and different parties. But as far as I know, the members of party B usually do not get to offer input into the list chosen by party A, and in some cases, only party insiders and not their rank and file, have any input into such decisions. For example, in 2003 the British Labour party expelled George Galloway for attacking Prime Minister Tony Blair over the Iraq war. Galloway then was only able to regain a parliament seat by forming his own political party. Some countries, e.g. Belgium, permit voting *both* for a party and for some scattered individuals, but that can sometimes have unintended effects. Another common problem with party list systems is officials "permanently in office."

[3] R.Loring tells me the following explicit historical example of how a voting system can "zero out" a minority. In North Carolina, the plurality rules effectively deny representation to African-Americans. They have enough voters to totally fill 2 election districts. However, they are a 25% minority scattered over 8 districts. So for 100 years they won no federal representation and many felt invisible as voters.

with all algorithms coded in the same computer language with a similar style.)

PB: Polynomially bounded runtime. The algorithm should run quickly with worst case input – preferably in a number of steps upper bounded by a *polynomial* in V and C.

Because computers can handle large brute force computations, the falsity of PB is not necessarily an insuperable obstacle. For example a computer could easily handle a $C = 25$-candidate election even using a 2^C-step algorithm, and indeed brute force algorithms often have the advantage of being simpler to describe. *But*, occasional very large elections might be infeasible for brute force to handle. For example, the 2003 California gubernatorial recall election had 135 candidates. A brute force algorithm which worked by considering, say, every possible 60-element "winner" subset of 135 candidates, would be completely beyond the reach even of all computers in the world running for centuries. Meanwhile a polynomial-time algorithm such as RRV (§7.5) could easily handle 1000s of candidates and winners.

Another reason PB could be less important than it seems, is the possibility of algorithms whose *typical* running time is much faster than any *worst-case* upper bound anybody is currently capable of *proving*. For example, no worst-case runtime bound better than 2^N is known for solving the N-city "traveling salesman problem"(TSP). But branch-and-bound algorithms are available based on very *good* bounds, which in practice seem to run far faster than that. Consequently it is now routine to find provably exact solutions of TSPs with thousands of cities [22]. Similarly, although the brute force implementation of our LPV_{0+} voting system (§7.9) requires examination of $\binom{C}{W}$ possible winner-sets, in practice I would conjecture that far faster runtimes would be achieveable via branch-and-bound search techniques. I speculate that in practice, this probably would make 100-winner LPV_{0+} elections feasible.

EP: Efficiently Parallelizable. When $V \gg C \geq 2$ there is an efficient way to perform the election in which each precinct only sends some kind of *subtotal* to the central tabulating agency, i.e. a much smaller amount of information[4] than sending *every vote* cast in that precinct individually. And all the necessary communication is one-way.

STV is an example of an election method which fails the EP criterion. In STV with one-way precinct→central communication only, either every vote must be sent, or $C!$ kinds of subtotals (since only the nonzero subtotals must be sent, actually only $\min\{C!, V\}$ numbers need to be sent), or, more cleverly, it is possible to send only $\min\{2^C, V\}$ numbers by sending the weighted-vote-totals for the weightings that would arise from each possible current-winners-subset, where note there are 2^C candidate-subsets. Actually if $W \ll C/2$ then the 2^C here even may be replaced by the substantially smaller upper bound $\sum_{K\leq W}\binom{C}{K}$, but even this is still exponentially large. In STV with *two*-way communication, only small amounts of information need to be sent (weighted-vote-totals in one direction; round winners and losers in the other direction). However, two-way communication can greatly increase election costs by possibly requiring all precinct workers to continue operating until every last detail has been worked out everywhere (which in the contemporary USA often can take months). In a worst-case scenario with near-ties in many rounds and adversairies trying to interfere with communications or precinct totals, STV elections could become a complete nightmare. We also remark that STV election winners are not necessarily easily guessed – it is easily possible to construct situations in which the countrywide STV winner differs from *every* precinct winner!

Despite all these complaints, Australia and Ireland have used STV for decades[5] and evidently have not collapsed. STV seems to behave much better typically than it does in rare worst-case scenarios. And in the UK it is standard procedure to send ballots to a central counting facility rather than count them at the precincts themselves – and the UK also has not collapsed. So EP failure is not an insuperable obstacle, and in the future brave new world of electronic voting, it may become even less of an obstacle since transmitting all ballots will become a trivality.

Nevertheless, it still seems somewhat desirable for an election method to obey EP since

1. centralized counting might make an election more vulnerable to a massive centrally-organized fraud, and
2. since it is pleasant to be able to summarize precinct results via some kind of subtotal.

RV: Reversal symmetry. If in a W-winner C-candidate election, each voter were to *reverse* the preference ordering of the candidates expressed in his vote, then in a $(C-W)$-winner election using the reversed ballots, the complement winner-set would be elected.

The RV criterion "feels good," but seems of no practical importance whatever.

4 Gender and Racial proportionality

New Zealand and Fiji, and perhaps other countries also, have racial quotas in various parts of their goverment in which certain seats are reserved for members of appropriate racial groups. I never liked that idea (one problem is *defining* somebody's race).

A more interesting idea, which is used by the Socialist Party of the USA for its internal elections, is to demand an exactly 50-50 split between men and women in all major positions (permitted to be off by one for certain positions of odd cardinality). For example there is a male co-chair and a female co-chair, plus exactly six male and six female members of the National Committee. The whole idea apparently was invented by Susan Dorazio.[6]

This effect is simply obtained by simply having two elections – one for the women's seats and one for the men's. I witnessed some SP-USA elections and the whole idea seemed to work quite well (despite numerous political problems suffered by the SP-USA in *other* aspects of their politics). Gender-based quotas have the advantage that one's gender is much

[4] To be precise, an amount of information upper bounded by polynomial in C and $\log V$.
[5] And in Australia, voting is obligatory.
[6] Socialist Party member and toddler teacher at the Nonotuck Massachusetts Community School.

more clearly defined than one's race, and sex demographics are unlikely ever to change much.

But we shall not consider this sort of idea further, since it does not belong to the domain of *mathematics*.

5 Impossibility theorems for multi-winner election systems

Theorem 1 (Reversal asymmetry [14]). *No proportional and reversal-symmetric multiwinner election system exists.*

Theorem 2 (NP-completeness). *Unless P=NP, no multiwinner election system with polynomially bounded runtime can enjoy representativeness.*

Proof: Let there be $V = 3W$ voters and C candidates where $W < C$. Suppose each candidate is approved by exactly 3 voters. Then this *question*: is there a subset of W among the C candidates, so that if those W were elected, then each of them would have been approved by exactly 3 voters, and each voter approves of exactly one among the W elected candidates ("his" personal representative)? is NP-complete, because it is merely a rephrasing in voting language of the EXACT COVER BY 3-SETS problem #SP2 in [8].

Indeed, more strongly, the the MINIMUM SET COVER and MINIMUM EXACT COVER problems are known to be APX-complete [2] and hence unless P=NP one cannot even *approximate* the minimum-cardinality set-cover to within a constant factor in polynomial time – indeed even more strongly it is known that a logarithmic approximation factor is best possible – and therefore one cannot approximate the minimum possible number of unrepresented voters to within an $O(\log V)$ factor in polynomial time. Q.E.D.

Theorem 3 (Fair$\Longrightarrow$ ¬Hyperproportional). *Hyperproportionality is not achievable by any "conventional" and "fair" multiwinner voting system in which the votes, and nothing else, determine the winner set. However hyperproportionality is achieved by "asset voting" [14] (an unconventional system in which the candidates, not just the voters, play a role) and in a weakened sense by party list systems (in which political parties, not just the voters, play a role).*

Proof: There are a number of subleties concerning hyperproportionality which perhaps we should have explained back when we defined it, but will do now.

Hyperproportionality *cannot be defined in a vacuum* – it can only be defined in the presence of *known party structures.* To see that, consider the following scenario. The voters fall into C equinumerous disjoint camps, each approving of exactly one candidate, and all the C candidates are in different camps, and $W = C/2$. In that case any W-element winner-set is as good as any other. Hence the "party" consisting of the $C/2$ non-winners would be totally shut out. In other words, in this scenario in the absence of knowledge of which $C/2$ of the candidates form one of the "political parties," it is simply impossible for any voting system to assure hyperproportionality, but in the presence of such knowledge, in this example achieving hyperproportionality is trivial.

Asset voting (§7.2), would, however, understand party membership if the candidates themselves understood it; while party list systems explicitly incorporate party membership information into the election system. Q.E.D.

Caveat: In §7.7 we shall discuss a way to (almost) evade theorem 3.

Theorem 4 (Representative$\Longrightarrow$ ¬Hyperproportional). *Hyperproportionality and representativeness are incompatible – no voting system can enjoy both properties.*

Proof: See footnote 17 to election example 8.6. Q.E.D.

Theorem 5 (Representative$\Longrightarrow$ ¬Proportional). *Proportionality and representativeness are incompatible – no voting system can enjoy both properties – at least with appropriate definitions of voter "approval" of candidates.*

Proof: Suppose there are 10^5 voters split approximately 50-50 into "red" and "blue" camps. This by proportionality forces there to be approximately 50% red and 50% blue winners. However, now let there be one each green, orange, white, and black candidate and one each green, orange, white, and black voter. All voters "approve" only of their-color candidate. In that case the only way to get a 6-winner set enjoying representativity is to have one winner of each color, but that would be a serious proportionality failure. Q.E.D.

Another construction: Suppose 99 of the 100 voters top-rank A and bottom-rank B – but the single remaining voter bottom-ranks A. In that case if bottom-ranked candidates are regarded as "disapproved," the only way to get a representative single winner is for it to be C. But this is highly disproportional. Q.E.D.

But we shall see later that the LPV_{0+} voting system enjoys both properties, i.e. evades this impossibility theorem. It accomplishes that by using a different definition of "approval," namely "receives a positive real number as vote."

6 Tabular comparison of properties of voting systems

system \ property	FA	MN	PR	HY	PB	EP	RE*	RV*	DC	#	$\#_W$
Asset[7]	yes	yes[8]	yes	yes	yes	yes	no	no	4	6	23
RRV	yes	yes	yes	no	yes	no	no	no	9	4	17
Birational	yes	yes	iff $C \le 2$	no	no	no	no	no	5	2 to 3	10 to 15
LPV_{0+}	yes	yes	yes	no	no	no	yes	no	5	3.5	16
Hare/Droop STV[9]	yes	no	yes	no	yes	no	no	no	20	3	12
BTR-STV	yes	no	yes	no	yes	no	no	no	22	3	12
W-votes	yes	yes	no	no	yes	yes	no	yes	3	4.5	15
mw-plurality(SNTV)	yes	yes	no	no	yes	yes	no	no	2	4	14
mw-range	yes	yes	no	no	yes	yes	no	yes	2	4.5	15
mw-approval	yes	yes	no	no	yes	yes	no	yes	2	4.5	15
mw-Borda	yes	yes	no	no	yes	yes	no	yes	3	4.5	15
mw-Condorcet	yes	yes	no	no	no	yes	no	yes	10	3.5	13
party lists	no	yes	yes[10]	yes	yes	yes	no	no	5	5	18
?	yes	yes	yes	no	yes	yes	no	no	?	5	19
(propty weight)	5	5	5	4	2	2	1	1	0		

Figure 6.1. Properties of multiwinner election systems. (Properties listed in decreasing order of subjective practical importance.) The systems in the top part of the table are new to the political science literature. The "mw-" systems are naive conversions of well known single-winner voting systems to multiwinner and are not advocated. In the "W-votes" system each voter names W candidates as his vote, and the W $(1 \le W < C)$ winning candidates are the most-named ones. This is very much like mw-approval except more complicated because of the W-cardinality constraint. The hypothetical "?" voting system may or may not exist (that is an open question; see the Conclusion) and in view of our impossibility theorems would be best possible for a "conventional" voting system. DC is approximate descriptive complexity, as a line-count, and is rather subjective. # is the number of properties the system satisfies, with *-properties (which seem less important) each counting only 1/2; meanwhile $\#_W$ is a *weighted* sum over properties, using the weights on the bottom line. Based on property count (#), asset voting is the best system here, and if asset's low DC is also taken into account, and/or if the reduced-strength sense in which party list satisfies PR and HY is taken into account, and/or if the reader agrees that the first six properties (or some subset of them) are more important than the others, then asset's lead widens. ▲

Based on these properties, the author's own "asset voting" system seems to be the best system among the 10 tabulated here.

7 Annotated descriptions of the systems

7.1 Naive multiwinner versions of common single-winner schemes

In the "plurality" voting system, each voter names a candidate as his vote, and the most-named candidate wins. One of course could convert this into a multiwinner system ("mw-plurality") by making the most-named, second-most-named,..., Wth-most-named candidates all win.

Similarly we could define "mw-approval voting": Each voter provides a list of the candidates he "approves" of, as his vote. The 1-Wth most approved candidates win.

Similarly we could define "mw-range voting": Each voter provides a numerical score from 0-99 for each candidate, as his vote. The candidates with the 1-Wth highest average scores, win.

Similarly we could define "mw-Condorcet voting": Each voter provides a rank ordering of the C candidates as his vote. (There are $C!$ possible votes.) Now draw the C-vertex directed graph where an arrow is drawn from candidate A to candidate B if A is preferred to B by the electorate (based on the votes with all candidates other than A and B ignored) and label that arrow with A's margin of victory in this pairwise election. Now select as the W winners the W candidates such that all arrows joining a winner and non-winner candidate point away from the winner – or if that is impossible, then such that the sum of the numbers ("penalties") on the wrong-way arrows is minimum possible.[11]

The most obvious problem with all of the above "naive conversions" of a single winner to multiwinner election system is that they yield extreme failures of proportionality. Namely: all of the above mw-systems can easily elect 100% Tory winners in a 2-camp party-line-vote Tory-Whig election with 51% Tory 49% Whig electorate.

We could define "mw-Borda voting": Each voter provides a score of $C + 1 - k$ to his kth-most favored candidate, for all $k = 1, 2, \ldots, C$, as his vote. The candidates with the 1-Wth highest total scores, win. This would also exhibit the same massive proportionality failure provided (1) the Whig and Tory candidates were equinumerous, (2) each Tory voter

[7] Asset is an unconventional "voting system" in which the candidates play an active, not merely a passive, role.

[8] Asset is "monotonic" *before* its "negotiation stage" *but* certain pre-declared negotiation strategies can introduce non-monotonic effects. See §8.

[9] Meek STV [9] has the same properties as Hare/Droop STV but even larger DC. Tideman & Richardson's "CPO-STV" [20] is even more complicated still, requires massive computer usage to run an election, and is extremely hard to describe.

[10] Only "proportional" if the political parties are the "disjoint camps"; this is a weaker notion of proportionality than the other proportional systems satisfy.

[11] This W-winner selection task unfortunately is easily shown to be NP-complete.

could pair himself with a Whig voter who ordered all the candidates in reverse order. Since this pathology seems unlikely in practice, it seems plausible mw-Borda would in practice behave somewhat better in this respect than the preceding mw-systems. However, mw-Borda would be massively vulnerable to "teaming"; the Whigs could assure a huge victory simply by running an enormous number of candidates.

For these reasons I consider all of these mw-systems to be unacceptable. Moral: naively converting a single-winner election system to become multiwinner, usually is a bad idea because the resulting system usually suffers from massive proportionality failures in common situations.[12] Nevertheless mw-plurality has been used in various governments. Japan used mw-plurality to elect district representatives and thus, not surprisingly, suffered under unbroken single-party rule from 1954 to 1994. This lesson was not learned when the newly created Afghan and Iraq constitutions were written in the early 2000s; they also both employ mw-plurality. See figure 7.1 p.68 of [18] for a plot demonstrating the disproportional nature of the Japanese voting system during 1963-1908; it consistently gave a greater fraction of the seats to the LDP than their fraction of the votes, about equal seats and votes fractions to the J.Socialist party, and fewer seats than votes to all the other parties. The Japanese eventually became so disgusted with LDP corruption that it fell from power during a brief interlude from 1994-1996, and then the first priority of the reformers was to change the election system [5]. The reformers evidently perceived the disproportional nature of mw-plurality since the new system was a hybrid of single-winner plurality districts and multi-winner districts that employ a party-list PR system (replacing mw-plurality). However, the net effect still was not enough to destroy single-party domination and the LDP resumed its stranglehold from 1996 onward. The USA presidential and vice-presidential election was originally a 2-winner mw-plurality election, but this was later abandoned in favor of a 1-winner plurality system among 2-candidate teams.

7.2 Asset voting [14]

procedure Asset-election

1. In a C-candidate asset election, each vote is a real C-vector, each entry of which is nonnegative and with all the entries summing to 1. For example a legal vote would be $(0.4, 0.3, 0, 0.3)$ in a 4-candidate election, since $0.4 + 0.3 + 0.3 = 1$.
2. Compute the sum-vector $\vec{s}$.
3. Now regard each s_n as the amount of an "asset" now owned by candidate n. The candidates now negotiate; any subset of them may redistribute their assets among themselves.
4. After all negotiations and redistributions end, the W wealthiest candidates win.

Asset voting is *proportional* and indeed *hyperproportional* in the sense that any coalition of candidates which has m quotas worth of assets under its control (where a "quota" means any fraction $> [W+1]^{-1}$ of the total assets) can by cooperating easily assure that at least m of its members will win seats – regardless of what the other candidates do.

Indeed, asset voting seems to enjoy a far *finer* degree of proportionality than any conventional voting system. Namely even candidates with too few asset-votes to be elected, still can have power and influence at the negotiating table, by donating (and/or threatening to donate) those assets. Since their assets will eventually be donated, none of them will go to waste, i.e. *no* voter's vote, even one cast for a no-hoper, is "wasted." Similarly, if a candidate acquires *unnecessarily many* asset-votes, those votes too need not get wasted.

And these advantages are achieved with an incredibly simple algorithm.

7.3 Hare/Droop STV (single transferable vote with reweightings)

Here is a pseudocode description of the Hare/Droop STV procedure [19][12][10][7]. Let there be C candidates, from whom V voters are to choose W winners ($0 < W < C$, $0 < V$).

procedure STV-election

1: Obtain from each voter a preference ordering (permutation) of the C candidates;
2: Associate each vote with a real "weight" w with $0 \le w \le 1$, where initially all weights are 1;
3: Compute the "Droop Quota" $Q = \lfloor V/(W+1) \rfloor + 1$;
4: **loop**
5:   **repeat**
6:     **for** $c = 1$ to C **do**
7:       Compute F_c, the sum, over all votes ranking candidate c first, of that vote's weight;
8:     **end for**
9:     $g = \operatorname{argmax} F_c$;
10:     {g is the "good" canddt with the most 1st-place votes}
11:     **if** $F_g \ge Q$ **then**
12:       Multiply the weight of each vote which ranks g first, by $(F_g - Q)/F_g$;
13:       Declare g to be a "winner" and eliminate g from all preference orderings;
14:     **end if**
15:     **exitwhen** W canddts have been declared winners;
16:   **until** $F_g < Q$
17:   $b = \operatorname{argmin} F_c$;
18:   {b is the "bad" canddt with fewest 1st-place votes}
19:   Declare b to be a "loser" and eliminate b from all preference orderings;
20: **end loop**

This is a fairly complicated procedure. Many variants of it, both less and more complicated, also exist. One of the best seems to be Meek's weighting scheme [9],[13] but it is much more complicated, e.g. involving a nonlinear multivariable iteration to convergence.

To explain the concepts inside Hare/Droop STV in English: there are two things going on: *elimination* of loser-candidates top-ranked by the fewest voters (in step 19), and *winner declarations* for candidates top-ranked by enough voters (exceeding

[12] One could similarly define mw-Nanson and mw-Woodall-DAC voting [23] (see also [17]), and STV without vote-reweighting (formerly used by Australia, before it adopted Hare/Droop reweighting in 1949) and observe that they too suffer from massive proportionality failures.

[13] According to [21], Meek-STV may now be used in some New Zealand elections.

the "Droop quota"). After either move, the winner or loser is eliminated from all votes. Votes then are *reweighted* (step 12) so that anybody who just voted for a winner will have smaller vote-weight in the next round.

The most important theorem about `STV-election` is the **proportionality theorem** [19][14]: This postulates that the voters and candidates consist of several disjoint types of people, and each voter of a given type always prefers each candidate of his same type, above every candidate of any other type. Let the number of voters of type t be V_t, the number of candidates of type t be C_t, and the number of winners of type t be W_t. Then: $W_t \ge \lfloor WV_t/Q \rfloor$ if $C_t \ge \lfloor CV_t/Q \rfloor$. The reweighting and quota formulas were carefully designed to force proportionality.

Unfortunately STV is nonmonotonic – top-ranking a candidate can actually cause him to lose – and can encourage bizarre dishonest voting strategies.[14]

7.4 BTR-STV, Nanson STV, and other STV variants

IRV, or "instant runoff voting" is the single-winner special case of STV voting: votes are candidate-orderings, the candidate with fewest top-rank votes is eliminated, and we continue eliminations until only a single candidate (the winner) remains. "BTR-IRV" (BTR stands for "bottom two ranks" and its advocates like to pronounce it "better") is the variant where the bottom *two* candidates (with the fewest top-rankings by voters) are considered, and whichever pairwise loses to the other (based on the votes with all the other $C - 2$ candidates erased) is eliminated. BTR-IRV was invented by Rob LeGrand. It has the advantage that it always elects a Condorcet ("beats all") winner if one exists. Jan Kok then suggested creating BTR-STV which is the same as Hare/Droop STV except that in line 19 the eliminated candidate again is the pairwise loser among the bottom two.

Fact: BTR-STV obeys the same proportionality theorem, with the same proof [14], as ordinary Hare/Droop STV. Not only that, but it is possible to eliminate *any* candidate (according to *any* rule, even crazy-sounding ones such as "eliminate the candidate with the *most* top-rankings") in line 19 and the proportionality theorem and its proof [14] still will work! That is because it is the reweightings, transfers, and Droop quotas that force proportionality; the elimination rule has no effect on it aside from forcing the whole procedure to terminate.

Hence there are an infinite number of variant-STV voting systems, all of which obey the same proportionality theorem. It is entirely unclear which among these infinity is the "best" but BTR and ordinary STV both seem plausible contenders. Another might be "Nanson-STV" which instead eliminates the candidate with the lowest Borda-score in line 19. In the single-winner case this reduces to Nanson's voting system, which, like BTR-IRV, always elects a Condorcet winner if one exists.

It is not difficult to check that Nanson, BTR, and regular STV all obey exactly the same subset of our properties. All have "no" entries in their row of table 6.1 in each position where Hare/Droop STV's row contains "no"s. (To see that note that the only two difficult cases, RE and RV failure, follow from proportionality and the impossibility theorems and proofs in §5.) Therefore, at least reckoned by our properties table, none of them are superior to ordinary Hare/Droop STV. However, they might be considered superior for some other reason.

7.5 RRV: Reweighted Range Voting [16]

Let there be C candidates from whom V voters are to select W winners, $0 < W < C$, $0 < V$.

procedure Reweighted-Range-Vote
1: Each voter k supplies a C-vector $\vec{x}_k$ as his vote, each entry of which is a real number in $[0, 1]$. The cth entry of this vector expresses that voter's opinion of candidate c (i.e. 1=great, 0.5=middling, 0=terrible);
2: Each C-vector vote has associated with it, a "weight" $w_k \in [0, 1]$.
3: **for** $r = 1$ to W **do**
4: **for** $k = 1$ to V **do**
5: Let the weight of vote k be $w_k = 1/(X+1)$, where the sum of vote $\vec{x}_k$'s winner-entries is X. (Thus, initially, there are no winners and all weights are 1.)
6: **end for**
7: Compute the weighted-vote-sum vector $\vec{s} = \sum_{k=1}^{V} w_k \vec{x}_k$ (actually, this step would be best programmed as combined into step 5, but we have written it separately to enhance clarity);
8: The candidate C with the largest $\vec{s}$-entry (among candidates who have not yet been declared "winners") is declared to be the rth winner.
9: **end for**

In the 1-winner case, RRV reduces to range voting [13], i.e., it simply adds up all the vote vectors $\vec{s} = \sum_{k=1}^{V} \vec{x}_k$ and then declares the winner to be the index of the largest entry in $\vec{s}$. The first RRV winner in fact is always the same as the range-voting winner, but the second RRV winner is not necessarily the same as the candidate range voting would say was in second place. That is because the reweightings cause the supporters of the first winner to have diminished influence on the choice of the second.

The reweighting formula again was carefully designed [16] to force proportionality. The point is that voters who have already supported many previous winners heavily, will have diminished vote-weights.

It can also confer additional benefits of "encouraging voter honesty." For example, a voter has incentive not to exaggerate his high opinion of some candidate Y too greatly, because then (if Y wins) that exaggeration will decrease the weight of the voter's vote in later RRV rounds.

Later in the examples (§8), we shall see some other ways RRV can encourage voter honesty. From an unsystematic perusal of these examples I got the impression that in multiwinner elections voters are typically more motivated to be more honest than they are in single-winner elections, i.e. honesty and strategy seem to coincide to a much greater extent in the multiwinner case, both in RRV, LPV_{0+}, and asset voting. *But* in STV, one gains the opposite impression – strategic voters

[14] One strategy known to occur among Australian STV voters is, instead of ranking their favorite candidate (whom they know to be popular) top, voters intentionally rank a less-popular less-favorite candidate top to try to prevent him from being eliminated.

tend to be comparatively honest in *single*-winner STV elections but feel less forced to be honest in multiwinner ones.

7.6 An argument suggesting the superiority of RRV over every STV variant

It was argued in [16] that RRV *obsoleted* STV; the latter seemed inferior for all elections except perhaps small hand-counted ones where STV might have the advantage of being easier to count. The superiority of RRV over STV variants also is suggested by our properties table 6.1.

We now offer a different such argument. It is not a rigorous proof of anything, but it is highly suggestive. The argument lives inside the following **mathematical model:**

Model all candidates and voters as being "colored" from one of a finite number of colors. (Example: Joe is "blue.") In addition, each candidate has a numerical utility for each voter. (Example: Mary's election has utility 2.354 for Joe.)

Voter behavior model: First priority is, each voter prefers each candidate of his own color to every candidate of any other color, and says so in his vote. Second, aside from that, voters may order candidates in any way they please (including dishonest-strategically) with the goal of electing a subset with highest utility-sum.

Reasoning: Now, suppose all voters know the colors of all candidates, and know the proportions of each color in the electorate. Therefore, in any Hare/Droop STV-type proportional representation multiwinner voting system (regardless of which "elimination rule" is used in line 19) they know exactly what the color composition of the winner set is going to be. The only question is who the winning members of each color-group are going to be.

Each "round" of the STV election we either elect a new winner (noticing he is above-quota) or eliminate somebody. Call each run of $r \geq 0$ successive elimination-rounds followed by an elect-winner round, a "cycle." If a cycle elects more than one winner, regard that as more than one cycle *but* some of the cycles contain $r = 0$ elimination rounds. Under this definition each cycle elects exactly *one* additional winner.

Each cycle (including cycles with $r = 0$) is a *single-winner election* using weighted votes. The voter-weights are fixed and known at the beginning of the cycle and do not change during the cycle.

Now here is the thing. With STV, this single-winner election is just an IRV election (with the asterisks that the voters are pre-weighted and that we terminate "early" i.e. once some candidate reaches quota, which may be well be *before* all but one candidate is eliminated). With BTR-STV, it is a BTR-IRV election. With other STV variants it is other IRV variants. With RRV, it is a range voting election.

Conclusion: Assuming we believe Range Voting is a better single winner method than any particular IRV repeated elimination variant, for either honest or for strategic voters (this conclusion is supported by computer simulations measuring "Bayesian Regret" [13]) – and disregarding inter-cycle interactions and "early winners" altering the IRV-variant's rules – we conclude RRV is a better multiwinner voting method because each cycle is handled better by it as measurd by expected-utility-sum.

So in this mathematical model this "proves," cycle by cycle, that RRV is a better multiwinner voting system than any STV variant, provided we start with the belief that range voting is better than any IRV variant as a single-winner election (and disregard certain effects).

7.7 Methods based on measures of candidate "similarity" – a bad idea

It has been suggested that impossibility theorem 3 might be evaded by making voters additionally provide an assessment of candidate "similarities," thus enabling the election system to know which pairs of candidates are very similar "clones." However, with C candidates there are a quadratically-large number $(C-1)C/2$ of candidate pairs, which seems an unacceptably large amount of information to demand of voters, especially in the occasional large election such as the 2003 California Governor-recall election with $C = 135$ and $(C-1)C/2 = 9045$. It only seems reasonable to demand at most a linear amount of information. (Also, any such request for similarity information would be very vulnerable to strategic lying by voters.)

In response to these objections, it was proposed that the election system could *deduce* similarity information from the ballots without explicitly asking for it: The percentage of times other candidates Z lie between X and Y in preference-ordering votes could be used to assess the "similarity" of candidates X and Y.

Unfortunately election methods based on between-count similarity measures will be manipulable by candidate cloning, resulting in deleterious "vote splitting" and/or its opposite ("teaming") effects. That would defeat the original purpose of hyperproportionality, which was supposed to be to get clone-immunity. E.g.: If Z gets cloned, then X and Y wrongly are reckoned as "less similar." Or if X and Y are clones (i.e. totally similar) then they will correctly be reckoned as totally similar by this approach – *except* if V, W, X, Y *all* are clones then this approach will output a different view of the similarity of X and Y (assuming voters order clones randomly).

But there is a way around the cloning objection: the *correlation* between coordinates j and k in a set of range-vector (or approval-vector) votes can be used as a measure of the similarity of candidates j and k, and this measure is *immune*[15] to manipulation through candidate cloning.

That all suggests the following general approach to performing a multiwinner election.

procedure Similarity-based-election
1: Acquire range-vector style votes (C-vectors with all coordinates in the real interval $[0, 1]$) from each voter in a C-candidate election. Approval-vector style votes (C-vectors with all coordinates in the 2-element set $\{0, 1\}$) also could be used.
2: For each pair ij with $1 \leq i < j \leq C$ find the (centered) correlation between the votes for candidate i and those for candidate j.

[15] At least, with honest voters – conceivably strategic voters would be motivated to vote differently for two clones.

3: Use some kind of clustering algorithm to somehow group the candidates into disjoint "camps" where candidates in the same camp tend to have high positive correlation.
4: Compute some measure of the total voter support for each camp.
5: Elect winners via some method (such as d'Hondt) which assures each camp gets a number of winners proportional to its total voter support.

There are two reasons not to recommend this sort of approach. First, *theoretically* speaking, this approach simply *fails* to give us hyperproportionality. The Whig-Tory counterexample 8.6 given at the end of §8 shows that: the votes for the Tory candidates have exactly zero correlation – the same as the correlation between Tories and Whigs – so that this similarity measure is simply incapable of detecting "parties."

Second, to understand why we believe this kind of approach would be unacceptable *in practice*, let us try to carry it out using two *real-world* data sets from the 2004 US presidential election. Consider the covariance data in tables 7.1 (a) and (b).

(a)	Bush	Kerry	Nader	Badnrk	Cobb	Pertka	Calero
Bush	**1357, 1614**						
Kerry	**–1299, –983**	**1332, 1599**					
Nader	–259, 46	–130, 170	**767, 1095**				
Badnrk	–81, 106	–150, 13	**76, 266**	**196, 430**			
Cobb	–209, 6	–14, 201	**166, 489**	**103, 265**	**264, 639**		
Pertka	**48, 253**	**–326, –92**	–25, 131	**35, 139**	**16, 111**	**164, 491**	
Calero	–22, 85	–97, 17	**41, 139**	**60, 175**	**76, 184**	**10, 110**	**55, 172**

Figure 7.1. Covariances of votes for candidate-pairs. (All data from [15].) That is, if x is a vote for one candidate (whose mean vote is $\overline{x}$) and y is the same voter's vote for another candidate (whose mean is $\overline{y}$) then their covariance is the average value of $(x-\overline{x})(y-\overline{y})$. Each covariance value is written as a *pair* a,b, meaning that with $\geq 90\%$ probability, its true value is $\geq a$, and with $\geq 90\%$ probability, its true value is $\leq b$. Ranges a,b with both a and b having the same sign, are in **bold font**. These covariances are for the combined DSG∪WDS 122-range-vote dataset. Below table same, except from JNQ's 656-approval-vote set:

(b)	Bush	Kerry	Nader	Badnrk	Cobb	Pertka
Bush	**2314, 2426**					
Kerry	**–2366, –2243**	**2327, 2435**				
Nader	**–331, –134**	**127, 325**	**1515, 1758**			
Badnrk	–13, 28	–27, 14	–8, 46	**31, 96**		
Cobb	**–109, –50**	**34, 96**	**37, 135**	–2, 31	**125, 277**	
Pertka	–16, 38	**–74, –8**	–15, 40	–1, 32	–4, 0	**31, 155**

The *correlation* between two candidates X and Y is given by correl $= C_{XY}/\sqrt{C_{XX}C_{YY}}$ where C_{XY} is the *covariance.*

Given this data, what would our similarity-based election procedure do?

If our criterion that candidates be in the "same camp" is that they have covariance$\geq$ 0, then the range-vote data tells us that there are two camps, namely

$$\text{Bush}\cup\text{Peroutka}\cup\text{Nader}\cup\text{Cobb}\cup\text{Calero}\cup\text{Badnarik} \quad \text{versus} \quad \text{Kerry} \tag{1}$$

which would cause Bush to be the clear winner once he teamed up with all those others in "his camp." However, in fact, Cobb (and the vast majority of his supporters) would almost certainly support Kerry rather than Bush. (See [15] for sketches of the political stances of all the candidates.) That is probably also true of Nader and Calero. So if this were the election system, Cobb supporters would quite justifiably feel robbed and that their votes had been "hijacked" by Bush, and Cobb himself would doubtless step forward to protest the unfair election system.

Now suppose instead that our criterion that candidates be in the "same camp" is that they have covariance$>$ 0. Or that they have fairly large correlation. In either case, no self-consistent camp-assignment *exists*, because, e.g. Cobb and Badnarik have large correlation, and Cobb and Nader have large correlation, but Nader and Badnarik do not (according to the range-vote based covariances).

If instead we employ the approval data, then if "large correlation" is the criterion, then all candidates is in their own isolated camp, with even Nader and Cobb not in the same camp. If instead covariance$>$ 0 is the criterion then the 4 camps would be

$$\text{Bush}, \quad \text{Kerry}\cup\text{Nader}\cup\text{Cobb}, \quad \text{Peroutka}, \quad \text{and} \quad \text{Badnarik} \tag{2}$$

and this assignment actually makes some amount of sense based on their political stances. But note that Approval and Range give different camp-results, which does not inspire great confidence in the whole similarity-camp idea.

Conclusion: Correlation-based similarity measures do not actually seem to work very well. Any election system based on them would probably come under heavy criticism if it were used in practice because whatever clustering method was used would make a lot of arbitrary-seeming decisions – in many cases decisions that the affected candidates themselves would publicly disagree with – that would have huge effects. Voters might try to compensate with a lot of dishonest and damaging "strategic voting." (And similarity measures based on "between counts" would probably work even worse.)

7.8 The birational system

We shall now describe a new voting system which arose from Smith misunderstanding what Simmons did. We call it the "birational" system. The birational system plays an important inspirational role as a mental stepping stone towards LPV_{0+}.

procedure Birational-voting

1: Each voter provides a range-vector vote, awarding each candidate a real number score in $[0, 1]$.

2: For each subset W of the candidate-set C, define the "lagrangian"

$$L(W) = \sum_{\substack{\text{vote vectors}\\ \vec{x}}} \sum_{w \in W} \sum_{s \in W} \frac{x_w}{1 + x_s} \tag{3}$$

3: Elect the winner-set W (from among all possible winner-subsets of the right cardinality) with maximum $L(W)$ value.

Note, here we have abused notation by using W and C to denote the *sets* of candidates and winners, not (as before) their cardinalities.

The brilliant property of the birational Lagrangian is that it causes 2-camp party-line-vote elections to satisfy proportionality. However, unfortunately, proportionality fails if there are ≥ 3 camps.

Theorem 6 (Birational proportionality and lack thereof). *In a scenario where all voters vote party-line approval style votes, (i.e. 1 for candidates in their party, 0 for all other candidates), and*

1. *assuming enough candidates from each party are running so that there is no proportionality failure due merely to running out of candidates,*
2. *and assuming there are exactly two parties,*

then the winner-set W which maximizes $L(W)$ will have the same proportions of winners from the two parties, as there are voters from the two parties. (I.e. "proportionality.") However, this theorem becomes false *if there are three or more parties.*

Proof: Let f_j be the fraction of the voters in party j. Let q_j be the fraction of the winners from party j. Then the q_j that happen are the ones that maximize

$$\sum_{j \in C} f_j q_j \left[\frac{q_j}{2} + 1 - q_j\right] \tag{4}$$

subject to $\sum_j q_j = 1$ and $q_j \geq 0$ for all j. When we solve this optimization problem by Lagrange multiplier technique we get that at the optimum, the q_j must obey

$$(1 - q_j) f_j = (\text{constant independent of } j). \tag{5}$$

If proportionality held then this would be true if $q_j = f_j$. However, $(1 - f_j) f_j$ is *not* a constant independent of j, in general, if there are ≥ 3 parties.

But if there are 2 parties, then this is just $f_1(1 - f_1) = f_2(1 - f_2) = f_1 f_2$ since $f_1 + f_2 = 1$, which *is* the same as $f_2 f_1$. And it is not hard to see that the unique optimum (ignoring round-to-integer effects) is $q = f_1$. E.g. replace $q_1 = q$ and $q_2 = 1 - q$ in EQ 4 to get

$$q f_1 \left[\frac{q}{2} + 1 - q\right] + (1 - q) f_2 \left[\frac{1 - q}{2} + q\right] \tag{6}$$

and optimize over q. Using ordinary 1-variable calculus, we set the q-derivative to 0 to find $q = f_1/(f_1 + f_2) = f_1$ as the location of the unique extremum with $0 < q < 1$; then we recognize that this is a *maximum* because the second derivative there is $-f_1 - f_2 = -1 < 0$. At this maximum, the function value is $(1 - f_1 f_2)/2$, which note is greater than the function values $f_2/2 = (1 - f_1)/2$ and $f_1/2 = (1 - f_2/2)$ at the respective endpoints $q = 0$ and $q = 1$ of the allowed q range. Therefore, we know this is the *global* maximum. Q.E.D.

Theorem 7 (birational monotonicity). *The birational voting system is monotonic.*

Proof: Follows from the fact that $x/(1 + x)$ is a monotone increasing function of $x > 0$ while $1/(x + 1)$ is monotone decreasing; hence increasing some x_w in a vote $\vec{x}$ cannot decrease w's chances of being in the winner set. Q.E.D.

Theorem 8 (NS is non-representative). *The birational voting system does not enjoy representativeness.*

Proof sketch. In an enormous election with a clear winner-set, adding just one more voter who does not approve of any of the winners, is not going to be enough to change the result since it will only cause a bounded additive change in the $L(W)$ values. Q.E.D.

7.9 LPV_κ: Logarithmic Penalty voting

If W is a subset of the full set C of candidates, define the "logarithmic penalty function"

$$L_\kappa(W) = \sum_{\substack{\text{vote vectors}\\ \vec{x}}} \sum_{j \in C} x_j \ln \frac{\kappa + |W|}{\kappa + \sum_{s \in W} x_s} \tag{7}$$

where we use the notation $|W|$ =cardinality(W) and where $\kappa > 0$ is a tunable constant parameter. (Incidentally, this expression looks much like Shannon's "entropy." Coincidence?)

This leads to the following new proportional voting system for multiwinner elections:

procedure Logarithmic-Penalty-Voting$_\kappa$

1: Voters submit range-vector style vector votes $\vec{x}$ in $|C|$-candidate election.

2: $L_\kappa(W)$ is computed by brute force for every possible acceptable winner-subset W of C, using EQ (7) above.

3: The winner-set is the W with the least $L_\kappa(W)$ value.

If $\kappa = 0$ then EQ 7 can divide by 0, yielding an infinite term. This can be regarded as a *good* thing in the sense that avoiding the infinity forces "representativity" – but still if no W-set exists which every voter approves (that is, gives a nonzero x_w to) then some infinite term is forced. That difficulty can be avoided by agreeing to *interpret* EQ 7 when $\kappa = 0$ as follows: after minimizing over W, take the limit as $\kappa \to 0+$. For short we shall call the resulting voting system "LPV_{0+}."

Theorem 9 (LPV_{0+} representativity). *LPV_{0+}, using the lagrangian in EQ 7, enjoys representativity in the common (over all W) limit $\epsilon \to 0+$. Indeed, LPV_{0+} will minimize the number u of unrepresented voters.*

Proof: $L(W)$ tends to infinity as $\kappa \to 0+$ asymptotically to $u|\ln\kappa|$ where u is the number of unrepresented voters, we see that LPV_{0+} with this interpretation in fact will not only assure each voter has a representative (if that is possible), but, more strongly, actually will *minimize* the number of unrepresented voters. If more than one winner-set accomplishes this minimization, then LPV_{0+} will break the tie by using the other, finite, terms inside the $L(W)$ formula. Q.E.D.

Theorem 10 (LPV_κ monotonicity). *For each $\kappa > 0$ (as well as $\kappa = 0+$) LPV_κ is a monotonic voting system.*

Proof: We assume $|W| \geq 2$. Consider the facts that the x-derivative of $x\ln(|W|/[s+x])$ is $\ln(|W|/[s+x]) - x/(s+x)$ the x-derivative of $\ln(|W|/[s+x])$ is $-1/(s+x)$, and the x-derivative of $x\ln(|W|/s)$ is $\ln(|W|/s)$. Hence

$$\frac{\partial L_\kappa(W)}{\partial x_w} = \begin{cases} \ln\frac{\kappa+|W|}{\kappa+\sum_{s\in W} x_s} - \sum_{j\in C}\frac{x_j}{\kappa+\sum_{s\in W} x_s} & \text{if } w \in W \\ \ln\frac{\kappa+|W|}{\kappa+\sum_{s\in W} x_s} & \text{if } w \notin W. \end{cases} \tag{8}$$

Increasing x_w in some vote $\vec{x}$ (while leaving all other x_s with $s \neq w$ unaltered) will increase $L_\kappa(W)$ in EQ 7 for each w-noncontaining set W by more than it increases any $L_\kappa(W)$ for a w-containing set W. Hence w's chances of being in the winner-set (minimizing $L_\kappa(W)$) cannot decrease. Q.E.D.

Theorem 11 (LPV_κ proportionality). *For any fixed $\kappa > 0$ including $\kappa \to 0+$, LPV_κ is proportional.*

Proof: (Based on an argument by F.W.Simmons.) We shall use the fact that $\ln(b/a) = \int_a^b \mathrm{d}q/q$.

Suppose (to avoid too much abstraction) that there are five political parties $P_1, P_2, \ldots, P_5$, and that their respective supporters number $V_1, V_2, \ldots V_5$, and that the total number of available seats up for grabs is $W = W_1 + W_2 + \ldots W_5 + 5\kappa$, where (for some constant of proportionality θ) it turns out that

$$V_j = (W_j + \kappa)\theta \;\text{ for }\; j = 1, 2, \ldots, 5. \tag{9}$$

In other words, W_j is party P_j's proportional fair share of the available seats in the legislature.

We need to show that among all possible allocations of the seats to various parties, this set of W_js is the one that maximizes the LPV_{0+} sum, assuming that members of each party approve all of their own (and only their own) candidates.

On this loyalty assumption the LPV_{0+} sum is

$$\sum_{j=1}^{5} \int_{W_j+\kappa}^{W+\kappa} V_j \frac{\mathrm{d}q}{q} \tag{10}$$

where

$$W_j = \sum_{s\in W} x_s^{(j)} \tag{11}$$

and $\vec{x}^{(j)}$ is the approval-style vote cast by voters in party P_j.

To see that this sum is minimal for the given number W of seats, first note that for each j, the integrand is decreasing in q, so that its greatest value is its leftmost value $V_j/(W_j+\kappa)$ which is just the constant of proportionality θ.

Now suppose that we were to increase one or more of the W_js at the expense of the others (to maintain their constant sum $W - 5\kappa$). This would replace some of the smaller-height chunks of the integral (less than θ) with an equal width chunk of larger-height (greater than θ), so the value of the LPV_{0+} sum would increase.

Although the above proof only holds if all the proportionalities are *exact*, we can easily extend it to show that at the minimum, the maximum excess of any W_j above its theoretical (generally non-integer) fair share, plus the maximum shortfall of any W_j below its fair share, does not exceed 1. (If it did, then the LPV_{0+} sum could be decreased by adding 1 to the smaller W_j and subtracting 1 from the larger, a contradiction.) Q.E.D.

8 Illustrative election examples

# voters	they approve	$\mathrm{RRV/LPV}_{0+}$	Asset vote	STV vote
32	both A and B	(1,1,0)	(1/2,1/2,0)	$A > B > C$ & $B > A > C$
28	both A and C	(1,0,1)	(1/2,0,1/2)	$A > C > B$ & $C > A > B$
20	only B	(0,1,0)	(0,1,0)	$B > A > C$ & $B > C > A$
20	only C	(0,0,1)	(0,0,1)	$C > A > B$ & $C > B > A$

Figure 8.1. Illustrative example by F.W.Simmons. 100 voters choose 2 winners from 3 candidates.

RRV: In RRV with approval-style votes ($\{0,1\}$-only), A (the most-approved candidate) wins, and then B is the second winner.

Birational: Under the birational system (using the RRV vote-tuples), A & B again win since

$$L(AB) = 32(\frac{1}{2}+\frac{1}{2})2 + 28(\frac{1}{2}+1) + 20(\frac{1}{2}+1) + 20(1+1) = 64+42+30+40 = 176$$

exceeds both

$$L(AC) = 32(\frac{1}{2}+1) + 28(\frac{1}{2}+\frac{1}{2})2 + 20(\frac{1}{2}+\frac{1}{2}) + 20(\frac{1}{2}+1) = 48+56+20+30 = 154$$

and

$$L(BC) = 32(\frac{1}{2}+1) + 28(\frac{1}{2}+1) + 20(\frac{1}{2}+1) + 20(\frac{1}{2}+1) = 48+42+30+30 = 150.$$

LPV_{0+}: But under LPV_{0+} (using the RRV vote-tuples), B and C win because

$$L(BC) = 32\ln 2 + 28\ln 2 + 20\ln 2 + 20\ln 2 = 100\ln 2$$

is less than both

$$L(AC) = 32\ln 2 + 28\ln 1 + 20\ln\frac{2}{\epsilon} + 20\ln 2 = 72\ln 2 - \ln\epsilon$$

and

$$L(AB) = 32\ln 1 + 28\ln 2 + 20\ln 2 + 20\ln\frac{2}{\epsilon} = 68\ln 2 - \ln\epsilon$$

for all ϵ with $0 < \epsilon < 2^{-32}$. Simmons argues this LPV_{0+} winner-set is superior because every voter approves of some election winner, i.e. nobody is left "unrepresented."

But it is not 100% convincing that $\{B, C\}$ *must* be the societally best 2-winner set because if A were far better than C, then society would be better off with $\{A, B\}$ since it would be worth paying the cost of leaving the C-loving voters unrepresented, in order to get A's wonderful services. While that scenario may be unlikely, it is possible and is compatible with the vote data.

Asset: Under asset voting, if A were truly superior and if the A-approving voters recognized that and hence gave the lion's share of each 2-man vote to A, then A would be a winner. However, asset voting would make B and C the winners if all the 2-man votes were evenly split as illustrated in the table.

STV: Assuming both kinds of possible STV votes are provided equinumerously, A is eliminated in the first round so that B and C win. But if A is regarded as superior by enough of those approving him, so that the first among the two STV votes in each of the first two rows predominate, then instead one of $\{B, C\}$ would be eliminated. However, in that circumstance the voters of the first two types would be highly motivated to strategically exaggerate their opinions of B and C to try to prevent their most-favored among these two from being eliminated. Depending on how much strategizing occurs, *any* of the 3 candidates could easily be the one eliminated, in which case the STV election outcome might look random. ▲

# voters	STV vote	RRV	Asset
25	$A > B > C$	(1, 0.9, 0)	(0.7, 0.3, 0)
35	$B > A > C$	(0.9, 1, 0)	(0.3, 0.7, 0)
40	$C > A > B$	(0.1, 0, 1)	(0, 0, 1)

Figure 8.2. 100 voters choose 1 winner from 3 candidates. In this example, candidates A and B are similar and also similar to 60% of the voters. The remaining 40% of the voters prefer C.

STV: A is eliminated in the first round, then B wins 60-to-40 over C. While it seems "fair" that one of $\{A, B\}$ should win the election, since 60% of the voters are in the AB camp, the C-camp's view that $A > B$ is completely disregarded, causing these voters to have zero influence. It seems wrong that B wins even though 65% of the voters think $A > B$. Had the C-voters predicted this outcome, they would have been tempted to dishonestly strategically vote $A > C > B$ instead – yielding a better outcome (A wins) from their point of view, but "betraying" their favorite C, causing him to wrongly appear to be top-ranked by nobody and ruining his future political career.

RRV=LPV_{0+}=range voting: (These three systems all are equivalent in the single-winner case. Birational voting is also equivalent assuming a pre-transformation $x \to x/(1+x)$ is made to all vote-vector extries x.) A wins.

Asset voting: The assets as we enter the negotiation round are $A = 25\times0.7+35\times0.3 = 28$, $B = 25\times0.3+35\times0.7 = 32$, $C = 40$. Assuming A and B agree they are both superior to C, then A can award his assets to B and assure B's victory. Although in that event C would be doomed to lose, C still would have "kingmaker" power: C could tell A "you do not need to give it up to B – I'll give you my assets so you can win." Of course in reality both A and B would probably try to make deals with C (even if they regard him as evil incarnate) in order to get his assets. Hence, despite losing, C could still have considerable power and influence under asset voting. ▲

# voters	STV vote	approvals	RRV
409	$\mathbf{A_1} > A_2 > A_3$	$A_1A_2A_3$	$A_1 = 1, A_2 = 0.9, A_3 = 0.8$, rest=0
1	$\mathbf{A_1}$	A_1	$A_1 = 1$, rest=0
326	$\mathbf{B_1} > \mathbf{B_2} > B_3$	$B_1B_2B_3$	$B_1 = 1, B_2 = 0.8, B_3 = 0.6$, rest=0
2	$\mathbf{B_1}$	B_1	$B_1 = 1$, rest=0
2*	$\mathbf{B_2}$	B_2	$B_2 = 1$, rest=0
259	$\mathbf{C_1} > C_2 > C_3 > \mathbf{B_1} > \mathbf{B_2} > B_3$	$C_1C_2C_3B_1B_2B_3$	$C_1 = 1, C_2 = 0.9, C_3 = 0.8, B_1 = 0.7, B_2 = 0.6, B_3 = 0.5$, rest=0
1*	$\mathbf{C_1}$	C_1	$C_1 = 1$, rest=0

Figure 8.3. 1000-voter 9-candidate 3-winner election scenario based on one by David Gamble.

STV: With STV voting in which truncated preferences are allowed, $\{A_1, B_1, C_1\}$ win in that order after elimination of B_3, C_3, C_2, A_2, A_3 and then B_2. This somehow seems "clearly the right result for society" because after elimination of the no-hopers B_3, C_3, C_2, A_2, A_3 all of whom have zero top-rank votes, we would get only the STV vote parts in bold print.

LPV_{0+} with approval-style voting: But translated into approval-style votes, under LPV_{0+} A_1, B_1, B_2 win. (Note that this leaves only the single last voter "unrepresented." STV would have left the 2 voters for B_2 unrepresented.) Judged from the STV votes, this result seems insane. This suggests that approval-style votes (and perhaps even STV-style votes) can be inadequate, in the sense that they simply do not provide enough information to be useful in multiwinner elections. Also, it is interesting and disturbing that the fate of this LPV_{0+} election arguably lies entirely in the hands only only the three fringe *-voters. If they had split 1:2 instead of 2:1 then C_1 would have been elected rather than B_2. Is this really fair? This problem is not specific to LPV_{0+}, but in fact would arise with *any* "representative" voting system. That is one sense in which the "representativeness" property has a dark side.

RRV: But RRV votes provide far more information than either approval- or STV-style votes, since in RRV, votes are vectors of real numbers in $[0, 1]$, i.e. range-vote style votes. With the range-style votes tabulated, the winner set is $\{B_1, A_1, B_2\}$, the same as LPV_{0+}/approval! B_1 is the first RRV winner (score$= 328 + 259 \times 0.7 = 509.3$ beating out A_1 with 410), then A_1 is the second winner (score=410 beating out B_2 with $326 \times 0.8/2 + 259 \times 0.6/1.7 = 221.81$) and finally B_2 is the third winner (score$= 221.81$ beating out $C_1 = 259/1.7 + 1 = 153.35$ and $A_2 = 409 \times 0.9/2 = 184.05$).

LPV_{0+} with range-style (RRV) voting: Same winner set as under approval-style voting (given the tabulated range-votes).

Asset: Assuming the same votes as under RRV except scaled to have unit sum [thus the first 409 voters each would vote $(0.333, 0.300, 0.267, 0, \ldots, 0)$] the As would get 410, the Bs 433.6, and the Cs 156.4 worth of assets. Then A_1 would surely be a winner. The Bs and Cs could by cooperating assure themselves the 2 remaining seats. One of them almost surely would be B_1, but it is not clear who would win the remaining seat, B_2 or C_1; it depends on how the negotiations proceed. The Cs carry considerable clout since they have the power to throw the remaining seat to the hated As. But the Bs have the power to win seats for *both* B_1 and B_2 *provided* the C's and As do nothing. Also, the As have enough clout that even after electing A_1, they still could either throw the third seat to the Cs or to the Bs, as they wish.

Summary: The STV view that "obviously" $\{A_1, B_1, C_1\}$ should be the 3 winners, actually maybe now looks wrong! Because the voters who wanted C_1 and the voters who wanted B_1 both also like B_2, so once A_1 and B_1 are chosen, it seems like a compromise more beneficial to all is to elect B_2 in preference to C_1.

This example illustrates the fact that it can be not at all clear who should be the winners in multiwinner elections. Some probably would say it should be C_1 and others B_2 in the winner's circle. ▲

# voters	opinions	approve
48	$A > B, C$	$A_1A_2A_3$
45	$B > A, C$	$B_1B_2B_3$
7	$C > B \gg A$	$C_1C_2C_3$

Figure 8.4. 100 voters choose 3 winners from 9 candidates from 3 parties A, B, C.

RRV: Under RRV if voters all provide votes of 1 to their 3 approved candidates (and 0 to all others), then two As and one B win seats. However, if the last 7 voters also provide 1-votes to the three B candidates, then two B and one A win seats, a better result for the party-C voters. This shows that it can be strategically desirable to approve of more than W candidates in W-winner RRV with approval-style voting.

STV: Two Bs and one A win seats. Probably the same thing will happen under **asset voting**.

LPV_{0+}: one A, one B, and one C win a seat, unless the 7 C voters foolishly approve of the Bs, in which case two Bs and one A win seats. ▲

# voters	STV vote	RRV/LPV_{0+}	Asset
24	$A > B > C = D$	(1, 0, 0, 0)	(1, 0, 0, 0)
24	$B > A > C = D$	(0, 1, 0, 0)	(0, 1, 0, 0)
52	$C = D > A = B$	(0, 0, 1, 1)	(0, 0, 0.5, 0.5)

Figure 8.5. 100 voters choose 2 winners from 4 candidates. In this example, candidates A and B are similar ("Tories") and also C and D are similar ("Whigs"). However, the Tory voters are choosier than the Whigs.

STV and LPV_{0+}: Under either system, one Tory and one Whig win. However LPV_{0+}[16] seems to perform slightly better than STV here, in the sense that STV totally ignores the Whigs' opinions about the Tories and the Tories' opinions about the Whigs, whereas LPV_{0+} (if a few voters had expressed such opinions by providing slightly nonzero vote-values) would take those opinions into account. (Thus LPV_{0+} – and RRV – can "encourage voter honesty": The Whig voters would have been better off providing a nonzero vote for their least-hated Tory, as opposed to exaggerating and pretending both Tory candidates were completely worthless.) Under STV, strategic Whig voters would therefore be tempted to dishonestly rank a Tory first (or, less effectively, second) to try to get some influence over the Tory seat. With enough voters playing strategic games of that sort, the STV election outcome might become less predictable...

RRV: A Whig (say D) wins, then there is a 3-way near-tie between A, B, and C to become the second winner, which in this example C narrowly wins. The Tory voters have paid a severe price for being choosier and are "zeroed out." This is an example of a hyperproportionality failure. Of course, if the Tories had been slightly less hardnosed and had voted like $(1, 0.2, 0, 0)$, that would have been enough to assure them a seat. That is one way in which RRV can "encourage voter honesty": Strategic Tory voters cannot go too far in exaggerating the relative virtues of A over B because then they risk paying a price.

Asset: Entering the negotiation round, the assets are $A = B = 24$, $C = D = 26$. Assuming candidates feel bound by party loyalty, one Whig and one Tory will win. But the two Whigs can use their 18.66 leftover assets (above the quota of 33.33 required to assure a seat) to play "kingmaker" and assure the election of the Tory they dislike least. Meanwhile simultaneously the two Tories can use their 16.66 leftover assets to play kingmaker and assure the election of the Whig *they* dislike least. ▲

Figure 8.6. Maximally dramatic hyperproportionality failure scenario:

We generalize the example in figure 8.5 to maximize drama, although admittedly at the cost of making it look unrealistic. Let there be n Tory candidates and n Whig candidates,

[16] And birational.

i.e. $C = 2n$ candidates total. Suppose our job is to choose $W = n + 1$ winners. Suppose the electorate consists of 50% Whigs and 50% Tories. Each Tory voter awards a vote of 1 to a single Tory candidate (whom we may regard as chosen at random independently by each voter) and 0 to all other candidates. Each Whig voter awards a vote of 1 to all n Whig candidates, and 0 to all Tory candidates.

LPV_{0+} exhibits a severe hyperproportionality failure, selecting the winner set to consist of n Tories and 1 Whig. (As it must by theorem 9 since this is the only cardinality $W = n+1$ candidate set leaving no voter unrepresented![17]) LPV_{0+} thus *benefits* the choosier Tory voters.

RRV: Meanwhile, RRV exhibits a severe hyperproportionality failure of the opposite sort: it *penalizes* the choosier Tories by selecting a winner set consisting of n Whigs and 1 Tory![18]

STV voting (assuming truncated preference orderings are allowed as votes and assuming the voting system treats remaining unmentioned candidates as effectively being ordered randomly but below the mentioned ones) would eliminate the Tory candidates one by one until only one remained; thus n Whigs and 1 Tory would be elected.

Birational voting also exhibits a dramatic pro-Whig bias.

Asset voting of course would elect 50% Whig and 50% Tory winners. ▲

Although the above examples of the effects of "choosiness" are artificial, it is easily possible to devise natural-seeming examples in which choosiness in combination with strategic "exaggerating" voting, creates problems.

# voters	STV vote
30	$A_1 > A_2 > A_3$
25	$A_1 > A_3 > A_2$
44	$B_1 > B_2 > B_3$
1	$B_2 > B_3 > B_1$

Figure 8.7. 100-voter 2-party 6-candidate election with 3 winners.

STV (truncated preferences allowed): A_1, A_2 and B_1 win.

RRV with approval-style votes: If all voters approve their top three choices then A_1, B_1, and either A_2 or A_3 win. But if the party-A supporters foolishly only approve of their top two choices (and we assume the final B-voter does the same) then that choosiness costs them: A_1, B_2, and B_3 are elected. In this example the final B-voter's decision not to approve of B_1 has singlehandedly prevented his election, contrary to the wishes of every other B-voter! ▲

How asset-voting can exhibit strange, e.g. non-monotonic, effects: In a 3-candidate single-winner election, suppose each candidate pre-announces, before the election begins: "My negotiation strategy will be as follows: If I have the fewest assets, then I will award all my assets to candidate X (just which X, depends on that candidate)" In that case the election will behave very similarly – indeed, under the right circumstances identically – to a so-called "instant runoff" (IRV[19]) election, and therefore can exhibit all sorts of "voting paradoxes" analogous to those suffered by IRV. A voter who increases his asset-vote for C at the expense of his second-favorite B, can actually cause C to *lose* by causing B to be eliminated in the first round so that B's assets transfer to A (whereas otherwise A would have been eliminated with his assets transferring to C).

This also serves as a counterexample to the idea that plurality-style "all assets to one" asset voting always is strategically best. (Many other such counterexamples also exist.)

Asset voting works better in this respect if the set of candidates is large and diverse enough so that a candidate exists who corresponds very closely to almost any voter. In that case the voter could pick a candidate whose pre-announced transfer strategy agreed with his own. Even then, funny things could still happen: it might be wise for the voter to decrease his vote for some disliked (by most) but liked (by that voter) candidate who, however, pledged to make desirable asset-transfers. It could be argued that the fact that Asset voting can "simulate" Hare/Droop-STV voting in some but not all cases, means Asset voting is "superior" to STV because if Asset could simulate STV in more cases, that would only have made Asset "worse" by allowing more nasty paradoxes. Although I feel this argument is more true than false, it certainly cannot be considered convincing. Indeed David Gamble has argued to me that in practice he suspects most asset voters wil award their entire vote to their favorite candidate, and candidates (in the negotiation stage) will transfer their entire assets in one block to other candidates. If so, Gamble argues, Asset voting will behave much like STV voting, except *worse* from the point of view of the voter, who under these conditions would not be able to control the transfers, but with STV, would.

9 Conclusion

Asset voting seems the best available system judged either from the property counts in table 6.1 or subjectively via the number of disturbing phenomena in the illustrative examples in §8. However, many people object to, or at least are uncomfortable with, asset's unconventional nature. (One said he did not consider it to be a "voting system" at all.) The prospect of mysterious "negotiations" and possibly-secret and/or corrupt deals is unsettling, and the remarks at the end of §8 show that strange phenomena could occur (whether to call them "disturbing" is a matter of opinion). I would like to see asset voting implemented in some real governments so that we can see how well it actually works. I do not think it is possible for mathematics alone to come to a clear conclusion about the "quality" (which is not even defined) of an election system involving many human beings interacting before, during, and after the election.

Both asset and party-list voting are simply unuseable in some abstract scenario where we are not electing sentient beings, but rather abstract choices such as pizza flavors.

[17] This example proves theorem 4.

[18] RRV would elect $n - 1$ Whigs, then there would be a multiway tie between all remaining candidates, and whichever party's candidate won the tie, that party would lose in the next and final RRV round.

[19] IRV is the single-winner special case of Hare/Droop STV.

If asset voting is rejected because of this, then no system clearly springs to the fore as the best among the known remaining ones. All the other systems surveyed suffer from disadvantages which can be very severe, at least in some circumstances to some people's minds.

Could there be some not-yet-discovered voting system which would save us from that? We have shown "impossibility theorems" indicating that nothing is ever going to avoid HY-failure, nothing obeying PR is ever going to obey RV, and nothing obeying RE is ever going to have a polynomial runtime bound PB[20] or obey HY. So the best we could possibly hope for from a conventional voting system would be the set of properties listed in the "?" line of table 6.1.

Open question: does such a voting system exist?

But even if it does exist, then (at least if we robotically employ the quality scores represented by the "#" columns in table 6.1) it would *still* be discernibly worse than asset voting.

Table 6.1 is somewhat misleading (as are property comparisons generally) in that it tells us a binary yes/no answer about whether a system disobeys some property, but does not tell us *how often* such failures occur nor *how serious* they are. For example, most of the PR-failures in table 6.1 are extremely severe and common, whereas I suspect the HY-failures in the top part of the table are usually not very severe and not very common. I suspect that all of the systems in the top part of the table are in fact quite good in practice.

Richard Katz [11] has argued that governments, by employing proportional representation, suffer greater extremism on the issues and can enter disfunctional "gridlock" states caused by "interparty bickering." However, asset voting would appear to encourage moderation, both because opposing parties constantly need to make deals with each other, and also since party A often has enough clout to force party B to elect the more A-like of its two candidates. Thus asset voting is a provably proportional method that may evade Katz's criticism.

Here are some **reassuring simulation arguments** that asset voting is at least as good as any party-list system.

If everybody asset-votes only for the King of the Democrats or the King of the Republicans, the net effect is to simulate a *party-list system* with each party list chosen by that party's king.

On the other hand if every Democrat voter chose to distribute his votes equally among all Democrat candidates, and ditto for Republican voters, the net effect would be to simulate a party-list system with each party list chosen by some kind of negotiation within that party in which all their candidates have equal power in that negotiation.

These simulation arguments seem a decent foundation for arguing that asset voting is at least as good a voting system as any party-list system. (But that argument is not really convincing; cf. the discussion of STV-simulation at the end of §8.) To go further and argue that asset would be not merely "at least as good" but in fact "better," we note that certain apparently-beneficial effects would happen under asset voting but could not happen in contemporary party-list systems:

1. A Democrat voter could in fact give some of his vote to those few Republicans that he likes, thus enabling them to get more power within the Republican party, etc.
2. The more-popular Democrats (electorate-wide) would obtain more power than the less popular ones within their own party, but not a total monopoly on power.

[20] Unless P=NP.

10 Acknowledgements

I would like to thank Forest W. Simmons, David Gamble, Jan Kok, James Cooper, and Abd ul-Rahman Lomax for their emails to me and/or the approval-voting yahoo email group (and other groups). The first two mentioned perhaps indeed should have been coauthors since they made substantial contributions (credited in the text), but they declined.

References

[1] M.Abramowitz & I.Stegun (eds.): Handbook of mathematical functions, Dover 1974.

[2] Pierluigi Crescenzi & Viggo Kann (editors): A compendium of NP optimization problems, online continually updated version at `www.nada.kth.se/~viggo/wwwcompendium/wwwcompendium.html`; the original version was an appendix to the book G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, M. Protasi: Complexity and Approximation, Springer 1999.

[3] J.J. Bartholdi III & J.B. Orlin: Single transferable Vote resists strategic voting, Social Choice & Welfare 8,4 (1991) 341-354.

[4] Steven J. Brams & Peter C. Fishburn: Some logical defects of the single transferable vote, Chapter 14, pp. 147-151, in Choosing an Electoral System: Issues and Alternatives (Arend Lijphart and Bernard Grofman, eds.) Praeger, New York 1984.

[5] Raymond V. Christensen: The new Japanese election system, Pacific Affairs, Spring 1996.

[6] G. Doron & R. Kronick: Single Transferable Vote: An Example of a Perverse Social Choice Function, American J. Political Science 21 (May 1997) 303-311.

[7] Henry R. Droop: On methods of electing representatives, J.Statistical Society London 44,2 (1881) 141-196; comments 197-202.

[8] M.R.Garey & D.S.Johnson: Computers and intractibility, a guide to the theory of NP-completeness, Freeman 1978.

[9] I.D.Hill, B.A.Wichmann, D.R.Woodall: Algorithm 123, single transferable vote by Meek's method, Computer Journal 30,3 (1987) 277-281. Available online at `http://www.bcs.org.uk/election/meek/meekm.htm`. Brian Meek's original work was published in French and an English translation is available in issue 1 (March 1994) of the ERS publication *Voting Matters*, available online at `www.mcdougall.org.uk/VM/MAIN.HTM`.

[10] Clarence G. Hoag & George H. Hallett: Proportional Representation, Macmillan, New York 1926; Johnson reprint corp. 1965.

[11] Richard S. Katz: A theory of parties and electoral systems, Johns Hopkins Univ. press 1980.

[12] W.J.M. Mackenzie: Free elections, George Allen & Unwin Ltd. London 1958.

[13] Warren D. Smith: Range voting, #56 at `http://math.temple.edu/~wds/homepage/works.html`.

[14] Warren D. Smith: "Asset voting" scheme for multiwinner elections, #77 at
http://math.temple.edu/∼wds/homepage/works.html.

[15] W.D. Smith, D.S.Greene, J.N.Quintal: What if the 2004 US presidential election had been held using Range or Approval voting? #82 at
http://math.temple.edu/∼wds/homepage/works.html.

[16] Warren D. Smith: Reweighted range voting – new multiwinner voting method, #78 at
http://math.temple.edu/∼wds/homepage/works.html.

[17] Warren D. Smith: Descriptions of voting systems, manuscript.

[18] R.Taagepeera & M.S.Shugart: Seats & Votes, Yale Univ. Press 1989.

[19] Nicolaus Tideman: The Single transferable Vote, J. Economic Perspectives 9,1 (1995) 27-38. (Special issue on voting methods.)

[20] Nicolaus Tideman & Daniel Richardson: Better Voting Methods Through Technology: The Refinement-Manageability Trade-Off in the Single Transferable Vote.

[21] S.W.Todd: STV in New Zealand, Voting matters 16 (Feb. 2003), www.electoral-reform.org.uk/publications/votingmatters/P6.HTM

[22] A fairly recent picture of humanity's ability to solve travelling salesman problems may be garnered from the http://www.tsp.gatech.edu/ web-site.

[23] Douglas R. Woodall: Monotonicity of single seat preferential election rules, Discrete Applied Maths. 77,1 (1997) 81-98.