

Assent-Maximizing Social Choice

Katherine A. Baldiga and Jerry R. Green*
Harvard University

July 8, 2011

Abstract

We take a decision theoretic approach to the classic social choice problem, using data on the frequency of choice problems to compute social choice functions. We define a family of social choice rules that depend on the population's preferences and on the probability distribution over the sets of feasible alternatives that the society will face. Our methods generalize the well-known Kemeny Rule. In the Kemeny Rule it is known a priori that the subset of feasible alternatives will be a pair. We define a distinct social choice function for each distribution over the feasible subsets. Our rules can be interpreted as distance minimization – selecting the order closest to the population's preferences, using a metric on the orders that reflects the distribution over the possible feasible sets. The distance is the probability that two orders will disagree about the optimal choice from a randomly selected available set. We provide an algorithmic method to compute these metrics in the case where the probability of a given feasible set is a function only of its cardinality.

*The authors would like to thank seminar and conference participants at Harvard University, the Decentralization Conference, the Canadian Economic Theory Conference, and the Risk, Utility, and Decision Conference for their useful feedback on this work. We are also very appreciative of the many helpful comments of two anonymous referees and an editor at Social Choice and Welfare.

1 Introduction

This paper attempts to unify social choice theory with decision theory. In the tradition of Arrow (1951), the objective in social choice theory is to aggregate the ordinal preferences of a population into a single ordinal preference that represents this population. Given any subset of the alternatives, society's choice from that subset is the best element, according to this representative order. In this way, society's choice function will appear to be rational. In decision theory, the focus of the analysis is that the choice problem to be faced is not known with certainty. Decision theory seeks rules to make decisions in all problems, with the understanding that a good rule will make good choices in most problems, or in frequently encountered problems, but may make bad choices in some, perhaps less likely problems. In this paper, the social ordering used to represent a population's preferences is calculated using the frequencies with which different problems will be faced – combining the problem of social choice with the problem of decision theory.

This approach to social choice theory and decision theory is a generalization of that taken by Kemeny. Kemeny (1959), (see also Kemeny and Snell (1962), Young (1988), and the axiomatization of Young and Levenglick (1978)), study the mapping that maximizes the probability that a randomly selected member of the population would agree with the choice made by the social ordering, when the feasible set is known to be a pair of alternatives and all pairs are equally likely. Our generalization is to allow all distributions over subsets of alternatives. Our criterion is the same as Kemeny's: maximizing the probability of agreement with the social order. Kemeny shows that the assent-maximizing order that solves his problem can be characterized as a distance-minimizer in the space of all orders. The distance used is the Kendall τ distance.¹ Kemeny, Kemeny-Snell and Young-Levenglick were motivated by Condorcet's wish to "break" cycles in pairs in the least intrusive way possible, reversing pairwise choices that are supported by weak majorities. In this paper we adopt the same approach, but we do not restrict the problems to pairs. We define a metric on the space of orders that plays the same role for us as the Kendall τ distance does for the Kemeny rule. Our metric depends on the probability distribution over the feasible sets. As in Kemeny's case, it can be interpreted as the probability that a given pair of orders will make the same choice from the feasible set.

We provide a computational technique to calculate our metrics. The Kendall τ distance is the minimal number of adjacent pairs of alternatives that have to be reversed to go from one preference to another. Our metrics are determined by similar counts of pairwise exchanges. However the position in the order at which these reversals are made is relevant for us, not merely the total number of reversals.

¹Under this metric, the distance between two orders is the number of the pairwise disagreements between them. It is also known as the Dodgson distance, the bubble sort distance, or the inversion distance.

The paper proceeds as follows. In Section 2 we define our social choice rules. We call them Assent-Maximizing because we construct them using metrics that measure the probability that a pair of orders will agree on the choice at a randomly selected feasible set. In Section 3 we show how these social welfare functions compare to Kemeny’s Rule and to other well-known families of social welfare functions such as the scoring rules. Through a series of examples we explore the dependence of the expected assent-maximizing order on the distribution of problems to be faced. We also study the special case in which, as in Kemeny’s model, the distribution is only a function of the cardinality of the number of available alternatives. We show that if the assent-maximizing social welfare function is invariant to all distributions over problems then the plurality winners at each problem are consistent with the assent-maximizing order at every subset. Conversely, if the plurality winners at every problem are consistent with an order over the alternatives then this order will be the assent-maximizing social welfare function at every distribution. However, when we restrict attention to the case where the probability of a given available set is determined by its cardinality, this equivalence does not hold: there are populations at which the assent-maximizing order is invariant over this smaller family of distributions but the plurality winners are inconsistent with any order.

2 Assent Maximizing Social Welfare Functions and the Measurement of Conflict

The space of **alternatives** is denoted $X = \{a_1, a_2, \dots, a_n\}$. A **preference** π over X is identified with both an ordering of the elements of X and with a permutation of the integers $\{1, \dots, n\}$. The set of all $n!$ preferences over X is denoted Π . Thus $\pi = (a_{\pi(1)} a_{\pi(2)} \dots a_{\pi(n)})$ represents the preference in which a_i is preferred to a_j if and only if $\pi^{-1}(i) < \pi^{-1}(j)$. We will say in this case that a_i **precedes** a_j in the order determined by π . We will also say that if $\pi = (\pi_1 \dots \pi_i \dots \pi_n)$ the alternative $\pi_i = a_{\pi(i)}$ is in **position** i .

We will refer to the **natural ordering** of the alternatives as the element $e = (a_1 a_2 \dots a_n) \in \Pi$ and the associated permutation is the identity $e(i) = i$ for all $i = 1, \dots, n$.

Let the set of all non-empty subsets of X be denoted \mathcal{X} . Typically we denote a set of **feasible alternatives** by $A \in \mathcal{X}$. Given a preference π let $c_\pi(A)$ be the element in A that precedes all other elements in A according to π . The function $c_\pi : \mathcal{X} \rightarrow X$ is the rational **choice function** generated by the preference π .

Let ν be a probability distribution over \mathcal{X} . We define our **measure of conflict between preferences** as

$$f(\pi, \pi'; \nu) = \nu\{A \in \mathcal{X} | c_\pi(A) \neq c_{\pi'}(A)\}$$

A **population** of preferences is a probability distribution λ over the set of all preferences Π . Let Λ be **set of all populations**. A **social welfare function**

is a correspondence that assigns a set of preferences to each population. We will define a set of social welfare functions that we call assent-maximizing. Each assent-maximizing social welfare function $\pi^*(\lambda, \nu)$ depends on the distribution ν over the set of choice problems that might be faced.

Given a population λ and a probability distribution ν , we can describe the level of expected conflict between a randomly-drawn member of the population and any potential candidate for a social ordering. Let the **frequency of expected conflict between a candidate social ordering π' and individual preferences in the population λ** be defined by

$$q(\pi', \lambda, \nu) = \sum_{\pi} f(\pi, \pi'; \nu) \lambda(\pi) \quad (1)$$

Using this notation the **assent-maximizing social welfare function** is defined as the set of social orderings that minimize this expected conflict

$$\pi^*(\lambda, \nu) = \arg \min_{\pi'} q(\pi', \lambda, \nu) \quad (2)$$

In our model the preference aggregation operates before the feasible set is determined, and depends upon the probability distribution over feasible sets. The rule selects a preference that will be operative at all feasible sets, in such a way that the largest fraction of the population, on average, agrees with the choice that this preference will make when a problem arrives.

The determination of $\pi^*(\lambda, \nu)$ is straightforward once $f(\pi, \pi'; \nu)$ is known. For general probability distributions ν , $f(\pi, \pi'; \nu)$ must be calculated at each pair of preferences. However, when $\nu(A)$ depends only on the cardinality of A , as in Kemeny's case where $\nu(a) = \frac{1}{n(n-1)}$ at every A with cardinality 2 and zero otherwise, f is a form of distance on Π . In this case we will say that ν is **exchangeable**. Exchangeable ν can be summarized by the distribution of the cardinality of the feasible set, or the **size distribution**, of the feasible set, $\mu = (\mu_1, \dots, \mu_n)$. For every probability vector μ there is one exchangeable distribution of the feasible set and in that distribution all sets A with cardinality i are equally likely and collectively have probability μ_i . We will denote Kemeny's case, in which it is sure that the feasible set will contain exactly two alternatives, by $\mu^K = (0, 1, 0, \dots, 0)$. If ν is exchangeable and we want to emphasize the dependence of social choice on the size distribution μ of the set of alternatives, we will write the assent-maximizing social welfare function as a function of μ , with the slight abuse of notation $\pi^*(\lambda, \mu)$.

In Appendix 1 we show that when ν is exchangeable we can give an explicit formula for f and an algorithm to compute it. We will use the exchangeable case in several examples, to contrast the properties of assent-maximizing social welfare functions with other rules and to study the dependence on ν more explicitly. In working out these examples, we will employ the computational methods described in the appendix. Our results and the spirit of our analysis, however, do not depend on the exchangeability of ν .

It is possible that two preferences will have the same average assent in the population, but this is a non-generic event in the space Λ of all probability

distributions over Π . Therefore, assuming $\mu_2 > 0$, each assent-maximizing social welfare function can be regarded as a partition of Λ into subsets on which $\pi^*(\lambda, \nu)$ is a constant with respect to λ .

When (1) is evaluated at the assent-maximizing social welfare function we have the **measure of conflict with assent-maximization**

$$Q(\lambda, \nu) = q(\pi^*(\lambda, \nu), \lambda, \nu) = \sum_{\pi} f(\pi, \pi^*; \nu) \lambda(\pi) \quad (3)$$

The measure Q corresponds to the "goodness of fit" of the social preference to the population of individual preferences.

Given a population λ , we can also define the **measure of internal conflict**

$$M(\lambda, \nu) = \sum_{\pi} \sum_{\pi'} \lambda(\pi) \lambda(\pi') f(\pi, \pi'; \nu) \quad (4)$$

The measure M does not depend on a social ordering. It is a direct measure of the diversity of preferences within the population.

3 Properties of Assent-maximizing Social Welfare Functions and Examples

We first discuss the relationship of the social ordering that is produced by our measures with that produced by the Kemeny method. In the Kemeny method, ν is exchangeable and concentrated on the family of all pairs. The size distribution of the feasible set is denoted μ^K . We will frequently contrast μ^K with another special case, where all sets of cardinality two or greater are equally-likely, and we denote this size distribution by μ^E . Similar examples can be constructed for many other size distributions. The important point is that larger potential feasible sets will generate a different pattern of agreement than if only pairs are possible.

Example 1 *Maximizing Assent to Plurality-Induced Choices*

This is our basic example where we show that the Kemeny method can lead to substantially less average assent than the assent-maximizing social welfare function determined by ν . This phenomenon is increasingly prevalent when larger feasible sets are likely. We take the case of equally-likely feasible sets for concreteness, with $n = 3$ and $\mu = \mu^E = (0, \frac{3}{4}, \frac{1}{4})$.

The population λ is concentrated on three of the six orderings in Π according to the distribution: $\lambda((a_1 a_2 a_3)) = .49$, $\lambda((a_3 a_2 a_1)) = .48$, $\lambda((a_2 a_3 a_1)) = .03$. This distribution exhibits the tension between selecting the ordering consistent with the transitive majority relation (Condorcet Consistency) and selecting the Plurality winner from the three-alternative set $A = X$. Using Plurality rule on every subset, the choice function would be

$$\begin{aligned}
c(\{a_1, a_2\}) &= a_2 \\
c(\{a_1, a_3\}) &= a_3 \\
c(\{a_2, a_3\}) &= a_2 \\
c(\{a_1, a_2, a_3\}) &= a_1
\end{aligned}$$

Condorcet Consistency requires that the social ordering be $(a_2 a_3 a_1)$, as the majority choices from the pairs are transitive. But this ordering does not describe the result of plurality rule when all three alternatives are available, as a_1 , not a_2 , is the plurality choice from the order three subset. This illustrates a well-known, important point: populations that satisfy Condorcet Consistency will not necessarily generate plurality choice functions that are consistent with an ordering. Thus, if the social choice procedure must produce an ordering, it will have to make trade-offs between matching plurality choices on subsets of various sizes. The Kemeny method selects the ordering $\pi = (a_2 a_3 a_1)$, which implies a choice of $c(\{a_1, a_2, a_3\}) = a_2$, supported by just 3% of the population. If $A = \{a_1, a_2, a_3\}$ never arises, then this choice from the triple does not create any expected dissent. If $A = \{a_1, a_2, a_3\}$ may arise, however, expected assent may be maximized by an ordering that has a_1 as its top element instead of a_2 . When all subsets with two or more elements are equally likely, we have $\pi^*(\lambda, \mu^E) = (a_1 a_2 a_3)$, implying a choice of a_1 from $A = \{a_1, a_2, a_3\}$ that is supported by 49% of the population. This gain of 46% in expected support for the social choice at $A = \{a_1, a_2, a_3\}$ more than outweighs the decrease of 1% or 2% in support that occurs when $\{a_1, a_2\}$ or $\{a_1, a_3\}$ are the feasible sets. Assent-maximization with μ^E sacrifices matching the majority preference on two of the three pairs problems in order to better match the population preferences in the case where all three alternatives are available. \square

Example 1 emphasizes the tension between assent-maximization and the ordering requirement. In these cases we need to consult the distribution ν when forming a social order. On the other hand, when the plurality choice from each A is explainable by an ordering, that is, when there exists $\hat{\pi}$ such that $c_{\hat{\pi}}(A)$ coincides with the plurality winner at every $A \in \mathcal{X}$, then this tension does not exist and, as the following Theorem states, the social ordering should be independent of ν – in fact it should be the same as the ordering $\hat{\pi}$.

Theorem 1 *If $\lambda \in \Lambda$ is a population of preferences such that the assent-maximizing social welfare function $\pi^*(\lambda, \nu)$ is independent of ν then plurality choice at each $A \in \mathcal{X}$ coincides with $\pi^*(\lambda, \nu)$. Conversely, if plurality choice on \mathcal{X} is a rational choice function, then this choice function is the assent-maximizing social welfare function for every distribution ν .*

Proof of Theorem 1

Define $\alpha(\pi, \lambda, A)$ to be the fraction of the population λ such that their most preferred element in A coincides with $c_{\pi}(A)$, the choice that would be made by the preference π .

We will prove the contrapositive of the first part of the theorem. Suppose that for $\lambda \in \Lambda$ the plurality choice at some $A \in \mathcal{X}$ does not coincide with the choice of the assent-maximizing social welfare function $c_{\pi^*(\lambda, \nu)}(A)$ for some ν . We need to show that the assent-maximizing social welfare function $\pi^*(\lambda, \nu)$ will not be invariant with ν . We know that $\alpha(\pi, \lambda, A) < \alpha(\pi', \lambda, A)$ for some $A \in \mathcal{X}$. We can then choose a new ν' such that $\nu'(A)$ is sufficiently large ($\nu'(A) = 1$ would always work) that π' will garner more assent than π .

To prove the other direction, suppose plurality choice on \mathcal{X} is a rational choice function. We will prove that this choice function is the assent-maximizing social welfare function for every distribution ν . Consider any $\pi' \in \Pi$, $\pi' \neq \pi$. Then, we have $\alpha(\pi, \lambda, A) > \alpha(\pi', \lambda, A)$ for every $A \in \mathcal{X}$. We have that $c_\pi(A)$ garners the most possible assent for every A , since it is the choice supported by the plurality, so π will clearly be the assent-maximizing ordering. There is no tension in this case between assent-maximization and the requirement that social choice be generated by an ordering. ■

The condition that the plurality choice function be derivable from an ordering is very strong. The simplest type of population which generates a rational plurality choice function contains a majority held preference, where $\lambda_i > .5$ for some π_i . It is obvious in this case that the choice function induced by π_i maximizes assent at every subset $A \in \mathcal{X}$, implying that π_i is assent-maximizing for any distribution ν . Populations with a majority-held preference are a strict subset of the populations which satisfy the premises of Theorem 1.

Note that if we restrict attention to only exchangeable distributions, Theorem 1 would not hold. It is possible that the assent-maximizing ordering is invariant over the smaller space of exchangeable distributions and yet the plurality choice function is not rational. We provide an example below.

Example 2 *Theorem 1 Fails when Restricted to Exchangeable Distributions*

Take a population λ at which plurality rule would generate cyclic choice over the pairs in a three-alternative problem: $\lambda(a_1a_2a_3) = \frac{1}{3} - \varepsilon$, $\lambda(a_2a_3a_1) = \frac{1}{3} + \varepsilon$, $\lambda(a_3a_1a_2) = \frac{1}{3}$, where ε is small and positive. The claim is that $\pi^* = (a_2a_3a_1)$ is the assent-maximizing order for all exchangeable distributions. Consider the following table of assents to choices from the pair subsets for each potential social ordering. Each column contains the assents to $c_\pi(A)$ for each A for the given π .

A	e	$a_1a_3a_2$	$a_2a_3a_1$	$a_2a_1a_3$	$a_3a_1a_2$	$a_3a_2a_1$
$\{a_1, a_2\}$	$\frac{2}{3} - \varepsilon$	$\frac{2}{3} - \varepsilon$	$\frac{1}{3} + \varepsilon$	$\frac{1}{3} + \varepsilon$	$\frac{2}{3} - \varepsilon$	$\frac{1}{3} + \varepsilon$
$\{a_2, a_3\}$	$\frac{2}{3}$	$\frac{1}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
$\{a_1, a_3\}$	$\frac{1}{3} - \varepsilon$	$\frac{1}{3} - \varepsilon$	$\frac{2}{3} + \varepsilon$	$\frac{1}{3} - \varepsilon$	$\frac{2}{3} + \varepsilon$	$\frac{2}{3} + \varepsilon$
Sum	$\frac{5}{3} - 2\varepsilon$	$\frac{4}{3} - \varepsilon$	$\frac{5}{3} + 2\varepsilon$	$\frac{5}{3}$	$\frac{5}{3}$	$\frac{4}{3} + 2\varepsilon$

For an exchangeable distribution, the assent to each pair is weighted equally in the maximization. Thus, the final row of sums captures the relevant information about assent to the pairs. We see $a_2a_3a_1$ strictly dominates the other

candidates on pairs. The total expected assent for any candidate for any exchangeable distribution will be a weighted sum of the pairs assent and the triple assent. Since a_2 is the assent-maximizing choice from the order three subset, $a_2a_3a_1$ weakly dominates the other candidates on the triple. Therefore $a_2a_3a_1$ is assent-maximizing at all exchangeable distributions. But, the plurality choice function of this population is not rational because it cycles on the pairs.

By Theorem 1, there are non-exchangeable distributions at which $a_2a_3a_1$ would not be assent-maximizing. For example, a non-exchangeable distribution with $\nu\{a_1, a_2\} > \frac{1}{3}$, $\nu\{a_2, a_3\} = \nu\{a_1, a_3\}$ and $\nu\{a_1, a_2, a_3\} = 0$ would not produce $a_2a_3a_1$ as a social ordering. \square

We now concentrate on the typical case in which there is a tension between assent-maximization and the ordering requirement. Through a series of five examples we will demonstrate some features of assent-maximizing social welfare functions. We will show that assent-maximization may select something other than a Condorcet winner, even when one exists. We show that the assent-maximizing ordering may lie outside the support of the preferences in the population. We show that the method of compromise inherent in assent-maximization is different from both the compromises made by positional methods such as scoring rules and the compromises made by methods that use only the numerical vote tallies from pairwise contests. We explore the relationship between the dispersion around the social choice and the measure of average disagreement across pairs in the population, showing the counterintuitive result that a more internally-conflicted population may have less expected dissent with its social ordering than a population with less internal conflict. Finally, we provide two examples of non-exchangeable problems.

Example 3 *Assent-maximization may not rank a Condorcet winner first*

This is a well-known four-alternative problem, that of choosing the capital city of Tennessee (see Young (1988), Moulin (1988)), which we adapt for the present discussion by allowing the feasible sets to be larger than pairs. For concreteness, we assume that all subsets are equally likely.

The four alternatives are the four largest cities in Tennessee:

Memphis = a_1 , Nashville = a_2 , Chattanooga = a_3 , Knoxville = a_4

The distribution of preferences λ is: $\lambda((a_1a_2a_3a_4)) = .42$, $\lambda((a_2a_3a_4a_1)) = .26$, $\lambda((a_3a_4a_2a_1)) = .15$, $\lambda((a_4a_3a_2a_1)) = .17$

We have $\pi^*(\lambda, \mu^E) = (a_1a_2a_3a_4)$ and $\pi^*(\lambda, \mu^K) = (a_2a_3a_4a_1)$. The difference between our method and the Kemeny method is based on a rather narrow difference in average assents to the social choice: $q((a_1a_2a_3a_4), \lambda, \nu^E) = .5282$, but $q((a_2a_3a_4a_1), \lambda, \nu^E) = .5173$. This is an admittedly small difference in average assent. Nevertheless, the induced choice functions are quite different because $(a_1a_2a_3a_4)$ selects Memphis whenever it is available, whereas $(a_2a_3a_4a_1)$ never puts the capital in Memphis at any feasible set.

Notice that Nashville (a_2) is a Condorcet winner, beating the other three alternatives in the three pairwise contests where it is one of the options. For that reason the Kemeny method ranks a_2 first. On the other hand, in every

available set with three or more alternatives in which a_1 is one of the options, it is chosen by the plurality of the population. Therefore as long as the probability is high enough that the actual available set will have three or more options, a_1 should be chosen from these sets. In these cases, therefore, a_1 should be first in the social ordering. \square

Example 4 *Compromising on a Second Best*

In both Example 1 and Example 3 the assent-maximizing π^* was one of the orderings in the support of λ . This is not always the case as the following four alternative example shows:

Let $\lambda((a_3a_1a_2a_4)) = .33$, $\lambda((a_2a_1a_3a_4)) = .34$, $\lambda((a_4a_1a_2a_3)) = .33$. For a variety of metrics, including the equally-likely size distribution μ^E , $\pi^*(\lambda, \mu) = (a_1a_2a_3a_4)$. This social ordering is somewhat of a compromise of the three orderings in the population; it places the alternative that everyone in the population ranks second, a_1 , at the top of the social ordering. Though this means that no one in the population will agree with the choice from the order four subset, this social ordering will perform better on average on subsets of order two and three than an ordering chosen to match the first-place choice of one segment of the population. The assent-maximization method trades off support on the order four subset in order to garner greater support on the more likely order two and three subsets. Note, however, that if the probability of facing the order four subset was sufficiently high, the social ordering would change to $\pi = (a_2a_1a_3a_4)$. \square

The social choice rules based on assent-maximization differ from the two main classes of methods that have been applied and characterized in the literature. These methods are the scoring rules, which produce an ordering by combining rank information from the individual preferences, and pairwise methods that produce an ordering as a function of either the majority tournament or the matrix of voting margins across all pairs of alternatives. Theorem 1 demonstrates assent-maximization is not equivalent to any scoring rule. The reason is that in populations where a majority shares the same preference, scoring rules will not always produce a social ordering coincident with that majority preference.² A minority with very different preferences will be able to influence the relative scores.³ We can also see that assent-maximization is not a member of the class of social welfare functions based only on pairwise comparisons and

²Daugherty, Eustis, Minton and Orrison (2009) show that the Borda scores can be recovered from the binary vote margin matrix. However, the ordering of alternatives by their Borda score will in general be different from the assent-maximizing order.

³A population at which a strict majority would not dominate the social ordering under a scoring rule is $\lambda((a_1a_2a_3a_4)) = .55$, $\lambda((a_2a_4a_3a_1)) = .25$, $\lambda((a_2a_3a_4a_1)) = .2$. For all ν our method reproduces the majority's preference $\pi^* = (a_1a_2a_3a_4)$ but Borda count produces $(a_2a_1a_3a_4)$ because the minority preferences give a_2 an advantage over their last-choice, a_1 .

satisfying neutrality. There can well be a tension between the way that pairwise methods break a Condorcet cycle, at its weakest link, and the plurality choice from a larger feasible set. The following example shows that the assent-maximizing ordering can differ from the unique recommendations of both of these classes of methods, and thus that the social welfare function it defines does not lie in either class.

Example 5 *Positional and Pairwise Methods and Assent-Maximization*

In this example we show that for a particular population, with preferences over just three alternatives, positional methods, pairwise methods, and our method for a range of μ can generate three distinct social orderings.

Consider the population λ defined by

$$\begin{aligned}\lambda((a_1a_2a_3)) &= .10 \\ \lambda((a_1a_3a_2)) &= .30 \\ \lambda((a_3a_2a_1)) &= .25 \\ \lambda((a_2a_1a_3)) &= .35\end{aligned}$$

Saari (2001) has shown that all positional methods will select $(a_1a_2a_3)^4$, while Kemeny's method (and all pairwise-based methods) will select $(a_2a_1a_3)$. Our rule, applied to exchangeable μ with $\mu_3 > \frac{2}{5}$, will yield $(a_1a_3a_2)$. Under our method, when μ_3 is relatively large, more consideration is given to the choice from the three-alternative set, so the assent-maximizing ordering has alternative a_1 in first-place, as it will be supported as the choice from the triple by a plurality of the voters. Alternative a_2 from the triple would be supported by just .35 of the population, and a choice of alternative a_3 is supported by just .25. Thus, with respect to the top-ranked alternative, our method looks similar to the positional methods. However, unlike the positional methods, our method respects the will of the majority on the ordering of the final two alternatives, placing a_3 before a_2 . Thus, in general, for a given λ , no positional method or pairwise-based method will agree with our methods over the entire range of μ . Nor is there a positional method that will agree with our methods for a given μ over the entire range of λ . \square

Example 6 *Conflict and Compromise*

Let λ and λ' be two populations such that for a given μ , λ is less internally conflicted than λ' : $M(\lambda, \nu) < M(\lambda', \nu)$. Will there be more assent to the assent-maximizing social ordering at λ than there will be at λ' ? In other words, is it easier to find a good compromise when the population is less conflicted?

⁴Saari's result can be seen as follows: Consider any scoring rule, standardized to the vector $(1, s, 0)$ where $s \in [0, 1]$. Then, we have the following scores for a_1, a_2, a_3 , respectively: $.4 + .35s, .35 + .35s, .25 + .3s$.

A simple example shows that this will not generally be the case. We will fix $\mu = \mu^E$ and let λ and λ' be given by:

$$\begin{aligned}\lambda((a_1a_2a_3)) &= .5 \\ \lambda((a_2a_1a_3)) &= .5\end{aligned}$$

$$\begin{aligned}\lambda'((a_1a_2a_3)) &= .5 \\ \lambda'((a_1a_3a_2)) &= .25 \\ \lambda'((a_2a_1a_3)) &= .25\end{aligned}$$

It is straightforward to show that λ' is more internally conflicted than λ .

$$\begin{aligned}M(\lambda, \nu) &= .0625 \\ M(\lambda', \nu) &= .0703\end{aligned}$$

But

$$\begin{aligned}Q(\lambda, \nu) &= .125 \text{ with } \pi^*(\lambda) = \{(a_1a_2a_3), (a_2a_1a_3)\} \\ Q(\lambda', \nu) &= .0938 \text{ with } \pi^*(\lambda') = \{(a_1a_2a_3)\}\end{aligned}$$

The reason for this difference in conflict measures can be traced to the fact that the space of preferences is discrete. The measure of conflict with the assent-maximizing ordering is like a second moment taken around a point in the space which is not a true mean – one that is chosen in asymmetric situations because it is an approximate compromise and lies within the discrete space from which compromises must be selected. The measure of internal conflict, on the other hand, deals directly with pairs of preferences chosen from the population and does not require any such point of reference. \square

Example 7 *Two Non-Exchangeable Models*

Situations in which the probabilities of available sets with a given cardinality are not all the same arise naturally. Our general framework extends to these cases as well. One such situation is where the individual alternatives are either available or not available, the set of available alternatives forming the feasible set for the social choice. If their availability is statistically independent and each has the same probability of availability, we are in the exchangeable case with a binomial distribution over the cardinality of the available set. Otherwise we have a non-exchangeable problem. A model of this form called the "unavailable candidate model" has been explored by Lu and Boutilier (2010).

A second family of non-exchangeable problems arises when there is a set of different underlying issues and within each issue there is a distinguished outcome, which we can think of as a status quo. An alternative is a selection

of outcomes, one for each issue – either the status quo or one of the other alternatives. For example, suppose that there are two issues b and h . Issue b can be decided in favor of a new policy, b_1 , or the society can remain with the status quo b_0 . Similarly issue h has a status quo policy h_0 , or a new policy h_1 . The members of the population have preferences which can exhibit substitutability or complementarity across the two new policies – they can prefer to do one of the new policies, or the other new policy, but not both (the substitutes case), or they can prefer to do both new policies but not either one without the other (the complements case). In the general notation of our model there are four "alternatives" which are composites of these four individual policies as follows:

Alternative	Corresponding choices on b and h
a_1	b_1h_1
a_2	b_0h_1
a_3	h_0b_1
a_4	b_0h_0

In this model the uncertainty about the available set is due to the set of the issues that are "on the table". If an issue is not on the table then the status quo is sure to be the outcome on that issue. In the example we are examining, the possibilities are that b can be on the table but not h , or h but not b , or both. Thus there would be three possible feasible sets of alternatives, $\{a_2, a_4\}$, $\{a_3, a_4\}$, and $\{a_1, a_2, a_3, a_4\}$, instead of the family of all subsets. The social preference will not be called upon to make a decision from a feasible set such as $\{a_1, a_3, a_4\}$ because if both issues on on the table then a_2 must be available as well. A model of this form has been studied by Ahn and Oliveros (2010).

As in the case of the unavailable candidate model, assent-maximization can be used to make the selection of the social order. Because of this factored structure of the alternative space, parameterized functional forms for the preferences and probabilities of combinations of active issues will result in a much simpler calculation of the assent-maximizing social order than would have been the case in general with the same number of alternatives.

We defer further exploration of this issue to the Appendix, where computational aspects of assent-maximizing social choice via choice-based metrics on preferences are discussed. \square

4 Conclusion

We have tried to unite social choice theory and decision theory by recognizing the effect of the distribution of potential feasible sets as a determinant of the social choice rule. We use a distance-minimizing approach to the selection of the choice rule, where the distance between preferences is the probability that these preference would make the same selection from a random feasible set. We give examples illustrating the dependence of the choice rule on the distribution of the feasible set, at particular populations. For the case in which the

decision problems are selected using an exchangeable distribution μ we give an algorithmic method to compute the distance between two preferences.

Two further topics to be explored as a result of this paper are the following:

1) Our theory suggests a new source of "data" for welfare economics. Instead of focusing exclusively on the choice function c which describes the choice made at each feasible set, one should develop a theory that takes into account the frequency of different problems. The data of choice theory would be (c, ν) instead of only c , where ν matches the empirical distribution of choice problems. By extending our metrics to the domain of choice functions, we could use assent-maximization to determine which π is closest to the observed choices under $F(\pi; \nu)$. This type of methodology is similar to the classic revealed preference approach but would allow for welfare assessments in the case of irrational choice data.

2) Our theory defines a mapping from Λ to Π corresponding to each distribution μ . Thus, given μ , the simplex Λ is divided into equivalence classes of populations that have the same social ordering. The algebraic structure of these partitions and their dependence on μ needs to be explored further. The recent work of Daugherty, Eustis, Minton and Orrison (2009) makes significant progress in that regard. Their group theoretic structure is an alternative to examining the geometric structure of the equivalence classes for scoring rules as studied by Saari (1994).

Appendix 1: Exchangeability and computing the measure of conflict:

In this appendix we discuss how our measure of conflict depends on the size distribution of the feasible subset of alternatives in the exchangeable case, where the size distribution $\mu = (\mu_1, \dots, \mu_n)$ determines the probability of each feasible set⁵. We will show that μ determines a semi-metric on Π , and when $\mu_2 > 0$ this semi-metric is indeed a metric on Π . The implication of this technical result is that the assent-maximizing social welfare functions are the minimizers of the expected distance to the preferences in the population, where the expectation is taken over the population distribution λ .

We propose a simple computational method for finding $f(\pi, e; \nu)$ that reveals its dependence on ν in the exchangeable case. The result is linear in the size distribution μ and the coefficients can be calculated algorithmically as a function of π . While the number of steps in the algorithm rises rapidly with n , the method is practical because, once accomplished for a given n , these coefficients are fixed and can be used at any population facing n alternatives. It does not have to be redone at each population. As a result of this computational argument, the assent-maximizing social welfare function can be found via integer programming methods: minimize the average distance between $\pi^*(\lambda, \mu)$ and π , using weights $\lambda(\pi)$; the integer constraints are needed to express the fact that π^* is an order. This integer programming description is well-known in the Kemeny case⁶ and the present section establishes how this construction can be extended to all exchangeable distributions.

Once we know the probability of disagreement between an arbitrary π and the natural ordering e the probability of disagreement between all pairs of orderings are determined by virtue of the equality of probabilities of all sets of alternatives with the same cardinality. The probability of disagreement between π and e depends on ν through the size distribution μ of the feasible set A , thus we can define it as

$$F(\pi; \mu) = f(\pi, e; \nu).$$

The elements of Π form a non-directed graph in which the **edges** are pairs of permutations obtainable from each other by a single transposition of adjacent elements. That is (π, π') is an edge if there is some $i \in \{1, \dots, n\}$ such that $\pi_i = \pi'_{i-1}$, $\pi_{i-1} = \pi'_i$ and $\pi_j = \pi'_j$ for all $j \neq i-1, i$. If this relation holds we say that π differs from π' by a **transposition at position i** , the index i being the larger of the two positions at which they differ.

If π differs from e only by a transposition at position i , then the family of feasible sets on which $c_\pi(A)$ differs from $c_e(A)$ is precisely those sets that contain a_{i-1} and a_i and do not contain any element a_j for $j < i-1$. For example, if

⁵We usually will set $\mu_1 = 0$. Choice problems with only a single feasible element are possible, but as there can be no disagreement they are trivial for our purposes. To simplify calculations we frequently assume that such problems do not arise at all.

⁶See Conitzer (2006).

π differs from e by a transposition at position n then the only feasible set on which they differ is $A = \{a_{n-1}, a_n\}$. If π differs from e by a transposition at position i , then there are 2^{n-i} sets A at which $c_\pi(A) \neq c_e(A)$ because any subset of $\{a_{i+1}, \dots, a_n\}$ when combined with $\{a_{i-1}, a_i\}$ will be such a feasible set. Therefore a transposition at position i generates a collection of changes in the choice function, and we know the number of sets of each cardinality k that are affected by this transposition.

A **path** ρ from π to π' is a list of preferences ρ_0, \dots, ρ_M such that (ρ_{m-1}, ρ_m) is an edge for all $m = 1, \dots, M$, and $\rho_0 = \pi, \rho_M = \pi'$. Clearly there are many paths between any two permutations.

There are many algorithms which can be applied to a permutation π that will connect it via a path to e . One of the best known of these sorting procedures is **bubble sort**.⁷ In bubble sort we make a series of passes through the alternatives, beginning each time with the best (left-most) alternative. At each step in the algorithm, we compare the element in position i with the element in position $i + 1$. If this pair of elements does not appear in its natural order, the elements are transposed. The algorithm continues this series of pairwise comparisons, moving from $i = 1$ to $i = n - 1$, at which point a pass is complete. The algorithm then begins again at $i = 1$, making another pass through the ordering. The algorithm terminates when a pass is completed with no transpositions made.

In this paper, we will use a very similar algorithm: **reverse bubble sort (RBS)**. RBS is exactly like bubble sort except that its passes move leftward from position $i = n$ until $i = 2$. Let us keep track of the positions at which RBS makes a transposition. In each pass, either there is or is not a transposition at position i ; there can never be more than one transposition at position i in any one pass because the pass continually steps through the positions. Let x_i be the number of passes at which there is a transposition at position i . We call the vector $x = (x_2, \dots, x_n)$ the **RBS signature of π** . (If it is necessary to be explicit we write $x(\pi)$ instead of x .)

The RBS relationships between the six orderings of the three-alternative case can be represented geometrically as a permutahedron (see Ziegler (1995) for more on these diagrams). Figure 1 features each ordering as a vertex of a regular hexagon. The RBS signature between each ordering and its neighbors, the two orders obtainable by a single transposition, appear on the edges of the hexagon. The element e appears in the upper left-hand corner. The arrows in the interior of the diagram are labeled with the RBS signatures for the move from the ordering e to the three non-adjacent orderings which require more than one transposition.

While bubble sort and reverse bubble sort both convert π into e and both make only transpositions of adjacent alternatives, they generate different sequences of choice functions along the way. We choose RBS because its associated sequence has an analytically-useful property that we exploit; ordinary bubble

⁷Bubble sort is not efficient as an algorithm. We are not concerned with computational efficiency here. As will be seen below, we are using the sorting procedure as an analytic tool.




Figure 1: The above figure illustrates the $n = 3$ case. Each preference ordering appears as a vertex of the regular hexagon. The RBS signatures between each ordering and its neighbors, the two orders obtainable by a single transposition, appear around the outside of the hexagon. The element e appears in the upper left-hand corner. The arrows in the interior of the hexagon are labeled with the RBS signatures of the moves from the ordering e to the three non-adjacent orderings, each of which requires more than one transposition. Diagrams of this type have a long history in group theory and social choice, see Kemeny (1959), Le Conte de Poly-Barbut (1990), Reuter (1996), and Ziegler (1995).

sort does not share this property. The following example is an illustration.

Example 8 *Bubble Sort and Reverse Bubble Sort*

Let $n = 3$ and $\pi = (a_3a_2a_1)$. As bubble sort and RBS each sort π , the following sequences of preferences is created. Each successive pair (ρ_{i-1}, ρ_i) , $i = 1, 2, 3$ is an edge because at each step only one transposition is made.

Bubble Sort	RBS
$\rho_0 = (a_3a_2a_1)$	$\rho_0 = (a_3a_2a_1)$
$\rho_1 = (a_2a_3a_1)$	$\rho_1 = (a_3a_1a_2)$
$\rho_2 = (a_2a_1a_3)$	$\rho_2 = (a_1a_3a_2)$
$\rho_3 = (a_1a_2a_3)$	$\rho_3 = (a_1a_2a_3)$

Using ordinary bubble sort to transform π into e induces the following series of changes in the choice function. The first link changes $c_\pi(\{a_1, a_2, a_3\})$ and $c_\pi(\{a_2, a_3\})$. The second link changes $c_\pi(\{a_1, a_3\})$. And the third link changes $c_\pi(\{a_1, a_2, a_3\})$ and $c_\pi(\{a_1, a_2\})$. Note that $c_\pi(\{a_1, a_2, a_3\})$ would be changed twice along this path, first from a_3 to a_2 , then from a_2 to a_1 .

By using RBS, the choice from each subset is changed exactly once along the path from π into e . The first link changes $c_\pi(\{a_1, a_2\})$. The second link changes $c_\pi(\{a_1, a_3\})$ and $c_\pi(\{a_1, a_2, a_3\})$. And the third link changes $c_\pi(\{a_2, a_3\})$. No set A has its associated choice changed more than once. The RBS signature is $x(a_3a_2a_1) = (1, 2)$ because the first and second passes make transpositions at position 3 but only the first pass makes a transposition at position 2. \square

Example 8 illustrates a general point that is central to our analysis. We will show below that RBS never changes the choice at a set A more than once. For the purpose of characterizing preferences π by the choice functions they generate, we can use RBS to sort them each into e and keep track of the positions at which RBS makes a transposition. The counts of how many transpositions RBS makes as it transforms π into e will form a sufficient statistic for the family of sets at which π makes a different choice than e . These ideas are formalized as follows:

A **lexicographically minimal path** from π to π' is a path ρ_0, \dots, ρ_M in which there is no set $A \in \mathcal{X}$ at which $c_{\rho_{m-1}}(A) \neq c_{\rho_m}(A)$ for more than one value of m . We call these paths lexicographically minimal because their associated signatures are minimal when compared lexicographically to signatures generated by other paths. Lexicographically minimal paths are useful because the set of decision problems on which π and π' disagree can be characterized by the lexicographically minimal paths between them.

Proposition 1 *The permutations successively reached as reverse bubble sort converts π to e constitute a lexicographically minimal path from π to e .*

Proof of Proposition 1

Let $\pi \in \Pi$ and let ρ be the path generated by RBS as it sorts π into e . Assume $\pi \neq e$ and that $c_\pi(A) \neq c_e(A)$. Then $c_\pi(A)$ precedes $c_e(A)$ in the ordering π .

Each step in each pass of RBS involves a possible transposition of alternatives. We will say that y is promoted beyond x if the alternatives x and y are transposed at a step and x initially precedes y . As a pass of RBS proceeds, the alternatives that are promoted are continually better according to e . Thus in each pass of RBS there are three possibilities:

(i) There is some step in the pass where $c_e(A)$ is promoted and it continues to be promoted until it is promoted beyond $c_\pi(A)$.

(ii) There is some step where $c_e(A)$ is promoted but some $x \notin A$ becomes its immediate predecessor, $x < c_e(A)$, and $c_\pi(A)$ is a predecessor of x at this step. In this case $c_e(A)$ is no longer promoted during this pass. At the next step in this pass, either x is promoted beyond its immediate predecessor, or there is an even better alternative, according to e that is promoted. As some point in this pass some alternative better than $c_\pi(A)$ according to e will be promoted beyond $c_e(A)$. However, this alternative cannot be a member of A because it must be better than $c_e(A)$ which is, by definition, the best alternative in A according to e . Thus no member of A can be promoted beyond $c_\pi(A)$ in this pass under the conditions of possibility (ii).

(iii) There is no step where $c_e(A)$ is promoted. In this case, as in the argument above, when $c_e(A)$ stops being promoted, its immediate predecessor is better than $c_e(A)$ according to e . As in that (ii), there is some x that is promoted beyond $c_e(A)$ at this pass, but, for the same reason as above, this x cannot be a member of A as that would contradict the definition of $c_e(A)$.

In case (i), at the step where $c_e(A)$ and $c_\pi(A)$ are transposed, the choice from A will change. By the definition of $c_e(A)$ it can never change again at any later step of the algorithm.

In cases (ii) and (iii) $c_\pi(A)$ will remain the highest-ranked member of A at all steps of this pass.

Thus on the first pass such that case (i) obtains, there will be a change in the choice from A . At all future passes $c_e(A)$ is the highest ranking alternative among those in A and these passes will be in either case (ii) or case (iii). Thus the path generated by RBS is lexicographically minimal. ■

Given π , $c_\pi(A) \neq c_e(A)$ for some family of sets A . All lexicographically minimal paths from π to e change the choice at each of these sets exactly once. RBS can be used as a computational method to form a list of all sets in this family and to compute how many there are of each cardinality.

Because of our exchangeability assumption, all sets of a given cardinality have equal probability under ν . Therefore, once we know how many sets of each cardinality are changed by RBS, the probability of disagreement will be the linear combination of these counts weighted by μ .

We will now show how the RBS signatures between preferences can be used to define a metric⁸ on Π . Proposition 2 shows that the RBS signature between

⁸More generally, in limiting cases, the RBS signatures define only a semi-metric, although

two preferences is uniquely determined at all lexicographically minimal paths between these preferences. This generalizes the diagrammatic situation from Figure 1 to cases of larger n , where there can be many lexicographically minimal paths. Proposition 3 then shows that the disagreement probabilities that we have discussed above define a metric on Π .

A converse of Proposition 3 is also valid, and we provide a proof in the Appendix 2. Any metric on Π that is exactly additive along the lexicographically minimal paths in Π can be viewed as a disagreement probability for some suitably chosen distribution over \mathcal{X} . The additivity property along lexicographically minimal paths reflects the choice-based nature of the metric. Theorem 2 (in Appendix 2) characterizes all such metrics, extending the axiomatization of Kemeny-Snell (1962) who required additivity along all paths. Taken together, Proposition 3 and Theorem 2 show that the assumption of assent-maximization is equivalent to the assumption that the procedure used for preference aggregation is a function of a metric structure on preferences.

Proposition 2 *Let ρ_0, \dots, ρ_M be any lexicographically minimal path from e to π and let y_i be the number of transpositions at position i that are made along this path. Let $y = (y_2, \dots, y_n)$. Then $y = x(\pi)$.*

Proof of Proposition 2

Let ρ and ρ' be two paths from π to e with RBS signatures x and x' , respectively. Let k be the lowest index such that $x_k \neq x'_k$. The number of $A \in \mathcal{X}$ with cardinality $n - k + 2$ at which the choice from A changes along these paths is determined by (x_2, \dots, x_k) and (x'_2, \dots, x'_k) because all transpositions at position $k + 1$ and higher affect only sets of cardinality $n - k + 1$ and smaller. Since $x_i = x'_i$ for all $i < k$, by the definition of k , these two paths cannot make the same number of changes at sets with cardinality $n - k + 2$. Since both of these paths transform π into e , one of the paths must change one set of cardinality $n - k + 2$ more than once. ■

Proposition 3 *For any μ , the probabilities $F(\pi; \mu)$ form a semi-metric on Π . If $\mu_2 \neq 0$, then $F(\pi; \mu)$ is a metric on Π .*

Proof of Proposition 3

We need only verify the triangle inequality. This follows from the fact that the concatenation of lexicographically minimal paths, one from π to π' and the other from π' to π'' may or may not be a lexicographically minimal path from π to π'' . If it is not lexicographically minimal, strict inequality in the triangle inequality will hold.

As to the assertion that F is a metric whenever $\mu_2 \neq 0$, note that in this case there is a positive probability that A will be any particular pair. Thus

the difference will not be operationally important below. In these limiting cases we cannot distinguish between two preferences based on the choices they induce. The semi-metric induces a metric on the equivalence classes of preferences.

if π and π' rank any pair differently, that pair could be the available set and the distance between these two orderings will be non-zero. (When $\mu_2 = 0$ then two orderings that differ only in their two lowest ranked alternatives will never disagree because some better alternative will be available and will be the choice of both.) ■

Having a metric on Π is an analytical convenience in the computation of social choice rules and other functions on the space of distributions over Π . We use it to find a central preference that best represents the preferences of a population and to measure the dispersion among preferences. Other instances of the use of metrics on the space of preferences in social choice theory for related purposes are Kemeny (1959), Craven (1996), Klamler (2008), and Barthelmy and Monjardet (1981).

We now turn to the computational use of the RBS signature in deriving the disagreement probabilities. We will compute the number, $r(k, i)$, of sets A with cardinality k such that, if a transposition is made at position i that converts π to π' then $c_\pi(A) \neq c_{\pi'}(A)$. The transposition changes the order of two alternatives, π_{i-1} and π_i . If A has k elements in all, then $k - 2$ of them must come from those $n - i$ alternatives that follow π_{i-1} and π_i in the order π . Thus

$$r(k, i) = \binom{n-i}{k-2}$$

Of course if $k - 2 < 0$ then $r(k, i) = 0$.

As all sets with cardinality k are equally likely under our assumptions, each has probability $\frac{\mu_k}{\binom{n}{k}}$. Therefore the probability w_{ik} that the feasible set is one of the sets with cardinality k where the choice is changed by the transposition at position i is

$$w_{ik} = \frac{\mu_k}{\binom{n}{k}} r(k, i)$$

Each transposition is a measure of disagreement between π and e . The algorithm changes the choice at a family of feasible sets from $c_\pi(A)$ to $c_e(A)$ which is precisely the family at which π and e disagree, and it never makes a change where $c_\pi(A) = c_e(A)$. The lexicographical minimality of the path assures that the total disagreement is partitioned among the edges with no double counting. Thus, the probability that the feasible set is such that π and e disagree can be computed by adding the disagreement probabilities at every step of the RBS algorithm. Adding the probabilities as computed above and weighting each of them by the number of transpositions required at each position we have:

Proposition 4 *If π has RBS signature $x(\pi)$ and ν is exchangeable with size distribution μ then the probability of disagreement between π and e is given by*

$$F(\pi; \mu) = \sum_{i=2}^n \sum_{k=2}^n w_{ik} x_i = \sum_{i=2}^n x_i(\pi) \sum_{k=2}^n \frac{\mu_k}{\binom{n}{k}} \binom{n-i}{k-2} \quad (5)$$

This Proposition follows from the above discussion of counts of sets by cardinality.

Formula (5) shows how π and the size distribution μ enter into the probability of disagreement. The effect of the preference π is completely summarized by its RBS signature $x(\pi)$. This disagreement probability is a linear function of the signature, with coefficients dependent on the size distribution of the feasible sets.

We now examine some special cases of exchangeable distributions and see what this formula tells us.

The case in which it is sure that the feasible set will contain exactly two alternatives is the problem studied by Kemeny (1959). In this case we define

$$\mu^K = (0, 1, 0, \dots, 0)$$

We know from the Kemeny (1959) and Young and Levenshick (1978) papers that in this case F is Kendall's tau: $w_{ik} = 1$ for all k , so the formula reduces to $\sum_{i=2}^n x_i$, the total number of pairwise exchanges needed to carry π into e .

In the case where it is certain that all alternatives will be available, so that $A = X$ with probability one, we define

$$\mu^X = (0, \dots, 0, 1)$$

Two preferences π and π' disagree if and only if $\pi_1 \neq \pi'_1$. Thus $w_{2n} = 1$ and $w_{ik} = 0$ for $i > 1$ or $k < n$. In this case F assigns zero distance between all preferences that agree with e on the top element, and a unit distance for all those that disagree. It is in cases such as this that F is a semi-metric instead of a metric, because $\mu_2 = 0$.⁹

Finally, in the case in which all $A \in \mathcal{X}$ with two or more elements are equally likely, we define

$$\mu^E = \frac{(0, \binom{n}{2}, \dots, \binom{n}{n})}{\sum_{j=2}^n \binom{n}{j}}$$

The scalar factor adjusts for the fact that we put zero probability on sets of order zero and one. In the case where all subsets are equally likely, a transposition at position i affects the optimum if and only if both π_{i-1} and π_i are in A and none of π_k are in A for $k < i - 1$. The probability of this event is $\frac{1}{2^i}$. Below, we work through an example of how to compute $F(\pi; \mu)$ in this case.

Example 9 *Computing $F(\pi; \mu)$*

⁹Craven (1996) examines the social choice rule associated several metrics on the space of orderings different from μ^K following the Kemeny's approach in other respects. One of his metrics is μ^X , corresponding to the case in which it is certain that all alternatives are available. The other metrics are among those described in Diaconis (1988). They do not lead to an interpretation as the probability of disagreement and are thus not "choice-based" in our terminology.

Suppose we want to compute $F(a_3a_1a_2; \mu^E)$.

Step 1: Use RBS to sort $a_3a_1a_2$ into e . The following table catalogs the number of transpositions made at each position.

	$a_3a_1a_2$	$a_1a_3a_2$	$a_1a_2a_3$	Total
Transposition at position 2	0	1	0	1
Transposition at position 3	0	0	1	1

We write the signature of $a_3a_1a_2$ as $x(a_3a_1a_2) = (1, 1)$.

Step 2: Use the formula given in Theorem to compute probability of disagreement.

$$\begin{aligned}
F(\pi; \mu) &= \sum_{i=2}^n \sum_{k=2}^n w_{ik} x_i = \sum_{i=2}^n x_i(\pi) \sum_{k=2}^n \frac{\mu_k}{\binom{n}{k}} \binom{n-i}{k-2} \\
&= \sum_{i=2}^3 x_i(\pi) \sum_{k=2}^3 \frac{\mu_k}{\binom{3}{k}} \binom{3-i}{k-2} \\
&= x_2(\pi) \left[\sum_{k=2}^3 \frac{\mu_k}{\binom{3}{k}} \binom{3-2}{k-2} \right] + x_3(\pi) \left[\sum_{k=2}^3 \frac{\mu_k}{\binom{3}{k}} \binom{3-3}{k-2} \right] \\
&= x_2(\pi) \left[\frac{\mu_2}{\binom{3}{2}} \binom{3-2}{2-2} + \frac{\mu_3}{\binom{3}{3}} \binom{3-2}{3-2} \right] + x_3(\pi) \left[\frac{\mu_2}{\binom{3}{2}} \binom{3-3}{2-2} + \frac{\mu_3}{\binom{3}{3}} \binom{3-3}{3-2} \right] \\
&= 1 \left[\frac{\frac{3}{4}}{\binom{3}{2}} \binom{3-2}{2-2} + \frac{\frac{1}{4}}{\binom{3}{3}} \binom{3-2}{3-2} \right] + 1 \left[\frac{\frac{3}{4}}{\binom{3}{2}} \binom{3-3}{2-2} + \frac{\frac{1}{4}}{\binom{3}{3}} \binom{3-3}{3-2} \right] \\
&= 1 \left[\frac{\frac{3}{4}}{3} \binom{1}{0} + \frac{\frac{1}{4}}{1} \binom{1}{1} \right] + 1 \left[\frac{\frac{3}{4}}{3} \binom{0}{0} + \frac{\frac{1}{4}}{1} \binom{0}{1} \right] \\
&= \left[\frac{1}{4} + \frac{1}{4} \right] + \left[\frac{1}{4} + \frac{1}{4} (0) \right] \\
&= \frac{3}{4} \quad \square
\end{aligned}$$

Appendix 2 Axiomatic foundation for metrics on preference orderings

In the paper we define the metrics on Π as probabilities of disagreement at a randomly chosen $A \in \mathcal{X}$. In this appendix we give a set of axioms that give rise to these metrics, without having to mention random feasible sets or to be specific about their distribution. Our axioms are very much like those of Kemeny and Snell (1962) and thus we will be brief. There is one key modification of the Kemeny-Snell axioms. It is this modification that allows for all the metrics in our family other than Kendall's tau.

We are concerned with a function

$$f : \Pi \times \Pi \rightarrow \mathbb{R}$$

Axiom 1 f is a semi-metric

That is, f is a non-negative, symmetric function satisfying the triangle inequality $f(\pi, \pi') + f(\pi', \pi'') \geq f(\pi'', \pi)$

Axiom 2 f is order-preserving under permutations

If $f(\pi, \pi') \geq f(\pi, \pi'')$ then for all permutations ρ , $f(\rho \circ \pi, \rho \circ \pi') \geq f(\rho \circ \pi, \rho \circ \pi'')$

Axiom 3 If π^i and π^j are obtained from e by a single adjacent transposition at positions i and j respectively, then $i < j$ implies $f(\pi^i, e) \geq f(\pi^j, e)$

Axiom 4 If $(\rho_0, \rho_1, \dots, \rho_M)$ is a lexicographically minimal path from π to π' then $\sum_{k=1}^M f(\rho_{k-1}, \rho_k) = f(\pi, \pi')$.

Theorem 2 Let $x(\pi)$ be the RBS signature of π and assume that f satisfies Axioms 1-4. There exists a non-increasing set of non-negative numbers d_k for $k = 2, \dots, n$ such that

$$f(\pi, e) = \sum_{k=2}^n x_k(\pi) d_k$$

The proof of this Theorem is exactly like the proof in Kemeny and Snell except that, because of the weakening of Axiom 4 to lexicographically minimal paths instead of all paths (which Kemeny and Snell called "lines"), we can construct the distance from π to e only along lexicographically minimal paths. By Theorem 2 above, all lexicographically minimal paths have the same RBS signature. Thus, we need only specify the distances between two permutations that differ in a single transposition at i for all values of i . These can be any non-increasing numbers, according to Axiom 3. Then our Axiom 4 gives us the values of f as linear combinations of these distances. Because all lexicographically minimal paths generate the same distance, f is well-defined.

Obviously the Kendall τ distance corresponds to $d_k = 1$ for all k . If $d_m = 0$ for any m then f will be a semi-metric but not a metric, as $f(\pi, \pi')$ can be zero for $\pi \neq \pi'$. When all $d_m > 0$, f will be a metric.

If $(\rho_0, \rho_1, \dots, \rho_M)$ is a path from π to π' but not a lexicographically minimal path, as required by Axiom 4, then $\sum_{k=1}^M f(\rho_{k-1}, \rho_k) \geq f(\pi, \pi')$. To see this in an example, let us reconsider Example 8. There are two paths from $(a_3 a_2 a_1)$ to e given in this example: One, the RBS path, is $(a_3 a_2 a_1)$, $(a_3 a_1 a_2)$, $(a_1 a_3 a_2)$, $(a_1 a_2 a_3)$. The other (which is obtained by the regular bubble sort algorithm) is $(a_3 a_2 a_1)$, $(a_2 a_3 a_1)$, $(a_2 a_1 a_3)$, $(a_1 a_2 a_3)$. The former is lexicographically minimal while the latter is not. Calculating the distances along these paths, the former sums to $d_2 + 2d_3$ but the latter sums to $2d_2 + d_3$. Clearly these are equal in the case of the Kemeny metric and the triangle inequality holds strictly in all other cases. The lexicographically minimal path defines the distance $f((a_3 a_2 a_1), e)$.

To find the exchangeable distribution on \mathcal{X} corresponding to a given non-increasing vector (d_2, \dots, d_n) we can solve the system of equations derived from the formula in Theorem 4 for μ .

For each $i = 2, \dots, n$:

$$d_i = \sum_{k=2}^n \frac{\mu_k}{\binom{n}{k}} \binom{n-i}{k-2}$$

- Ahn D, Oliveros S (2010) Combinatorial Voting. *Econometrica* (conditionally accepted)
- Apestequia J, Ballester MA (2009) Choice by Sequential Procedures
- Arrow, Kenneth J (1951) *Social Choice and Individual Values*. Wiley, New York
- Barthelemy JP, Monjardet B (1981) The Median Procedure in Cluster Analysis and Social Choice Theory. *Mathematical Social Sciences* 1:235-268
- Conitzer V, Davenport A, Kalagnanam J (2006) Improved Bounds for Computing Kemeny Rankings. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06)*, Boston
- Craven J (1996) Best Fit Social Choices: An Alternative to Arrow. *The Economic Journal* 106(438):1161-1174
- Daugherty, Z., Eustis, A., Minton. G and M. Orrison, (2009) Voting, the symmetric group, and representation theory, *The American Mathematical Monthly* 116(8), 667-687.
- Diaconis P (1988) *Group Representations in Probability and Statistics*. Lecture Notes Monograph Series, volume 11, Hayward Institute of Mathematical Statistics
- Kemeny JG (1959) Mathematics Without Numbers. *Deadalus* 88: 577-591
- Kemeny JG, Snell JL (1962) *Mathematical Models in the Social Sciences*, Chapter 2, MIT Press, Cambridge
- Klamler C (2008) A Distance Measure for Choice Functions. *Social Choice and Welfare* 30: 419-425
- Le Conte de Poly-Barbut, C (1990) Automorphismes du permutoeedre et votes de Condorcet, *Math. Inf. Sci. hum.* 111, 73-82.
- Lu T, Boutilier C (2010) The Unavailable Candidate Model: A Decision-Theoretic View of Social Choice. *Proceedings of the Eleventh ACM Conference on Electronic Commerce*, 2010
- Moulin H (1988) *Axioms of Cooperative Decision Making*. *Econometric Society Monograph* number 15, Cambridge University Press, Cambridge
- Reuter K (1996) *Linear Extensions of a Poset as Abstract Convex Sets*. *Mathematisches Seminar*, Universitat Hamburg
- Saari D (2001) *Chaotic Elections*. American Mathematical Society
- Saari D (1994) *The Geometry of Voting*. Springer-Verlag, Berlin, New York
- Young HP (1988) Condorcet's Theory of Voting. *American Political Science Review* 82(2): 1231-1244
- Young HP, Levenglick A (1978) A Consistent Extension of Condorcet's Election Principle. *SIAM Journal on Applied Mathematics*, 35(2): 285-300
- Ziegler GM (1995) *Lectures on Polytopes*, Springer-Verlag, Graduate Texts in Mathematics, 152