On the Discriminative Power of Tournament Solutions

Felix Brandt and **Hans Georg Seedig**
Technische Universität München
{brandtf, seedigh}@in.tum.de

Abstract

Tournament solutions constitute an important class of social choice functions that only depend on the pairwise majority comparisons between alternatives. Recent analytical results have shown that several concepts with appealing axiomatic properties such as the Banks set or the minimal covering set tend to not discriminate at all when the tournaments are chosen from the uniform distribution. This is in sharp contrast to empirical studies which have found that real-world preference profiles often exhibit Condorcet winners, i.e., alternatives that all tournament solutions select as the unique winner. In this work, we aim to fill the gap between these extremes by examining the distribution of the number of alternatives returned by common tournament solutions for empirical data as well as data generated according to stochastic preference models such as impartial culture, impartial anonymous culture, Mallows mixtures, spatial models, and Pólya-Eggenberger urn models.

Introduction

A key problem in social choice theory is to identify functions that map the preference relations of multiple agents over some abstract set of alternatives to a socially acceptable alternative. Whenever the social choice function is required to be impartial towards alternatives and voters, it may be possible that several alternatives qualify equally well to be chosen. Depending on the rationalization of the social choice function, this might be a rare exception or a common phenomenon. Since alternatives are generally assumed to be mutually exclusive, it is typically understood that ties will eventually be broken by some procedure that is independent of the agents' preferences. This can for example be achieved by using a lottery (thus achieving *ex ante* fairness), letting a chairman (with known or unknown preferences) pick his most-preferred of the remaining alternatives, or simply assuming that a single alternative will be chosen according to a procedure that is completely unknown to the agents (see, e.g., Gärdenfors, 1979; Brandt and Brill, 2011; Brandt, 2011a). The uncertainty the agents face when it comes to the final selection process can be used as a powerful tool to satisfy certain formal criteria (such as impartiality, consistency, or strategyproofness) that would otherwise be impossible to attain. On the other hand, the uncertainty can also be viewed as a burden because a social choice function that leaves too much ambiguity may be unacceptable to society. In general, it seems desirable to narrow down the choice as much as possible based on the preferences of the voters alone. The goal of this paper is to study the discriminative power of various social choice functions—i.e., how many tied alternatives are returned—when preferences are drawn from common distributions that have been proposed in the literature.

An important class of social choice functions only depends on the pairwise majority relation between alternatives. When the pairwise majority relation is asymmetric, as is the case when there is an odd number of agents with linear preferences, these functions are known as *tournament solutions*. The tradeoff between discriminative power and axiomatic foundations is especially evident for tournament solutions as many of them can be axiomatically characterized as the *most discriminating* functions that satisfy certain desirable properties.[1] Tournament solutions are known to return rather large choice sets and are therefore particularly well-suited for an analysis of their discriminative power. We are considering all common tournament solutions in this paper: the top cycle TC, the uncovered set UC, the Banks set BA, the iterated uncovered set UC^∞, the minimal covering set MC, the tournament equilibrium set TEQ, the bipartisan set BP, the Slater set SL, the Copeland set CO, and the Markov set MA. All tournament solutions return a *Condorcet winner*—i.e., an alternative that is preferred to every other alternative by some majority of voters—whenever one exists. Moreover, all tournament solutions except CO, SL, and MA return a single alternative *if and only if* there is Condorcet winner.

The set-theoretic relationships between tournament solutions are well-studied (Laslier, 1997). For example, it is known that BA is contained in UC which in turn is contained in TC. Recently, Brandt et al. (2013c) gave a number of instructive examples showing that these inclusions are strict even when the number of alternatives is relatively small.

Analytical results about the discriminative power of tournament solutions for realistic distributions of preferences are very difficult to obtain. To the best of our knowledge, all existing papers explicitly or implicitly consider a uniform

[1] For example, TC is the most discriminating tournament solution satisfying expansion-consistency. Similar characterizations are known for UC, BA, MC, and BP (see, e.g., Brandt et al., 2013b, Chapter 6, Section 2.2.2)

distribution over all tournaments of a fixed size. Under this assumption, it was shown by Fey (2008) that BA almost always selects all alternatives as the number of alternatives goes to infinity. By the above-mentioned inclusion relationship this implies the same statement for UC and TC. Later, an analogous result was shown for MC by Scott and Fey (2012). More precise results for BP have been given by Fisher and Reeves (1995) who identified the whole distribution of $|BP|$ for any fixed number of alternatives m. They found that the probability that BP returns exactly k alternatives is $2^{-(m-1)}\binom{m}{k}$ if k is odd and zero otherwise. This directly implies that on average, BP returns half of the alternatives for odd $|T|$. In fact, for large tournaments, BP almost always chooses close to half of the alternatives (Scott and Fey, 2012).

These analytical results stand in sharp contrast to empirical observations that Condorcet winners exist in many real-world situations (see, e.g., Feld and Grofman, 1992), implying that tournament solutions very frequently return singletons. Simulations with stochastic preference models have been used for the analysis of several problems in (computational) social choice. For example, Laslier (2010) generated voting instances to derive estimates for the frequency of Condorcet winners and to compare the results of different voting rules such as plurality, Borda, approval voting, and Copeland's rule to each other. In his work, he has used a Rousseauist culture, capturing the idea of a pre-existing truth, as well as spatial and redistributive cultures. Earlier, McCabe-Dansted and Slinko (2006) have used computational experiments to obtain a hierarchical clustering of voting rules. To this end, they considered the number of times two voting rules coincide on a sample set as a measure for their similarity. They used the same setting as Shah (2003) with 5 alternatives and 85 voters and employed the Pólya-Eggenberger urn model by Berg (1985) to generate preferences. In comparison, we consider tournaments of larger sizes because several tournament solutions are known to always coincide when there are only few alternatives (Brandt et al., 2013c).

The remainder of this paper is structured as follows. First, we introduce terminology for preferences and tournament solutions and briefly define the tournament solutions considered in this paper. Next, we define and discuss the stochastic preference models employed in our analysis. In the section on experimental results, we describe our methodology, visually present the data obtained, and discuss some conclusions.

Methodology

Preference Profiles and Tournament Solutions

Let A be a set of alternatives with $|A| = m$ and $N = \{1, \dots, n\}$ a set of voters. The preferences of voter $i \in N$ are represented by a complete and antisymmetric *preference relation* $R_i \subseteq A \times A$. The set of all preference relations will be denoted by $\mathcal{R}$. The interpretation of $(a, b) \in R_i$, usually denoted by $a\ R_i\ b$, is that voter i values alternative a at least as much as alternative b. When individual preferences are transitive, we also speak of *rankings*. A *preference profile* $R = (R_1, \dots, R_n)$ is an n-tuple containing a preference relation R_i for each agent $i \in N$. The majority relation $\succ_R$ of a given preference profile is defined as

$$a \succ_R b \;\Leftrightarrow\; |\{i \mid a\ R_i\ b\}| > |\{i \mid b\ R_i\ a\}|.$$

A *tournament* T is a pair $(A, \succ)$, where $\succ$ is an asymmetric and complete (and thus irreflexive) binary relation on A. Whenever the number of voters n is odd, $(A, \succ_R)$ constitutes a tournament. If there is an alternative a such that $a \succ_R b$ for all $b \in A \setminus \{b\}$, a is a *Condorcet winner* according to R. A *tournament solution* is a function that maps a tournament to a nonempty subset of its alternatives, the *choice set*, and uniquely chooses a Condorcet winner whenever one exists. The simplest tournament solution is $COND$ which chooses the set of all alternatives whenever there is no Condorcet winner. The other tournament solutions considered in this paper are

- the *Copeland set* (CO), consisting of all alternatives with maximum outdegree,
- the *top cycle* (TC), consisting of the unique smallest set of alternatives such that every element of the set dominates all alternatives not in the set,
- the *uncovered set* (UC), consisting of all alternatives that reach all other alternatives in at most two steps,
- the *iterated uncovered sets* (UC^{∞}), the result of iteratively computing the uncovered set of the subtournament obtained from the restriction to the uncovered set,
- the *bipartisan set* (BP), consisting of all alternatives that are in the support of the unique Nash equilibrium of the corresponding tournament game,
- the *Markov set* (MA), consisting of the alternatives with maximum probability in the unique stationary distribution of the tournament when it is interpreted as a Markov chain,
- the *Banks set* (BA), consisting of all alternatives that are maximal in an inclusion-maximal transitive subtournament,
- the *Slater set* (SL), consisting of those alternatives who are maximal in a linear order, that shares as many edges as possible with the tournament,
- the *minimal covering set* (MC), consisting of the unique inclusion-minimal set of alternatives that is UC-stable, and
- the *tournament equilibrium set* (TEQ), consisting of the union of inclusion-minimal sets of alternatives, that are TEQ-retentive.

For definitions and discussions on most of these concept, we refer to the excellent overview by Laslier (1997). Stability and retentiveness which are properties of tournament solutions used to define MC and TEQ are discussed in detail by Brandt (2011b) and Brandt et al. (2013a), respectively. Computational issues concerning tournament solutions are discussed by Brandt (2009) and Hudry (2009).

The set-theoretic relationships of these concepts are depicted in Figure 1. BA and TEQ are the only tournament

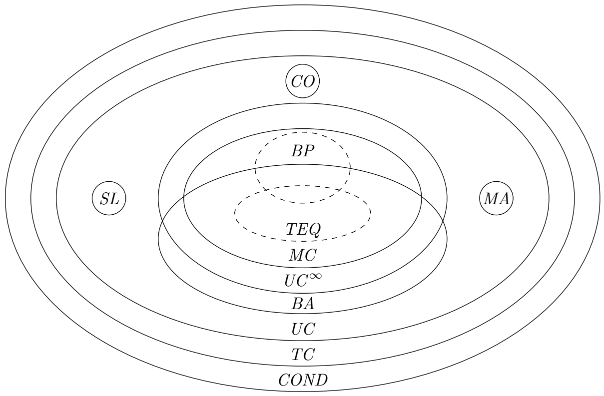

Figure 1: Set-theoretic relationships between tournament solutions. If the ellipses of two tournament solutions S and S' intersect, then $S(T) \cap S'(T) \neq \emptyset$ for all tournaments T. If the ellipses for S an S' are disjoint, however, this signifies that $S(T) \cap S'(T) = \emptyset$ for some tournament T. The exact location of BP in this diagram is unknown but it is contained in MC and intersects with TEQ in all known instances (Laslier, 1997). TEQ is contained in BA, but the inclusion in MC is uncertain.

solutions that are capable of discriminating in regular tournaments, i.e., tournaments in which all alternatives have the same degree.

It is important to realize that virtually all desirable properties of tournament solutions are properties of the choice set rather than its individual elements. For example, TC satisfies a certain notion of group-strategyproofness (Brandt, 2011a). Even though a single alternative will eventually be picked from TC, CO—which is always contained in TC—does *not* satisfy group-strategyproofness.

Empirical Data

In the preference library PrefLib (Mattei and Walsh, 2013), scholars have contributed data sets from real world scenarios ranging from preferences over movies or sushi via Formula 1 championship results to real election data. At the time of writing, PrefLib contained 354 tournaments induced from pairwise majority comparisons. Out of these, all except 9 exhibit a Condorcet winner. The remaining tournaments are still very structured as the uncovered set never contains more than 4 alternatives (even in the largest of the remaining tournaments with 242 alternatives). This is in line with earlier observations that real-world majority relations tend to be close to linear orders and often have Condorcet winners (Regenwetter et al., 2006).

Stochastic Models

As the available empirical data does not allow to draw conclusions about the differences in discriminative power of tournament solutions, we now consider stochastic models to generate tournaments of a given size m (usually by sampling preference profiles and considering the induced majority tournament).[2] In this section, we will briefly introduce the considered models and describe how to sample from them whenever this is not obvious.

The *uniform random tournament* model was used in the previous analysis of the discriminative power of tournament solutions (Fisher and Reeves, 1995; Fey, 2008; Scott and Fey, 2012). It assigns the same probability to each *labeled* tournament of size m, i.e.,

$$\Pr(T) = \frac{1}{2^{\binom{m}{2}}} \text{ for each } T \text{ with } |T| = m.$$

In all of the remaining models, we sample preference profiles and work with the tournament induced by the majority relation. The term *culture* has been coined for probabilistic preference models where the draws for each voter are independent from each other. Cultures are defined by the probabilities they put on each possible preference ranking. The most widely-studied model of this kind is the *impartial culture model* (IC), where every possible ranking of the alternatives has the same probability of $\frac{1}{m!}$. IC is a member of the family of *dual cultures*, defined by the property that each ranking has the same probability as its inverse. Dual cultures have been criticized for being too unrealistic (Regenwetter et al., 2006). Nevertheless, they are relevant for their susceptibility to analytical methods that helped to improve the understanding of voting phenomena (see, e.g., (DeMeyer and Plott, 1970)). If we add anonymity by having indistinguishable voters, the set of profiles is partitioned into equivalence classes. In the *impartial anonymous culture* (IAC), each of these equivalence classes is chosen with equal probability. Note that this is not a culture in the sense mentioned above.

A very different kind of model is the *spatial model*. Here, alternatives and voters are uniformly at random placed in a multi-dimensional space and the voters' preferences are determined by the (Euclidian) distanced to the alternatives. The spatial model has played an important role in political and social choice theory where the dimensions are interpreted as different aspects or properties of the alternatives (see, e.g., Ordeshook, 1993; Austen-Smith and Banks, 2000).

There are several models who assume a pre-existing truth in the form of reference rankings such that each agent reports a noisy estimate of said truth as his preferences. For these models, Laslier has introduced the term *Rousseauist cultures* (Laslier, 2010). Such models are usually parameterized by a homogeneity parameter that scales the noisiness of individual perceptions. In its arguably simplest form, every agent provides possibly intransitive preferences R where each pairwise preference $a\ R\ b$ is 'correct', i.e., coincides with the reference ranking R_0 with a probability p where $0.5 \leq p \leq 1$. We will call this the *Condorcet noise* model.

[2]This guided our choice of models for this paper. There are several stochastic models on rankings such as Thurstonian models, other Babington Smith models, or multi-stage ranking models that we do not consider in this work (Critchlow et al., 1991; Marden, 1995). These models are often very versatile but depend strongly on the choice of a rather large number of parameters. Also, sampling from a general Babington smith model is a very tedious task.

This is the only model we consider in which individual preferences can be intransitive. An interesting aspect of this model is that all pairwise majority comparisons are independent of each other and can be computed by

$$\Pr(a \succ_R b \mid a\ R_0\ b) = \sum_{v=\frac{n}{2}+1}^{n} \binom{n}{v} p^v (1-p)^{n-v}.$$

For $p = 0.5$, the Condorcet noise model with any number of voters coincides with the model of uniform random tournaments.

In *Mallows-ϕ model* (Mallows, 1957), the distance to a reference ranking is measured by means of the Kendall-tau distance which counts the number of pairwise disagreements. Let R_0 be the reference ranking. Then, the Kendell-tau distance of a preference ranking R to R_0 is $\tau(R, R_0) = \binom{m}{2} - |R \cap R_0|$. According to the model, this induces the probability of a voter having R as his preferences to be $\Pr(R) = \phi^{\tau(R,R_0)}/C$ where C is a normalization constant and $\phi \in (0, 1]$ is a dispersion parameter. Small values for ϕ put most of the probability on rankings very close to R_0 whereas for $\phi = 1$ the model coincides with IC.

Obviously, one can define a number of such distance-based models. Besides the Kendell-tau distance, Spearman-rho distance has been considered (resulting in Mallows-θ model), as well as the distance measures named after Cayley, Hammond, and Ulam. See (Critchlow et al., 1991) for a discussion.

A property that makes distance-based models less appealing for this particular study is their bias towards transitive majority relations which makes the issue of choosing trivial. In fact, Mallows-ϕ even satisfies *strong unimodality* as defined in Critchlow et al. (1991) since a single preference ranking has maximum probability and ranking probabilities are non-increasing as we move along a path of rankings, where in each step two adjacent alternatives are swapped causing an increase in the Kendell-tau distance to the modal ranking.

To overcome this unimodality of the preference distribution, *mixtures* of models have been considered. A mixture model consists of several ordinary models with a probability distribution over them. While this idea could theoretically be applied to any set of models that may just differ in their parameterization or even belong to different model families, it has been considered the most with respect to the Mallows-ϕ model. For simplicity and to reduce the number of free parameters, we consider uniform mixtures over k Mallows-ϕ with a shared parameter ϕ and refer to this as *Mallows k-mixtures*. Sampling from Mallows-ϕ (or Mallows mixtures) is conveniently possible by a repeated insertion model (Doignon et al., 2004; Lu and Boutilier, 2011).

In the Pólya-Eggenberger *urn model*, each possible preference ranking is thought to be represented by a ball in an urn from which individual preferences are drawn. After each draw, the chosen ball is put back and $\alpha \in \mathbb{N}_0$ new balls of the same kind are added to the urn (Berg, 1985). This models the effect of an interdependence of multiple voters' preferences as the next voter chooses from a modified distribution. The urn model subsumes both IC ($\alpha = 0$) and IAC ($\alpha = 1$).

Experimental Results and Discussion

In our experimental setup, we generated tournament instances according to the aforementioned models and computed the different choice sets for them. For the sampling step, we built on the implementations of Mallows-ϕ and an urn model from Mattei and Walsh (2013) to generate preference profiles of which we considered the majority relation. The computation of the various tournament solutions was done via counting (CO), matrix multiplication (UC, UC^∞), depth-first-search (TC), linear programming (BP, MC), eigenvalue decomposition (MA), branch-and-bound (SL), or tailored algorithms (BA, TEQ). Computing BA, TEQ, and SL is NP-hard wheres the remaining tournament solutions can be computed efficiently.

First, we considered the frequency of majority relations with a Condorcet winner of each stochastic model. In case of the Condorcet noise model, they can be computed directly: let $p_M = \Pr(a \succ_R b \mid a\ R_0\ b)$. Then,

$$\Pr(\succ_R \text{ has Cond. winner}) = \sum_{i=1}^{n} {p_M}^{\,n-i-1} (1-p_M)^{\,i}.$$

For the other models, we resorted to computational experiments. Some of these results are shown in Figure 2. The distributions of IAC, Mallows-ϕ, and urn models with a fixed α are not displayed but are very similar to IC. For all models, we see that increasing the number of alternatives generally makes Condorcet winners less likely. Note, however, that in case of the Mallows-4-mixture with $\phi = 0.9$, there is a non-monotonic region for $n \geq 25$ voters where Condorcet winners become more frequent again when the number of alternatives is sufficiently large. A similar phenomenon is visible in Figure 3 for the Mallows-4-mixture with $\phi = 0.95$. We currently have no conclusive explanation of this effect.

For our choice of parameters, we briefly mention the effects of the other parameters. In all Rousseauist cultures, increasing the number of voters increases the likeliness of Condorcet winners. The same holds for parameter changes that increase homogeneity such as increasing p in the Condorcet noise model, decreasing ϕ in Mallows-ϕ model, or increasing α in the urn model. For the spatial model, we have not found the dimension to have a large impact on the results as long as it is at least 2.

Secondly, we examined the ability of the various solutions to rule out alternatives. Our informal measure for discriminative power of a tournament solution on a specific model is the distance of its average choice set size to the average size of $COND$ which, by definition, is the least discriminative tournament solution. In our comparisons, we provide $COND$ not only as a baseline but also as an indicator for the frequency of tournaments with a Condorcet winner. We examined the average choice set sizes of the aforementioned tournament solutions for a fixed number of voters $n = 51$. The results are shown in Figure 3.

Due to the large number of Condorcet winners in these samples, the standard deviations of the measured choice set sizes are usually about as large as the reported values themselves. In cases of very high or very low average choice set sizes as in the spatial or in the uniform random tournament

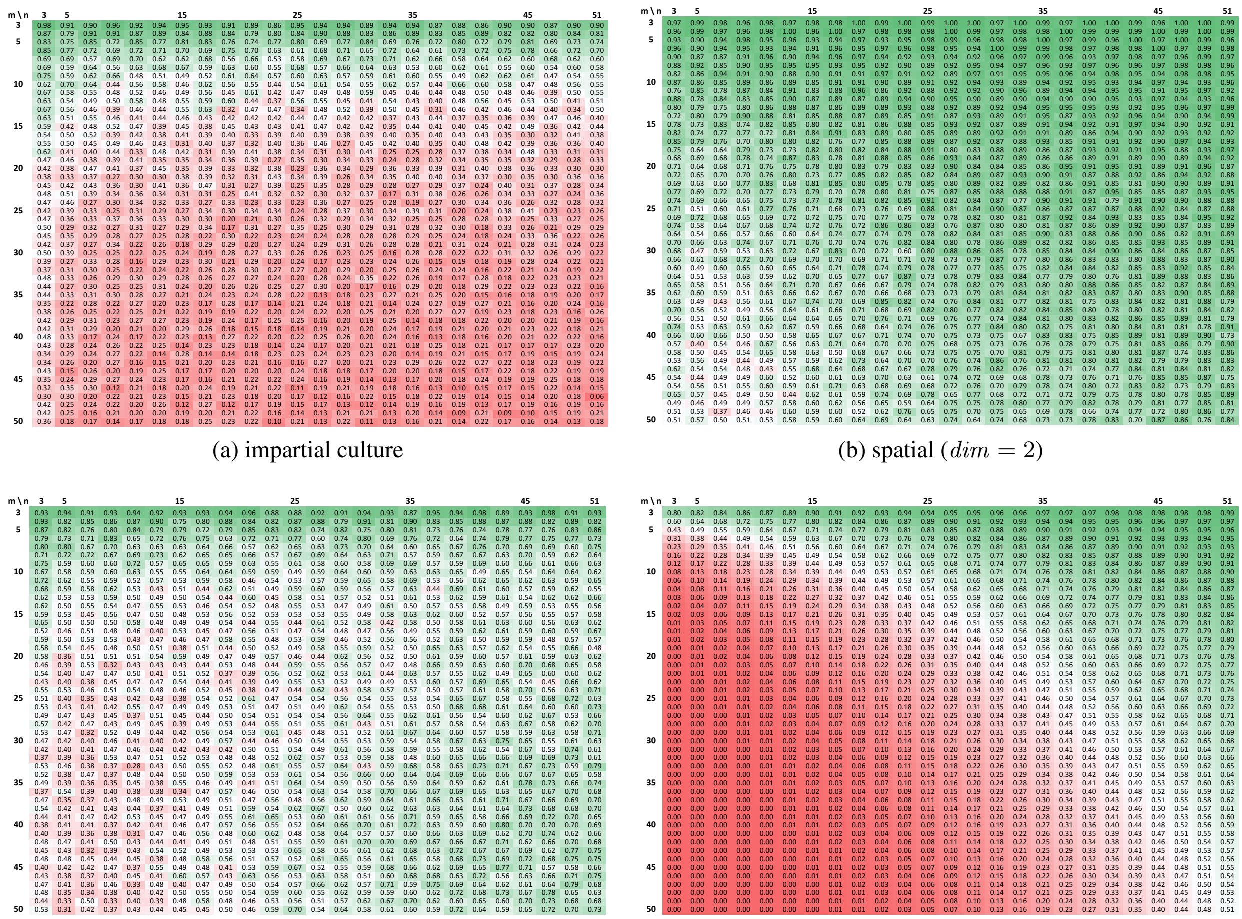

(a) impartial culture

(b) spatial ($dim = 2$)

(c) Mallows 4-mixture ($\phi = 0.9$)

(d) Condorcet noise ($p = 0.65$)

Figure 2: Frequencies of tournaments with a Condorcet winner as a heat map for four stochastic models. The number of voters ranges from 3 to 51 on the horizontal axis, the number of alternatives from 3 to 50 on the vertical axis. Displayed are the proportions of tournaments with a Condorcet winner. The color green corresponds to a high frequency of such tournaments, the color red indicates that few tournaments sampled from this model exhibit a Condorcet winner. For the first three models, the values are taken over 100 samples, whereas the probabilities for the Condorcet noise model were computed directly.

model, the standard deviation is, of course, low. A notable exception from this behavior is BP in case of the uniform random tournament model and the similar Condorcet noise model with $p = 0.55$. There, BP on average chooses less than half of the alternatives with low standard deviation.

The following conclusions can be drawn from our results.

- TC is almost as undiscriminating as $COND$.
- All other tournament solutions are much more discriminating than the analytical results for uniform random tournaments suggest. In fact, for all reasonable parameterizations of the considered models with transitive individual preferences and at least 10 alternatives (including impartial culture) all tournament solutions except TC discarded at least 75% of the alternatives on average.
- All tournament solutions except TC behave similarly in terms of discriminative power. One may conclude that the decision which one to use in practical applications should not be based on discriminative power, but rather on axiomatic properties.
- Using a more fine-grained analysis, tournament solutions can be divided into five clusters based on their discriminative power. The first cluster merely consists of TC. The second cluster contains UC and BA. UC^∞, MC, and TEQ are contained in the third cluster. BP forms a cluster of its own. Finally, tournament solutions based on scoring (SL, CO, and MA) are much more discriminating than all other tournament solutions and form the fifth cluster. Out of these, MA stands out as the most selective one. It is almost always unique.
- UC^∞ (and thereby also MC) discriminates more than BA. This observation could not be deduced from the set-theoretic relationships between tournament solutions.
- BP is not only remarkably discriminating in uniform random tournaments (which already follows from the analytical results), but even more discriminating in the Condorcet noise model with $p = 0.55$. Within the group of tournament solutions with appealing characterizations, it discriminates the most (and is efficiently computable).

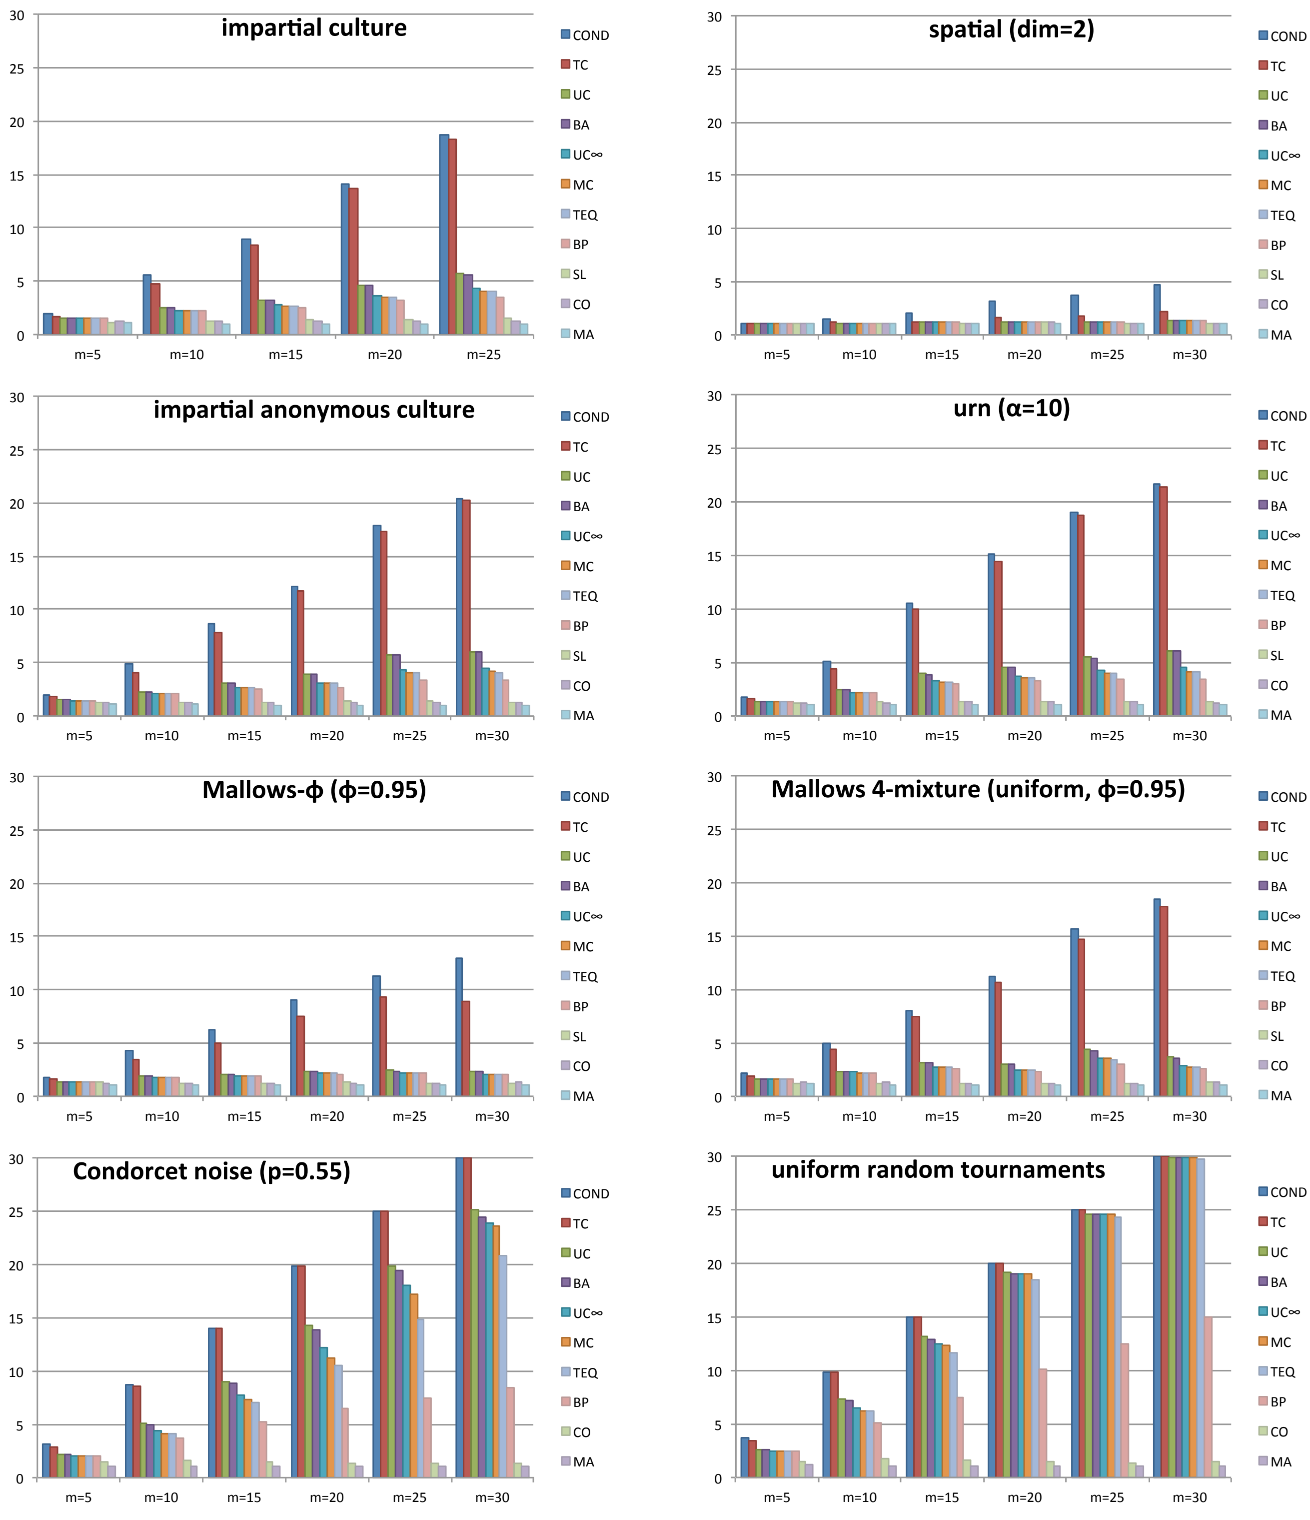

Figure 3: Comparison of average absolute choice set sizes for various stochastic preference models. The number of alternatives is on the horizontal axis, the number of voters is $n = 51$. Averages are taken over 100 runs. The Slater set (SL) is omitted whenever its computation was infeasible.

References

D. Austen-Smith and J. S. Banks. *Positive Political Theory I: Collective Preference*. University of Michigan Press, 2000.

S. Berg. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice*, 47:377–387, 1985.

F. Brandt. *Tournament Solutions – Extensions of Maximality and Their Applications to Decision-Making*. Habilitation Thesis, Faculty for Mathematics, Computer Science, and Statistics, University of Munich, 2009.

F. Brandt. Group-strategyproof irresolute social choice functions. In T. Walsh, editor, *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 79–84. AAAI Press, 2011a.

F. Brandt. Minimal stable sets in tournaments. *Journal of Economic Theory*, 146(4):1481–1499, 2011b.

F. Brandt and M. Brill. Necessary and sufficient conditions for the strategyproofness of irresolute social choice functions. In K. Apt, editor, *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 136–142. ACM Press, 2011.

F. Brandt, M. Brill, F. Fischer, and P. Harrenstein. Minimal retentive sets in tournaments. *Social Choice and Welfare*, 2013a. Forthcoming. Also appeared at AAMAS-2010.

F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*, chapter 6, pages 213–283. MIT Press, 2nd edition, 2013b.

F. Brandt, A. Dau, and H. G. Seedig. Bounds on the disparity and separation of tournament solutions. 2013c. Working paper.

D. E. Critchlow, M. A. Fligner, and J. S. Verducci. Probability models on rankings. *Journal of Mathematical Psychology*, 35:294–318, 1991.

F. DeMeyer and C. R. Plott. The probability of a cyclical majority. *Econometrica*, 38(2):345–354, 1970.

J.-P. Doignon, A. Pekeč, and M. Regenwetter. The repeated insertions model for rankings: missing link between two subset choice models. *Psychometrika*, 69(1):33–54, 2004.

S. L. Feld and B. Grofman. Who's afraid of the big bad cycle? Evidence from 36 elections. *Journal of Theoretical Politics*, 4:231–237, 1992.

M. Fey. Choosing from a large tournament. *Social Choice and Welfare*, 31(2):301–309, 2008.

D. C. Fisher and R. B. Reeves. Optimal strategies for random tournament games. *Linear Algebra and its Applications*, 217:83–85, 1995.

P. Gärdenfors. On definitions of manipulation of social choice functions. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.

O. Hudry. A survey on the complexity of tournament solutions. *Mathematical Social Sciences*, 57(3):292–303, 2009.

J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.

J.-F. Laslier. In silico voting experiments. In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, chapter 13, pages 311–335. Springer-Verlag, 2010.

T. Lu and C. Boutilier. Learning Mallows models with pairwise preferences. In L. Getoor and T. Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pages 145–152, 2011.

C. L. Mallows. Non-null ranking models. *Biometrika*, 44 (1/2):114–130, 1957.

J. I. Marden. *Analyzing and Modeling Rank Data*. Number 64 in Monographs on Statistics and Applied Probability. Chapman & Hall, 1995.

N. Mattei and T. Walsh. PrefLib: A library for preference data http://www.preflib.org. In P. Perny, M. Pirlot, and A. Tsoukiàs, editors, *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, volume 8176 of *Lecture Notes in Computer Science (LNCS)*, pages 259–270. Springer, 2013.

J. C. McCabe-Dansted and A. Slinko. Exploratory analysis of similarities between social choice rules. *Group Decision and Negotiation*, 15(1):77–107, 2006.

P. Ordeshook. The spatial analysis of elections and committees: four decades of research. Mimeo, 1993.

M. Regenwetter, B. Grofman, A. A. J. Marley, and I. M. Testlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, 2006.

A. Scott and M. Fey. The minimal covering set in large tournaments. *Social Choice and Welfare*, 38(1):1–9, 2012.

R. Shah. Statistical mappings of social choice rules. Master's thesis, Stanford University, 2003.