Strategy in Range Voting and COAF voting systems

Warren D. Smith*
WDSmith@fastmail.fm

November 18, 2005

***Abstract* — (1) We define "range voting." (2) More generally we define the wide class of "COAF voting systems," (3) We reach an understanding of optimum voter strategy in COAF voting systems, at least in a certain probabilistic model, the "Gaussian model," of how the other voters behave, and in the limit $V \to \infty$ of a large number of voters. (4) This understanding also works for Condorcet and IRV voting (which are not COAF), albeit in a more limited way. (5) We find that range voting is the uniquely best COAF system, in a certain sense: Roughly speaking, range voting is the only COAF system which allows voters to be maximally expressive without being strategically stupid. We also find (for the first time) an infinite number of nontrivial voting systems which satisfy Mike Ossipoff's "favorite betrayal criterion" (FBC), and show that range voting is the unique FBC-obeying COAF system with maximum voter expressivity. Most of these results were due to me in 1999-2000 but I did not write them up very well at that time.**

1 Range and Approval Voting

In a C-candidate election conducted with *range voting* your vote is an assignment of a score in the interval $[0, 1]$ to each candidate. For example a legal vote in a 4-candidate election would be $(0, 1, 0.4, 1)$. The candidate with the greatest score-sum wins.

Approval voting [2] is the same, except that only the two endpoints of the $[0, 1]$ interval are allowed as scores. Giving a score of 1 is said to represent your "approval" of that candidate.

2 COAF Voting Systems (Compact set based, One-vote, Additive, Fair)

Let there be C candidates and V voters, C fixed, V large.

Definition 1. A "COAF Voting System" is a single-winner election that works as follows. Each voter chooses, from a fixed **C**ompact set $S \subset \mathbf{R}^C$ of "allowed votes," **O**ne C-vector. The vectors are **A**dded. The maximum entry in the summed C-vector corresponds to the winner. Such a system is "**F**air" if S is invariant under the group of $C!$ permutations of the C coordinates of $\mathbf{R}^C$.

Voting system	The Compact Set of Allowed Votes
Plurality	The C permutations of $(1, 0, 0, 0, \dots, 0)$
Approval	The 2^C vectors of form $(\pm1, \pm1, \pm1, \dots, \pm1)$
Dabagh "vote and a half"	The $(C-1)C$ permutations of $(2, 1, 0, 0, 0, \dots, 0)$
Used on Nauru Island	The $C!$ permutations of $(1,\ 1/2,\ 1/3,\ 1/4, \dots, 1/C)$
Borda[1]	The $C!$ permutations of $(C-1,\ C-2,\ C-3, \dots, 1,\ 0)$
"Vote for-and-against"	The $(C-1)C$ permutations of $(+1, -1, 0, 0, 0, \dots, 0)$
Anti-plurality	The C permutations of $(-1, 0, 0, 0, \dots, 0)$
Continuum "cumulative voting"	The $\vec{x}$ with $\sum_{j=1}^{C} x_j = 1$ and $x_j \ge 0 (\forall j)$
Boehm's "signed voting"	The $2C$ permutations & sign changes of $(\pm1, 0, 0, 0, \dots, 0)$
L_2-sphere voting	The $\vec{x}$ with $\sum_{j=1}^{C} x_j^2 = 1$
Range Voting	The unit C-dimensional hypercube $[0, 1]^C$

Figure 2.1. Some interesting COAF voting systems. Some well-known voting systems which are *not* COAF include: Condorcet systems, Instant Runoff (transferable vote), Bucklin, Woodall-DAC [4]; in all four of these cases the votes are *preference orderings* of the C candidates. ▲

*Non-electronic mail to: 21 Shore Oaks Drive, Stony Brook NY 11790.
[1] Some people also allow "truncated preference" vote-vectors, such as $(3, 2, \frac{1}{2}, \frac{1}{2})$, in (their version of) Borda.

2.1 Conditions which cause optimum range voting strategy to be approval-style

Theorem 2 (Approval strategy I). *Assume the candidates can be pre-ordered in decreasing likelihood of election chances, and that each's chance is hugely greater than the next one's. (Note: this assumption is valid in the Gaussian model with $V \to \infty$ that we shall describe in §5; assuming generic real parameters the probability ratios tend exponentially to ∞.) Then "approval style" range voting (where you give every candidate either the maximum- or minimum-possible score) is strategically optimum (or at least, the utility of the best approval-style vote only falls below that of the truly-best range vote by a negligibly small amount).*

Proof: In this case you can assign your scores to the candidates one by one in order, only basing your decision on a candidate's score on information about him and the previous candidates – the later ones are irrelevant to that decision. Then, for each candidate X you score, his score cannot affect the winning chances of the previously-scored ones if X does not win. X's score then can only affect the chance X wins, not the chances for previous candidates, unless X wins – and it does so in a monotonic manner. Hence the decision on X should be purely based on whether X is superior or inferior in utility to the expected utility among the previous winners given your previous votes: if superior, give X the maximum, otherwise the minimum score. □

Theorem 3 (Approval strategy II). *Assume the probability of an election-leading near-tie (which your vote can hope to alter) between two of the candidates is very unlikely and the probability of a 3-or-more-way near-tie is even more, indeed neglectibly, unlikely. Also assume your probability of breaking an AB tie is proportional to your vote's $A-B$ score difference. Then: "approval style" range voting is strategically optimum (or at least, the utility of the best approval-style vote only falls below that of the truly-best range vote by a neglectibly small amount).*

Proof: Give maximum score to A and minimum to B where AB is the most important (utility-wise, i.e. probability of that tie times candidate-utility difference in maximum) possible tie. There is no reason to want any other vote: You might want to lower A's score by some amount Δ if by so doing you could increase the vote difference for some other pair (or pairs) by Δ (in total); but no – since that other pair (or pairs) has (have) *less* utility, that move would be bad for you. So your best vote (maximizing expected utility) then always maxes-out or mins-out every score you can. □

Theorem 4 (Full knowledge). *For a voter with* complete knowledge *of all the other votes, an optimum range vote is always trivial to determine, and wlog is an "honest approval-style" (or one could even demand it be a "plurality-style"), vote. E.g. give the maximum score to the best candidate who can win and to all better candidates, and the minimum score to all others.*

3 Examples of non-approval-style optimum strategy in range voting

The assumptions in theorem 2 and 3, while plausible in large real elections, usually are false in small ones. We shall give incomplete-knowledge election examples in which all approval-style range votes are strategically sub-optimal.

3.1 First example: ignorance

In a 3-voter (or more voter) [0,1]-range-voting election, assume the totals of the votes of the other two (or more) voters are either

$$(1,\ 1.1,\ 1) \qquad \text{or} \qquad (0,\ 1.1,\ 1.2) \tag{1}$$

but you do not know which. Call the candidates A,B, and C. Your candidate-utilities:

$$U_A = 10,\ \ U_B = 5,\ \ U_C = 0. \tag{2}$$

- If you vote (1, 0, 0), then A wins in scenario #1 but C wins in scenario #2.
- If you vote (1, 1, 0), then B wins in scenario #2 but B wins in scenario #1.
- If you vote (1, 1, 1), then B wins in scenario #1 and C wins in scenario #2.

But if you vote (1, 0.5, 0), then A wins in scenario #1 *and* B wins in scenario #2, which is the best you can hope for in either case (and is strictly better than any of the above votes) so that this "honest" vote here also happens to be the strategic vote giving you the best possible expected utility.

3.2 Second example: you, a known, and a random voter

In the following 3-voter election, call the candidates A, B, and C. Let R_j denote independent uniform random numbers in $[0, 1]$, so that voter#2 is regarded as completely unknown to you, while voter#1 is regarded as completely known.

Voter	A	B	C
voter#1	0	0.798	0.618
voter#2	R_1	R_2	R_3
you(voter#3)	1	V	0

Assume the election-utilities (for you) of the three candidates are

$$U_A = 1, \ \ U_B = 0.832, \ \ U_C = 0 \tag{3}$$

Then here are the expected utilities of various possible votes you could make (result of 25 million Monte Carlo experiments):

Vote V	**Utility**
1(exaggerated)	0.8354
0.83(honest)	0.8436
0.58(best?)	**0.8512**
0(exaggerated)	0.8359

As you can see, neither the honest range vote, nor either of the two "approval style" exaggerated ones, is strategically best.

For those unsatisfied with inexact results arising from Monte-Carlo numerical integrations, we remark that it is in fact possible to find the *exact rational* utility answers by evaluating the volumes of certain 3-dimensional polyhedra. Specifically, voter #2's vote is a "unit cube" and the vote-total from all the voters is then a point in an appropriately translated unit cube. The probabilities of the various election outcomes are then the volumes of the intersections of this translated unit cube with the "payoff regions" (such as the set where $A > B$ and $A > C$). These intersections each are certain interior-disjoint convex polyhedra, whose volumes may be evaluated by dividing them into tetrahedra and then computing appropriate determinants.

3.3 Third example: you, a known, and two random voters

In the following 4-voter election, call the candidates A, B, and C. Again let R_j denote independent uniform random numbers in [0,1].

Voter	**A**	**B**	**C**
voter#1	0	0.896	0.219
voter#2	R_1	R_2	R_3
voter#3	R_4	R_5	R_6
you(voter#4)	1	V	0

Assume the election-utilities (for you) of the three candidates are

$$U_A = 1, \ \ U_B = 0.940, \ \ U_C = 0 \tag{4}$$

Then here are the expected utilities of various possible votes you could make (result of 25 million Monte Carlo experiments):

Vote V	**Utility**
1(exaggerated)	0.9435
0.94(honest)	0.9441
0.46(best?)	**0.9478**
0(exaggerated)	0.9409

Again, neither the honest range vote, nor either of the two "approval style" exaggerated ones, is strategically best.[2]

4 Review: Tricks with Gaussians

A *Gaussian probability density* in C-dimensional space is one proportional to

$$\exp\left(-\vec{x}M\vec{x}\right) \tag{5}$$

where M is a positive definite symmetric $C \times C$ matrix. If M is proportional to the $C \times C$ *identity* matrix, the Gaussian is "spherical," otherwise it is "ellipsoidal."

Gaussian densities are easy to integrate over halfspace or "slab" regions (bordered by two parallel hyperplanes). Simply perform the linear transformation $\vec{x} \to M^{1/2}\vec{x}$ to cause the Gaussian to become spherical with characteristic matrix I. This maps the halfspace (or slab), to a new halfspace (or slab). The spherical Gaussian integrated over a slab with distances of its bounding hyperplanes to the origin A and B, is simply $Z(B) - Z(A)$ where Z is the cumulative density function of the standard 1-dimensional Gaussian (with $M = 1$).

In particular [using $\int_{-\infty}^{\infty} \exp(-x^2)\mathrm{d}x = \sqrt{\pi}$] the integral of EQ 5 over *all* C-space is

$$\pi^{C/2}\ |\mathrm{det}M\ |^{-1/2} \tag{6}$$

[2] The idea that examples such as these might exist, was pointed out by Boris Alexeev, although his specific attempts to produce such an example failed due to incorrect arithmetic. The probabilities in this example again can be regarded as volumes of certain polyhedra – albeit this time *four* dimensional polyhedra – and hence again in principle could be computed as exact rational numbers. All Monte Carlo experiemnts have been done with several random number generators and feature self-agreement up to the final deciaml place.

and the reciprocal of this is, of course, the appropriate normalizing factor to make this Gaussian become a probability distribution.

Gaussians are also easy to deal with inside subspaces (whether translated or not) because they are (up to altering the normalizing factor) still Gaussians in the subspace.

The *central limit theorem* tells us that the sum of a large number V of bounded independent random C-vector variables with smooth densities, will, if translated so its mean lies at the origin $\vec{0}$ and then scaled by $V^{-1/2}$, approach a Gaussian density in the limit $V \to \infty$.

Consequently, instead scaling that sum by V^{-1} (which is what happens when computing the candidates' average scores in range voting) leads, in the $V \to \infty$ limit, to a Dirac delta function "spike" located at the mean.

These spike-like distibutions are even easier to integrate over convex regions R (asymptotically in the limit $V \to \infty$). That is because the integral is the same, to within asymptotically negligible error, as the integral within an arbitrarily small neighborhood of the *single point* in R where the density is greatest. If that **magic point** is in fact the centerpoint of the Gaussian, then a neighborhood of diameter$\gg V^{1/2}$ will suffice (we now assume a convention of *no* rescaling) but if it is any other point, then merely diameter$\gg 1$ is all that is required.

Incidentally, the only reason I am restricting myself to *convex* R is that this forces there to be at most one maximum of the Gaussian within R (one "magic point"). With nonconvex R more maxima could occur, including, possibly, coequal maxima. That would complicate things, but usually not intractibly so. And even for many nonconvex regions it often still is obvious that at most one maximum matters, the rest (if any) being exponentially neglectibly tiny in comparison when $V \to \infty$.

This "magic point method" makes performing all sorts of probability-integration tasks extremely easy in the $V \to \infty$ limit. Furthermore, its validity does not even depend on exactly having a Gaussian. Remember, the central limit theorem actually does not say anything particularly useful about the "tails" of our Gaussians. So one might worry about the validity of applying the magic point method inside tail regions. Let us examine that question. Consider the simplest (1D) case: the sum of V independent "Bernoulli" (binary-outcome coin flip $\pm 1/2$) random variables. As was well known since the days of de Moivre, this leads to the "binomial distibution"

$$\text{Prob}(S) = \frac{V!}{(V/2 - S)!\,(V/2 + S)!}, \quad |S| \le V/2 \tag{7}$$

which is approximated by a normal distribution when V is large:

$$\frac{V!}{(V/2 - S)!\,(V/2 + S)!} \approx \frac{2^V}{\sqrt{\pi V/2}} \exp \frac{-2S^2}{V} \tag{8}$$

As you can see, the true binomial distribution falls off somewhat faster in its tails than the normal approximation, but shares with that approximation the property of unimodality. This faster falloff causes our "magic point method" actually to work *better* for the true than the normal-approximate distribution in the tails, because it allows replacing our demand that the neighborhood size be $\gg 1$, with a neighborhood size of order 1 (or even some $o(1)$ sizes will do) centered at the magic point, *still* yielding asymptotic validity in the $V \to \infty$ limit.

5 Optimum voting strategy in a Gaussian probability model with a large number of voters

We are going to consider the question of what your strategically-optimum (expected-utility-maximizing) vote is, in any COAF voting system, in the following **Gaussian probabilistic model:**

1. You know your utility-value U_X for each of the C candidates (here candidate X).
2. You have seen a pre-election poll of a random voter subsample, which tells you that the total vote from the other voters is a Gaussian random variable and tells you the mean C-vector (centerpoint) and the covariance matrix of that Gaussian. (Or perhaps no covariance matrix is given, in which case we shall assume the Gaussian is spherical.)

We shall assume all of the real numbers in this data are "**generic**," that is, satisfy no algebraic relationships. (This allows us to avoid worrying about "ties.") Finally, we shall assume that the number V of voters is made very large ($V \to \infty$) so that their mean-vote Gaussian random variable becomes spike-like.

The "payoff regions" of C-space are the regions of vote-vector-total space in which one candidate wins, for example A wins if $A > B$, $A > C$, and $A > D$ in a 4-candidate election. These payoff regions are unbounded convex polytopes.

The **7-step procedure to find the strategically-best vote-vector** is as follows.

Step 0. Order the candidates in decreasing order of a priori likelihood of winning the election.

Every probability (including these) that we are going to care about is going to be a Gaussian integral over a convex polytopal region R.

These integrals may be rapidly estimated using the "magic point method." If the magic point (maximizing the Gaussian within R) happens to be on the *boundary* of R, e.g. a *corner*, then some difficulties can ensue because, e.g. the solid-angle factor represented by the corner can be difficult to evaluate. However, this does not matter because these factors are enormously outweighed in the $V \to \infty$ limit by the exponentially huge differences in size among the Gaussian heights at the magic points for various regions. So (assuming generic reals so there are no height *equalities*) only these Gaussian heights matter, nothing else matters, when it comes to deciding which integral of two is the larger.

Also, we should point out that *finding* the "magic point" at which a given Gaussian is maximized (or equivalently, where its logarithm – a concave-$\cap$ quadratic function of $\vec{x}$ – is maximized) inside a convex polytope R, is just an instance of "convex quadratic programming" and hence is soluble by a polynomial time algorithm.

Assuming (as we have) generic real data, each candidate's election chances are exponentially enormously huger than the next's in the $V \to \infty$ limit (justifying the assumption in theorem 2).

Step 1. Find the most-likely lead-near-tie between two candidates.

The likelihood of a tie-within-Δ between A and B is the integral of the Gaussian over the slab-like region $|A - B| \le \Delta$ intersected with the convex polytope $A > C$, $A > D$, $B > C$, $B > D$. Again, this is a Gaussian integral over a convex polytopal region.

Step 2. Of those two candidates, give the one with the most utility the maximum, and the other the minimum, vote-score allowed by the rules of that voting system (and given that the previously-assigned scores are fixed).

Note: for many voting systems, e.g. plurality and L_2-ball voting, at some early point in this **FOR**-loop, the rest of the vote often becomes uniquely determined, i.e. forced by the rules of that voting system given the preceding scores. In that case we can terminate the algorithm early.

Step 3. Now proceed as follows among the remaining candidates X in most-likely to least-likely order. **FOR** each X:

Step 4. Determine the most likely way in which X can be involved in a lead-tie to within Δ.

This lead-near-tie is going to be among a set of candidates, e.g. A, B, C, X, corresponding to the point of the linearly-transformed space in which the Gaussian is spherical and $\vec{0}$-centered, in which X is behind the lead by at most Δ, and which (among all such X-nearly-leads points) lies the closest to the origin $\vec{0}$ (i.e. which maximizes the Gaussian in the original space). Unfortunately (in the sense that it makes calculations harder) this candidate-set is not necessarily of cardinality=2.

Step 5. By integrating the Gaussian over the full set of such X-nearly-leads points, find the a priori (conditioned on X nearly leading) probabilities P_Y that each non-X candidate Y would win that election (without X), and compute $E = \sum_Y P_Y U_Y$.

The simplest case is if there is only one non-X candidate Y involved in the lead-near-tie, in which case $P_Y = 1$ without any calculation needed. Otherwise, the P_Y calculations unfortunately can involve difficult multiple integrations of Gaussians over polytopal regions - "difficult" because the magic point method is inadequate because a finer degree of accuracy is required. (The difficulty is that *all* the eligible non-X candidates involved in the tie share the *same* "magic point" therefore our usual technique of just dismissing one as comparatively negligible, is not applicable – we actually need to calculate.) In practice, though, Monte-Carlo approximate integrations are quite adequate, and we do not need to find each P_Y in its own integration; it suffices to evaluate the expected utility $E = \sum_Y P_Y U_Y$ in a *single* integration. (Another way of saying the same thing: *all* the P_Y may be found in a single run of Monte-Carlo experiments.)

Step 6. If $U_X > E$ then award X the maximum possible score as X's vote-score; if $<$ then the minimum (and $=$ is not possible by our genericity assumption).

Step 7. END FOR.

We call this the "moving weighted average" method because each candidate X's vote-score is chosen by contrasting his utility U_X with a weighted average of the utilities of certain other (previously-scored) candidates, and with the weights and the candidate-set possibly changing as we go.

Theorem 5 (Optimum strategy in COAF voting systems). *The above algorithm determines the unique strategically optimum (maximizing expected utility of the election result) vote, in any given COAF voting system, in the Gaussian probability model of the other voters, in the $V \to \infty$ limit, assuming generic real parameters and assuming no "ties" occur at any stage of the algorithm.*

Proof: By theorem 2 and its proof, it is generically valid to choose the vote-scores for each candidate in decreasing order of their a priori election chances, ignoring the later ones when making the candidate-X decision. In that case each successive decision must be made to maximize expected utility, which is precisely what we do. □

6 Examples of optimum voting strategy in the Gaussian $V \to \infty$ model

We here shall consider voting systems other than range voting. Range voting will be addressed next section.

6.1 Borda voting – typical example

Suppose the candidates are A, B, C, D, E in decreasing order of winning-chances. In fact (to be more precise) suppose the pre-election poll (which you regard as being a random subsample of the Borda votes) gives totals for (A, B, C, D, E) proportional to (20, 16, 13, 11, 10), which 5-vector we regard as the mean of a spherical Gaussian. Suppose your utilities for the election of each are

$$U_A = 40, \quad U_B = 60, \quad U_C = 25, \quad U_D = 75, \quad U_E = 10. \tag{9}$$

Then your honest and optimally-strategic Borda votes are:

$$\begin{array}{cc} \text{Honest:} & \text{Strategic:} \\ (A, B, C, D, E) = (2, 3, 1, 4, 0) & (A, B, C, D, E) = (0, 4, 1, 3, 2) \end{array} \tag{10}$$

Observe that these two votes bear little relation to one another. This example is quite typical of the way that strategic Borda votes often seem "random" and almost entirely unrelated to the honest ones.

To outline how the strategic vote was found: realize that (with this Gaussian-mean) each candidate, successively, is, assuming it is going to be nearly-tied for the lead, most likely to do so by being tied with A alone. Therefore, we go through the candidates X in order awarding each either the least or grestest still-available score (under the rules of Borda voting and considering the previously-assigned scores) depending on whether $U_X < U_A$ or not.

6.2 Plurality and Continuum Cumulative voting

Generically, the optimum strategy is always the following: Vote for the candidate, among the two pre-election-poll-leaders, with the greatest utility. This has been called the "lesser of two evils."

In continuum cumulative voting, the best strategy is the same (vote everything for the best-utility among the two pre-election-poll frontrunners, and zero for everybody else). If all voters do that, the same "lesser of two evils" winner as with plurality is elected.

6.3 Anti-Plurality voting

Generically, the optimum strategy is always the following: Vote against the candidate, among the two pre-election-poll-leaders, with the least utility. Note that if all voters do that, then the winner of the election (assuming $C \geq 3$) is assured to be "dark horse" who was *not* one of the pre-election poll leaders. I call voting systems such as this which tend (with strategic voters) to invalidate their own pre-election polls, "suicidal."

6.4 Borda voting – DH3 pathology

Suppose the candidates are A, B, C, D in decreasing order of winning-chances, and that you regard A, B, C as roughly equally likely (probabilities 1/3 each) to win. (We here are departing from our usual genericity assumption.) Also suppose the "dark horse" D is universally perceived as having no chance to win – because Society seems divided into three roughly-equinumerous camps favoring A, B, and C respectively.

If you are in the A-camp and hence for you $U_A > U_B \approx U_C \gg U_D$, then your most-strategic vote is $(A, B, C, D) = (3, 1, 0, 2)$ or $(3, 0, 1, 2)$; choose one of these two at random by a coin flip. If everyone acts this way, then the universally-agreed-to-be-worst candidate D will win the election. This is quite suicidal (for every meaning of the word)!

Incidentally: if all voters act this way then the winner in any *Condorcet* voting system will also be D. That may be a strong indictment of all Condorcet voting systems.

6.5 Borda voting – another pathology

Suppose the candidates are A, B, C, D, E in decreasing order of winning-chances based on the pre-election poll. Suppose roughly half the populace thinks $U_A > U_B$ and half thinks $U_B > U_A$. Then roughly half the strategic votes will be of the form $(A, B, C, D, E) = (4, 0, [3, 2, 1])$ and the other half will be of the form $(A, B, C, D, E) = (0, 4, [3, 2, 1])$ where the square brackets denote "some permutation of." If all voters act this way, then some member of $\{C, D, E\}$ will (generically) win the election. This again is "suicidal" voting system behavior and again can force the election of a horrible "dark horse."

These kinds of pathological behaviors indeed were immediately observed on Kiribati, the world's only government to employ Borda voting [3]. Kiribati then abandoned Borda and switched to plurality voting. Nauru's system, in contrast, is immune to both DH3 and this pathology.

6.6 Dabagh voting

Generically, the optimum strategy is always the following: Vote 2 for the candidate (call him X), among the two pre-election-poll-leaders (call them X and Y), with the greatest utility and 0 for the other. Now go though the remaining candidates in decreasing order of their winning chances. For the first one Z you find who has greater utility than the expected utility of the winner among the candidates more-likely than Z to win (conditioned on Z being in a near-tie with that winner), give Z your 1-vote. If no such Z exists then award the 1 to the least-likely candidate to win.

Generically, if all voters act this way then the same winner (lesser of two evils) will occur as with strategic plurality voting, but with the twist that the least-likely "darkest horse" candidate simultaneously will receive an astoundingly large vote total.

6.7 For-and-against and L_2-sphere voting

With L_2-sphere voting, your best vote is of the form $(+2^{-1/2}, -2^{-1/2}, 0, 0, 0, \dots, 0)$ where your positive and negative votes are for the two pre-election-poll frontrunners. This is the same as the best strategic vote with for-and-against voting. Generically the same "lesser of two evils" winner will then arise as in strategic plurality voting; however, if those two are closely matched, then a small fraction of other kinds of voters can cause a hard-to-predict winner.

6.8 Boehm's signed voting

Is interesting. It can be possible for the usual lesser-of-two-evils candidate to win, *or* a hard to predict candidate. Our algorithm actually is inadequate in that it produces two "tied" possible votes for the two frontrunner candidates, namely $(+1, 0)$ or $(0, -1)$. Deciding which among these to use, actually does depend on using the "later" candidates to perform a "utility tiebreak" (this would be an enhanced version of our algorithm).

7 Range Voting strategy in the Gaussian model

7.1 Three results about range votign that are valid fully generally – i.e. without need of the Gaussian model

Theorem 6 (Monotonicity). *Increasing your vote-score for X cannot decrease X's chances of winning; decreasing your vote-score for X cannot increase X's chances of winning.*

Theorem 7 (No favorite-betrayal). *It is* never *strategically forced to give X less than the maximum possible score, if X is your true favorite (or one of your true favorite) candidates.*

Proof: Increasing X's score will not change the winner (except if it changes to X). □

Similarly, there is no strategic reason to give your most-hated candidate anything other than the minimum-possible score.

Theorem 8 (No order reversal in 3-candidate elections). *It is* never *strategically forced to misorder $X > Y$ if you truly regard $Y > X$, in your range vote,* if *there are ≤ 3 candidates.*

Proof: Call the three candidates X, Y, and Z. By the preceding result we can assume wlog your favorite and most-hated candidates are respectively given the maximum- and minimum-possible vote-scores. It then is not possible for you to vote $X > Y$ if you honestly think $U_Y > U_X$. □

But if there are *four* candidates, then we shall see, via either of two examples, that dishonesty can be strategically forced.

7.2 Strategic dishonesty in range voting I: correlation (highly ellipsoidal Gaussian)

There are two liberals L_1, L_2 and two conservatives C_1, C_2 running. You believe that L_1 and L_2 will get a near-equal number of the other people's votes, and ditto for C_1, C_2, but don't know whether the liberals or conservatives will be ahead[3] Then your best strategy is to vote in the style

$$L_1 = 1, \quad L_2 = 0, \quad C_1 = 1, \quad C_2 = 0 \tag{11}$$

even if you prefer both L's over both C's (or both C's over both L's). (We assume you slightly prefer the candidates with smaller subscripts.)

(This example also can be generalized to $2N$ candidates falling into N ultra-correlated pairs. It apparently originally traces to S.J.Brams and was shown to me by Rob LeGrand.)

[3] This kind of situation can be modeled with a highly ellipsoidal Gaussian distribution being returned by a pre-election poll in which the correlation and covariance matrices are published, not just the estimated means. We admit that the sort of pre-election polls we have seen in newspapers and television never do publish covariance information, so this is somewhat unrealistic. However, this scenario as a whole nevertheless is tolerably realistically plausible.

7.3 Strategic dishonesty in range voting II: even without correlations (exactly spherical Gaussian)

Let the 4 candidates be A, B, C, D. Let the pre-election poll results say the vote totals (mean of the Gaussian) are $A = 100$, $B = 80$, $C = 70$, $D = 10$, (well more precisely, proportional to these; the totals are these times a number of order V) say, so A appears most likely to win, then B, then C, then D. We assume these 4 vote totals are samples from a spherically symmetric 4-dimensional Gaussian, i.e. with no correlations or anticorrelations among candidates. (It is easier to make one of our examples if correlations are assumed – as we just saw – but our point here is they are not needed.) There are V voters where we assume $V \to \infty$ (we work in this *limit*) and the Gaussian has mean of order V and peakwidth (standard deviation) of order $\sqrt{V}$. Let your true utilities be:

$$U_A = 0, \;\; U_B = 20, \;\; U_C = \text{HUGE}, \;\; U_D = 25. \tag{12}$$

Here HUGE denotes some sufficiently large constant.

Theorem 9 (Dishonesty example). *In this scenario in the Gaussian model, the unique strategic range vote (which maximizes your expected election-result-utility) is* $(A, B, C, D) = (0, 1, 1, 0)$. *This is dishonest for B and D.*

Proof: The top-two pair is AB – most likely to be tied for lead – other ties are exponentially less likely when $V \to \infty$ – so we give $A = 0$ and $B = 1$ in the vote to give us maximum chance to favorably break an AB tie. (Since $U_B > U_A$.)

Now AC is the next pair (second most likely to tie for the lead, but in the $V \to \infty$ limit, it is exponentally far less likely to be AC than it is to be AB; but AC is exponentially more likely than any other non-singleton set) so give $C = 1$ to make C win over A in a situation where AC tied for lead and B and D's winning chances may be neglected. (Since $U_C > U_A$.)

Finally among sets involving D tied for the lead, the most likely to result in a lead-tie is A, B, C, D 4-way near-tie near 260/4=65 each. That is because any other location in the 4-space having D maximal has a larger sum-of-squares distance from (100,80,70,10) and hence is exponentially less likely in the $V \to \infty$ limit.

Now even if you don't believe in the 4-way-near-tie, fine. More precisely, all we need to believe, is it is going to come down to either a DC, DB, or DA battle (conditioned on the assumption that D is involved in the lead-tie) with some fixed positive conditional probabilities for each, not necessarily $1/3, 1/3, 1/3$, in the $V \to \infty$ limit. Call these probabilities P_A, P_B, P_C. Basically what happens is as you move away from the magic $(65, 65, 65, 65)$ point, but preserving the assumed fact that D is maximal: as soon as you go a distance of order 1, i.e. order 1 vote worth, that causes a factor of order $e^{O(1)}$ falloff in probability. Hence when $V \to \infty$ there is negligible probability (conditioned on D leading or co-leading) that we are more than a distance of order $(\log V)/V$ away from the magic $65, 65, 65, 65$ point. You can see that using EQ 5 defining the Gaussian (which we assume is still valid out here in the tail, perhaps somewhat unrealistically, but it is ok since we took the whole Gaussian model as an *assumption*, and anyhow as I've argued it apparently would be valid *enough* even out here in the tail, even without such an assumption) using standard deviation of order $\sqrt{V}$ and distance-to-center of order V. That is the underlying reason for the positive P_A, P_B, P_C existing.[4]

So in this final case where we assume D has a chance (which is an assumption needed to make it worth deciding on a vote about D at all) we should vote $D = 0$ provided that

$$25 = U_D < P_A U_A + P_B U_B + P_C U_C \tag{13}$$

which by making U_C HUGE enough (since the values of P_A, P_B, P_C exist and do not depend on U_A, U_B, U_C), always happens. So we must vote $D = 0$. □

The thinking in this example actually can be generalized to provide a way to determine your strategically optimum range-vote in *any* Gaussian-model scenario in the $V \to \infty$ limit, namely the algorithm in §5.[5]

8 Axiomatic characterizations of range voting

8.1 Favorite-betrayal

Mike Ossipoff has emphasized the importance of the following criterion for single-winner voting systems in order that they lead to a good democracy.

[4] You may here worry that perhaps a *3-way* or 4-way tie could be important, in which case the strategically best vote might not even be approval-style at all. However, if we agree in this example to make our spherical Gaussian have *variance* which is a sufficiently large constant, then the conditional probability of a more-than-2-way tie can be made arbitrarily small, so this is not a worry. Arbitrarily large variances arise in practice if, e.g, the votes are postulated to come in blocs of order 1000 highly correlated votes. Also, a different reason why the optimum vote here must be approval-style, regardless of worries about 3-way near-ties, was given in theorem 2, which makes this whole issue moot.

[5] But Mike Ossipoff comments: I want to re-emphasize that, although in extremely unlikely examples, reversing a preference can be optimal in Approval, one should never vote someone over one's favorite. One should always vote top for one's favorite in Approval or Range voting by theorem 7; no exceptions.

One author familiar with the subject said that the kind of probability knowledge that would be needed to make preference-reversal optimal in Approval is so rare that such situations can be disregarded.

And (WDS further comments) even with such knowledge... these situations still (empirically) seem quite rare. And even when they do occur they usually have a small probability of changing the winner, hence do not matter much. My computer simulations with 4- and 5-candidate elections indicated a typical utility change due to true-strategic voting versus "honest approval" (threshhold strategy) voting of order 1 part in 1000.

Favorite-Betrayal Criterion (FBC): Voters should have no incentive to vote someone else over their favorite.

(Cf. theorem 7.) While FBC may sound like an obvious criterion, very few of the usual election method proposals comply with it. For example, the following election suffices to demonstrate FBC-failure for most of the strict-rank-order-based election methods (e.g. Nanson, Condorcet, Instant Runoff, Borda, Woodall DAC) one runs across:

#voters	Their Vote
8	$B > C > A$
6	$C > A > B$
5	$A > B > C$

Most methods agree that the winner is B, but if the six $C > A > B$ voters insincerely switch to $A > C > B$, betraying their favorite C, then A becomes the winner, which those voters prefer.

For most methods with rank-ordering votes *allowing equalities*,

#voters	Their Vote
3	$A = B > C$
3	$C = A > B$
3	$B = C > A$
2	$A > C > B$
2	$B > A > C$
2	$C > B > A$

is a 3-way tie which breaks in favor of A if the $A = B > C$ voters change their vote to $A > B > C$, thus betraying their true co-favorite B but improving the election result in their view. (This latter example is due to Kevin Venzke.)

Theorem 10 (FBC-obeying COAF systems). *The COAF voting systems (aside from trivial ones where no voter choice is possible) which obey FBC, are precisely the ones whose convex-hull's extreme points are the 2^C approval-votes. (These include approval voting itself, and range voting.)*

Proof: If your favorite candidate is not one of the $k \geq 2$ pre-election-poll frontrunners, then it is easy to see in the Gaussian model of §5 that your vote-score for him will necessarily be below the strategic vote-score you give to one of those frontrunners, in some situations, *unless* the set of allowed C-vector votes, scaled and translated to just fit inside the C-dimensional unit cube, includes every one of the 2^C corners of that hypercube (i.e. the 2^C approval votes). □

Remark. This leads to an infinite number of FBC-obeying COAF voting systems, arising from any subset of the unit C-dimensional cube which includes all of its vertices. E.g. for a mere 2-parameter infinity of examples, consider range voting where the voter is required to score at least one candidate above a and at least one candidate below b, for any given a, b with $0 < a < b < 1$.

Theorem 11 (Range Voting characterization I). *Range voting is: the unique FBC-obeying COAF voting system with maximum C-dimensional volume for its allowed-vote-set (i.e. maximum "expressive freedom" for the voter) given that that set is pre-scaled to have some fixed diameter or to have some fixed maximum−minimum score-range for some given candidate.*

8.2 Expressivity for strategic voters

For any COAF voting system, there is a set of possible allowed votes you can cast. We now enquire about **generically strategically accessible subsets (SSAs)** of those possible votes. Specifically, in the Gaussian model of §5, only certain of the allowed votes are actually possible outputs of the algorithm that finds strategically optimum votes as a function of both the Gaussian's parameters and your candidate-utility C-vector (assuming both consist of generic reals).

If the Gaussian parameters g are regarded as fixed, then in general there is an even smaller g-dependent set SSAGF(g) of strategically accessible votes.

Let $k \geq 2$. Let SSAGF$(g, u)_k$ be the subset of allowed votes that still remain strategically accessible given g and your utilities u for the k candidates most likely (according to g) to win (both regarded as made of generic reals), but with your remaining candidate-utilities as-yet unspecified. The larger the SSAGF$(g, u)_k$s are, the more "expressive freedom" the *strategically-aware* voter may be said to have.

Theorem 12 (Range Voting characterization II). *Range voting is the unique COAF voting system (for a given diameter of[6] the allowed-vote compact set) in which, for each $k \in \{0, 2, 3, 4, \ldots, C\}$, the set SSAGF$(g, u)_k$ generically has nonzero, and maximum possible, $(C - k)$-dimensional measure.*

Proof: This may be proven by induction on the number of dimensions (number of candidates). The point is that, inductively, after k candidates have been removed from the picture by pre-assigning their vote-scores, the remaining vote scores will necessarily be range votes for $C - k$ candidates. The basis of the induction (1-3 candidate cases) is easy.

In fact, in the 3-candidate case, what we are demanding is that the intersection of the allowed-vote set (pre-scaled to have fixed L_p diameter) with the plane ($X = 1$ say) giving the maximum score to one candidate, and the plane ($Y = 0$, say) giving

[6] Here "diameter" may be measured in any L_p metric, $p \geq 1$.

the minimum score to another, always be a 1-dimensional set of maximum possible total measure (length), no matter which two axes "X" and "Y" were.

Now wlog we may choose X and Y to be such that the allowed-vote set has maximum widths in the X and Y directions, and minimum in the Z direction, if all such widths are not identical. That already would invalidate the theorem if these widths were non-identical as we can see by considering the voting system whoxe allowed-vote set is a "brick." Given that the wdiths are identical, it is plain that the intersection of the two planes and the vote-set can be a subset of a vertical line whose total measure is at most equal to the Z-width, and this is always achieved if and only if the entire "edge of the cube" is in the set. So we may restrict ourselves to voting systems which include the entire "wire frame" [7] of the unit cube defining the maximum 1D widths, in their allowed vote set, in the 3-candidate case. To now maximize 3-volume we now are forced to demand the allowed vote set be the entire cube. That provides the basis (when $C - k = 3$) of the induction.

This basis case may be restated as follows: "Range voting is the unique 3-candidate COAF voting system in which it always remains possible to vote honestly and maximally expressively (i.e. one's vote score for Q is completely unrestricted between the maximum and minimum allowed scores for any candidate) about the least-popular-in-the-pre-election-polls candidate Q, given that this voter has already chosen the scores for the two more-popular candidates by consideration of overriding strategic considerations." □

9 Optimum voting in Condorcet and other non-COAF systems

Condorcet voting systems share some features with COAF systems – specifically, one may first "add up" a set of votes (more precisely, one adds up pairwise election totals of various allowed types), and then deduce the winners from the totals. Hence it might be hoped that optimum strategy for Condorcet systems would also be easily comprehensible. However, I doubt that.

In a 3-candidate election let X, Y and Z denote the number of $A > B$, $B > C$, and $C > A$ votes respectively (or $-X$ is the number of $B > A$ votes if X is negative). Then 3-candidate Condorcet votes are actually the same thing as 3-candidate approval votes (with the "stupid" $+++$ and $---$ approval votes disallowed), *but* the winner is determined quite differently. Specifically (ignoring ties), A wins if and only if

$$(X > 0 \wedge Z < 0) \vee (X > 0 \wedge Y > 0 \wedge Z > 0 \wedge Z < X \wedge Z < Y) \vee (X < 0 \wedge Y < 0 \wedge Z < 0 \wedge Z < X \wedge Y < X). \tag{14}$$

More generally, all of the usual Condorcet election schemes may be thought of as being just like COAF methods *but* in a $(C-1)C/2$-dimensional space where C is the number of candidates, and *except* with the winner determined by membership in certain *non*convex polytopal sets. (This was only allowing strict preference orderings as votes. If equal rankings are allowed in votes, then $(C-1)C$-dimensional space.) These nonconvex polytopes can be expressed as boolean combinations of convex ones (as here) but those boolean expressions can become very large and complicated, e.g. with an exponential-in-C number of terms in them.

The non-convexity and high logical complexity of the payoff regions means that determining the most-likely way for some candidate to win (or nearly-win), is no longer necessarily easy. And that in turn leads me to suspect that determining strategically optimum Condorcet votes might be very difficult if C is allowed to be large.

For several voting systems, including "Instant-Runoff Voting," it has been proven [1] to be NP-hard – in situations with a large number C of candidates – to vote optimally.

However, if C is *fixed*, strategic voting in both IRV and all the usual Condorcet systems, becomes generically feasible. One simply subdivides the payoff regions into a fixed number of simplices (which are, of course, convex) and then uses *them*. The usual procedure then works: pre-order the candidates in decreasing likelihood of their election chances (each is exponentially hugely more likely than the next to win in the $V \to \infty$ limit); for each candidate X compute the most likely way it can get into a near-win; and then alter the still-available degrees of freedom in the vote in order to maximize or minimize X's winning chances in that scenario (depending on U_X versus a weighted average U for the previous candidates). Keep doing this until the vote is uniquely specified.

Because the polytopes in C dimensions with C fixed necesaarily have bounded complexity, and because subdividing any C-dimensional polytope in to simplices is a polynomial time task if C is *fixed*, [8] and finally because convex programming is in polynomial time, we conclude

Theorem 13 (Polynomial time strategizing). *If the number C of candidates is fixed, then in the Gaussian statistical model there is polynomial-time algorithm to find your strategically-optimal vote (maximizing expected utility for you of the election winner), in both Instant Runoff Voting and all the usual Condorcet Voting Systems. If C is not fixed, i.e. allowed to become large, this is still a valid algorithm, although it might no longer be in polynomial time.*

This has never been stated previously. Note that for many COAF systems, e.g. range voting, we get a polynomial-time algorithm even with *no constraint* on C, but that is not the case here.

[7] Actually, only half of the wires are demanded, and a specific half, but this does not matter.

[8] Chop up everything into *convex* polytopes using the facial hyperplanes, then divide them into simplices by "coning off" to an extra central vertex, treating the faces recursively in one dimension lower.

References

[1] J.J.Bartholdi, J.A. Orlin: Single transferable vote resists strategic voting, Social Choice and Welfare 8 (1991) 341-354.

[2] S.J.Brams & P.C.Fishburn: Approval Voting, American Political Science Review 72,3 (September 1978) 831-847; also book with same title, Birkhäuser Boston 1983.

[3] Benjamin Reilly: Social Choice in the South Seas: Electoral Innovation and the Borda Count in the Pacific Island Countries, Int'l Political Science Review 23,4 (2002) 355-372.

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