A Failure of Representative Democracy*

Katherine Baldiga
Harvard University

August 30, 2011

Abstract

We compare direct democracy, in which members of a population cast votes for alternatives as choice problems arrive, and representative democracy, in which a population elects a candidate whose ordering of alternatives serves as a binding, contingent plan of action for future choice problems. While direct democracy is normatively appealing, representative democracy has practical advantages and is the more common institution. The key question, then, is whether representative democracy can successfully implement the choices that would be made under direct democracy. We perform a best case analysis, constructing models of direct and representative democracy that are as similar as possible and considering only those populations where majority rule over alternatives is consistent with an ordering. We show that even in this best case, where the normative recommendation of direct democracy is clear, representative democracy may not select the ordering consistent with majority rule over alternatives. Finally, we propose a population restriction that serves as a sufficient condition for consistency between direct and representative democracy.

*A special thank you to Jerry Green for many helpful conversations about this work. This author also gratefully acknowledges the comments and suggestions of Drew Fudenberg, Alvin Roth, Kenneth Shepsle, and several seminar participants at New York University and Harvard University.

1 Introduction

Direct democracy is a fundamental principle of collective decision-making. When a choice problem arrives, individuals should have the opportunity to express preferences over the alternatives. A good decision-making rule then aggregates these preferences into a choice that reflects the will of the group. While many aspects of social decision-making have been debated, this individual right to direct participation has remained a normative ideal from both a theoretical and popular standpoint.

Yet, despite its normative appeal, direct democracy is relatively rare in practice. Just over half of the states in the United States allow for recalls and/or referendums and no provisions for direct democracy exist at the federal level. Direct democracy plays a similarly minor role in the governments of countries around the world, with the well-studied exception of Switzerland (Frey [7]). Most institutions instead take the form of representative democracies, under which elected officials make decisions on behalf of the electorate. From a practical perspective, representative democracies have an edge over direct democracies as they dramatically reduce transaction costs and shift the burden of decision-making to a small group of well-informed leaders. The question, then, is whether representative democracy, with its practical advantages, can successfully implement the choices that would be reached under the more normatively attractive direct democracy.

This paper tackles this question from a theoretical perspective. The first and most crucial step of our analysis is to build models of both of these forms of governance. We begin with a population of individuals with strict preferences over a finite set of alternatives. Then, we develop and apply theories of how outcomes are reached under each system: how alternatives are chosen under direct democracy and how candidates are elected under representative democracy. With these models in place, we ask whether the choices of the population under direct democracy are preserved under representative democracy.

Our goal is to construct models of direct and representative democracy that are as similar as possible, so that we may study the case that is most likely to lead to positive results. As we describe below, both models depend only on individuals' preferences over pairs of alternatives. Thus, our theories map the same preference information from the population into two different sets of outcomes.

We use the majority rule tournament to model direct democracy. Under this system, alternatives are compared pairwise, with the majority preference dictating the winner of each contest. We focus on populations whose majority rule tournaments are consistent with an ordering of the alternatives. In these cases, it is easy to identify the normative standard for representative democracy. In our model of representative democracy, members of the population vote over candidates rather than alternatives. To simplify our setting, we define a candidate as an ordinal ranking of alternatives: a binding, contingent plan of action for future choice problems. We assume that when a choice problem of alternatives arrives, the social decision is made according to the ordering of the elected candidate.

The key modeling assumption is how individuals vote over candidates. We map individuals' preferences over alternatives into preferences over candidates by assuming that an individual votes for the candidate with whom she is most likely to agree about the choice from a randomly-selected pair of alternatives. While other models of candidate selection could be used, this model parallels the direct democracy case in an unique way. In this version of representative democracy, individuals form preferences over candidates solely based upon the decisions those candidates would make in choice problems that consist of pairs of alternatives. These are exactly the choice problems faced by individuals in the direct democracy tournament. Therefore, individuals' pairwise preferences completely determine both their votes over alternatives and over candidates.

These models map the preferences of the population into two distinct tournaments: a tournament over alternatives (direct democracy) and a tournament over orderings (representative democracy). Each tournament is a complete, asymmetric binary relation. The direct democracy tournament on the alternative space is well-studied. Our model of representative democracy generates a new type of tournament, a tournament over the orderings. One candidate beats another in this tournament relation if it earns a majority of the population's votes. We explore the basic properties of this new type of tournament and exploit its structural properties to prove our results.

The challenge in tournament theory is to identify the set of tournament winners. For our direct democracy tournaments, we do not need to employ a specific solution concept, as we focus exclusively on populations whose majority preferences are consistent with an ordering. The tournaments over orderings generated by these populations, however, may cycle. Identifying the winners of these tournaments is non-trivial. We use the uncovered set as our solution concept. The uncovered set has several nice properties: it is a subset of the top cycle, it is Condorcet consistent, and it contains no Pareto dominated alternatives (Miller [14]). Perhaps most importantly, Miller and Shepsle and Weingast have shown that many familiar voting methods under both sincere and strategic voting implement alternatives in the uncovered set (Miller [14] , Shepsle and Weignast [18]). In this way, the uncovered set best characterizes the likely outcomes under most pairwise forms of electoral competition. The use of this solution concept provides an institution-free setting, allowing us to derive results that generalize to many well-known voting rules.

With this machinery in place, we can address the question of whether representative democracy achieves the same outcomes as direct democracy. For example, suppose a population has the transitive majority preference $a_1 \succ a_2 \succ a_3$. In this case, we expect the winning candidate under representative democracy to be the ordering $a_1 \succ a_2 \succ a_3$, since this candidate would make the same choices as the population does under direct democracy.

We find that for problems with a small number of alternatives, representative democracy does succeed in electing candidates that implement the choices made under direct democracy. But, for general problems, this result does not hold.

We show that even for these most well-behaved populations, where majority preferences over the alternatives are consistent with an ordering, representative democracy may not elect the candidate with this ordering. In this way, representative democracy can fail in settings when the normative ideal is clear.

Importantly, the failure of representative democracy that we identify here fundamentally differs from many of the existing negative results in the fields of social choice and political economy. We prove our negative result for a non-cyclic population whose choices under direct democracy are perfectly consistent with an ordering, typically a non-problematic environment. And, in our counterexample, the ordering selected by the majority rule tournament over alternatives is also the ordering recommended by any positional method, such as the Borda count or plurality voting. Therefore, it is clear that this failure of representative democracy does not stem from a tension between matching majority preferences over alternatives and producing an ordering, nor does it stem from a tension between majoritarian and positional voting methods.

In the last section, we discuss restrictions on the distribution of preferences that are sufficient to guarantee that representative democracy implements the same choices as direct democracy. The sufficiency condition we propose can be interpreted as an analog to the well-known single-peakedness result proposed by Black and Arrow, as it describes a class of populations for which this form of majority rule (majority rule over candidates) functions well (Black [4] and Arrow [1]).

The rest of the paper proceeds as follows. Section 2 introduces notation and defines tournaments over orderings, our model of representative democracy. In Section 3, we explore the structure of tournaments over orderings and prove that McGarvey's Theorem [12] does not hold in this context. That is, given an arbitrary tournament relation, it may not be possible to find a population which will generate this relation as the outcome of a tournament over orderings. Section 4 considers the relationship between direct and representative democracy, proving consistency results for the case of $n = 3$ but negative results for more general problems. Section 5 concludes.

2 Notation and Model

2.1 Notation

First we introduce some notation. Our space of **alternatives** is finite and discrete, denoted $X = \{a_1, a_2, ..., a_n\}$. A **preference**, denoted π, is an ordering of the alternatives, where π corresponds to a permutation of the integers $\{1, ..., n\}$; given the preference $\pi = (a_{\pi(1)}a_{\pi(2)}...a_{\pi(n)})$, a_i is preferred to a_j if and only if $\pi^{-1}(i) < \pi^{-1}(j)$. It will be useful to write e to represent the alphabetical ordering of the alternatives, $e = (a_1a_2...a_n) \in \Pi$. The set of all $n!$ preferences over X is Π.

A **population**, λ, is a distribution over Π. Let Λ be the set of all distributions over Π.

We model a **candidate** as a strict ordering of the alternatives in X. We will write candidate π to denote the candidate with ordering π.

A **choice problem**, $A \subseteq X$, is a non-empty subset of alternatives; the set of all choice problems is $\mathcal{X}$. To each preference (and candidate), π, we can associate a **rational choice function**, $c_\pi : \mathcal{X} \to X$, where for each A, $c_\pi(A)$ is the element in A that is preferred to all other elements in A according to π. Our model assumes that an elected candidate, π, implements choices according to c_π.

We model these political environments using tournament theory. A tournament is a complete, asymmetric binary relation. Our analysis considers two types of tournaments: tournaments on the alternative space and tournaments on the candidate space. We use $\Gamma(X)$ to denote a **tournament on the space of alternatives**; we reserve the traditional T to refer to a **tournament on the space of candidates**, $T(\Pi)$. In both cases, a tournament depends upon the preferences of the population; therefore, we write Γ^λ or T^λ to denote the tournaments generated by the population λ.

2.2 Models of Direct and Representative Democracy

Under direct democracy, members of the population vote over alternatives. We use the majority tournament to model these decisions. In this tournament, the relationship between any pair of alternatives is defined by majority rule: $a_i \Gamma^\lambda a_j$ if $\sum_{\pi \in \Pi} (\lambda(\pi) | a_i \succ a_j) > \frac{1}{2}$. We write $a_i \Gamma^\lambda a_j$ if a_i beats a_j in the tournament on X, denoted $\Gamma^\lambda(X)$. In the populations we consider below, majority rule over alternatives will be consistent with an ordering, which we will denote π^*. We study this case for two reasons. Firstly, when the majority rule tournament is consistent with an ordering, the normative recommendation for representative democracy is clear: it should select precisely this ordering that is consistent with the direct democracy choices. Thus, this case serves as the most straightforward test of consistency between the two methods. Secondly, we hypothesize that this case has the highest likelihood of leading to positive results. Our negative result for this case suggests we would derive similar negative results in environments where there is additional tension between matching direct democracy choices and producing an ordering.

Building a model of representative democracy requires a theory of how individuals choose to vote over candidates. That is, we must map individuals' preferences over alternatives into preferences over candidates. Our goal is to construct a model of representative democracy that is as similar as possible to our model of direct democracy, as this provides the strictest test of our negative result. Because our direct democracy model depends only on preferences over pairs of alternatives, we build a model of representative democracy that relies upon only this same preference information.

We assume that an individual votes for the candidate with whom she is most likely to agree about the choice from a randomly-selected pair of alternatives. To present this theory formally, we use choice-based metrics. Baldiga and Green define choice-based metrics on the space of ordinal preferences as follows (Baldiga and Green [2]). First, we define a **probability distribution,** ν, over $\mathcal{X}$. Then, the distance between a pair of orderings can be described as the probability that the orderings disagree about the optimal choice from a randomly-selected subset of alternatives.

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

Our model of representative democracy will use the choice-based metric associated with the distribution over choice problems where each choice problem is a pair, and all pairs are equally likely. This particular choice-based metric is equivalent to the Kemeny distance (also known as the bubble sort distance and the Kendall distance), which defines the distance between π and π' as the number of pairs of alternatives the two orderings rank differently. Following the notation of Baldiga and Green, we denote it this choice-based metric by $f(\pi, \pi'; \mu^K)$.

We use the Kemeny distance to define individuals' preferences over candidates: voters prefer candidates whose orderings are closer to their own. Given two candidates π and π', a voter with preference π'' will prefer π if and only if $f(\pi'', \pi; \mu^K) < f(\pi'', \pi'; \mu^K)$. The Kemeny distance is perhaps the most familiar metric on the space of ordinal preferences. It has been employed in a similar manner in previous work in social choice. For instance, in their investigation of the strategy-proofness of social welfare functions, Bossert and Storcken assume, as we do, that individuals' preferences over orderings are determined by relative proximity under the Kemeny distance (Bossert and Storcken [5]). In Figure 1, we provide an illustration of the Kemeny distances between the six orderings in Π for the $n = 3$ case.

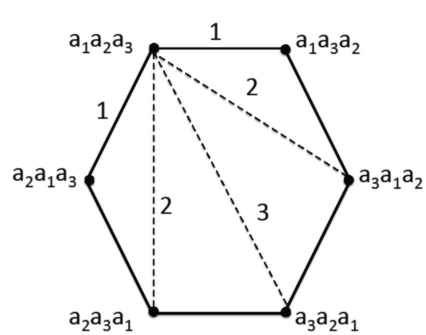

Figure 1: We use a hexagon to demonstrate the space of orderings for the $n = 3$ case, with each vertex representing an ordering. Above, we illustrate the Kemeny distances from $e = a_1a_2a_3$ to each of the other orderings in the space.

We use voters' preferences over candidates to generate a tournament on Π, denoted T^{λ,μ^K}. We say that π' **attracts a majority against** π if there exists a subset of preferences Π_1 with $\lambda(\Pi_1) > \frac{1}{2}$ such that $f(\pi'', \pi'; \mu^K) < f(\pi'', \pi; \mu^K)$ for all $\pi'' \in \Pi_1$.We write $\pi' T^{\lambda,\mu^K} \pi$ if π'attracts a majority against π. In the case where $f(\pi'', \pi'; \mu^K) = f(\pi'', \pi; \mu^K)$, we assume that the votes of those with preference π'' are split evenly between π and π'.[1] It straightforward to check that this binary relation is complete and asymmetric.

The tournament over orderings induced by $f(\pi, \pi'; \mu^K)$ is the most similar to the tournaments over alternatives induced by majority rule. Tournaments over orderings induced by other choice-based metrics take into account preferences over subsets other than pairs, information that is absent from the majority rule analysis of tournaments over alternatives. The tournaments defined by $f(\pi, \pi'; \mu^K)$ depend only on pairwise preferences. This tournament serves as a "best case" analysis; it is the environment in which we would expect the most consistency between direct and representative democracy. We prove negative results in this framework, illustrating that even in the case when these forms of direct and representative democracy have the same informational basis, their solutions are not consistent with one another.[2]

2.3 Tournament Solutions

We consider only the simplest direct democracy tournaments: those populations in which majority rule over alternatives is consistent with an ordering. In these cases, there is a clear winner of the tournament: the alternative which defeats all others in the tournament relation, the Condorcet winner. And in fact, it is straightforward to rank each of the alternatives in these tournaments, as there

[1]This assumption means we employ *relative* majority rule. We could alternatively define our tournaments in terms of *absolute* majority rule. Under this assumption, we would have $\pi' T^{\lambda,\mu^K} \pi$ only if a majority of voters are strictly closer to π' than π. For absolute majority rule, if neither candidate is closer to more than half of the population, the two candidates would tie in the tournament relation.

In our framework, voters will often be indifferent between candidates. Each voter has a set of indifference curves: a voter most prefers the candidate with his own ordering, then he equally prefers all candidates with whom he disagrees about the choice from one pair of alternatives, and next he equally prefers all candidates with whom they disagree about the choice from two pairs of candidates, etc.

Given the large amount of indifference in our population, choosing to use absolute majority rule would result in a large number of ties in our tournaments over orderings. These ties would disregard the information we have on the voters who are not indifferent. For instance, in a tie between candidate π and π', we may have that 30% of the voters are indifferent between candidate π and candidate π', 49% of the voters prefer π to π', and only 21% of voters prefer π' to π. Despite the large disparity in the number of voters that strictly prefer π to π', these two orderings would tie in the tournament relation. By using relative majority rule, we use this information on strict preference, even in the cases where large subsets of voters are indifferent between the two candidates.

[2]In the appendix, we explore a model of representative democracy based upon expected utility maximization. The inclusion of cardinal information is a sharp departure from the ordinal nature of the direct democracy tournament; thus, we would expect even less consistency between direct and representative democracy in this expected utility framework.

are no cycles.

However, for general tournaments, including the tournaments over orderings we study below, identifying the winners is non-trivial. In the simplest tournaments, there may be a Condorcet winner, but usually, the tournament relation will cycle. Therefore, we need a tournament solution, a mapping $S : T \to 2^{\Pi} \backslash \phi$ that will determine the best elements given an arbitrary tournament structure. We use the uncovered set as our tournament solution. Miller ([14]) provides the first characterization of the uncovered set. Here, we follow the definition given by Laslier ([11]), applied to our tournaments over orderings. First, define the covering relation of T. For a given T, we say π_i **covers** π_j if and only if:

(a) $\pi_i T \pi_j$, *and*

(b) $\forall \pi_k \in \Pi$, $\pi_j T \pi_k \Rightarrow \pi_i T \pi_k$

The **uncovered set of** T is the set of maximal elements of the covering relation: $\pi_i \in UC(T)$ iff $\nexists \pi_j \in \Pi$ such that π_j covers π_i.

As Miller ([14]) first described, the uncovered set has a number of appealing properties. The uncovered set is always a non-empty subset of the top cycle. And, unlike the top cycle, it contains only Pareto undominated orderings. It is Condorcet consistent: if a Condorcet winner exists, it will be the sole member of the uncovered set. Miller also noted the following game-theoretic interpretation of the uncovered set. Consider the following two-player zero sum game. The feasible strategies are the orderings, $\pi_1, ..., \pi_n \in \Pi$. If player i and player j play (π_i, π_j), then player i receives a payoff of 1 if $\pi_i T \pi_j$ and a payoff of -1 otherwise. Then the set of undominated strategies in this game is exactly the uncovered set of the tournament $T(\Pi)$.

In addition to having appealing normative properties, the uncovered set also characterizes the outcomes under a variety of familiar voting rules. Miller ([14]) and Shepsle and Weingast ([18]) have shown that a number of voting procedures under both sincere and sophisticated voting implement elements of the uncovered set. We discuss three broad classes of those procedures here.

The first is the amendment procedure under sophisticated voting. The amendment procedure begins with a majority vote over a pair of alternatives. The loser of the pairwise contest is eliminated from contention; the winner of that pairwise contest advances to face another alternative in the next stage. This process continues through $n-1$ stages, with the surviving alternative selected as the winner. Miller proves that the amendment procedure under sophisticated voting must select a member of the uncovered set as a winner. Further, Shepsle and Weingast show that an alternative $a_i \in X$ can be the sophisticated voting outcome of an amendment procedure agenda containing a_j in the first stage if and only if a_j does not cover a_i. Therefore, the uncovered set characterizes the set of implementable outcomes under the amendment procedure with sophisticated voting.

Cooperative models of voting also lead to the selection of elements of the uncovered set. Miller shows that when voters may form coalitions and play cooperatively, communicating and making binding agreements to vote together, then any majoritarian procedure will implement an element of the uncovered set.

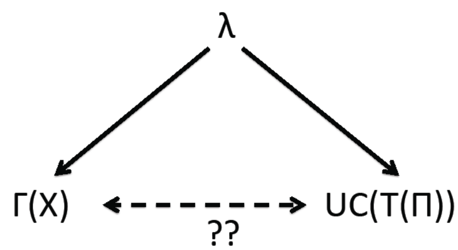

Figure 2: Our Approach

Finally, Miller discusses a model of strategic electoral competition that chooses elements of the uncovered set. In his model, there are two political parties, A and B. Each party must select a "platform" from the set of possible alternatives, $a_i \in X$. Voters are assumed to have rational preferences over the alternatives; they have a dominant strategy to vote sincerely for the party whose platform they prefer. Under these assumptions, this model is a symmetric, 2-person, zero-sum game for the parties A and B. Miller shows that the set of undominated strategies is exactly the uncovered set of X.[3]

These results motivate our use of the uncovered set. By working with this tournament solution, we avoid making specific institutional assumptions. Instead, we identify the likely winners of the tournament more generally. The negative result we prove in this framework can be applied to any voting procedure which leads to the election of a member of the uncovered set.

2.4 Our Approach

We are now ready to pose our question more formally. We consider a population with preferences over a finite set of alternatives. We look at the choices this population would make under direct democracy by computing the tournament over alternatives induced by majority rule. Then, we determine which candidates this population would elect under representative democracy by computing the tournament over orderings induced by the Kemeny distance. We ask whether the ordering consistent with the majority preferences over alternatives is contained in the uncovered set of the tournament over orderings. This approach is summarized in Figure 2.

Below we discuss our three main results. First, we analyze the structure

[3]McKelvey [13] proves similar results in the context of multidimensional choice spaces and quasiconcave preferences.

of T^{λ,μ^K} and show how the underlying majority preference over alternatives imposes constraints on the space of tournaments over orderings that can be generated. We conclude that the mapping $\lambda \to T^{\lambda,\mu^K}$ is not onto. That is, given an arbitrary tournament, T, we cannot in general find a population λ which generates this binary relation as a tournament over orderings induced by $f(\pi, \pi'; \mu^K)$. This is in contrast to the famous theorem of McGarvey, which states that given an arbitrary tournament over a set of n alternatives X, a population of strict preferences exists which will generate this tournament as the outcome of majority rule (McGarvey [12]).

We then turn our attention to the uncovered sets of tournaments over orderings. We start in the case of $n = 3$ and show that for these problems, the results of representative democracy are consistent with the results of direct democracy. We prove that if majority rule on X is consistent with an ordering, then this ordering is the sole member of the uncovered set of T^{λ,μ^K}. This consistency result, however, does not hold for more general problems. In fact, we show that for general n, even if majority rule on X is consistent with an ordering, this ordering may not be a member of the uncovered set of T^{λ,μ^K}. This result is surprising in light of the fact that both tournaments depend only of the population's pairwise preferences over alternatives. Perhaps more importantly, this negative result arises in the case where social choice is typically least problematic: when there are no cycles in majority rule. The wedge between direct and representative democracy stems not from a tension between matching majority preferences over alternatives and producing a social ordering, but rather from a more pervasive tension between the ordering that represents majority preferences and the ordering committed to ex ante.

3 The Structure of Tournaments over Orderings

One of the most influential results in social choice is **McGarvey's Theorem.** It states that given an arbitrary tournament over a set of n alternatives X, a population of strict preferences exists which will generate this tournament as the outcome of majority rule (McGarvey [12]). Thus, when it comes to majority rule over alternatives, anything goes; the structure of the tournament space over alternatives is unconstrained.

This result does not hold for tournaments over orderings defined by the Kemeny distance. Consider an arbitrary binary relation over the orderings. Then, holding the metric $f(\pi, \pi'; \mu^K)$ fixed, it is not generally possible to find a population λ that will generate this binary relation as the outcome of the tournament over orderings. That is, the mapping $\lambda \to T^{\lambda,\mu^K}$ is not onto. This result is interesting from a mathematical standpoint and it is also computationally helpful, as restrictions on the space of tournaments over orderings reduce the number of cases to consider in the analysis that follows.

The additional structure on the space of tournaments over orderings stems from the tournament's reliance on the underlying majority relation over alternatives. We illustrate this dependence with a simple example for the $n = 3$

case. Recall that the distance between π and π' as measured by the metric $f(\pi, \pi'; \mu^K)$ is proportional to the number of pairwise disagreements between the two orderings. Therefore, when determining whether a voter with ordering π'' will vote for candidate π or π' in T^{λ,μ^K}, it is enough to count the number of pairs on which π'' disagrees with each of the two orderings. The candidate with more pairwise choices in common with π'' will win the support of voter π''. As a result, we can use majority rule over pairs of alternatives to pin down the relationship between certain pairs of orderings in the population.

Consider a population whose majority preference over the pair (a_1, a_2) is $a_1 \succ a_2$, which yields $a_1 \Gamma^\lambda a_2$. This determines the relationship between pairs of candidates whose orderings disagree only about the relative ranking of a_1 and a_2. For example, consider the pair of candidates $a_1a_2a_3$ and $a_2a_1a_3$. Since the majority of the population prefers $a_1 \succ a_2$, we must have that a majority of the population is closer to $a_1a_2a_3$ than $a_2a_1a_3$ according to $f(\pi, \pi'; \mu^K)$. Similarly, we can deduce that a majority of the population is closer to $a_3a_1a_2$ than $a_3a_2a_1$. If the majority preference on (a_1, a_2) were reversed, the relationship between both of these pairs of orderings would also be reversed. Therefore, we have:

$$\begin{aligned} a_1 \Gamma^\lambda a_2 \quad &\Leftrightarrow \quad a_1a_2a_3 T^{\lambda,\mu^K} a_2a_1a_3 \\ a_1 \Gamma^\lambda a_2 \quad &\Longleftrightarrow \quad a_3a_1a_2 T^{\lambda,\mu^K} a_3a_2a_1 \\ a_1a_2a_3 T^{\lambda,\mu^K} a_2a_1a_3 \quad &\Longleftrightarrow \quad a_3a_1a_2 T^{\lambda,\mu^K} a_3a_2a_1 \end{aligned}$$

This makes it clear that there is no population such that $a_1a_2a_3 T^{\lambda,\mu^K} a_2a_1a_3$ but $a_3a_2a_1 T^{\lambda,\mu^K} a_3a_1a_2$.

Because of the close tie between a population's majority preferences over pairs of alternatives and the tournament relation T^{λ,μ^K}, it will be helpful to classify all $n = 3$ populations according to their majority preferences. The first class of populations has a transitive majority preference of the type $a_1 \Gamma^\lambda a_2$, $a_2 \Gamma^\lambda a_3$, and $a_1 \Gamma^\lambda a_3$. There are six possible majority relations in this class, one corresponding to each of the six possible orderings. We will show that for any one of these transitive majority relations, there are only two possible tournaments over orderings T^{λ,μ^K} that can be generated. The second class of populations has a cyclic majority preference; there are only two possible cyclic majority relations: $a_1 \Gamma^\lambda a_2$, $a_2 \Gamma^\lambda a_3$, $a_3 \Gamma^\lambda a_1$ and $a_2 \Gamma^\lambda a_1$, $a_1 \Gamma^\lambda a_3$, $a_3 \Gamma^\lambda a_2$. In the appendix, we show that for either one these cyclic majority preferences, there are six possible tournaments over orderings T^{λ,μ^K} that can be generated. Thus, we can conclude that only 24 of the 32,768 (2^{15}) possible asymmetric binary relations can be generated by a tournament over orderings induced by a population λ and the metric $f(\pi, \pi'; \mu^K)$.

Before we move to our results, we discuss a few useful observations about the general structure of T^{λ,μ^K}.

Proposition 1 *Consider a population where majority preferences are consistent with e. Take any ordering π that has at least one pair of adjacent alternatives ordered according to e. Obtain π' by performing one transposition of* *adjacent alternatives that appeared in the natural order in π. Then, we have $\pi T^{\lambda,\mu^K}\pi'$ for all such π'. Furthermore, we can take any of these π' that have at least one pair of adjacent alternatives ordered according to e and obtain π'' by performing one transposition of alternatives that appeared in the natural order in π and π'. Then, we have $\pi T^{\lambda,\mu^K}\pi''$.*

Proof: Intuitively, when π and π' agree on all but a single pair of alternatives, π'' will be closest to whichever of these orderings it agrees with on the pair in the question. Since π agrees with the majority preference on the pair in question, more than half of the population must be closer to π, yielding $\pi T^{\lambda,\mu^K}\pi'$.

We can take this logic one step further to prove the claim for cases where π and π'' agree on all but two pairs of alternatives. Those orderings which agree with π on exactly one pair will be equidistant from π and π''. Those subsets of the population which agree with π on both of the pairs in question will be closer to π; those which disagree with π on both of the pairs in question will be closer to π''. We can prove the claim by formally analyzing these subsets of the population. Denote by λ_α the subset of the population that agrees with π on the first pair in question (the one transposed to obtain π'), denoted α. Denote by λ_β the set that agrees with π on the second pair in question (the one transposed to obtain π''), denoted β. Because π is consistent with the majority preference on both of these pairs, we know that $\lambda_\alpha > .5$ and $\lambda_\beta > .5$. Therefore:

$$2 \times (\lambda_\alpha \cap \lambda_\beta) + \lambda_\alpha \backslash \lambda_\beta + \lambda_\beta \backslash \lambda_\alpha > 1$$

We can rearrange this expression to show that:

$$\lambda_\alpha \cap \lambda_\beta > 1 - (\lambda_\alpha \backslash \lambda_\beta + \lambda_\beta \backslash \lambda_\alpha + \lambda_\alpha \cap \lambda_\beta)$$

The left-hand side of this equation is the fraction of the population that agrees with π on both of the pairs in question and will be closer to π. The right-hand side of this equation is the fraction of the population that disagrees with π on both of the pairs in question and will be closer to π''. The rest of the population will be equidistant. Thus, the equation tells us that a larger fraction of the population will be closer to π than π'', yielding $\pi T^{\lambda,\mu^K}\pi''$. ■

The following corollary is a straightforward implication of Proposition 1 and will prove useful in the following sections.

Corollary 1 *An ordering that is consistent with the majority preferences of a population λ, call this π^*, must dominate all orderings that are fewer than two transpositions away from it; that is, we must have $\pi^* T^{\lambda,\mu^K}\pi$ for all π such that $f(\pi^*, \pi; \mu^K) \leq \frac{2}{2^n-(n+1)}$.*

4 Consistency Results for Direct and Representative Democracy

4.1 Defining Consistency

With this knowledge of the structure of tournaments over orderings, we turn now to computing solutions. For our populations, in which the majority preferences over alternatives are consistent with an ordering, there is a clear test of whether representative democracy implements the choices made under direct democracy. We ask whether the ordering consistent with the majority preferences over alternatives is a member of the uncovered set of the tournament over orderings.

Definition 1 *Strong Order Consistency: If majority preferences over alternatives are consistent with an ordering, then this ordering is the sole member of the uncovered set of* T^{λ,μ^K}.

Or, we might require a weaker condition, which simply requires inclusion of this ordering in the uncovered set of the tournament over alternatives:

Definition 2 *Order Consistency: If majority preferences over alternatives are consistent with an ordering, then this ordering is a member of the uncovered set of* T^{λ,μ^K}.

In the analysis that follows, we show that order consistency holds for the case of $n = 3$ but fails for larger problems with $n > 3$. Then, we provide a counterexample which illustrates that for $n \geqq 10$, $UC(T^{\lambda,\mu^K})$ does not satisfy order consistency. In the appendix, we address the case in which majority preferences over alternatives cycle. There, we discuss notions of consistency for these cyclic populations and prove that for the case of $n = 3$, cyclic populations satisfy reasonable notions of consistency.

4.2 Consistency for *n=3*

We exploit results about the structure of tournaments over orderings to prove results for the $n = 3$ case with transitive majority preferences consistent with an ordering. Here, we work through the case where a_1 is the Condorcet winner of the tournament with $a_1\Gamma^\lambda a_2$, $a_2\Gamma^\lambda a_3$, and $a_1\Gamma^\lambda a_3$, proving that a population with this underlying majority relation can generate only two possible tournaments T^{λ,μ^K}. We show that simply knowing the majority preference over the three pairs of alternatives is enough to determine 13 of the 15 pairwise relationships between candidates and the uncovered set of T^{λ,μ^K}; the remaining two relationships, which do not impact the uncovered set, are determined jointly by the relative sizes of the majorities.

Our first step is to map majority preferences over pairs onto T^{λ,μ^K}. The majority preference over each of the three pairs (a_1, a_2), (a_1, a_3), and (a_2, a_3)

determines the relationship between three pairs of orderings. To see this, we work out the full set of implications of $a_1\Gamma^{\lambda}a_2$.

Let $a_1\Gamma^{\lambda}a_2$. Then, we know:

$$\sum_{\pi_i\in\Pi}(\lambda(\pi_i)|a_1 \succ_{\pi_i} a_2) > \sum_{\pi_i\in\Pi}(\lambda(\pi_i)|a_2 \succ_{\pi_i} a_1)$$
$$\lambda(a_1a_2a_3) + \lambda(a_1a_3a_2) + \lambda(a_3a_1a_2) > \lambda(a_2a_1a_3) + \lambda(a_2a_3a_1) + \lambda(a_3a_2a_1)$$

This inequality implies a set of T^{λ,μ^K} relations. As made clear by the upper left-hand diagram in Figure 3, the orderings on the left-hand side of the inequality are closer to $a_1a_2a_3$ than $a_2a_1a_3$, closer to $a_1a_3a_2$ than $a_2a_3a_1$, and closer to $a_3a_1a_2$ than $a_3a_2a_1$. Therefore:

$$a_1a_2a_3T^{\lambda,\mu^K}a_2a_1a_3$$
$$a_1a_2a_3T^{\lambda,\mu^K}a_2a_1a_3$$
$$a_3a_1a_2T^{\lambda,\mu^K}a_3a_2a_1$$

We can do this for each of the three pairs, pinning down nine of the relationships in T^{λ,μ^K}. Figure 3 illustrates this process.

Next, we can map pairs of majority preferences onto T^{λ,μ^K}. For instance, since we have $a_1\Gamma^{\lambda}a_2$ and $a_1\Gamma^{\lambda}a_3$, we know:

$$\sum_{\pi_i\in\Pi}(\lambda(\pi_i)|a_1 \succ_{\pi_i} a_2) + \sum_{\pi_i\in\Pi}(\lambda(\pi_i)|a_1 \succ_{\pi_i} a_3) >$$
$$\sum_{\pi_i\in\Pi}(\lambda(\pi_i)|a_2 \succ_{\pi_i} a_1) + \sum_{\pi_i\in\Pi}(\lambda(\pi_i)|a_3 \succ_{\pi_i} a_1)$$
$$\lambda(a_1a_2a_3) + \lambda(a_1a_3a_2) + \lambda(a_3a_1a_2) + \lambda(a_1a_2a_3) + \lambda(a_1a_3a_2) + \lambda(a_2a_1a_3) >$$
$$\lambda(a_2a_1a_3) + \lambda(a_2a_3a_1) + \lambda(a_3a_2a_1) + \lambda(a_2a_3a_1) + \lambda(a_3a_1a_2) + \lambda(a_3a_2a_1)$$
$$2\lambda(a_1a_2a_3) + 2\lambda(a_1a_3a_2) > 2\lambda(a_2a_3a_1) + 2\lambda(a_3a_2a_1)$$
$$\lambda(a_1a_2a_3) + \lambda(a_1a_3a_2) > \lambda(a_2a_3a_1) + \lambda(a_3a_2a_1)$$

The orderings on the left-hand side of the inequality are closer to $a_1a_2a_3$ than $a_2a_3a_1$ and closer to $a_1a_3a_2$ than $a_3a_2a_1$. Since the other orderings, $a_2a_1a_3$ and $a_3a_1a_2$ are equidistant from each pair of orderings, this is enough to determine the relationship between each of these pairs.

$$a_1a_2a_3T^{\lambda,\mu^K}a_2a_3a_1$$
$$a_1a_3a_2T^{\lambda,\mu^K}a_3a_2a_1$$

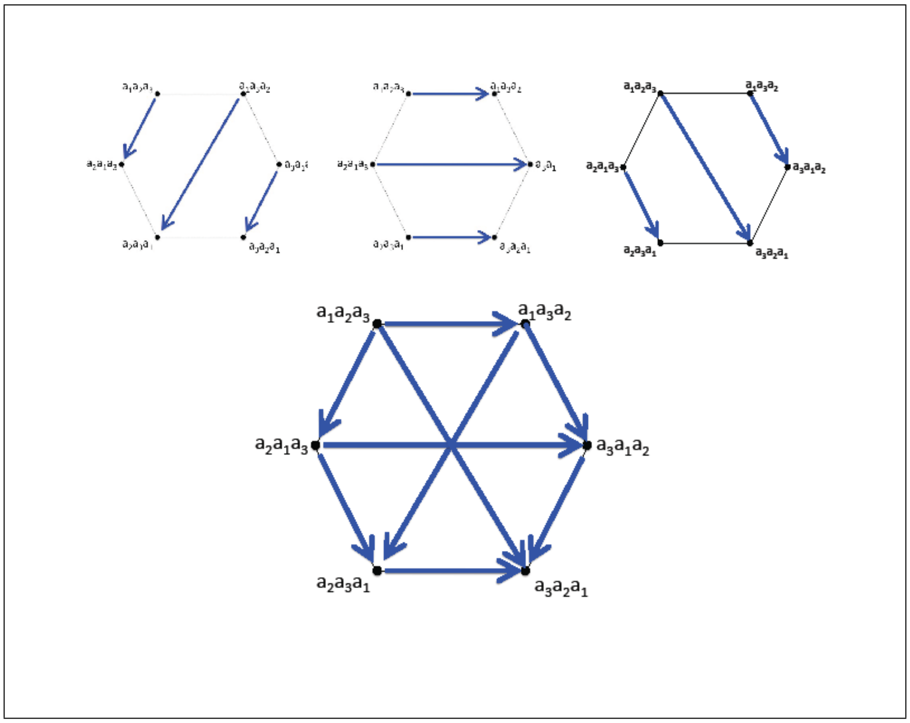

Figure 3: We use a hexagon with an ordering on each vertex to represent the structure of the $n = 3$ space of orderings. Throughout this paper, we'll use an arrow to illustrate a T relationship, with the arrow directed toward the defeated ordering. In the upper left-hand corner, the three blue arrows illustrate T^{λ,μ_K} relationships determined by $a_1\Gamma^{\lambda}a_2$. In the middle of the top row, the three blue arrows illustrate T^{λ,μ_K} relationships determined by $a_2\Gamma^{\lambda}a_3$.In the upper right-hand corner, the three blue arrows illustrate T^{λ,μ_K} relationships determined by $a_1\Gamma^{\lambda}a_3$.The diagram on the bottom row summarizes the nine T^{λ,μ_K} relationships determined by majority preferences over alternatives.

We can draw a similar set of conclusions working from the majority preference of $a_1\Gamma^\lambda a_3$ and $a_2\Gamma^\lambda a_3$, deducing:[4]

$$a_2a_3a_1T^{\lambda,\mu^K}a_3a_2a_1$$
$$a_1a_2a_3T^{\lambda,\mu^K}a_3a_1a_2$$

Figure 4 summarizes these relationships:

Regardless of how the final two links are resolved, we see that $a_1a_2a_3$ is the Condorcet winner of the tournament of orderings, beating every other ordering directly. Therefore, we have $UC(T^{\lambda,\mu^K}) = \{a_1a_2a_3\}$.

Relating this result to our definitions of consistency above, we can state the following proposition:

Proposition 2 *For $n = 3$, $UC(T^{\lambda,\mu^K})$ satisfies strong order consistency.*

4.3 Inconsistency for General Problems

First, we present an example that shows that for $n = 4$, strong order consistency fails. This proves that for $n > 3$ there exist populations whose majority preferences are consistent with an ordering but whose uncovered sets of T^{λ,μ^K} are multi-valued. Consider the following population:

π	$\lambda(\pi)$
$a_1a_2a_3a_4$	.399
$a_2a_4a_1a_3$	.2
$a_1a_4a_3a_2$	.2
$a_3a_4a_1a_2$	.201

It is clear that the majority preferences of this population are consistent with $e = a_1a_2a_3a_4$. But, there is no Condorcet winner of the tournament over alternatives. We can show that there exists an ordering, denoted $\hat{\pi}$ such that $\hat{\pi}T^{\lambda,\mu^K}e$. Consider $\hat{\pi} = a_4a_1a_3a_2$. We have $f(\hat{\pi}, a_1a_4a_3a_2) < f(e, a_1a_4a_3a_2)$ and $f(\hat{\pi}, a_3a_4a_1a_2) < f(e, a_3a_4a_1a_2)$, and we have $f(\hat{\pi}, a_2a_4a_1a_3) = f(e, a_2a_4a_1a_3)$. So,

$$\sum_{\pi\in\Pi}[\lambda(\pi)|f(\hat{\pi},\pi) < f(e,\pi)] > \sum_{\pi\in\Pi}[\lambda(\pi)|f(e,\pi) < f(\hat{\pi},\pi)]$$

As a result, $\hat{\pi}T^{\lambda,\mu^K}e$. And, in fact, we can show that $\hat{\pi} \in UC(T^{\lambda,\mu^K})$, with $UC(T^{\lambda,\mu^K}) = \{e, a_1a_2a_4a_3, a_1a_4a_2a_3, \hat{\pi}\}$. We are able to extend this $n = 4$ example to a problem with an arbitrary number of alternatives by preserving the structure above for the first four alternatives and simply appending additional alternatives in their natural order to the right end of each of the four orderings above. This leads to the following proposition.

[4]Note that the majority preference over $a_1\Gamma^\lambda a_2$ and $a_2\Gamma^\lambda a_3$ only tells us that $\lambda(a_1a_2a_3) > \lambda(a_3a_2a_1)$, which alone is not enough to determine the T^{λ,μ_K} relationship between any pair of orderings.

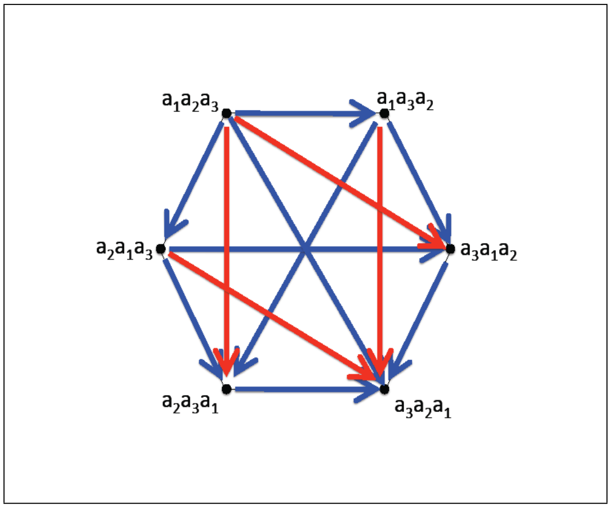

Figure 4: The four red arrows illustrate the additional T^{λ,μ_K} relationships that are determined by pairs of majority preferences over alternatives.

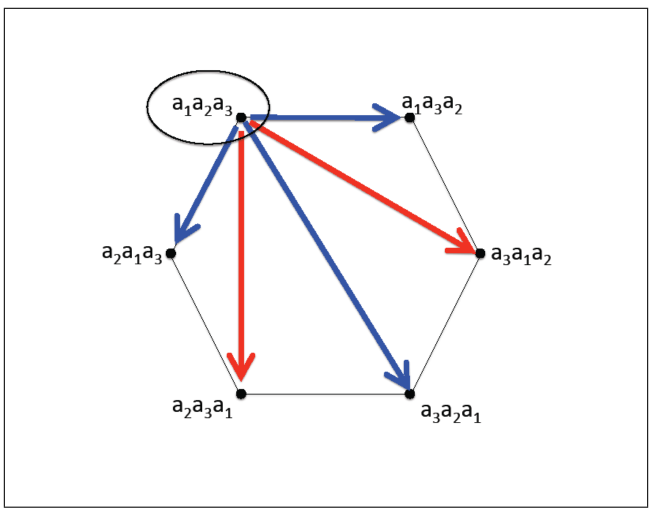

Figure 5: The diagram above illusrates the uncovered set of the transitive population for $n = 3$. Because $a_1a_2a_3$ beats all other orderings, it is the Condorcet winner of the tournament and the sole member of the uncovered set.

Proposition 3 *For $n \geq 4$, $UC(T^{\lambda,\mu^K})$ fails strong order consistency.*

This is the strongest result that we can achieve for $n = 4$, as it is straightforward to demonstrate that for populations that are consistent with a majority preference, this ordering must be a member of its uncovered set. Assume majority preferences are consistent with e. We can show that there is no ordering that can cover e. The key step is to recognize that we can apply Proposition 1 to prove that any ordering fewer than five transpositions from e cannot cover e: $eT^{\lambda,\mu^K}\pi$ for any π within two transpositions, so they cannot cover e, and for those three or four transpositions away, even if they beat e, they will be defeated by at least one ordering one or two transpositions from e (which e beats). So, the only orderings that could potentially cover e are five or six transpositions away from e: $\{a_4a_3a_1a_2, a_4a_2a_3a_1, a_3a_4a_2a_1, a_4a_3a_2a_1\}$. In the appendix, we rule these out one at a time, proving that order consistency must hold for $n = 4$ populations, as $e \in UC(T^{\lambda,\mu^K})$ for all populations whose majority preferences are consistent with e:

For $n \geq 10$, we are able to prove a stronger result. We identify populations in which the ordering consistent with the majority preferences is covered. For smaller problems like the one above, we can find a $\hat{\pi}$ such that $\hat{\pi}T^{\lambda,\mu^K}e$, but this $\hat{\pi}$ does not beat everything that e beats. That is, we can apply Proposition 1 to find a chain of the type $eT^{\lambda,\mu^K}\pi T^{\lambda,\mu^K}\hat{\pi}$ where π is just two transpositions away from both e and $\hat{\pi}$. In order to prevent these types of chains, $\hat{\pi}$ must be at least five transpositions from e. We can do this most simply for problems with at least 10 alternatives. However, as we discuss after the proof, it may be possible to find other types of populations that violate order consistency with fewer than 10 alternatives.

Theorem 1 *For $n \geq 10$, $UC(T^{\lambda,\mu^K})$ fails order consistency.*

Proof: We prove this through a general counterexample. First, we construct the population. Let $\pi^*(\lambda, \mu^K) = e$ for n alternatives. The majority preferences of our population λ will be consistent with e on all pairs of alternatives. Select j pairs of alternatives, where $5 \leq j \leq \frac{n}{2}$. Each pair should consist of two adjacent elements in the natural ordering, and all pairs should be mutually exclusive. For example, it would be permissible to select $\{a_1, a_2\}$ and $\{a_3, a_4\}$ as two of the pairs, but one could not select $\{a_1, a_2\}$ and $\{a_2, a_3\}$, or $\{a_1, a_3\}$ and $\{a_4, a_6\}$. It will be helpful to have notation for the j pairs; let them be denoted $p_1, p_2, ..., p_j$. Note that since the pairs consist of adjacent and mutually exclusive alternatives, it is always possible to find an ordering π that agrees with the majority preference on any particular subset of the pairs $\{p_1, p_2, ..., p_j\}$ exactly.

We will associate with each of the j pairs a particular ordering, π_{p_i}, where π_{p_i} agrees with the majority preference on pair p_i, disagrees with the majority preferences on the other $j - 1$ pairs in the set $\{p_1, p_2, ..., p_j\}$, and agrees with the majority preference on all other pairs of alternatives.

Allocate the population as follows. Let $\lambda(e) = \frac{1}{2} - \varepsilon$, where $\varepsilon < \frac{1}{2j}$. Divide the rest of the population evenly among the orderings $\{\pi_{p_1}, \pi_{p_2}, ..., \pi_{p_j}\}$, creating j equal masses of size $\frac{\frac{1}{2}+\varepsilon}{j}$.

It is straightforward to check that this population produces majority rule that agrees with e. For each pair not included in $\{p_1, p_2, ..., p_j\}$, the population unanimously prefers a_i to a_{i+1}. For each of the pairs in $\{p_1, p_2, ..., p_j\}$, we have that $(\frac{1}{2} - \varepsilon) + (\frac{\frac{1}{2}+\varepsilon}{j}) > \frac{1}{2}$ agrees with e.

Now we will show that for this population $\pi^*(\lambda, \mu^K) \notin UC(\lambda, \mu^K)$. We do so by identifying an ordering which covers π^*. Consider the ordering which disagrees with the majority preferences on all of the pairs $\{p_1, p_2, ..., p_j\}$ and agrees with the majority preference on all other pairs. Denote this order $\hat{\pi}$. Since we are working with the metric generated by μ^K, the distance between any two orderings is, up to a scale factor, the number of pairs over which the two orderings disagree. For simplicity, we'll scale our distances below to the number of pairwise disagreements.

First we'll show that $\hat{\pi} T \pi^*$. We have $f(\pi_{p_i}, \hat{\pi}) = 1$ and $f(\pi_{p_i}, \pi^*) = j - 1$ $\forall i \in \{1, ..., j\}$. Therefore, we have $\frac{1}{2} + \varepsilon$ of the population that is closer to $\hat{\pi}$ than π^*, so $\hat{\pi} T \pi^*$.

In order to prove that $\hat{\pi}$ covers π^*, we must show that there cannot exist a π' such that $\pi^* T \pi'$ but $\pi' T \hat{\pi}$. Suppose there did exist such a π'.

We have that $\pi^* T \pi'$. This implies that we have $f(\pi^*, \pi_{p_i}) \leq f(\pi', \pi_{p_i})$ for at least some $i \in \{1, ..., j\}$. Because we know $f(\pi^*, \pi_{p_i}) = j - 1$, this implies that $f(\pi', \pi_{p_i}) \geq j - 1$.for at least some $i \in \{1, ..., j\}$.

We also know that $\pi' T \hat{\pi}$. Then we must have that $f(\pi', \pi_{p_k}) \leq f(\hat{\pi}, \pi_{p_k})$ for at least some $k \in \{1, ..., j\}$. We know that $f(\hat{\pi}, \pi_{p_k}) = 1$ for all $k \in \{1, ..., j\}$,which implies that $f(\pi', \pi_{p_k}) \leq 1$ for at least some $k \in \{1, ..., j\}$.

Finally, we know that $f(\pi_{p_i}, \pi_{p_k}) \leq 2$ for any $i, k \in \{1, ..., j\}$.

This creates the following violation of the triangle inequality: $f(\pi', \pi_{p_k}) \leq 1$, $f(\pi_{p_i}, \pi_{p_k}) \leq 2$, and $f(\pi', \pi_{p_i}) \geq j - 1$, where $j \geq 5$.This is a contradiction.

Therefore, there can exist no π' such that $\pi^* T \pi'$ but $\pi' T \hat{\pi}$. As a result, we can conclude that $\hat{\pi}$ covers π^*. ■

This proof describes a method for constructing populations for which direct and representative democracy yield different choices. It shows that when there are 10 or more alternatives, it is always possible to construct a population such that the majority preferences are consistent with an ordering but for which this ordering is not a member of the uncovered set of T^{λ, μ^K}. These populations have a rather intuitive interpretation. Let's think about a population constructed by the method above for the case of $n = 10$. First, we note the distinction between the five "contested" choices $(\{a_1, a_2\}, \{a_3, a_4\}\{a_5, a_6\}\{a_7, a_8\}\{a_9, a_{10}\})$ and the other 40 pairwise choices which are decided unanimously. The largest mass of voters, just under half of them, have the preference e.Let's call these our "mainstream" voters. The remaining voters are divided evenly among five smaller minority preferences. Each minority agrees with the mainstream preference on just one of the contested issues; on the other hand, each minority block agrees with every other minority block on three of the five contested issues. In this way, the minority preferences are all more similar to one another than to the mainstream voters.

Let's think about the choices this population would make under direct and representative democracy. When voting directly over the alternatives, the population implements choices consistent with e. Most of the choices are unanimous; and, for the five contested pairs, the mainstream voters and one of the minority groups form a majority. Though the minority groups have similar preferences, when voting issue-by-issue, they never vote all together on a contested issue. As a result, the mainstream voters are able to implement their preferred choices. We can contrast this with the dynamics under representative democracy. In this setting, candidate e cannot be elected. Though e attracts the mainstream voters, there exist candidates which all five minority groups prefer to e. Consider candidate $\hat{\pi}$, which agrees with the unanimous choices of the population but disagrees with the majority preference on the five contested pairs. All five minority preferences prefer $\hat{\pi}$ to e; together, they consist of a majority of the population and can succeed in electing this compromise candidate. The choices of the selected candidates in representative democracy are more closely aligned with the minority preferences than the choices under direct democracy. The ability of the minority groups to compromise under representative democracy produces outcomes that are at odds with the choices implemented under direct democracy.

While our model is quite different, our results echo the findings of Besley and Coate, who study the question of whether citizens' initiatives, which allow citizens to cast votes directly over issues, improve upon the outcomes reached under electoral competition among representatives (Besley and Coate [3]). Their model consists of a two party political system, where the population makes decisions in a two-dimensional policy space. Using this framework, they show that the elected candidates may implement policies that are odds with the majority preferences of the population. They attribute these errors to the bundling of issues that is inherent in the election of a representative. As in our model, when issues are decided upon concurrently, via the choice of a representative, decisions may diverge from those made when citizens are able to vote directly over issues, one at a time.

The mechanics behind this result also reveal a connection to the political science literature on vote trading, or log rolling (see Tullock [19], Riker and Brams [15]). A vote trade may occur when two voters (or parties) are on opposite sides of two issues, with each being a pivotal voter in the majority on exactly one the two. In these circumstances, both voters may stand to gain by agreeing to vote against their true preference on the issue in which they are pivotal, in exchange for the other doing the same. Through vote trades, voters may succeed in implementing their preferred choices on the issues they value most, when they would otherwise have been out-voted by the majority. A sizable literature exists documenting the potential for this type of vote trading to alter outcomes under majority rule. As Schwartz ([17]) observes, in cases in which majority rule does not cycle and thus would lead to a unique, stable outcome, the outcome under vote trading can be different. While the specifics of the two models are rather different, the underlying intuition is similar. In the model presented here, representative democracy, like vote trading, provides a way for small parties to impact choices through joint action. By each sacrificing their preferred choice on one issue, the minority parties are able to agree upon a compromise candidate which is an improved outcome for all of them.

While our result is similar in spirit to the vote trading literature, we point out that it is rather different from most other negative results about majority rule, and social choice in general, because it does not stem from a problem of cyclical collective preferences. It would be less surprising to find inconsistencies between direct and representative democracy when there was no ordering consistent with the direct democracy choices. But in our example, there is no tension between producing an ordering and matching the majority preferences of the population. And in fact, the divergence noted here is distinct from the type of conflict that has been identified between pairwise methods and scoring rules in the work of Saari and others (see Saari [16]). For the populations in the counterexample above, the ordering generated by applying a scoring rule to the preferences over alternatives is exactly the order that is consistent with majority rule. Thus, the prediction of representative democracy in this setting is at odds with the predictions of most familiar forms of electoral competition and proposed theoretical voting rules, which all agree on the "correct" ordering for this type of population. This makes the result more surprising, and perhaps more troublesome. We have proven the result for a relatively well-behaved population; we expect that for populations with cyclical majority preferences or inconsistencies between scoring rules and pairwise methods, the choices under direct and representative democracy will diverge to a greater extent. Better understanding the tournaments over orderings associated with more general classes of populations is a topic for future study.

A natural question to ask in this context is whether we can impose restrictions on the distribution of preferences that would guarantee order consistency. This approach has been adopted by many social choice theorists in attempts to rule out other paradoxical outcomes; perhaps most classic is the single-peakedness restriction pioneered by Black ([4]) and Arrow ([1]). Their preference restriction considers populations that are unimodal in the sense that all members of the population, for any particular triple of alternatives, can agree on an alternative that is not worst. Assuming the number of voters is odd, this condition is sufficient for transitive majority rule. Clearly, this restriction is not enough to assure order consistency, as the class of populations we consider in our counterexample above are indeed single-peaked. However, we can use a similar idea, that of restricting the number of modes in the distribution, in order to derive a sufficient condition for order consistency in our framework. The class of populations with transitive majority rule consistent with an ordering, denoted π^*, can be thought of as having a "peak" or cluster of weight around π^*. Our sufficiency condition says that as we move away from π^*, we must not encounter another cluster of similar orderings. In order to state this condition more formally, it will be useful to introduce some new terminology. When referring to a population with transitive majority rule consistent with π^*, we will call any pairwise disagreement with π^* an **error**. For example, we will say that an ordering π that is m transpositions from π^* contains m errors. We can state our sufficiency condition in terms of these errors.

Condition 2 *Consider the class of populations with transitive majority rule consistent with π^*. Then, order consistency holds if for any set of m errors, $m \geqq 5$, we have*

$$\sum \lambda(\pi|\pi \text{ contains at least } \frac{1}{2} \text{ of these } m \text{ errors}) < \frac{1}{2}$$

Proof. Suppose $\pi^* \notin UC(T^{\lambda,\mu^K})$. We will show there must exist a set of m errors, $m \geqq 5$, such that $\sum \lambda(\pi|\pi$ contains at least $\frac{1}{2}$ of these m errors) $< \frac{1}{2}$. Since $\pi^* \notin UC(T^{\lambda,\mu^K})$, we know there exists $\hat{\pi}$ such that $\hat{\pi}$ covers π^*. Let $\hat{\pi}$ contain m errors; we know $m \geqq 5$ in order for $\hat{\pi}$ to cover π^*. Since $\hat{\pi}$ and π^* agree on all pairwise choices outside of the m errors, we know that the distance between any π and $\hat{\pi}$ and π^* is determined only by how many of these m errors π shares. Those that have less than $\frac{1}{2}$ of the m errors will be such that $f(\pi, \pi^*; \mu^K) < f(\pi, \hat{\pi}; \mu^K)$. So, suppose the set of orderings that had at least $\frac{1}{2}$ of these m errors in common with $\hat{\pi}$ had mass less than $\frac{1}{2}$. Then, we would have $\sum \lambda(\pi|f(\pi, \pi^*; \mu^K) < f(\pi, \hat{\pi}; \mu^K)) > \frac{1}{2}$. This would imply $\pi^* T^{\lambda,\mu^K} \hat{\pi}$, contradicting $\hat{\pi}$ covers π^*. ■

This sufficiency condition has a straightforward intuition. If we encounter a population that contains a mass of orderings that are both (a) relatively distant from π^*, and (b) relatively close to one another, then we may have the type of counterexample presented above. This type of cluster of similar orderings far from π^* may be able to agree upon a compromise candidate which covers π^*, but only if together they constitute a majority. The condition rules out this possibility by assuring that no such cluster of mass greater than $\frac{1}{2}$ exists.

Like in the case of single-peakedness, this is a sufficient but not necessary condition for $\pi^* \in UC(T^{\lambda,\mu^K})$. A gray area exists between the class of populations described in our counterexample above and the class of populations described by this sufficiency condition. For populations that fall into neither of these classes, perhaps the best answer to whether π^* is a member of $UC(T^{\lambda,\mu^K})$ is it depends. For some populations that fail the condition above, the distribution of mass on orderings far from π^* may be too dispersed to agree upon an ordering like $\hat{\pi}$ which could beat everything that π^* beats. One might ask whether we could improve the sufficiency condition by restricting this set of distant orderings to fall within a certain radius of one another. For an example that illustrates why this strategy fails, please see the Appendix.

We summarize our results for populations with order consistent majority preferences in the table below:

	$n=3$	$n=4$	$5 \leq n \leq 9$	$n \geq 10$
strong order consistency	√	x	x	x
order consistency	√	√	?	x

Order consistency results for $5 \leq n \leq 9$ remain an open question. Our strategy here focuses on generating populations by manipulating preferences over mutually exclusive, non-overlapping pairs. Preferences over one pair of alternatives are independent of preferences over a mutually exclusive non-overlapping pair. As a result, we are able to combine arbitrary preferences over these pairs into an ordering. This makes generating a population with particular majority preferences much simpler. This approach has been fruitful in constructing populations for which order consistency fails, though it may require more alternatives than might be necessary with other approaches. An important open question is what is smallest number of alternatives necessary to have order consistency break down.

We have chosen to use the uncovered set as the solution concept for our tournaments. How heavily does our result depend on this choice? One of the attractive features of the uncovered set is that many other popular tournament solution concepts are subsets of the uncovered set (Laslier [11]). Therefore, it is possible to extend our negative result to many other solution concepts. This includes the basic refinements of the uncovered set, the iterated uncovered set and the minimal covering set. It also includes the Banks solution and the Bipartisan set. Another well-studied method for identifying tournament winners is ranking the members of the tournament based upon their victories and losses within the tournament. The most popular of these ranking methods include the Copeland solution, the Markov solution, and the Slater solution, each of which is also a refinement of the uncovered set.

However, if we expanded our solution concept to the top cycle, which contains the uncovered set, we coud prove a positive result. For any population with majority preferences consistent with an ordering, this ordering is a member of the top cycle of the tournament over orderings. We can prove this using Proposition 1. By performing one transposition at a time, each transposing a pair of alternatives that were ordered according to the majority preference, we can construct a chain, $\pi^* T^{\lambda,\mu^K} \pi_i T^{\lambda,\mu^K} \pi_j T^{\lambda,\mu^K} ... T^{\lambda,\mu^K} \pi$, from π^* to any other ordering π. of the type, $\pi^* T^{\lambda,\mu^K} \pi_i T^{\lambda,\mu^K} \pi_j T^{\lambda,\mu^K} ... T^{\lambda,\mu^K} \pi$. Thus, we must have $\pi^* \in TC(T^{\lambda,\mu^K})$. But for populations like the one above, the top cycle is large; it will contain many other orderings, including our $\hat{\pi}$. Thus, while order consistency would hold, strong order consistency would fail.

5 Conclusion

In this paper, we introduce a tournament over orderings as a natural model of representative democracy. Instead of looking at pairwise comparisons of alternatives, as traditional tournaments do, our tournaments are built from pairwise comparisons of candidates, modeled as orderings of alternatives. We assume that when comparing two candidates, an individual votes for the candidate with whom she is most likely to agree about the choice from a randomly-selected pair of alternatives.

We find that when applied to the same problem, tournaments over alternatives and tournaments over orderings often predict different outcomes. In fact, even when the choices of the population under direct democracy are consistent with an ordering, this ordering may not be selected under representative democracy. We conclude that the choice of institution plays a significant role in determining outcomes. While this is a familiar theme in much of the existing social choice and political economy literature, our result is novel because of the context in which we prove it. We consider a majoritarian method and find that even in settings when a Condorcet winner exists, a candidate may be elected who will not implement this alternative in certain choice problems.

Many open questions remain. One avenue to pursue would be considering different methods for defining tournaments over orderings. Because we are interested in comparing our results to the outcome of tournaments defined by majority preferences over pairs of alternatives, we focus on tournaments over orderings defined by the Kemeny distance. However, it would be interesting to consider tournaments generated by other choice-based metrics, corresponding to different distributions over choice problems. We could ask how the likelihood of a choice problem relates to the probability of representative democracy implementing the outcome recommended by direct democracy. That is, are representative and direct democracy more likely to agree on choice problems which are very likely to arrive?

We could also explore ways of adapting this framework to more applied settings. In this work, we study a very abstract model of representative democracy. In most institutional settings, representative democracy has more structure than what we've imposed here: there may be party systems which shape which candidates are available, restrictions on how many issues or alternatives can be decided upon at once, priority given to status quo alternatives or candidates. And, the process is dynamic. Considering these modifications would enrich the model, leading to sharper predictions and results that speak to existing forms of representative democracy.

6 Appendices

6.1 Expected Utility Tournaments

In the analysis above, we consider a model of representative democracy generated by the choice-based metric, $f(\pi, \pi'; \mu^K)$. Alternatively, we could use an expected utility framework to model representative democracy. To create a cardinal model, we must enrich our basic model in a few key ways. We introduce utility functions over the alternatives, $U : X \to R$; a utility function, u, assigns a real number to each alternative in X. Then, we redefine a population, λ_u, as a distribution over utility functions rather than a distribution over orderings. We can easily generate a model of direct democracy, $\Gamma^{\lambda_u}(X)$, via majority preference in this setting: $a_i \Gamma^{\lambda_u} a_j$ if $\sum_{u \in U} [\lambda_u(u) | u(a_i) > u(a_j)] > \frac{1}{2}$.[5]

To define preferences over the candidates, we build up from our basic utility function over alternatives. Consider a member of the population with utility function, u. She evaluates a candidate π according to an expected utility function, $\omega_{u,\nu}$:

$$\omega_{u,\nu}(\pi) = \sum_{A \in \mathcal{X}} \nu(A) \times u(c_\pi(A))$$

She evaluates candidate π by considering the choices π would make from each possible choice problem. She takes a weighted average of the utility of each choice weighted by the likelihood of that choice problem arriving, which is given by ν. From here, it is straightforward to define a rule for voting over candidates. We say that π' **attracts a majority against** π if there exists a subset of utility functions, U_1, with $\lambda_u(U_1) > \frac{1}{2}$ such that $\omega_{u,\nu}(\pi) < \omega_{u,\nu}(\pi')$ for all $u \in U_1$. In the case where $\omega_{u,\nu}(\pi) = \omega_{u,\nu}(\pi')$, we assume that the votes of those with utility function u are split evenly between π and π'. It is clear that this rule also generates a tournament over the orderings.

A natural question to ask is what is the relationship between the ordinal, metric-based model and the expected utility model here. Consider a tournament over orderings generated by a population λ and a choice-based metric, $f(\pi, \pi'; \nu)$.Is it possible to find a population of utility functions, λ_u, that generates this same tournament under $\omega_{u,\nu}$? In general, the answer to this question is no. And in fact, for general n, there is no expected utility tournament that corresponds to the tournament over orderings induced by $f(\pi, \pi'; \mu_K)$. However, if we restrict our attention to the $n = 3$ case and those ν that only put positive probability on choice problems of order two, we can find such a mapping.

Proposition 4 *Consider a tournament over orderings generated by a population λ and $f(\pi, \pi'; \mu^K)$. For $n > 3$, it is not generally possible to find a population of utility functions, λ_u, that generates this same tournament under $\omega_{u,\nu}$.*

Proof: We can show this with a quick counterexample. Let $n = 4$. Consider a member of the population whose preference is e. We will show that we cannot find a utility function u such that his preferences over orderings according to $\omega_{u,\nu}(\pi)$ are the same as his preferences over orderings induced by $f(e, \pi; \mu^K)$. We know that for this agent $\omega_{u,\nu}(\pi)$ must be the same for all π such that $f(e, \pi; \mu^K) = \frac{1}{9}$ (these are all the π that are exactly one transposition from e). Thus, we must have that $u(a_i) - u(a_{i+1}) = u(a_{i+1}) - u(a_{i+2})$ for $i \in \{1, 2\}$. For convenience, let us label this distance $u(a_i) - u(a_{i+1}) = \alpha$. Now consider the following pair of orderings: $a_3a_1a_2a_4$ and $a_2a_1a_4a_3$. For both of these π, we have $f(e, \pi; \mu^K) = \frac{2}{9}$. But, in terms of expected utility according to $\omega_{u,\nu}(\pi)$, when choices are made according to $a_3a_1a_2a_4$, an agent with preference e has

[5]For simplicity, assume u assigns unique values to each $a_i \in X$. This avoids the need to define tie-breaking rules.

a utility loss of 2α (for choosing a_3 rather than a_1 from $\{a_1a_3\}$) plus a utility loss of α (for choosing a_3 rather than a_2 from $\{a_2a_3\}$). When choices are made according to $a_2a_1a_4a_3$, an agent with preference e has a utility loss of α (for choosing a_2 rather than a_1 from $\{a_1a_2\}$) plus a utility loss of α (for choosing a_4 rather than a_3 from $\{a_3a_4\}$). So, $\omega_{u,\nu}(a_2a_1a_4a_3) > \omega_{u,\nu}(a_3a_1a_2a_4)$ even though $f(e,\pi;\mu^K) = \frac{2}{9}$ for both of these π. Therefore, we cannot find a utility function that induces the same preferences over orderings as $f(e,\pi;\mu^K)$ does. ■

Proposition 5 *Suppose $n = 3$. Consider a tournament over orderings generated by a population λ and a choice-based metric, $f(\pi,\pi';\nu)$, with ν such that $\nu(A) \succ 0$ only if $|A| = 2$. Then it is possible to find a population of utility functions, λ_u, that generates this same tournament under $\omega_{u,\nu}$.*

Proof: Suppose that there are three alternatives, $\{a_1, a_2, a_3\}$, and the likelihood of the four choice problems are given by $\nu : \nu(a_1a_2) = p, \nu(a_2a_3) = q, \nu(a_1a_3) = 1 - p - q$, where $p, q \in [0, 1]$. Without loss of generality, consider a member of the population with ordering $\pi = a_1a_2a_3$. Then we have:

$f(\pi,\pi;\nu) = 0$
$f(\pi, a_1a_3a_2;\nu) = q$
$f(\pi, a_2a_1a_3;\nu) = p$
$f(\pi, a_2a_3a_1;\nu) = 1 - q$
$f(\pi, a_3a_1a_2;\nu) = 1 - p$
$f(\pi, a_3a_2a_1;\nu) = 1$

We will show that we can find a utility function, u, to assign to this member of the population such that his preferences over the orderings induced by $\omega_{u,\nu}$ will be the same as his preferences induced by the metric $f(\pi,\pi';\nu)$. Without loss of generality, let $u(a_1) = x$, $u(a_2) = y$, $u(a_3) = 0$, with $x > y > 0$. Then, the expected payoff from each ordering is:

$\omega_{u,\nu}(\pi) = [1 - q]x + qy$
$\omega_{u,\nu}(a_1a_3a_2) = [1 - q]x$
$\omega_{u,\nu}(a_2a_1a_3) = [p + q]y + [1 - p - q]x$
$\omega_{u,\nu}(a_2a_3a_1) = [p + q]y$
$\omega_{u,\nu}(a_3a_1a_2) = px$
$\omega_{u,\nu}(a_3a_2a_1) = py$

Straightforward calculations reveal that $f(\pi,\pi';\nu)$ and $\omega_{u,\nu}$ induce the same preference over orderings for 11 of the 15 possible pairwise comparisons. That is, for 11 of the 15 pairs of orderings, we have $f(\pi,\pi';\nu) < f(\pi,\pi'';\nu)$ if and only if $\omega_{u,\nu}(\pi') > \omega_{u,\nu}(\pi'')$. There are four remaining relationships to be determined: $(a_1a_3a_2, a_2a_1a_3)$, $(a_1a_3a_2, a_2a_3a_1)$, $(a_2a_1a_3, a_3a_1a_2)$, and $(a_2a_3a_1, a_3a_1a_2)$.

The four relationships in question are governed by the following inequalities, shown simplified here.

$$a_1a_3a_2 \succ a_2a_1a_3 : px > (p + q)y \tag{1}$$

$$a_1a_3a_2 \succ a_2a_3a_1 : (1 - q)x > (p + q)y \tag{2}$$

$$a_2a_1a_3 \succ a_3a_1a_2 : (p+q)y + (1-p-q)x > px \tag{3}$$

$$a_2a_3a_1 \succ a_3a_1a_2 : (p+q)y > px \tag{4}$$

It is clear that for all values of p, q, x, and y 4 implies 3 and that 1 implies not 4. There are five possible cases to consider. For each case, we list the relevant ranges of p and q. Then, we list the corresponding preference over orderings induced by ν for this range of probabilities.

(1) $1-p > p > q$: 1, 2, and 3 hold and 4 does not

(2) $p > 1-p$, $p > q, (1-q) > p$: 1 and 2 hold and 3 and 4 do not

(3) $p > 1-p$, $p > q, (1-q) < p$: 1 holds and 2, 3, and 4 do not

(4) $1-q > q > p$: 1 does not hold and 2, 3, and 4 do hold

(5) $q > p$, $q > 1-q$: 1 and 2 do not hold and 3and 4 do hold

For each case, we must prove that we can find a u (values of x, y) such that we deduce the same preferences over orderings when looking at $\omega_{u,\nu}$.

Case 1: Need to show $\exists\ x, y$ such that 1, 2, and 3 hold and 4 does not. In this case, since $p < (1-q)$, 1 implies 2. Therefore, we only need to check:

$$\begin{aligned}
(p+q)y &< px < (p+q)y + (1-p-q)x \\
0 &< px - (p+q)y < (1-p-q)x \\
0 &< p - (p+q)\frac{y}{x} < (1-p-q) \\
-p &< -(p+q)\frac{y}{x} < (1-2p-q) \\
\frac{p}{p+q} &< \frac{y}{x} < \frac{2p+q-1}{p+q}
\end{aligned}$$

Since the left-hand side is weakly less than 1, we just need to check that $p < 2p+q-1$. This holds as long as $p < 1-q$, which is true in this case. $\square$

Case 2: Need to show $\exists\ x, y$ such that 1 and 2 hold and 3 and 4 do not. In this case, since $p < (1-q)$, 1 implies 2. And, not 3 implies 1. Therefore, we only need to check:

$$\begin{aligned}
px &> (p+q)y + (1-p-q)x \\
p &> (p+q)\frac{y}{x} + (1-p-q)x \\
2p+q-1 &> (p+q)\frac{y}{x} \\
\frac{y}{x} &< \frac{2p+q-1}{p+q}
\end{aligned}$$

So, we just need to have $0 < 2p+q-1$. This holds as long as $p > 1-p-q$, which is true in this case since $p > 1-p$. $\square$

Case 3: Need to show $\exists\ x, y$ such that 1 holds and 2, 3, and 4 do not. In this case, not 3 implies 1. Therefore, we only need to check:

$$\begin{aligned}(1-q)x &< (p+q)y < px - (1-p-q)x \\ (1-q) &< (p+q)\frac{y}{x} < 2p+q-1 \\ \frac{(1-q)}{p+q} &< \frac{y}{x} < \frac{2p+q-1}{p+q}\end{aligned}$$

Since $p > 1-q$, the left-hand side is less than 1. We just need to check that $(1-q) < 2p+q-1$. This holds as long as $p > 1-q$, which is true in this case. $\square$

Case 4: Need to show $\exists\ x, y$ such that 1 does not hold and 2, 3, and 4 do hold. In this case, 4 implies 3 and not 1. Therefore, we only need to check:

$$\begin{aligned}px &< (p+q)y < (1-q)x \\ p &< (p+q)\frac{y}{x} < 1-q \\ \frac{p}{p+q} &< \frac{y}{x} < \frac{1-q}{p+q}\end{aligned}$$

Since the left-hand side is weakly less than 1, we just need $(1-q) > p$, which is true in this case. $\square$

Case 5: Need to show $\exists\ x, y$ such that 1 and 2 do not hold and 3and 4 do hold. In this case, 4 implies 3 and not 1. And, since $(1-q) > p$, we have not 2 implies 4. Therefore, we only need to check:

$$\begin{aligned}(p+q)y &> (1-q)x \\ \frac{y}{x} &> \frac{1-q}{p+q}\end{aligned}$$

We just need $(1-q) < p+q$, which is true in this case since $q > 1-q$. $\square$

By exhausting these five cases, we've completed our proof. We've shown that for any $n = 3$ tournament over orderings generated by a population λ and a choice-based metric, $f(\pi, \pi'; \nu)$, with ν such that $\nu(A) \succ 0$ only if $|A| = 2$, it is possible to find a population of utility functions, λ_u, that generates this same tournament under $\omega_{u,\nu}$. $\blacksquare$

Despire the differences between the choice-baed metric models and the expected utility models, the counterexample presented in Section 4 does work for a family of expected utility generated tournaments. Restrict all members of the population to the same scale of cardinalization; that is, each utility function assigns the alternative in the ith spot the same real number $\forall i \in 1, ..., n$. Then if $u(a_k) - u(a_{k-1}) = u(a_{k-1}) - u(a_{k-2})$, the counterexample is still valid. Furthermore, holding fixed any particular utility function, common to all members of the population, it is likely possible to construct counterexamples that would allow us to prove an expected utility analog to 4.3. This is a topic for future work.

6.2 Cyclic Populations

6.2.1 Defining Consistency

Here, we consider a few reasonable notions of consistency between direct and representative democracy for cyclic populations. In the next subsection, we'll use these concepts to evaluate cyclic populations for $n = 3$.

Our first notion of consistency makes an effort to adapt the idea of order consistency to the case of cyclic populations. While majority preferences over alternatives are not generally consistent with an ordering, the Kemeny rule provides a way to identify an ordering that is "closest" to the population's preferences (Kemeny [8], Kemeny and Snell [9]). The Kemeny ordering is the ordering which minimizes the expected distance to a randomly-drawn member of the population, where distance is given by $f(\pi, \pi'; \mu^K)$:

$$\pi^*(\lambda, \mu^K) = \arg\min_{\pi} \sum_{\pi' \in \Pi} f(\pi, \pi'; \mu^K)\lambda(\pi')$$

In the case where the population's majority preferences are consistent with an ordering, this ordering is the Kemeny ordering. This suggests a reasonable pair of consistency conditions for our tournament comparison:

Definition 3 *Strong Kemeny Order Consistency: For any population,* $\pi^*(\lambda, \mu^K) = UC(T^{\lambda,\mu^K})$.

Definition 4 *Kemeny Order Consistency: For any population,* $\pi^*(\lambda, \mu^K) \in UC(T^{\lambda,\mu^K})$.

Requiring the membership of the Kemeny ordering in the uncovered set of the tournament over alternatives is one intuitive operationalization of consistency in this context. It asks that the ordering that best fits the population's majority preferences be selected as a winner of the tournament over orderings. Given that we define our tournaments over orderings by $f(\pi, \pi'; \mu^K)$, it seems plausible to expect that the Kemeny ordering be a member of $UC(T^{\lambda,\mu^K})$.

We know, however, that the Kemeny method is just one way of producing a representative social ordering from a population of individual rational preferences. As Arrow's Theorem and the rich literature that followed has shown, there is no clear answer as to what ordering best represents the majority preferences of the population (Arrow [1]). Therefore, when developing a notion of testable consistency for our framework, it may be appealing to employ a more flexible concept. Instead of looking for any particular ordering in the uncovered set, we could ask that the orderings in $UC(T^{\lambda,\mu^K})$ simply rank highly the winners from the tournament over alternatives. To formalize this concept, we define weak consistency:

Definition 5 *Weak Consistency: Suppose* $a_i \in UC(\Gamma(X))$ *and* $a_j \notin UC(\Gamma(X))$. *Then, for any* $\pi \in UC(T^{\lambda,\mu^K})$*,we do not have* a_j *precedes* a_i.

This axiom requires that members of $UC(T^{\lambda,\mu^K})$ rank alternatives in $UC(\Gamma(X))$ above elements that are not in $UC(\Gamma(X))$. In the case where there is a Condorcet winner among alternatives, $a_i^* = UC(\Gamma(X))$, weak consistency requires that a_i^* be ranked first in any ordering in $UC(T^{\lambda,\mu^K})$. Note that in contrast to the notions of consistency proposed above, weak consistency does not make any statements about the ranking of alternatives within $UC(\Gamma(X))$. In this sense, weak consistency is more in keeping with the rationale of the uncovered set. Because any alternative within the uncovered set can be considered a winner of direct democracy, we have no grounds for requiring a specific ranking of these elements in the orderings of $UC(T^{\lambda,\mu^K})$.

6.2.2 Results

For populations with cyclic majority preferences, $UC(\Gamma(X)) = \{a_1, a_2, a_3\}$. It is immediately clear that weak consistency will be vacuous in this case, as every alternative is a member of the uncovered set. But we can evaluate our other notions of consistency by looking at the structure of these tournaments. For illustration, we study the case of the cyclic majority preference $a_1\Gamma^\lambda a_2$, $a_2\Gamma^\lambda a_3$, and $a_3\Gamma^\lambda a_1$ and prove that this majority preference can generate six possible tournaments with identical consistency properties.

We proceed in the same fashion as above. First, we map majority preferences over pairs onto T^{λ,μ^K}. Again, the majority preference over each of the three pairs (a_1, a_2), (a_1, a_3), and (a_2, a_3) determines the relationship between three pairs of orderings. As shown in Figure 6, the only wrinkle is the reversal of the relationship between the pairs of orderings that depend on the majority preference of (a_1, a_3).

Because we no longer have that $a_1\Gamma^\lambda a_3$, we cannot employ the technique of combining information about pairs of majority preferences over pairs. Instead, we have to rely on information about the strength of majorities in order to flesh out the rest of the tournament relation. We know that for a given cyclic majority preference, we can rank the strength of the majority for each pair; so, we can identify a strongest link (the pair for which the size of the majority is largest) and a weakest link (the pair for which the size of the majority is smallest).

Suppose in our example that the weakest link is $a_3\Gamma^\lambda a_1$. The fact that $a_3\Gamma^\lambda a_1$ is supported by a smaller majority than $a_1\Gamma^\lambda a_2$ yields the following inequality:

$$[\lambda(a_3a_1a_2) + \lambda(a_3a_2a_1) + \lambda(a_2a_3a_1)] - [\lambda(a_1a_2a_3) + \lambda(a_1a_3a_2) + \lambda(a_2a_1a_3)] <$$
$$[\lambda(a_1a_2a_3) + \lambda(a_1a_3a_2) + \lambda(a_3a_1a_2)] - [\lambda(a_3a_2a_1) + \lambda(a_2a_3a_1) + \lambda(a_2a_1a_3)]$$
$$2[\lambda(a_3a_2a_1) + \lambda(a_2a_3a_1)] < 2[\lambda(a_1a_2a_3) + \lambda(a_1a_3a_2)]$$
$$[\lambda(a_3a_2a_1) + \lambda(a_2a_3a_1)] < [\lambda(a_1a_2a_3) + \lambda(a_1a_3a_2)]$$

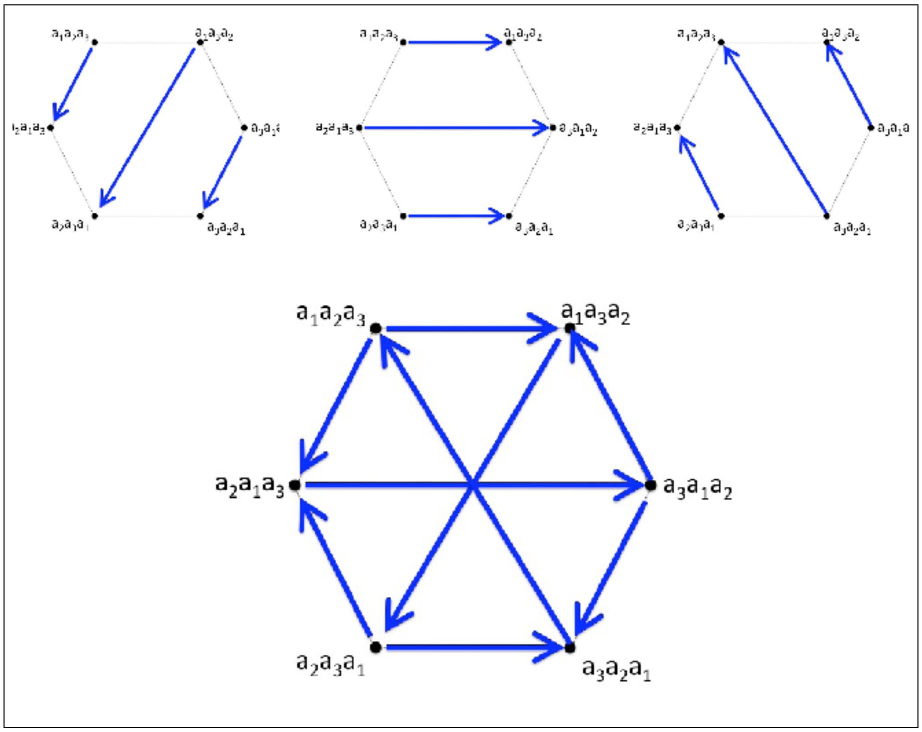

Figure 6: In the upper left-hand corner, the three blue arrows illustrate T^{λ,μ_K} relationships determined by $a_1\Gamma^{\lambda}a_2$. In the middle diagram of the top row, the three blue arrows illustrate T^{λ,μ_K} relationships determined by $a_2\Gamma^{\lambda}a_3$.In the upper right-hand corner, the three blue arrows illustrate T^{λ,μ_K} relationships determined by $a_3\Gamma^{\lambda}a_1$.The figure in the bottom row summarizes the nine T^{λ,μ_K} relationships determined by majority preferences over alternatives.

We can determine two pairs of T^{λ,μ^K} relationships from this inequality: $a_1a_2a_3T^{\lambda,\mu^K}a_2a_3a_1$ and $a_1a_3a_2T^{\lambda,\mu^K}a_3a_2a_1$. Similarly, the fact that $a_3\Gamma^\lambda a_1$ is supported by a smaller majority than $a_2\Gamma^\lambda a_3$ tells us that $[\lambda(a_3a_2a_1)+\lambda(a_3a_1a_2)] < [\lambda(a_1a_2a_3)+\lambda(a_2a_1a_3)]$, which in turn determines $a_1a_2a_3T^{\lambda,\mu^K}a_3a_1a_2$ and $a_2a_1a_3T^{\lambda,\mu^K}a_3a_2a_1$. See Figure 7 for diagrams of these relationships.

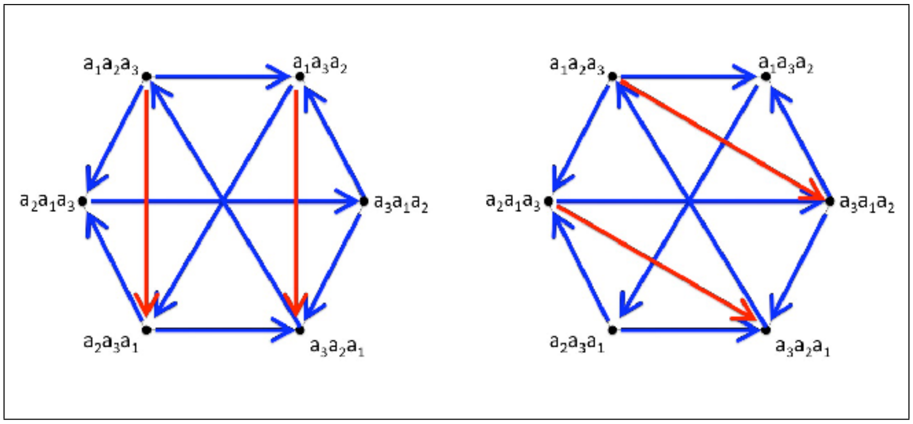

Figure 7: On the left, the red arrows are determined by the fact that $a_3\Gamma^\lambda a_1$ is supported by a smaller majority than $a_1\Gamma^\lambda a_2$.On the right, the red arrows are determined by the fact that $a_3\Gamma^\lambda a_1$ is supported by a smaller majority than $a_2\Gamma^\lambda a_3$.

We are left with the same two relationships to determine as we had in the transitive case; here, however, the uncovered set of the tournament over orderings will depend on the resolution of these links, as $a_1a_2a_3$ is no longer a Condorcet winner. Again we can turn to the relative strength of the majorities in favor of (a_1, a_2) and (a_2, a_3). Suppose that the strength of the majority for preference for $a_1 \succ a_2$ is greater than the strength of the majority preference for $a_2 \succ a_3$. Then we deduce that $[\lambda(a_1a_3a_2) + \lambda(a_3a_1a_2)] > [\lambda(a_2a_1a_3) + \lambda(a_2a_3a_1)]$, and we must have $a_1a_3a_2T^{\lambda,\mu^K}a_2a_1a_3$ and $a_3a_1a_2T^{\lambda,\mu^K}a_2a_3a_1$. If we instead have that the strength of the majority preference for $a_2 \succ a_3$ is greater than the strength of the majority preference for $a_1 \succ a_2$, then both of these T^{λ,μ^K} relationships would be reversed.

Therefore, for a given cyclic majority relation, we can have six possible tournaments. The differences stem from what pairs are supported by the weakest and strongest majorities: there are three possible weakest links, and each weakest link can be paired with two possible strongest links.

We can check these uncovered sets for Kemeny order consistency and see that it is always satisfied. The Kemeny winner of each tournament is the ordering which breaks the cycle at its weakest majority: in the top row, when the smallest majority is $a_2 \succ a_3$, $\pi^*(\lambda, \mu^K) = \{a_3a_1a_2\}$, in the middle row, when the smallest majority is $a_3 \succ a_1$, $\pi^*(\lambda, \mu^K) = \{a_1a_2a_3\}$, and in the last row, when the smallest majority is $a_1 \succ a_2$, $\pi^*(\lambda, \mu^K) = \{a_2a_3a_1\}$. Thus, in

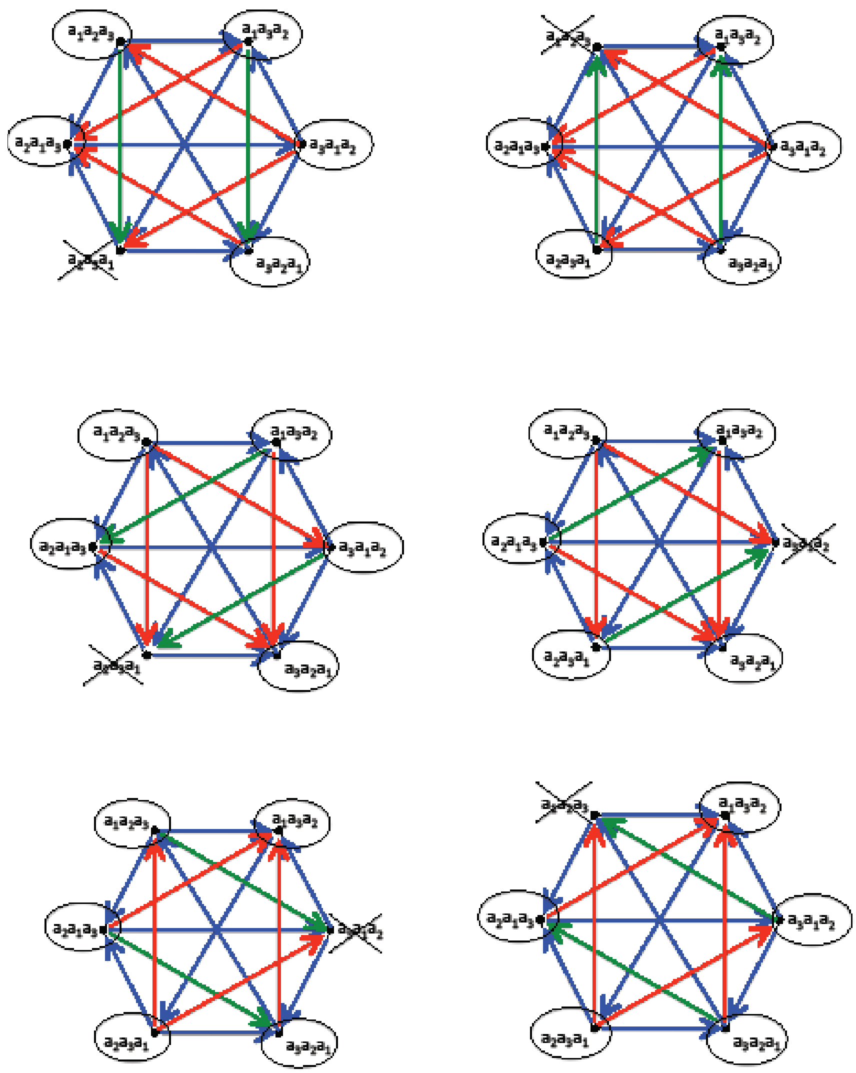

Figure 8: The above diagram illustrates the uncovered sets for the six possible tournaments in the cyclic case. A circe around the ordering indicates that is a member of the uncovered set; an X through the ordering indicates that it is covered. The top row features the two tournaments generated when $a_2\Gamma^\lambda a_3$ is supported by the smallest majority: on the left is the tournament generated when $a_1\Gamma^\lambda a_2$ is strongest, one the right is the tournament generated when $a_3\Gamma^\lambda a_1$ is strongest. In the second row, $a_3\Gamma^\lambda a_1$ is supported by the smallest majority ($a_1\Gamma^\lambda a_2$ strongest on the left, $a_2\Gamma^\lambda a_3$ on the right). In the third row, $a_1\Gamma^\lambda a_2$ is supported by the smallest majority ($a_2\Gamma^\lambda a_3$ strongest on the left, $a_3\Gamma^\lambda a_1$ on the right). While some of these tournaments have the same uncovered sets, no two of these tournaments are identical.

every case we have $\pi^*(\lambda, \mu^K) \in (T^{\lambda,\mu^K})$ and we see that Kemeny order consistency is always satisfied for any $n = 3$ population. But, strong Kemeny order consistency fails for cyclic tournaments, as the Kemeny ordering is not the sole member of the uncovered set for the populations above. In all of the cyclic cases, $|UC(T^{\lambda,\mu^K})| = 5$. Interestingly, note that the ordering covered in any particular case above agrees with the majority preference on two of the three pairs. The three orderings that are always in the uncovered set: $a_1a_3a_2, a_2a_1a_3$,and $a_3a_2a_1$ only agree with the majority preference on one of the three pairs. Thus, while the closest ordering to the majority preferences as defined by the Kemeny method is always a member of the uncovered set, agreement with the majority preferences is not in general a good predictor of membership in $UC(T^{\lambda,\mu^K})$.

We can summarize the results for the $n = 3$ case with the following proposition:

Proposition 6 *For $n = 3$, $UC(T^{\lambda,\mu^K})$ satisfies strong order consistency, Kemeny order consistency, and weak consistency.*

Therefore, for problems with only three alternatives, our results indicate that direct and representative democracies yield similar policy outcomes. In the case of transitive majority preferences in the population, our model predicts that a candidate selected via representative democracy will implement the same policy choice that would have been agreed upon under direct democracy for any pairwise choice problem. In the case of a cyclic majority preference, the policy predictions under direct democracy are less clear - the uncovered set predicts any of the possible policies could potentially be a winner. Correspondingly, the tournament over orderings identifies a larger set of electable candidates. Importantly, this set always includes the ordering closest to the majority preferences, the Kemeny winner. Thus, for the $n = 3$ case, the choice of mechanism does not impact predicted policy outcomes: both direct and representative democracy yield choices consistent with the majority preferences over alternatives.

6.3 Additional Results for *n=4*

Proposition 7 *For $n = 4$, $UC(T^{\lambda,\mu^K})$ satisfies order consistency.*

Proof: Assume majority preferences are consistent with e. We show that there is no ordering that can cover e. The key step is to recognize that we can apply 1 to rule out any ordering fewer than 5 transpositions from e: $eT^{\lambda,\mu^K}\pi$ for any π within two transpositions, so they cannot cover e, and for those three or four transpositions away, even if they beat e, they will be defeated by at least one ordering one or two transpositions from e (which e beats). So, the only orderings that could potentially cover e are five or six transpositions away from e: $\{a_4a_3a_1a_2, a_4a_2a_3a_1, a_3a_4a_2a_1, a_4a_3a_2a_1\}$. We rule these out one at a time:

– We cannot have $a_4a_3a_1a_2$ covers e, since $eT^{\lambda,\mu^K}a_2a_3a_1a_4T^{\lambda,\mu^K}a_4a_3a_1a_2$ for any population with majority preferences consistent with e.1 proves $eT^{\lambda,\mu^K}a_2a_3a_1a_4$.

We cannot have $a_2a_3a_1a_4T^{\lambda,\mu^K}a_4a_3a_1a_2$ since all of the orderings closer or equidistant to $a_4a_3a_1a_2$ than $a_2a_3a_1a_4$ have a_4 precedes a_2. Thus, if more than half the population were closer to or equidistant to $a_4a_3a_1a_2$, we would not have $a_2 \succ a_4$ in the majority preference.

– We cannot have $a_4a_2a_3a_1$ covers e, since we must have $eT^{\lambda,\mu^K}a_4a_2a_3a_1$. All of the orderings closer or equidistant to $a_4a_3a_1a_2$ than e have a_4 precedes a_1. Thus, if more than half the population were closer to or equidistant to $a_4a_2a_3a_1$ than to e, we would not have $a_1 \succ a_4$ in the majority preference.

– We cannot have $a_3a_4a_2a_1$ covers e, since $eT^{\lambda,\mu^K}a_1a_4a_2a_3T^{\lambda,\mu^K}$ $a_3a_4a_2a_1$ for any population with majority preferences consistent with e.1 proves $eT^{\lambda,\mu^K}a_1a_4a_2a_3$. We cannot have $a_3a_4a_2a_1T^{\lambda,\mu^K}a_1a_4a_2a_3$ since all of the orderings closer or equidistant to $a_1a_4a_2a_3$ than $a_3a_4a_2a_1$ have a_3 precedes a_1. Thus, if more than half the population were closer to or equidistant to $a_3a_4a_2a_1$, we would not have $a_1 \succ a_3$ in the majority preference.

– We cannot have $a_4a_3a_2a_1$ covers e, since $eT^{\lambda,\mu^K}a_1a_3a_2a_4T^{\lambda,\mu^K}a_4a_3a_2a_1$ for any population with majority preferences consistent with e.1 proves $eT^{\lambda,\mu^K}a_1a_3a_2a_4$. We cannot have $a_4a_3a_2a_1T^{\lambda,\mu^K}a_1a_3a_2a_4$ since all of the orderings closer or equidistant to $a_4a_3a_2a_1$ than $a_1a_3a_2a_4$ have a_4 precedes a_1. Thus, if more than half the population were closer to or equidistant to $a_4a_3a_2a_1$, we would not have $a_1 \succ a_4$ in the majority preference. ∎

6.4 Population Restrictions

Example 1 *Why Tightening the Sufficiency Condition by Restricting the Radius of the Outlier Orderings Does Not Work*

Consider the following population, a slight variant from the counterexample presented in Section 4:

π	$\lambda(\pi)$
e	$\frac{1}{2} - 2\varepsilon$
$a_1a_2a_4a_3a_6a_5a_8a_7a_{10}a_9$	$\frac{1}{5}(\frac{1}{2} - \varepsilon)$
$a_2a_1a_3a_4a_6a_5a_8a_7a_{10}a_9$	$\frac{1}{5}(\frac{1}{2} - \varepsilon)$
$a_2a_1a_4a_3a_5a_6a_8a_7a_{10}a_9$	$\frac{1}{5}(\frac{1}{2} - \varepsilon)$
$a_2a_1a_4a_3a_6a_5a_7a_8a_{10}a_9$	$\frac{1}{5}(\frac{1}{2} - \varepsilon)$
$a_2a_1a_4a_3a_6a_5a_8a_7a_9a_{10}$	$\frac{1}{5}(\frac{1}{2} - \varepsilon)$
$a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1$	3ε

Using the strategy from the proof above, we can show that $\hat{\pi} = a_2a_1a_4a_3a_6a_5a_8a_7a_{10}a_9$ covers e, the ordering consistent with majority preferences. We need to show that (a) $\hat{\pi}T^{\lambda,\mu^K}e$ and (b) $\forall\pi' \in \Pi$, $eT^{\lambda,\mu^K}\pi' \Rightarrow \hat{\pi}T^{\lambda,\mu^K}\pi'$. First we will show that $\hat{\pi}T^{\lambda,\mu^K}e$. For the five orderings in population with weight $\frac{1}{5}(\frac{1}{2} - \varepsilon)$, we have $f(\pi, \hat{\pi}) = 1$ and $f(\pi, e) = 4$. And, we know $a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1$ is closer to $\hat{\pi}$ than e, since it is maximally distant from e. Thus, $\frac{1}{2} + 2\varepsilon$ of the population is closer to $\hat{\pi}$ than e, so $\hat{\pi}T^{\lambda,\mu^K}e$.Now we need to show there cannot exist π' such that $eT^{\lambda,\mu^K}\pi'$ but $\pi' T^{\lambda,\mu^K}\hat{\pi}$. Suppose there did exist such a π'. Then, $eT^{\lambda,\mu^K}\pi'$ implies that for at least one of the orderings π in population other than e, $f(e,\pi) \leq f(\pi',\pi)$. We know there cannot exist a π', $\pi' \neq e$, such that $f(e, a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1) \leq f(\pi', a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1)$. So, it must be that this is true for one of the remaining five orderings. Since for any of these orderings $f(e,\pi) = 4$, we must have $f(\pi',\pi) \geq 4$ for at least one of those five orderings π. And, the fact that $\pi' T^{\lambda,\mu^K}\hat{\pi}$ implies that we have at least one of the following two cases:

1. For at least one of the orderings with weight $\frac{1}{5}(\frac{1}{2}-\varepsilon)$, we have $f(\pi',\pi) \leq f(\hat{\pi},\pi)$.
2. For both e and $a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1$, we have $f(\pi',\pi) < f(\hat{\pi},\pi)$.

For case 1, we know $f(\hat{\pi},\pi) = 1$, so this would imply, $f(\pi',\pi) \leq 1$ for one of the orderings with weight $\frac{1}{5}(\frac{1}{2}-\varepsilon)$. This leads to the same violation of the triangle inequality that we reached above, since for any two orderings with weight $\frac{1}{5}(\frac{1}{2}-\varepsilon)$, we have $f(\pi_i,\pi_j) \leq 2$. For case 2, $f(\pi',e) < f(\hat{\pi},e)$ implies $f(\pi',e) < 5$. And, $f(\pi', a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1) < f(\hat{\pi}, a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1)$ implies $f(\pi', a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1) < 40$. But, $f(\pi',e) < 5$ and $f(\pi', a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1) < 40$ cannot both hold, since the first implies π' has fewer than five errors and the second implies it has more than 5 errors. This leads to a contradiction. Thus, order consistency fails for this population.

This example illustrates the difficulty we encounter if we attempt to tighten the sufficiency condition for $\pi^* \in UC(T^{\lambda,\mu^K})$ by imposing a radius around the orderings with common errors. Taking the basic counterexample from above, where the minority orderings all lie relatively close to another, we can move some weight to $a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1$ and still arrive at $\hat{\pi}$ covers e.[6] Thus, it is not always true that we need the minority orderings to be relatively close to one another in order to have order consistency fail.

References

[1] Arrow, Kenneth J. 1951. *Social Choice and Individual Values.* Wiley, New York.

[2] Baldiga, Katherine A. and Jerry R. Green. Assent-Maximizing Social Choice. Forthcoming in Social Choice and Welfare.

[3] Besley, T. and S. Coate. 2008. Issue Unbundling via Citizens' Initiatives. Quartely Journal of Political Science 3(4), pp. 379-397.

[4] Black, Duncan. 1948. On the Rational of Group Decision Making. Journal of Political Economy, 56.

[6] In fact, we can move even more weight than we have here.

[5] Bossert, Walter and Tom Storcken. 1992. Strategy-proofness of social welfare functions: The use of the Kemeny distance between preference orderings. Social Choice and Welfare, 9, 345-360.

[6] Cox, Gary W. 1987. The Uncovered Set and the Core. American Journal of Political Science 31(2), 408-422.

[7] Frey, B. 1994. Direct Democracy: Politico-Economic Lessons from the Swiss Experience. American Economic Review, 84(2), 338-342.

[8] Kemeny, John G. 1959. Mathematics Without Numbers. Daedalus 88. 577-591.

[9] Kemeny, John and J. Laurie Snell. 1962. Mathematical Models in the Social Sciences, chapter 2, Cambridge: MIT Press.

[10] Koehler, David H. 1975. Vote Trading and the Voting Paradox: A Proof of Logical Equivalence. The American Political Science Review 69(3), 954-960.

[11] Laslier, J.F. 1997. *Tournaments Solutions and Majority Voting.* Berlin, New York: Springer-Verlag.

[12] McGarvey, David C. 1953. A Theorem on the Construction of Voting Paradoxes. Econometrica 21, 608-610.

[13] McKelvey, Richard D. 1986. Covering, Dominance, and Institution-Free Properties of Social Choice. American Journal of Political Science 30(2), 283-314.

[14] Miller, Nicolas R. 1980. A New Solution Set for Tournaments and Majority Voting: Further Graph-Theoretical Approaches to the Theory of Voting. American Journal of Political Science 24(1), 68-96.

[15] Riker, William H. and Steven J. Brams. 1973. The Paradox of Vote Trading. The American Political Science Review 67(4), 1235-1247.

[16] Saari, Donald G. 1995. *Basic Geometry of Voting.* Berlin: Springer-Verlag.

[17] Schwartz, Thomas. 1977. Collective Choice, Separation of Issues and Vote Trading. The American Political Science Review 71(3), 999-1010.

[18] Shepsle, Kenneth A. and Barry R. Weingast. 1984. Uncovered Sets and Sophisticated Voting Outcomes with Implications for Agenda Institutions. American Journal of Political Science 28(1), 49-74.

[19] Tullock, Gordon. 1970. A Simple Algebraic Logrolling Model. American Economic Review 60, 419-426.

[20] Young, H. Peyton. 1988. Condorcet's Theory of Voting. American Political Science Review. 82(2). 1231-1244.

[21] Young, H. Peyton and Arthur Levenglick. 1978. A Consistent Extension of Condorcet's Election Principle, SIAM Journal on Applied Mathematics, 35(2), 285-300.