Alice and Bob Will Fight: The Problem of Electing a Committee in the Presence of Candidate Interdependence

Joel Uckelman[1]

Abstract. The problem of electing a committee which satisfies voters is one for which good solutions are scarce. Extending single-winner voting rules to the multi-winner case works well only when voters have no preferential dependencies among candidates for the committee. (Our motivating example is a voter who believes that Alice and Bob are the best candidates, but also that the worst possible committee is one with *both* Alice and Bob.) In order to tackle the interdependence problem, we propose a voting rule called the Goalbase Summation Rule (GSR), which uses goalbases (sets of weighted propositional formulas) as ballots. Using goalbases as ballots lets voters express complex preferences in a compact fashion, while the computational complexity of finding winning committees remains reasonable when the number of seats is fixed. Additionally, the GSR is able to simulate and extend many existing voting rules.

1 The Problem

Suppose that we are electing a three-seat committee from the five candidates Alice, Bob, Charlie, Dave, and Elaine (a, b, c, d, e). Suppose further that a voter Valerie believes that

1. Alice and Bob are the best candidates, so any committee with one of them is better than any committee with neither, and
2. Alice and Bob will fight if they are on the committee together, so any committee with both is worse than any committee with neither.

These two constraints generate Valerie's preference order:

$$acd, ace, ade, bcd, bce, bde > cde > abc, abd, abe.$$

Given her preference, how should Valerie vote? If Valerie votes for both Alice and Bob in hopes that only one will win, she risks electing one of her least-favored committees. If Valerie votes for neither Alice nor Bob in hopes that other voters will prefer one over the other, she risks electing her second choice committee cde. If Valerie votes for Alice but not Bob, or Bob but not Alice, she has no principled way to choose between these options, as she prefers Alice-committees and Bob-committees equally.

Valerie's problem arises for a wide variety of voting rules:

Definition 1 (m-vote, ℓ-support, top-k)**.** *Call a voting rule* m-vote *if each voter may cast at most m votes;* ℓ-support *if each voter may cast at most ℓ votes per candidate; and* top-k *if the k candidates receiving the most votes are the winners.*

Each tuple $\langle k, \ell, m \rangle$ where $k, \ell, m \in \mathbb{N}$ and $\ell \leq m$ defines a voting rule. Standard, single-winner plurality is the 1-vote 1-support top-1 rule, while approval voting is the n-vote 1-support top-1 rule. Cumulative vote is m-vote m-support top-1, while Borda is $\frac{n(n-1)}{2}$-vote $n-1$-support top-1 with the additional constraint that support must be given exactly in the amounts $n-1, n-2, \ldots, 1, 0$.

All m-vote ℓ-support top-k rules have the following flaw when $k > 1$: Voters for whom the candidates are not independent may have no principled way to vote.

Proposition 2. *Let C be a set of candidates, with $a, b, x, y \in C$ all distinct, $\vec{c} \subseteq C \setminus \{a, b, x, y\}$, and $|C| \geq 4$. Then no m-vote ℓ-support top-k rule has a ballot compatible with the partial order $ab\vec{c} < xy\vec{c} < ax\vec{c}$, $by\vec{c}$, for $k \geq 2$.*

Proof. Fix a ballot $c_1, \ldots, c_n$, where c_i denotes the number of votes given that candidate. Suppose that the ballot induces the order $ab\vec{c} < xy\vec{c} < ax\vec{c}$, $by\vec{c}$. Wlog, let $\vec{c} = \emptyset$. So $a + b < x + y < a + x$, $b + y$. But then $a + b < a + x \implies b < x$, and $x + y < b + y \implies x < b$, contradiction. □

Thus, if Valerie is faced with any m-vote ℓ-support top-3 rule, her preferences are inexpressible. (Similar things can be said for distance-based rules, as found in [2].)

2 The Goalbase Summation Rule

Rather than electing a committee by voting for individual candidates, we might consider voting for *properties* of the winning committee. What Valerie wants is to indicate her preference for "Alice-XOR-Bob-ness" in the winning committee. Now we present a balloting framework in which this is expressible.

First, we require some definitions, following [5]:

Definition 3 (Goalbases and Utility Functions)**.** *A* goalbase *G is a set of* weighted propositional formulas *(φ, w), where φ is formed from Boolean connectives and propositional symbols from the set $\mathcal{PS}$, and $w \in \mathbb{R}$. A* utility function *$u\colon 2^{\mathcal{PS}} \to \mathbb{R}$ maps* models *$M \subseteq \mathcal{PS}$ to their values. Each goalbase G generates a utility function u_G such that $u_G(M) = \sum_{(\varphi, w) \in G, M \models \varphi} w$.*

Imagine a voting rule where voters submit goalbases with binary weights as ballots. Valerie might then wish to cast

$$G = \{(\neg a \vee \neg b, 1), (a \dot{\vee} b, 1)\}$$

as her ballot. The utility function u_G which this ballot generates,

$$u_G(M) = \begin{cases} 2 & \text{if } a \in M, b \notin M \text{ or vice versa,} \\ 1 & \text{if } a, b \notin M, \\ 0 & \text{if } a, b \in M, \end{cases}$$

---

[1] Institute for Logic, Language and Computation, University of Amsterdam. Email: `j.d.uckelman@uva.nl`

induces an ordering on the committees which coincides exactly with Valerie's preference ordering. The models acd and bde score 2, ahead of cde which scores 1, ahead of abc which scores 0.

Now we formulate a general voting rule using goalbases as ballots, in which winners correspond to the optimal states over the sum of all voters' ballots:

Definition 4 (Goalbase Summation). *If G, G' are goalbases, then*

$$G \oplus G' = \left\{ \Big(\varphi, \sum_{(\varphi,a)\in G} a + \sum_{(\varphi,b)\in G'} b\Big) \;\middle|\; \varphi \in \mathrm{For}(G \cup G') \right\}$$

is their formula-wise sum, where $\mathrm{For}(G) = \{\psi \mid (\psi, w) \in G\}$.

Definition 5 (Optimal k-Sized Models). *The set* $\mathrm{opt}_k(G)$ *is the set of optimal k-sized models according to the goalbase G:*

$$\mathrm{opt}_k(G) = \operatorname*{argmax}_{\substack{M \subseteq \mathcal{PS} \\ |M| = k}} u_G(M).$$

Definition 6 (k-Seat Goalbase Summation Rule). *The* k-Seat Goalbase Summation Rule (GSR) *is the voting rule where voters $v \in \mathcal{V}$ submit goalbases G_v as ballots and the winning committees are the members of* $\mathrm{opt}_k(\bigoplus_{v\in\mathcal{V}} G_v)$.

Our first observation about the GSR is that when $k = 1$, we have a voting rule which generalizes a wide variety of standard single-winner rules in the sense that it accepts all of their ballots (after a bijective mapping to goalbases) as valid ballots itself. For example, plurality ballots correspond to goalbases $\{(c, 1)\}$ where c is the voter's top choice; approval ballots correspond to goalbases $\{(c, 1) \mid c \in \mathcal{A}\}$, where $\mathcal{A}$ is the set of the voter's "approved" candidates; cumulative ballots correspond to goalbases $\{(c, w_c)\}_{c\in C}$ where the $w_c \in \mathbb{N}$ sum to no more than a fixed point limit. In fact, there is an obvious mapping for all m-vote ℓ-support top-k voting rules:

Proposition 7. *Let $\mathcal{B}$ be a set of m-vote ℓ-support ballots. For each ballot $b = v_{c_1}, \ldots, v_{c_n} \in \mathcal{B}$, represented as a vector of votes for each candidate, define the goalbase $G_b = \{(c_i, v_{c_i}) \mid 1 \le i \le n\}$, collecting all $G_b \in \mathcal{G}$. Then k-GSR with ballots $\mathcal{G}$ will choose the same winners as the m-vote ℓ-support top-k rule with ballots $\mathcal{B}$.*

On its face, this is just a baroque way of rewriting the ballots of existing voting rules, so is not interesting by itself; rather, the interest is in what it points to: Voting rules we can *simulate* using GSR we can also *extend*. The problem we identified in Section 1 was that multi-winner extensions of single-winner voting rules have too little expressive power when voters have non-independent preferences. The ballots used by each of these rules are a proper subset of all possible goalbase ballots; therefore, relaxing the restrictions on allowable ballots is a direct way of increasing the expressivity of such rules.

Proposition 8. *Deciding winners for k-GSR may be done in time polynomial in the number of candidates, for any fixed k.*

Proof. Deciding winners for k-GSR is the decision-problem version of finding members of $\mathrm{opt}_k(\bigoplus_{v\in\mathcal{V}} G_v)$—exactly the problem k-MAX-UTIL, which involves deciding whether a model yielding at least some given amount of utility exists [4, Definition 7.5.2]. For a fixed k and $n = |C|$, there are only $\binom{n}{k} \le \frac{n^k}{k!}$ models to check. □

This makes *whatever* goalbase language we want for representing our voters' ballots computationally tractable (though not necessarily trivial) so long as the number of seats and candidates is not too large.

3 Discussion & Future Work

Many paths are yet to be explored. To use goalbase ballots for multi-winner voting, we must restrict the weights available to voters. (E.g., if voters may select weights from $\mathbb{N}$ then voting becomes a game of who can write the largest number.) Having established that restrictions are needed, we face the problem of selecting some—it is not presently obvious which restrictions are most suitable. One candidate is to limit the sum of weights in any goalbase: $\sum_{(\varphi,w)\in G} w \le K$. This is a limit on the input space. Another approach is to restrict the output space: For example, we might limit the utility of any admissible state: $u_G(M) \le K$ for all $M \subseteq \mathcal{PS}$ where $|M| = k$.

There are advantages and disadvantages to each method. If our voter is a person, he will find it easier to cast a valid sum-limited ballot than a valid state-limited one. Input limits are not uncommon: For example, in the U.S., the State of Illinois used single-winner cumulative voting with a 3-point limit (i.e., 3-vote 3-support top-1) within districts for electing members of its House of Representatives, 1870–1980 [3, 6]. Corporate boards of directors are usually elected using cumulative voting, where the point limit for each voter is the number of shares he owns. We know of no uses of output limits. Presumably this is because it is hard to see when working in the input space whether output limits are being respected; output limits expect too much of the average voter. However, output limits on elections of the size human voters are likely to face will not be difficult for machines to enforce, so might be practical if voters use a computer-aided voting system. This is a user-interface issue.

The basic question here seems to be how to set the value of preferences which are not over single states against those which are. What is an appropriate measure of voting power here? Input limits seem to favor voters with top-heavy orders, output limits seem to favor bottom-heavy voters. One way of quantifying the effect that a proposed weight limit could have is by considering the *efficacy* of voters with different preferences under that weight limit. (The efficacy of a ballot for a voter is a measure of how often that voter will be pivotal if he casts that ballot.) Ideally, all voters would have equally efficacious ballots to cast. Brams and Fishburn [1, Chapter 5] calculate the efficacy of ballots for approval voting and find that not all ballots are equally effective. A similar analysis could be done for GSR, with an eye to which weight restrictions treat voters most equitably.

ACKNOWLEDGEMENTS

This work is not entirely disjoint with [4, Chapter 7], which was supported by a GLoRiClass fellowship funded by the European Commission (Early Stage Research Training Mono-Host Fellowship MEST-CT-2005-020841).

REFERENCES

[1] S. J. Brams and P. C. Fishburn, *Approval Voting*, Springer, 2nd edn., 2007.

[2] S. J. Brams, D. M. Kilgour, and M. R. Sanver, 'A minimax procedure for electing committees', *Public Choice*, **132**(3–4), 401–420, (2007).

[3] B. F. Moore, *The History of Cumulative Voting and Minority Representation in Illinois, 1870–1908*, volume III.3 of *The University Studies*, University Press, Urbana-Champaign, Illinois, 1909.

[4] J. Uckelman, *More Than the Sum of Its Parts: Compact Preference Representation Over Combinatorial Domains*, Ph.D. dissertation, University of Amsterdam, 2009. ILLC Publication DS-2009-12.

[5] J. Uckelman, Y. Chevaleyre, U. Endriss, and J. Lang, 'Representing utility functions via weighted goals', *Mathematical Logic Quarterly*, **55**(4), 341–361, (2009).

[6] 'Alternative voting systems as remedies for unlawful at-large systems', *The Yale Law Journal*, **92**(1), 144–160, (1982).