Reprinted from Information Pooling and Group Decision Making. B. Grofman and G. Owen (eds.) Greenwich, Conn.: JAI Press, 1986.

OPTIMAL RANKING AND CHOICE FROM PAIRWISE COMPARISONS

H. P. Young

1. INTRODUCTION

In his famous essay on the mathematical theory of voting (*Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*), Condorcet (1785) examined voting schemes in terms of the likelihood that they yield "correct" results. He posited that voters act as honest but imperfect judges of the objective *merit* of different candidates or the truth of different propositions. The question he asked is: what voting procedure yields the candidate (or proposition) most likely to be the best or correct one? The answer, when there are exactly two alternatives, is simple majority rule—assuming that individual voters vote independently and are correct in their judgments more than half of the time. Moreover, the probability of the majority opinion being correct increases the more voters there are.

Condorcet goes on to pose a more complex question: what is the analog of simple majority rule—in the sense of the most likely to produce a correct result—when there are more than two alternatives? His well-known answer is: the candidate who obtains a majority over every other candidate (the so-called *majority* candidate) *assuming that such a candidate exists.* In case a majority candidate does not exist due to voting cycles, it is generally thought that Condorcet stated no clear method for resolving the inconsistencies in the voters' opinions. In fact he did propose a very natural method, but then confused his readers by suggesting a plausible heuristic algorithm for computing the rule that does not always give the answer he apparently had in mind.

In this paper we examine Condorcet's method as a statistical rule of inference. From a maximum likelihood standpoint, the method gives the ranking of all candidates that is most likely to be correct. When the objective is to choose the single best or correct candidate from among n candidates, however, Borda's method seems to be a superior rule of inference. The difficulty with Condorcet's approach for estimating the strongest single candidate is that the top-ranked candidate in a most likely ranking need not be the candidate that is most likely to beat a randomly chosen opponent.

Even if we reject the specific probabilistic model by which this conclusion is reached, there are still strong a priori grounds for asserting that the rules of Condorcet and Borda are the optimal rules for ranking and choice respectively. Indeed they can be uniquely characterized as such in terms of certain simple *axioms of inference* not predicated on *any* probabilistic model.

The concluding section discusses extensions of the results to situations in which inferences are drawn from sources whose a priori reliabilities differ.

2. MAXIMUM LIKELIHOOD METHODS

The problem of choosing a "best" or "correct" alternative, or a "correct" ordering of all alternatives based on a sample of pairwise comparisons, will be treated as a problem of statistical inference using maximum likelihood estimation. The objective is to estimate the most probable true state (i.e., the correct alternative or the correct ranking) consistent with the outcomes of a sequence of binary trials such as votes, contests, taste tests, weighings, and so on, assuming no prior information on the probability of the states.

To fix ideas, let E be a finite set of m possible *events* $E = \{1, 2, \ldots, m\}$ and Θ a set of possible *states*. Let $g(e|\theta)$ be the probability of observing $e \in E$ given that $\theta \in \Theta$ is the true state. An *outcome* is an m-dimensional vector $\mathbf{x}$ such that for each $e \in E$, x_e represents the number of times e occurred in a sequence of $N = \sum_E x_e$ independent observations. The probability of observing $\mathbf{x}$ given that the true state is θ is given by the *likelihood function*:

$$L(\mathbf{x}; \theta) = \frac{N!}{x_1!x_2!\ldots x_m!} \prod_{e \in E} (g(e|\theta))^{x_e}. \tag{1}$$

The *maximum likelihood decision rule* is to choose the state(s) θ that maximize (1):

$$f(\mathbf{x}) = \{\theta \in \Theta: L(\mathbf{x}; \theta) \geq L(\mathbf{x}; \theta'), \text{for all } \theta' \in \Theta\}.$$

Taking logarithms, this is equivalent to

$$f(\mathbf{x}) = \left\{\theta \in \Theta: \sum_{e \in E} x_e \log g(e|\theta) \geq \sum_{e \in E} x_e \log g(e|\theta') \text{ for all } \theta' \in \Theta\right\}. \quad (2)$$

3. BORDA AND CONDORCET AS MAXIMUM LIKELIHOOD RULES

Let $\{a_1, \ldots, a_n\} = A$ be n distinct *alternatives* (objects, candidates, contestants) that differ in some objective attribute. For the sake of specificity, assume that the alternatives are objects that differ slightly in *weight* and can be compared pairwise by means of a balance. The balance, however, is faulty: it only gives the correct comparison with a fixed probability p that is assumed to be unknown but larger than $\frac{1}{2}$.

Two cases will be considered. In one there is a single distinguished object, for example, one that is *heavier* than the others, and the goal is to correctly identify it after a series of independent "weighings." In the second situation, the objects are assumed to be slightly different in weight, and the problem is to correctly *rank* them. In the latter case it is also of interest to find the one most likely to be the heaviest.

Thus we have before us two distinct problems: to infer a correct ranking of alternatives (by weight) or to choose the "best" (weightiest), based on the outcome of binary comparisons. A rule of inference that does the former is called a *binary ranking rule*; one that does the latter is a *binary choice rule.* It turns out that the optimal rules for these two situations may differ. This point, almost obvious intuitively, is often obscured by the circumstantial fact that most choice rules implicitly give a ranking; conversely, a ranking rule suggests choosing the top-ranked alternative. This is a specious manner of reasoning that does not necessarily give optimal results.

A series of pairwise comparisons on a set of n alternatives A can be represented by an $n(n-1)$-dimensional vector $\mathbf{x} = (x_{ij})$, $1 \leq i \neq j \leq n$, having nonnegative integer coordinates, where x_{ij} represents the *number* of weighings in which a_i was "heavier" than a_j. For simplicity assume that there are no draws and that every alternative is involved in the same number of comparisons c. All pairwise trials are assumed to be independent, and in every trial there is some fixed (but unknown) probability of success $p > \frac{1}{2}$ that the heavier of the two objects will be selected; if the objects are identical, each will be selected with probability exactly $\frac{1}{2}$. The assumption of a fixed p is most relevant to situations in which the objects are similar and difficult to distinguish. Indeed, the more dissimilar objects are, the easier it becomes to distinguish them, so a comparison of a top-ranked with a bottom-ranked alternative would have a significantly greater probability of success than a comparison of two nearby ones.

Consider first the case in which there is exactly one distinguished alternative. Then there are n distinct states and the likelihood that $\mathbf{x}$ occurs given that a_k is the distinguished alternative is proportional to

$$g(\mathbf{x}|a_k) = p^{\sum_{i\neq k} x_{ki}}(1-p)^{\sum_{i\neq k} x_{ik}}. \tag{3}$$

Since for all $i \neq j$, $x_{ij} + x_{ji} = c$ by assumption, the maximum likelihood decision rule is

$$f(\mathbf{x}) = \left\{a_k \in A: \ \sum_{i\neq k} x_{ki} \geq \sum_{i\neq j} x_{ji}, \text{ for } 1 \leq j \leq n\right\}. \tag{4}$$

In other words, the optimal rule is to choose the alternative(s) with the maximum number of wins minus losses over all pairwise comparisons. This is a straightforward generalization of a method first suggested by Jean-Charles de Borda (1781) and later expressed in essentially the above form by Duncan Black (1958). Note the important fact that *the solution is independent of the particular value of* p, *so long as* $p > \frac{1}{2}$.

Suppose now that the objects have slightly different weights and the objective is to correctly *rank* them. A *ranking* is a linear order R on the alternative set A such that $a_i R a_j$ means a_i ranks before a_j (a_i is "heavier" than a_j). For any outcome $\mathbf{x}$ and ranking R, define $\tau(\mathbf{x}, R)$ as the sum of all x_{ij} such that $a_i R a_j$. The probability of observing $\mathbf{x}$ given that the true ranking is R is proportional to

$$L^*(p, R) = p^{\tau(\mathbf{x},R)}(1-p)^{\binom{n}{2}c-\tau(\mathbf{x},R)}. \tag{5}$$

Maximizing $L^*(p, R)$ with respect to R is equivalent to maximizing $\tau(\mathbf{x}, R)$ with respect to R, that is, to choosing the ranking(s) consistent with the greatest number of successes:

$$f(\mathbf{x}) = \left\{R: \tau(\mathbf{x}, R) = \sum_{(i,j):a_i R a_j} x_{ij} = \max\right\}. \tag{6}$$

Although the rule given by (6) has been rediscovered many times since, Condorcet in fact proposed it in his 1785 *Essai.*[1] Unfortunately, he gave the following heuristic algorithm for computing it that, while plausible, is not always optimal. Begin with the set P of all ordered pairs (a_i, a_j) such that $x_{ij} > x_{ji}$. If P contains cycles, find a pair (a_i, a_j) that is contained in a cycle and has smallest majority in its favor. Replace it by (a_j, a_i). Continue in this manner until P contains no cycles, and hence accords with some transitive ranking R.

Finding good algorithms for the linear optimization posed by (6) is an interesting mathematical problem. It amounts to finding a maximum-weight acyclic subgraph of the complete directed graph on n vertices $1, 2, \ldots, n$, in which each directed edge (i, j) has weight x_{ij}. For a discussion of computational aspects of Condorcet's rule, see Barthelemy and Monjardet (1981) and Young (1985).

Condorcet's rule was proposed as a method of social choice in a 1959 paper by John Kemeny (without being attributed to Condorcet). It is closely related to M. G. Kendall's measure of rank correlation, since (6) amounts to finding the ranking R that is maximally rank-correlated with a given outcome **x**.

Condorcet advocated his method not only as a rule for determining the most likely ranking but for determining the most likely best choice. Strictly speaking this is consistent with a maximum likelihood interpretation. But there are situations where Borda's rule may give a better estimate of the true top candidate. Consider the example shown in Figure 1. According to Condorcet's rule, (A, B, C) is the most likely ranking, and A is the best choice. Note that A has a majority over B and also over C. However, the Borda scores for A, B, and C are 13, 14, and 6 respectively, so Borda's method would choose B. We claim that there is good reason to choose B instead of A. Indeed, B has the highest probability of beating a randomly chosen opponent in any single pairwise trial. B's probability of beating C is 10/11, and of beating A is 4/11, which means it beats a random opponent

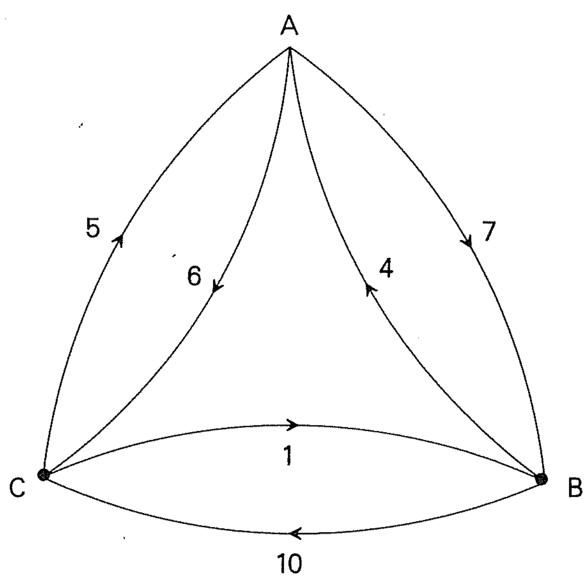

Figure 1. Borda chooses B while Condorcet chooses A.

with probability 7/11. On the other hand, A beats a random opponent with a probability of only 13/22.

There is a further argument that lends support to Borda's rule as a general method for inferring the most likely best candidate in certain cases. Let us view the likelihood function $L^*(p, R)$ as a posterior estimate of the relative probabilities of all rankings R. (This implicitly assumes a uniform prior distribution on the rankings.) Then the relative probability that a particular candidate is top-ranked is the sum of the likelihoods of all rankings in which it is top-ranked. If $\mathcal{R}_k$ denotes the set of $(n-1)!$ rankings in which a_k is top-ranked, let

$$L^*(p, a_k) = \sum_{R\in\mathcal{R}_k} L^*(p, R) = \sum_{R\in\mathcal{R}_k} p^{\tau(x, R)}(1-p)^{\binom{n}{2}c-\tau(x,R)}. \qquad (7)$$

Unlike the preceding two cases, the alternative a_k that maximizes (7) depends on p. Unfortunately, p is unknown. One possibility would be to estimate the most likely value of p from the data itself. Another would be to assume complete ignorance of p and integrate over a uniform prior distribution for p, $\frac{1}{2} \le p \le 1$. Neither of these approaches seems to lead to a result of practical value.

For p close to 1, the maximum of (7) occurs when $\tau(\mathbf{x}, R)$ is a maximum, that is, when R is selected by Condorcet's rule. For p close to $\frac{1}{2}$, the maximum likelihood decision rule implied by (7) is Borda's rule, as can be seen by taking the derivative of $L^*(p, a_k)$ with respect to p and evaluating it at $p = \frac{1}{2}$:

$$\left.\frac{dL^*(p, a_k)}{dp}\right|_{p=\frac{1}{2}} = \sum_{R\in\mathcal{R}_k} \left(2\tau(\mathbf{x}, R) - \binom{n}{2}c\right)\left(\tfrac{1}{2}\right)^{\binom{n}{2}c-1}$$

$$= \alpha \sum_{i\neq k} x_{ki} + \beta, \qquad (8)$$

where α, β are constants. By (8), the derivative of $L^*(p, a_k)$ at $p = \frac{1}{2}$ is proportional to the Borda score of a_k; therefore, since at $p = \frac{1}{2}$ all $L^*(p, a_k)$ are equal an alternative with maximum Borda score maximizes the likelihood function in a neighborhood of $p > \frac{1}{2}$.

Condorcet himself noticed that the maximum likelihood *choice* may not be the majority alternative, and illustrated this with an example (Condorcet, 1785, pp. 122–123). After considerable circumlocution he abandoned this observation in favor of finding a maximum likelihood ranking, but then later reverted to claiming that this ranking gives the best choice. As we have seen, there is room to doubt the latter conclusion. In the next section we shall show why straightforward reasoning suggests that Borda's method is better for inferring the single best alternative.

4. OPTIMAL RANKING AND CHOICE WHEN THERE IS NO OBJECTIVE STANDARD

The preceding analysis assumes that there actually is some measurable, objective standard that differentiates the alternatives. Such cases are common enough: a guilty party among innocents; a correct answer among false ones; differences in weight, length, time of occurrence, and the like. Kendall (1955) uses the nice example of "hardness" (A scratches B) to illustrate a measurable *ordinal* quality that may not be *cardinally* measurable. In many practical instances, however, there may be some quality "which we believe to be measurable but cannot measure for practical or theoretical reasons" (Kendall, 1955, p. 2). Common examples are "intelligence" or "ability." Where a cardinal measure is lacking, the use of pairwise comparisons to construct ordinal rankings or to make choices is particularly natural—as in ranking players or teams based on seasonal records.

Do the Borda and Condorcet rules have natural interpretations as rules of inference in this *relativistic* context? The answer is affirmative, and highlights in yet another way the distinction between making optimal choices and constructing optimal rankings. The difference can be illustrated by defining what is meant by a "strongest" or "best" player given the outcome of a sequence of pairwise contests (wins and losses). A natural definition of a *strongest* player would be one that beats a randomly selected opponent with greatest probability. Letting p_k be the *true* probability that player a_k will beat a randomly selected player a_j ($j \neq k$), then with no prior information, the player that actually does beat his opponents most often in a series is certainly the most likely estimate of the strongest player. This of course is Borda's rule, as given by (4).

But to use Borda's rule as a relative ranking of all the players involves a fallacy. Namely, the second-ranked alternative by Borda's rule is the one second most likely to be first, not the second in a most likely ranking. A more satisfactory definition of a best ranking in relative terms is one that maximizes the probability that a higher-ranked player will beat a lower-ranked one in a random series of pairwise trials. It is easy to verify that the best estimate of such a ranking, with no prior information, is given by Condorcet's rule.

5. GENERAL PRINCIPLES OF INFERENCE AND AXIOMATIC CHARACTERIZATIONS OF BORDA AND CONDORCET

The foregoing arguments are couched in terms of specific probabilisitic models of pairwise selection. How sensitive are the optimal decision rules to different assumptions? Put another way, are there more general *principles* of inferential choice that point to these same two rules independently of a specific statistical model? In this section we suggest a general principle of inference that, together with several other natural requirements, uniquely characterizes the Condorcet and Borda rules as the canonical methods for ranking and choice respectively.

Let f be any rule that infers a certain state (or set of states) in Θ given a sample *outcome* $\mathbf{x}$, assuming no prior information on the distribution of $\theta \in \Theta$. The *reinforcement principle* (or *axiom of confirming evidence*) states that if f infers θ from $\mathbf{x}$ and f also infers θ from $\mathbf{y}$, and $\mathbf{x}$ and $\mathbf{y}$ are independent outcomes, then f infers θ from the combined outcome $\mathbf{x}+\mathbf{y}$.

More generally, to allow for ties, the axiom states that

$$f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) \cap f(\mathbf{y}) \quad \text{whenever } f(\mathbf{x}) \cap f(\mathbf{y}) \neq \varnothing. \tag{9}$$

For example, (9) says that if f infers θ and θ' with equal likelihood from $\mathbf{x}$ but infers only θ from $\mathbf{y}$, then θ is most likely given $\mathbf{x}+\mathbf{y}$. This axiom certainly applies to all maximum likelihood rules as defined in Section 2. It has strong implications for the form a rule may take when further properties of the outcomes $\mathbf{x}$ and the state set Θ are taken into account.

Let $\mathbf{x}$ represent the outcomes of independent binary trials on a set A of n alternatives. f is a *binary choice rule* if $f(\mathbf{x})$ is a nonempty subset of A; f is a *binary ranking rule* if $f(\mathbf{x})$ is a nonempty subset of the n! rankings of A. Normally $f(\mathbf{x})$ produces a unique ranking or choice unless there are ties.

A binary choice or ranking rule f is *unbiased* (*neutral*) if for any permutation of the alternatives in A, and corresponding permutation of the coordinates of $\mathbf{x}$, $f(\mathbf{x})$ is permuted in like fashion.

A binary choice rule is *unanimous* is $f(\mathbf{x}) = \{a\}$ whenever all pairwise comparisons favor a; a binary ranking rule f is *unanimous* if $f(\mathbf{x}) = \{R\}$ whenever all pairwise outcomes agree with a fixed ranking R.

The following theorem (Young, 1974) shows why Borda's rule can be regarded as the canonical binary choice rule.

THEOREM 1. Borda's rule is the unique binary choice rule that is unbiased, unanimous, and satisfies reinforcement.

To characterize ranking rules, one more key idea is needed that amounts to a weakening of Arrow's famous "independence of irrelevant alternatives" condition. It might be called the *independence of remote alternatives* condition. Suppose that R is a ranking corresponding to $\mathbf{x}$: $R \in f(\mathbf{x})$. If a_i and a_j are *adjacent* in R, then their relative order in R should depend only on the data comparing a_i versus a_j alone (i.e., only on x_{ij} and x_{ji}) because they can always be switched without disturbing the order of the other more "remote" alternatives. This property is an instance of a very general principle of fair division which says that any restriction of a fair division should be fair. This principle also has important applications in the apportionment of representation (Balinski and Young, 1982) and bargaining theory (Lensberg, 1983).

For ranking functions on exactly two alternatives, unanimity, reinforcement, and nonbias uniquely characterize simple majority rule (Young, 1974, 1975). Hence, together with these other properties, independence of remote alternatives implies that a binary ranking function f on n alternatives satisfies the following: if $R \in f(\mathbf{x})$ and a_i precedes a_j in R, then $x_{ij} \geq x_{ji}$; moreover, if $x_{ij} = x_{ji}$, then also $R' \in f(\mathbf{x})$, where R' is the ranking obtained from R by interchanging a_i and a_j.

This implies the following result, which is a strengthening of a theorem proved by Young and Levenglick (1978).

THEOREM 2. Condorcet's rule is the unique binary ranking rule that is unbiased, unanimous, and satisfies reinforcement and independence of remote alternatives.

6. EXTENSIONS

The foregoing ideas may be extended to comparisons which are made by instruments or individuals having differing a priori reliabilities. (These generalize results of Shapley and Grofman, 1984, and Nitzan and Paroush, 1982, for the case of exactly two alternatives.) Consider first the problem of estimating the most likely distinguished object from a group of n objects, where pairwise comparisons are made by T different *types* of instruments or individuals, $t = 1, 2, \ldots, T$. The probability of success in a trial of type t is $\frac{1}{2} < p_t < 1$, which is assumed to be known a priori. An *outcome* is an $n(n-1)T$-dimensional vector $\mathbf{x}$ whose component x^t_{ij} represents the number of times a_i was selected over a_j in comparisons by instruments of reliability class t. Let c_t be the total number of comparisons of type t, for every pair of alternatives. Let $\mathbf{p} = (p_1, \ldots, p_T)$. The likelihood function takes the form

$$L(\mathbf{p}, a_k) = \prod_{t \in T} p_t^{\sum_{i \neq k} x^t_{ki}} (1 - p_t)^{\sum_{i \neq k} x^t_{ik}}.$$

This is maximized for those a_k which maximize

$$\log L(\mathbf{p}, a_k) = \sum_{t \in T} \left(\sum_{i \neq k} x^t_{ki} \right) \log\left(\frac{p_t}{1 - p_t} \right). \qquad (10)$$

This is a generalized form of Borda's rule in which the ordinary Borda scores are weighted by the log of the odds ratio for each reliability type. A similar generalization of Condorcet's rule determines the most likely ranking. Namely, choose those rankings R that maximize the expression

$$\sum_{t \in T} \left(\sum_{\substack{(i,j) \\ a_i R a_j}} x_{ij}^t \right) \log\left(\frac{p_t}{1 - p_t}\right). \tag{11}$$

ACKNOWLEDGMENT

This work was supported by the National Science Foundation under Grant No. MCS-82-07672.

NOTES

1. While Condorcet's meaning is quite clear taken in context, his style is so obtuse, and the argument so long and straggling, that many scholars since have taken Duncan Black at his word when he intoned with Nanson, "the general rules for the case of any number of candidates as given by Condorcet are stated so briefly as to be hardly intelligible . . . and as no examples are given it is quite hopeless to find out what Condorcet meant" (Black, 1958). A translation of Condorcet's statement and a discussion of Black's source of confusion are given in detail in Young (1985).