Towards a “Borda count” for judgment aggregation

William S. Zwicker[a]

[a]*Department of Mathematics, Union College, Schenectady, NY 12308*

Abstract

In social choice theory the Borda count is typically defined in terms of point scores awarded to individual alternatives, but in the broader context of judgment aggregation we cannot presume any underlying set of alternatives. To define a Borda judgment aggregator, then, we might seek some equivalent reformulation of Borda preference aggregation, sufficiently abstract to avoid reference to alternatives. We suggest here an approach based on the method of *orthogonal projection.* Loosely, a profile is viewed as a vector $\mathbf{v}$ decomposed into a first *infeasible* part and a second part orthogonal to the first:

$$\mathbf{v} = \mathbf{v}_{infeas} + \mathbf{v}_{\perp infeas}.$$

In the preference aggregation context, $\mathbf{v}_{infeas}$ is the *cyclic* component (tendency towards a majority cycle) while $\mathbf{v}_{\perp infeas}$ is *purely acyclic* (with zero tendency towards a cycle). The Borda outcome then turns out to be the ranking determined by the second component alone.

For judgment aggregation we can try something similar: discard the information in the infeasible component, basing the outcome on the *purely feasible* component alone. However, does this component retain enough information to yield anything useful or interesting – or by discarding $\mathbf{v}_{infeas}$ are we throwing the baby out with the bath water? The answer seems to depend on the particular feasibility constraints at hand.

Keywords: Borda count, judgment aggregation, orthogonal decomposition, purely feasible component

Extended abstract I: Borda preference aggregation

We'll begin with an analogy to Physics. Figure 1 shows a block B at rest on an inclined plane P, with $\mathbf{F}_G$, the force on the block due to gravity, decomposed as the sum of vectors $\mathbf{F}_{||P}$ (parallel to P) and $\mathbf{F}_{\perp P}$ (normal to P). That these components are orthogonal to each other suggests that they are independent in

Email address: `zwickerw@union.edu` (William S. Zwicker)

their effects. Of course they also interact, in that they have opposing implications for movement of the block. While $\mathbf{F}_{||P}$ tends to make the block slide, $\mathbf{F}_{\perp P}$ tends to make it stick, with the battle won by the "larger" of the two; B slides or sticks according to the sense of the inequality $||\mathbf{F}_{||P}|| \gtreqless C||\mathbf{F}_{\perp P}||$.

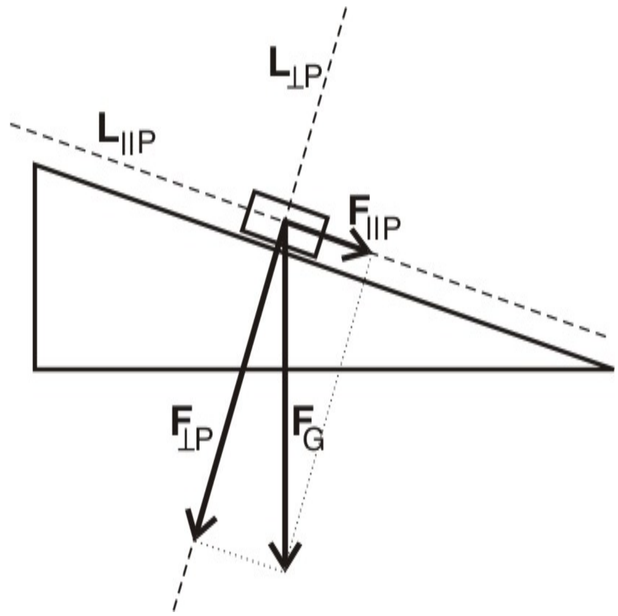

Figure 1: The orthogonal decomposition of the force due to gravity for a block resting on an inclined plane.

Here size of a component is measured the way we expect, via magnitude (Euclidean norm), and the constant C depends on the (static) coefficient of friction for the particular substances, as well as the contact area.

Now let's suppose that the origin of our coordinate system is located at the center of the block. The figure looks 2-dimensional, so we'll assume we are working in $\mathbf{R}^2$. We can then decompose $\mathbf{R}^2$ as the sum of two orthogonal subspaces (represented by the dashed lines), each of which is 1-dimensional in this example. Then $\mathbf{F}_{||P}$ is obtained as the *orthogonal projection* of $\mathbf{F}_G$ onto the subspace $\mathbf{L}_{||P}$ while $\mathbf{F}_{\perp P}$ is the corresponding projection of $\mathbf{F}_G$ onto $\mathbf{L}_{\perp P}$.

More generally, the situation for $\mathbf{R}^k$ is as follows. Take two subspaces $\mathbf{L}_0$ and $\mathbf{L}_1$ that are *orthogonal complements*, meaning that $v \in \mathbf{L}_j$ if and only if $v \perp w$ $(v \cdot w = 0)$ holds for each $w \in \mathbf{L}_{1-j}$. It follows that the dimensions of $\mathbf{L}_0$ and $\mathbf{L}_1$ sum to k, that each $v \in \mathbf{R}^k$ can be decomposed in one and only one way as $v = v_0 + v_1$ with $v_0 \in \mathbf{L}_0$ and $v_1 \in \mathbf{L}_1$, and that v_j is the orthogonal projection of v onto $\mathbf{L}_j$ for $j = 0, 1$.

To illustrate how the idea applies to preference aggregation, we'll look at an example – a profile P for $n = 30$ voters and $m = 4$ alternatives:

10	$a > b > c > d$
10	$b > c > d > a$
6	$b > a > c > d$
4	$a > b > c > d$

The Profile P

Notice that 20 voters prefer a to c and 10 prefer c to a, so we'll say that the net margin by which voters prefer a to c is $20 - 10 = 10$ (and the net margin by which voters prefer c to a is $10 - 20 = -10$). For each pair (x, y) of alternatives, the *net pairwise margin of x over y* is given by

$$Net_P(x > y) = [\text{the number of voters who prefer } x \text{ to } y] \\ -[\text{the number of voters who prefer } y \text{ to } x].$$

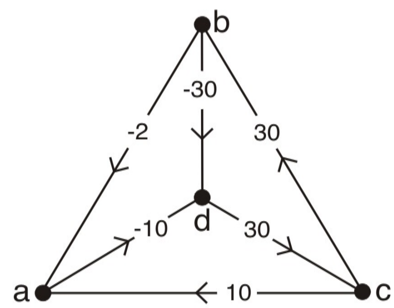

Figure 2: The net pairwise margins for profile P

Figure 2 shows the net pairwise margins for P as a flow on a directed graph G. Here G's vertices are the alternatives, and each pair of vertices is joined by a directed edge, whose direction is fixed arbitrarily. These directions serve to bookkeep the signs of the margins: the rule is that a directed edge $x \leftarrow y$ to x from y is labeled with $Net_P(x > y)$. This explains why the edge to c from d is labeled with the positive number $30 = Net_P(c > d)$ while that to a from b is labeled with the negative number $-2 = Net_P(a > b)$. (The directions are necessarily arbitrary because they stay fixed as P varies.) If we choose some enumeration of G's 6 edges, say

$$b \to a,\ b \to d,\ c \to b,\ a \to d,\ d \to c,\ c \to a$$

then by enumerating the edge labels in the corresponding order we can identify the flow of net pairwise margins with a vector in $\mathbf{R}^6$:

$$\mathbf{v}_P = (-2, -30, 30, -10, 30, 10).$$

In what follows we will blur the distinction between a profile P, the induced flow on G_P, and the enumerated vector $\mathbf{v} = \mathbf{v}_P$ of net pairwise margins, dropping the P subscript when the context allows.

Referring to Figure 2 one can quickly identify both the Condorcet and Borda outcomes for our profile P. We observe only positive labels on edges directed into b and only negative labels on edges directed out from b; all edges incident to vertex b thus indicate pairwise margins in b's favor. Hence b is a Condorcet alternative for P. Extending this reasoning, we see that pairwise majority rule yields the transitive ranking $b > a > c > d$. If we calculate the *net flow into* b by summing the labels on the b-incident edges (being careful to reverse the sign for each label on an edge directed *out* from b) we obtain b's *Borda score*, which is 62.[1] For our particular profile P, the *Borda outcome*, or ranking by Borda score, is the same as the Condorcet outcome: $b > a > c > d$.

Next, we consider the appropriate decomposition. A *basic cycle* is a directed unit flow around a simple closed loop of G. Figure 3 shows two basic cycles. Note that when the flow along some edge is zero we omit the numerical label. Also, the edge label is -1 when the direction of an edge on the loop is opposed to the direction in which the loop is taken. A *cycle* is any linear combination of basic cycles and the cycle subspace $\mathbf{V}_{cycle}$ is the subspace of all cycles – the linear span of the basic cycles. Thus $\mathbf{v} \in \mathbf{V}_{cycle}$ if and only if we can decompose $\mathbf{v}$ as a sum of scalar multiples of basic cycles.[2]

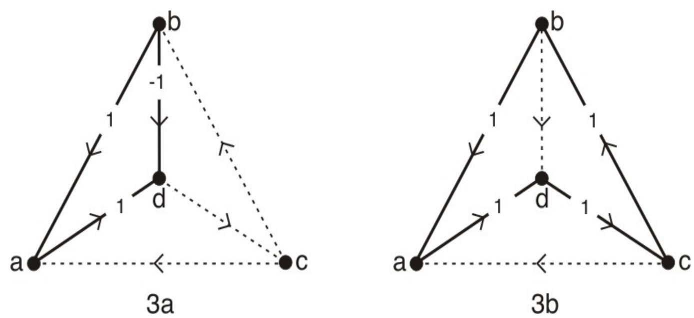

Figure 3: Two basic cycles

Let x be any alternative. The *basic cocycle for* x is a flow that labels each $x \leftarrow y$ edge with 1, each $x \rightarrow y$ edge with -1, and each edge not incident to x with 0. Figure 4 shows the basic cocycle for alternative b. A *cocycle* is any linear combination of basic cocycles and the cocycle subspace $\mathbf{V}_{cocycle}$ is the linear span of the basic cocycles. Thus $\mathbf{v} \in \mathbf{V}_{cocycle}$ if and only if we can decompose $\mathbf{v}$ as a sum of scalar multiples of basic cocycles.[3]

[1] With 4 alternatives, summing an alternative's margins in this way yields the Borda score via the vector $(3, 1, -1, -3)$ of scoring weights.

[2] This decomposition is not unique; the basic cycles are not linearly independent.

[3] Once again the decomposition is not unique.

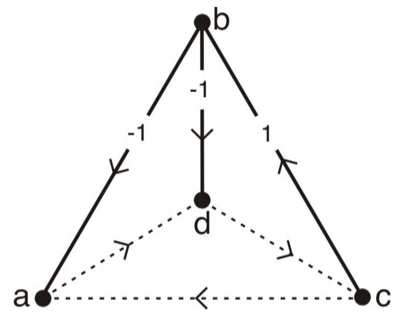

Figure 4: A cocycle

Let $\mathbf{v}_b = (-1,-1,1,0,0,0)$ denote the basic cocycle of Fig 4 and $\mathbf{v}_{3a} = (1,-1,0,1,0,0)$ denote the basic cocycle shown on the left in Fig 3. Then

$$\mathbf{v}_b \cdot \mathbf{v}_{3a} = (-1)(1) + (-1)(-1) + (1)(0) + (0)(1) + (0)(0) + (0)(0) = 0,$$

indicating that these vectors are orthogonal. It is not hard to see that every cocycle is similarly orthogonal to every basic cycle. It follows (from bilinearity of the dot product) that every linear combination of basic cocycles is orthogonal to every linear combination of basic cycles – that is, $\mathbf{V}_{cocycle} \perp \mathbf{V}_{cycle}$.

However, to show that $\mathbf{V}_{cocycle}$ and $\mathbf{V}_{cycle}$ are orthogonal complements in $\mathbf{R}^6$ (or, more generally, in $\mathbf{R}^{\binom{m}{2}}$) requires one more thing; we need to rule out the existence of any vector $\mathbf{w}$ outside $\mathbf{V}_{cocycle} \cup \mathbf{V}_{cycle}$ with $\mathbf{w} \perp (\mathbf{V}_{cocycle} \cup \mathbf{V}_{cycle})$. This will follow if we show that the dimensions of these two subspaces sum to that of the ambient space $\mathbf{R}^6$. With m alternatives, one can show that $dim(\mathbf{V}_{cocycle}) = m-1$, $dim(\mathbf{V}_{cycle}) = \frac{(m-1)(m-2)}{2}$, and these sum to the correct dimension $\binom{m}{2}$ (correct because with m alternatives, our graph G has $\binom{m}{2}$ edges, whence the dimension of the vector space of all flows on G is likewise $\binom{m}{2}$).[4]

Our sample profile P has $m = 4$ alternatives, so $dim(\mathbf{V}_{cocycle}) = 3$ and $dim(\mathbf{V}_{cycle}) = 3$.[5] If we take the vector $\mathbf{v}_P$ depicted in Fig 2 and project orthogonally onto the subspaces $\mathbf{V}_{cocycle}$ and $\mathbf{V}_{cycle}$ we obtain the decomposition $\mathbf{v}_P = \mathbf{v}_{cycle} + \mathbf{v}_{cocycle}$ shown in Figure 5. Figure 6 shows $\mathbf{v}_{cycle}$ written as the sum of three flows, each of which is a scalar multiple of a basic cycle.

---

[4]It is easy to show that any $m-1$ cocycles are linearly independent, and there is an easy inductive construction producing $\frac{(m-1)(m-2)}{2}$ linearly independent cycles (each of length 3); this shows the dimensions sum to $\binom{m}{2}$ or more, and "more" is not possible.

[5]It is also illustrative to consider the case of 3 alternatives, wherein the ambient space is $\mathbf{R}^3$, $dim(\mathbf{V}_{cocycle}) = 2$ and $dim(\mathbf{V}_{cycle}) = 1$, so that we can visualize the decomposition in 3-space. In this case, G is a triangle with 3 edges. If we orient these edges so that they "cycle" around the triangle, $\mathbf{V}_{cycle}$ becomes the one-dimensional line L through the origin with equation $x = y = z$ while $\mathbf{V}_{cocycle}$ becomes the 2-dimensional plane M perpendicular to L and through the origin, with equation $x + y + z = 0$.

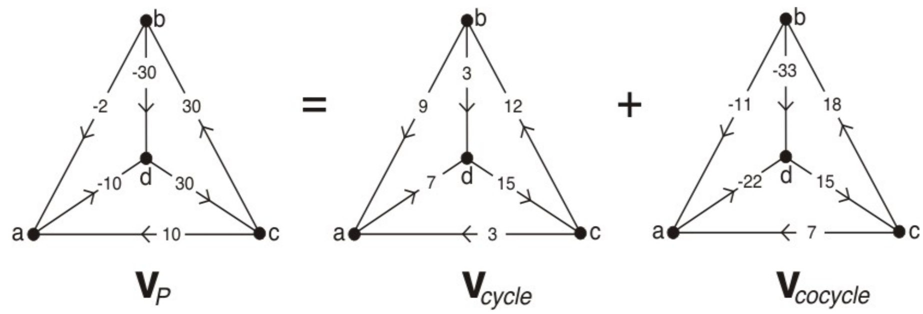

Figure 5: The decomposition $\mathbf{v}_P = \mathbf{v}_{cycle} + \mathbf{v}_{cocycle}$

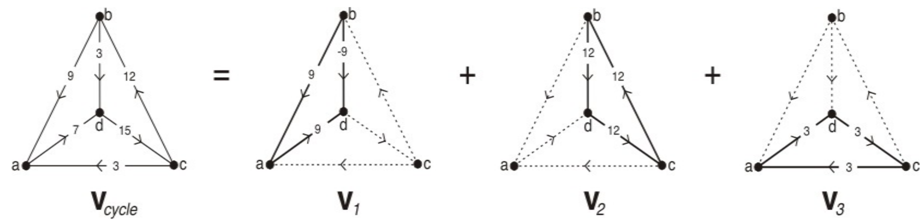

Figure 6: Expressing $\mathbf{v}_{cycle}$ as a linear combination of basic cycles

Why should we care about any of this?

1. Borda scores depend only on the $\mathbf{v}_{cocyclic}$ component of the net pairwise margins vector. In fact, the vector of Borda scores and $\mathbf{v}_{cocyclic}$ are essentially the same thing.
2. If we interpret the $\mathbf{v}_{cocyclic}$ flow as a tournament on the alternatives (by ignoring the magnitudes of the numerical labels, while noting only their signs along with the initial orientations assigned to G's edges) then the result is always a transitive relation. Thus, using the Borda count is tantamount to applying pairwise majority rule after suppressing $\mathbf{v}_{cycle}$.
3. Much as $\mathbf{F}_{||P}$ and $\mathbf{F}_{\perp P}$ have opposing implications for movement of the block on the inclided plane, $\mathbf{v}_{cycle}$ and $\mathbf{v}_{cocycle}$ have opposing implications for transitivity of the majority preference relation. If $\mathbf{v}_{cocycle}$ is large enough (relative to $\mathbf{v}_{cycle}$) then $\mathbf{v}_{cocycle}$ dominates, so that the Borda and Condorcet outcomes are identical (as was the case for our sample profile P).[6] If $\mathbf{v}_{cocycle}$ is too small, so that $\mathbf{v}_{cycle}$ dominates, then the majority preference relation is intransitive – a Condorcet cycle exists. Depending on the particular profile, however, there may be an intermediate situation for which including $\mathbf{v}_{cycle}$ *changes* the majority preference relation that

[6]The analogy with the inclined plane is not perfect, in that we are using "large enough" as a loose metaphor here – the precise version is not posed in terms of the Euclidean norm, as it was for the inclined plane. See [6] for details.

would result from $\mathbf{v}_{cocycle}$ alone, but without throwing it into a cycle. This is what happens when the Condorcet alternative does not prevail in a Borda election, for example.

In terms of point 1, note (Fig 5) that the Borda scores arising from $\mathbf{v}_P$ (calculated via net flow into each vertex, as we illustrated earlier) are exactly the same as those from $\mathbf{v}_{cocycle}$. In fact, it is easy to see why: any basic cycle contributes 0 to the Borda score of every alternative. Before we adapt these methods to the more general context of judgment aggreagation, we illustrate point 3, above, by adding multiples of a 4-voter Condorcet cycle to our initial profile P.

1	$a > b > c > d$
1	$b > c > d > a$
1	$c > d > a > b$
1	$d > a > b > c$

The Profile Q

It is easy to see that $\mathbf{v}_Q = 2\mathbf{v}_{3b}$; that is, to get $\mathbf{v}_Q$, just double all the edge labels in the basic cycle of Figure 3b. Now consider the profile $P + kQ$, obtained by adding k disjoint groups of voters (each with 4 voters, and each having Q as its profile) to the 30-voter electorate for P.

What is the effect of gradually increasing k on the Borda outcome ... or on the pairwise-majority relation? As the cocyclic component of Q is $\mathbf{0}$, Q contributes 0 to the Borda score of each alternative, and the effect on the Borda outcome is nil; regardless of the k value, the Borda outcome is $b > a > c > d$. The situation for the pairwise majority outcome is quite different, however (see Figure 7).

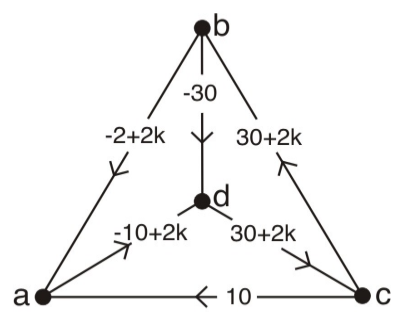

Figure 7: The net pairwise margins for $\mathbf{v}_{P+kQ}$

With $k = 0$, the two outcomes are the same. With $k = 2$ the sign of the edge label on $a \leftarrow b$ switches from $-$ to $+$, and the Condorcet outcome changes from $b > a > c > d$ to $a > b > c > d$, reflecting the intermediate situation described in point 3 (previous page); any Condorcet extension (Black's rule, for example) now declares a to be the winner. If we let k grow to the point where $k \geq 6$ then the sign of the edge label on $d \leftarrow a$ also becomes $+$, and the pairwise majority relation becomes intransitive; at this point Black's rule switches back to declaring b as the winner – a failure of *reinforcement* (aka "consistency").

Extended abstract II: Borda judgment aggregation

We'll assume that each of several voters (or experts) in the finite set N renders a binary judgment on each of k issues, so that a ballot is a k-tuple of $+1$s (or *yes* votes) and -1s (*no* votes). A *profile* P is a vector of ballots, with one ballot P_j per voter j. An outcome $F(P)$ is likewise a a k-tuple of $+1$s and -1s, as determined by some aggregation rule F, and we require that $F(P)$ lie inside some specified set $Feas \subseteq \{0, -1, +1\}^k$ of *feasible* outcomes, where we assume that the "completely tied" vector $\mathbf{0}$ of k 0s is feasible. For reasons that will become clear, we'll start by placing no corresponding restrictions on ballots – anything in $\{+1, +1\}^k$ is allowed.

Given a profile P, consider the issue-wise sum ΣP of ballots P_j for $j = 1, 2, \ldots, n$. (We are summing n vectors in $\{+1, -1\}^k$ to obtain a single vector in $\mathbf{R}^k$.) The *strictly issue-wise* rule $\mathcal{IW}$ outputs a $= 1, 0, -1$ for the i^{th} issue according to whether the i^{th} component of ΣP is positive, zero, or negative respectively; i.e., $\mathcal{IW}$ employs majority rule separately on each issue. Of course, $\mathcal{IW}(P)$ may be infeasible. We'll say that an aggregation rule G is *issue-wise* if $G(P) = \mathcal{IW}(P)$ holds for each profile P such that $\mathcal{IW}(P)$ is feasible.

We'll compare the behavior of such an issue-wise aggregator to that of the *Borda judgment aggregator* $\mathcal{B}$ in the context of three closely related examples:

Example 1a, 1b, 1c
There are 3 issues, so that each ballot is an ordered triple. For $1a$ we'll assume each issue is a proposal to spend money on a project: build a swimming pool, build a senior center, repave the streets in some neighborhood. However, budgetary constraints mean the town can afford to undertake at most two of the three projects. Thus, $(+1, +1, +1)$ is infeasible as an outcome, while the other 7 vectors of $+1$s and -1s are feasible. (We are postponing a careful discussion of 0s in the outcome.)

For $1b$ the voters are assessing the truth value of three propositional *wffs*: $p \wedge (\neg q \vee \neg r)$, q, $r \vee (\neg p \wedge \neg q)$. A truth table shows that there exist no truth assignments (to p, q, and r) making all three *wffs* true, or making all three false, but any other combination of true-false can be realized by some such assignment. Thus the infeasible outcomes are $(+1, +1, +1)$ and $(-1, -1, -1)$.

Example $1c$ will represent preference aggregation. The issues are assertions about pairwise preference among 3 alternatives: $a > b$, $b > c$, $c > a$. We'll assume, for example, that a vote of -1 on the first comparison indicates $b > a$ (we won't allow individual voters to express indifference). The feasibility constraints are those ruling out the two strict cycles, so the infeasible outcomes are $(+1, +1, +1)$ and $(-1, -1, -1)$.

Notice that $1b$ and $1c$ are formally identical, differing only in the interpretation we place on the issues, while $1a$ is slightly different in that $(-1, -1, -1)$

is feasible. However, $(-1,-1,-1)$ is a scalar multiple of $(+1,+1,+1)$, and so this difference disappears when we form the linear span $\mathbf{L}_{infeas}$ of the infeasible outcomes; for all parts of Example 1 we obtain a common one-dimensional subspace of $\mathbf{R}^3$:

$$\mathbf{L}_{infeas} = \{(x,x,x)|x \in \mathbf{R}\}$$

Also common to all of Example 1, the orthogonal complement $\mathbf{L}_{purelyfeas}$ of $\mathbf{L}_{infeas}$ in $\mathbf{R}^3$ is a (two-dimensional) plane through the origin (see footnote 5):

$$\mathbf{L}_{purelyfeas} = \{(x,y,z)|x+y+z=0\}.$$

Any vector $v \in \mathbf{R}^3$ now has a unique decomposition as a sum

$$\mathbf{v} = \mathbf{v}_{infeas} + \mathbf{v}_{\perp infeas}$$

with $\mathbf{v}_{infeas} \in \mathbf{L}_{infeas}$ and $\mathbf{v}_{\perp infeas} \in \mathbf{L}_{purelyfeas}$. Moreover, these components are obtained as orthogonal projections onto the corresponding subspaces: $\mathbf{v}_{infeas} = Proj_{infeas}(\mathbf{v})$ and $\mathbf{v}_{\perp infeas} = Proj_{purelyfeas}(\mathbf{v})$.

Given a profile P, the *Borda judgment aggregator* $\mathcal{B}(P)$ is then defined by the following steps:

- Take each individual ballot P_i cast by a voter i and replace it by its orthogonal projection $Proj_{purelyfeas}(P_i)$ onto $\mathbf{L}_{purelyfeas}$ (discarding P_i's infeasible component).
- Form the sum $\sum_{i=1}^n Proj_{purelyfeas}(P_i)$ of these projected ballots.
- Output $1, 0$, or -1 for the i^{th} issue according to whether the i^{th} component of $\sum_{i=1}^n Proj_{purelyfeas}(P_i)$ is positive, zero, or negative respectively.

One can get the general idea of what is going on from the following orthogonal decompositions for three of eight possible ballots. The purely feasible component precedes the infeasible component on the right side of each equation:

$$\begin{aligned}
(+1,+1,-1) &= (+\frac{2}{3},+\frac{2}{3},-\frac{4}{3}) + (-\frac{1}{3},-\frac{1}{3},-\frac{1}{3}) \\
(+1,-1,-1) &= (+\frac{4}{3},-\frac{2}{3},-\frac{2}{3}) + (+\frac{1}{3},+\frac{1}{3},+\frac{1}{3}) \\
(+1,+1,+1) &= \quad (0,\ 0,\ 0) \qquad + \ (+1,\ +1,\ +1)
\end{aligned}$$

Consider a profile $P^\star$ in which the number of voters casting a common ballot of $(+1,+1,+1)$ is A, the number casting ballot $(+1,+1,-1)$ is B, etc. (recalling that we are allowing individual ballots to be infeasible). The table below shows, in the third column, the total contribution made by each such batch of voters to the issue-wise sum $\sum_{i=1}^n P_i$, and the last column shows the corresponding contribution after projecting onto $\mathbf{L}_{purelyfeas}$.

Ballot	No. Votes	Issue-wise Contribution	Borda Contribution
$(+1,+1,+1)$	A	$(+A,+A,+A)$	$(0,0,0)$
$(+1,+1,-1)$	B	$(+B,+B,-B)$	$(+\frac{2B}{3},+\frac{2B}{3},-\frac{4B}{3})$
$(+1,-1,+1)$	C	$(+C,-C,+C)$	$(+\frac{2C}{3},-\frac{4C}{3},+\frac{2C}{3})$
$(-1,+1,+1)$	D	$(-D,+D,+D)$	$(\frac{-4D}{3},+\frac{2D}{3},+\frac{2D}{3})$
$(+1,-1,-1)$	E	$(+E,-E,-E)$	$(+\frac{4E}{3},-\frac{2E}{3},-\frac{2E}{3})$
$(-1,+1,-1)$	F	$(-F,+F,-F)$	$(-\frac{2F}{3},+\frac{4F}{3},-\frac{2F}{3})$
$(-1,-1,+1)$	G	$(-G,-G,+G)$	$(-\frac{2G}{3},-\frac{2G}{3},+\frac{4G}{3})$
$(-1,-1,-1)$	H	$(-H,-H,-H)$	$(0,0,0)$

Table 1

Thus the *strictly issue-wise* outcome $\mathcal{IW}(P^\star)$ uses the signs of the components in the vector of *issue-wise sums* (obtained as the sum of the third column):

$$\begin{aligned}\sum_{i=1}^{n} P_i &= (\mathcal{IW}_1, \mathcal{IW}_2, \mathcal{IW}_3)\\ &= ([A+B+C+E]-[D+F+G+H],\\ &\qquad [A+B+D+F]-[C+E+G+H],\\ &\qquad\quad [A+C+D+G]-[B+E+F+H])\end{aligned}$$

The *Borda judgment* outcome $\mathcal{B}(P^\star)$ uses instead the signs of the components from the vector of *Borda sums*, obtained as the fourth column sum and scaled by a factor of $\frac{3}{2}$ to clear fractions without changing signs:[7]

$$\begin{aligned}\frac{3}{2}\sum_{i=1}^{n} Proj_{purelyfeas}(P_i) &= (\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3)\\ &= ([B+C+2E]-[2D+F+G],\\ &\qquad [B+D+2F]-[2C+E+G],\\ &\qquad\quad [C+D+2G]-[2B+E+F])\end{aligned}$$

We'll make two observations comparing the vectors $(\mathcal{IW}_1, \mathcal{IW}_2, \mathcal{IW}_3)$ and $(\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3)$:

1. A and H terms appear in $(\mathcal{IW}_1, \mathcal{IW}_2, \mathcal{IW}_3)$, but not in $(\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3)$.
2. $(\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3) = \frac{2}{3}(\mathcal{IW}_1, \mathcal{IW}_2, \mathcal{IW}_3) + (\beta, \beta, \beta)$, where (β, β, β) translates $\frac{2}{3}(\mathcal{IW}_1, \mathcal{IW}_2, \mathcal{IW}_3)$ so that the mean of the three coordinates is 0.

[7]It would be misleading to use "Borda scores" in place of "Borda sums" here, because in the preference aggregation context $(\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3)$ does not correspond exactly to the vector of Borda scores, but rather to the cocyclic component of the orthogonal decomposition of $(\mathcal{IW}_1, \mathcal{IW}_2, \mathcal{IW}_3)$. For example, with $m = 4$ alternatives the vector of Borda scores is (of course) a 4-tuple, while the vector of Borda sums is a 6-tuple. However, in a sense that can be made quite precise, these two vectors are *isomorphic* to each other.

What do these observations imply about the behavior of issue-wise aggregators VS the Borda aggregator?

Observation 1 tells us that while votes of $(+1, +1, +1)$ or $(-1, -1, -1)$ have no effect on the Borda outcome, these ballots do count for issue-wise aggregators. Let's consider the issue-wise effect of such ballots for each of the distinct contexts of Example 1. In $1a$ (three projects with budget limit), a voter who finds the overall town tax rate to be already too high might wish to vote $(-1, -1, -1)$; with enough such votes, the collective decision might be to fund none of the three projects. Alternately, the $(-1, -1, -1)$ ballots might tip the balance, so that only one project gets funded rather than two. In that case, the budget-slashing voter may not know, in advance, which of the three projects her vote will defeat, so $(-1, -1, -1)$ may be her wisest choice. Although a $(+1, +1, +1)$ by itself represents an infeasible outcome, some of the same rationale applies; a voter who likes all three projects (perhaps feeling that the tax-cutters are destroying the town's quality of life) might wish to cast the infeasible $(+1, +1, +1)$, so as to maximizes the chance that one or two of the projects be approved (rather than none, or only one).

In $1b$ a somewhat parallel argument can be applied. If some expert has separately and independently judged each of the three propositions and decided each to be true (or decided each to be false), she will know that at least one of her three judgments must be mistaken, but may have reason to believe that the other two are probably correct (without knowing, of course, *which* two these are). Then, in the spirit of the Condorcet Jury Theorem, her ballot might make a net positive contribution to the probability that the group decision is correct, even though that ballot is known to be partly flawed. Whether this scenario can actually occur would seem to depend on the probability distribution at hand.[8]

In preference aggregation (Example $1c$) the $(-1, -1, -1)$ and $(+1, +1, +1)$ ballots represent cycles, and are typically disallowed. That may seems appropriate in a political context (an election for mayor, say) because our earlier rationale for infeasible ballots seems implausible – most of us do not rank candidates by making separate and independent judgments on the corresponding pairwise comparisons. Suppose the context is different, however – experts are determining the optimal time order in which to perform a sequence of operations, for example. If each expert weighs the potential advantages of performing operation i before j against those that favor j before i then it again seems plausible that some exerts might wish to cast ballots they know to be infeasible..

Turning now to the Borda aggregator, as votes of $(+1, +1, +1)$ or $(-1, -1, -1)$ are ignored by $\mathcal{B}$, such ballots are "wasted," and in the interests of openness and honesty we might choose to disallow their use when $\mathcal{B}$ is used. We might also take the above arguments for allowing these two ballots as constituting a

[8] I have not checked whether there exist plausible probability distributions for which casting a $(+1, +1, +1)$ is better (for correctness of the group decision) than would be abstaining, for this hypothetical voter. However, something similar seems easy to arrange with a larger number of issues when the logical constraints rule out only the two outcomes $(+1, +1, \ldots, +1)$ and $(-1, -1, \ldots, -1)$.

rationale against the Borda aggregator.

Before reaching any final conclusions, however, let's consider the implications of observation 2. Figure 8 shows the values of $\mathcal{IW}_1, \mathcal{IW}_2, \mathcal{IW}_3, \mathcal{B}_1, \mathcal{B}_2$, and $\mathcal{B}_3$ for some hypothetical (and unspecified) profile $P^\dagger$. They have been plotted on two vertical copies of the real line (with positive end up). What is common is the order of finish (the third issue accruing the greatest support, and the second issue the least), as well as the relative sizes of the gaps; for both the strictly issue-wise and Borda aggregators, issue 1 leads issue 2 by $\frac{7}{4}$ as much as the amount by which issue 3 leads issue 1.

What is different is the location of the 0. For the Borda sums, we know that the average value of $\mathcal{B}_1, \mathcal{B}_2$, and $\mathcal{B}_3$ is 0, and this determines the location of 0 relative to these sums. Figure 8 depicts a profile for which the middle sum is above the midpoint of the top and bottom sums, so 0 must appear between the middle and bottom sums. Thus issues 1 and 3 are collectively approved, with issue 2 rejected. But if $\mathcal{B}_1$ were moved down a bit, so that it fell below the $\mathcal{B}_2$, $\mathcal{B}_3$ midpoint, then 0 would lie between $\mathcal{B}_1$ and $\mathcal{B}_3$. Only the top-scoring issue 3 would then be approved. To summarize, in examples 1a, b, and c the Borda aggregator always approves the issue achieving greatest support, rejects the issue receiving least support, and approves or rejects the middle issue according to whether its support level comes closer to that of the approved issue or of the disapproved issue.

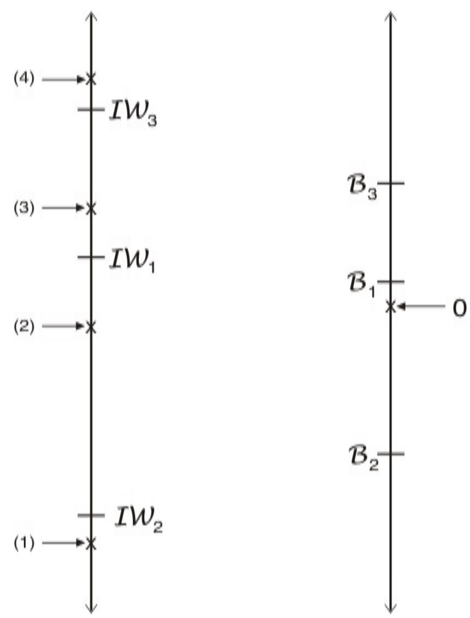

Figure 8: Issue-wise VS Borda aggregation
for Example 1, with profile $P^\dagger$

For the issue-wise sums, however, the 0 can appear in any position relative to the individual sums – for example, in any of the numbered positions shown on the left side of Fig 8. It all depends on the component of $P^\dagger$ ignored by $\mathcal{B}$. The issue-wise 0 location might lead to an infeasible outcome, or to a feasible outcome that differs from that of the Borda aggregator. For example, one can cook up a profile P that yields a version of Figure 8 for which the 0 on the $\mathcal{IW}$ scale is at position (3); any issue-wise aggregator then approves issue 3 only, while the Borda aggregator approves both issues 1 and 3.[9]

Extended abstract III: Conclusions, speculation

Based on the the examples and analysis presented here, should we think of the Borda aggregator as being only *slightly* different from issue-wise aggregators, or significantly different? Does that difference favor the use of the Borda aggregator, or argue against it?

At this point, we've only begun to look at Borda aggregation, and of course the examples considered here are extremely limited; in particular, the infeasibility subspace is only 1-dimensional for all parts of Example 1. So of course any answers will be quite tentative. Nevertheless, there is some reason to suspect that the answer to both questions depends quite a bit on the context.

If we consider Example 1c (preference aggregation) then whether one approves or disapproves of the middle-ranked issue (which, recall, is a pairwise comparison of two alternatives, *not* an individual alternative) can determine whether or not the Condorcet alternative wins the election. This difference is not considered to be slight, regardless of one's stance on scoring rules VS Condorcet extensions.

Perhaps in Example 1a (three projects with budget restriction) one might argue that the difference in approach is not so very great, as Borda and issue-wise approaches agree in their rankings of projects by level of support. Borda aggregation seems somewhat unappealing for this context, however, as one crucial difference is that Borda will always fund at least one project. Thus there is a feasible outcome that the Borda aggregator never chooses, and in particular the "no new taxes" group would certainly object to using $\mathcal{B}$. There are hints, however, that Borda aggregation might do some interesting things when there are a larger number of projects with varying costs, with a budget constraint over the total cost rather than the number of projects.[10]

Borda judgment aggregation seems potentially quite attractive, however, for Example 1b (experts judging the truth of three *wffs*, when consistency rules out that all three be true and that all three be false, but rules out nothing else).

---

[9] Moreover, such profiles P exist even if one disallows $(+1, +1, +1)$ and $(-1, -1, -1)$ ballots.

[10] Imagine that the voters cast a ballot that includes a real number (interpreted as the voter's ideal value for the total budget available for new projects), and this component is aggregated via the median (which is strategy-proof), to obtain a collective cap T on new spending. They simultaneously vote separately on each proposed project. The projects are funded in descending order of popularity, and are supported up to, but not including, the point at which T is exceeded. This procedure seems to be closely related to a modified version of Borda aggregation.

Imagine that each of the three *wff*s is judged to be true by more than half of the experts, leading to the infeasible $(+1, +1, +1)$ outcome. We might take that as evidence of some underlying bias towards $+1$ over -1. Perhaps each expert votes randomly and independently on each wff, with probabilities of $\frac{1}{2} + \epsilon$ of being correct, and $\frac{1}{2} - \epsilon$ of being incorrect (as in the standard assumptions for the Condorcet jury theorem), to which some (positive or negative) bias δ has been added: the probability of voting $+1$ is incremented by δ and that of voting -1 is decremented by the same amount. If we do not know δ's value in advance, but do know the probability distribution governing δ's selection, then we might use the issue-wise outcomes to help estimate the selected value of δ, and then compensate by translating the issue-wise scores accordingly. This seems to describe roughly what the Borda aggregator is doing.

The genesis of these ideas

The decomposition into cyclic and cocyclic components of a flow on a directed graph has been well-known to algebraic topologists for some time, and represents a special case of the boundary map of homology theory (see Croom [1], for example). Its application to the flow of electric current in a circuit is related to Kirchoff's Laws (Harary [2]). This same decomposition was applied to social choice in Zwicker [6], where it leads to a refinement of the *value restriction* analysis of Sen [5]. Applications to social choice were further developed by Saari [4].

At the December 2010 summer workshop held under the auspices of the *Centre for Mathematical Social Science* of the University of Auckland, I heard Clemens Puppe present a recent paper with Klaus Nehring and Marcus Pivato [3]. In his talk, Clemens mentioned the problem of extending the Borda count to the more general context of judgment aggregation.

[1] F. Croom, Basic concepts of algebraic topology, Springer Verlag, New York (1978).

[2] F. Harary, Graph theory and electrical networks, IRE Trans. Circuit Theory CT-6, 95-109 (May, 1958).

[3] K, Nehring, M. Pivato, C. Puppe, Condorcet admissibility: Indeterminacy and path-dependence under majority voting on interconnected decisions, preprint (July, 2011).

[4] D. Saari, Geometry of voting, Studies in Economic Theory 3, Springer Verlag, Berlin (1994).

[5] A. K. Sen, A possibility theorem on majority decisions, Econometrica 34 (2), 491-499 (1966).

[6] W. S. Zwicker, The voters' paradox, spin, and the Borda count, Mathematical Social Sciences 22, 187–227 (1991).