3

Tournament Solutions

Felix Brandt[a], Markus Brill[b], and Paul Harrenstein[c]

3.1 Introduction

Perhaps one of the most natural ways to aggregate binary preferences from individual agents to a group of agents is *simple majority rule*, which prescribes that one alternative is socially preferred to another whenever a majority of agents prefers the former to the latter. Majority rule intuitively appeals to democratic principles, is easy to understand and—most importantly—satisfies some attractive formal properties. As seen in Chapter 2 (Zwicker, 2016), May's Theorem shows that a number of rather weak and intuitively acceptable principles completely characterize majority rule in settings with two alternatives (May, 1952). Moreover, almost all common voting rules satisfy May's axioms and thus coincide with majority rule in the two-alternative case. It would therefore seem that the existence of a majority of individuals preferring alternative a to alternative b signifies something fundamental and generic about the group's preferences over a and b. We will say that alternative a *dominates* alternative b in such a case.

As is well known from Condorcet's paradox (see Chapter 2 (Zwicker, 2016)), the dominance relation may contain cycles. This implies that the dominance relation may not admit a maximal element and the concept of maximality as such is rendered untenable. On the other hand, Arrow writes that "one of the consequences of the assumptions of rational choice is that the choice in any environment can be determined by a knowledge of the choices in two-element environments" (Arrow, 1951, p. 16). Thus, one way to get around this problem—the one pursued in this chapter—is to take the dominance relation as given and define alternative concepts to take over the role of maximality. More precisely, we will be concerned with social choice functions (SCFs) that are based on the dominance relation only, i.e., those SCFs that Fishburn (1977) called $C1$ functions. Topics to be covered in this chapter include McGarvey's Theorem, various tournament solutions (such as Copeland's rule, the uncovered set, the top cycle, or the tournament equilibrium set), strat-

[a] Institut für Informatik, Technische Universität München, Germany
[b] Department of Computer Science, Duke University, USA
[c] Department of Computer Science, University of Oxford, UK

To appear in: *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (eds.), Cambridge University Press, 2016

egyproofness, implementation via binary agendas, and extensions of tournament solutions to weak tournaments. Particular attention will be paid to the issue of whether and how tournament solutions can be computed efficiently.

In this chapter, we will view tournament solutions as $C1$ SCFs. However, for varying interpretations of the dominance relation, tournament solutions and variants thereof can be applied to numerous other settings such as multi-criteria decision analysis (Arrow and Raynaud, 1986; Bouyssou et al., 2006), zero-sum games (Fisher and Ryan, 1995; Laffond et al., 1993a; Duggan and Le Breton, 1996), and coalitional games (Brandt and Harrenstein, 2010).

3.2 Preliminaries

We first introduce and review some basic concepts and notations used in this chapter. Let $N = \{1, \ldots, n\}$ be a set of voters, A a set of m alternatives, and $R = (\succsim_1, \ldots, \succsim_n)$ a vector of linear orders over A. $\succsim_i$ is the *preference relation* of voter i and R is called a *preference profile*. The *majority relation* $\succsim$ *for* R is defined such that for all alternatives a and b,

$$a \succsim b \quad \text{if and only if} \quad |\{i \in N : a \succsim_i b\}| \geq |\{i \in N : b \succsim_i a\}|.$$

See Figure 3.1 for an example preference profile and the corresponding majority relation. A *Condorcet winner* is a (unique) alternative a such that there is no other alternative b with $b \succsim a$ (or in other words, an alternative a such that $a \succ b$ for all $b \in A \setminus \{a\}$, where $\succ$ is the asymmetric part of $\succsim$). By definition, the majority relation is complete, i.e., $a \succsim b$ or $b \succsim a$ for all alternatives a and b. Apart from completeness, the majority relation has no further structural properties, i.e., every complete relation over a set of alternatives can be obtained as the majority relation for some preference profile. This result is known as *McGarvey's Theorem*.

Theorem 3.1 (McGarvey, 1953) *Let A be a set of m alternatives and $\geq$ a complete relation over A. Then, there is a preference profile $R = (\succsim_1, \ldots, \succsim_n)$ over A with $n \leq m(m-1)$ such that $\geq = \succsim$.*

Proof Denote the asymmetric part of $\geq$ by $>$. For every pair (a, b) of alternatives with $a > b$, introduce two voters, i_{ab} and j_{ab}, i.e., $N = \{i_{ab}, j_{ab} : a > b\}$. Define the preference profile R such that for all $a, b \in A$,

$$a \succ_{i_{ab}} b \succ_{i_{ab}} x_1 \succ_{i_{ab}} \cdots \succ_{i_{ab}} x_{m-2} \quad \text{and}$$
$$x_{m-2} \succ_{j_{ab}} \cdots \succ_{j_{ab}} x_1 \succ_{j_{ab}} a \succ_{j_{ab}} b,$$

where $x_1, \ldots, x_{m-2}$ is an arbitrary enumeration of $A \setminus \{a, b\}$. It is easy to check that the majority relation $\succsim$ for R coincides with $\geq$. By asymmetry of $>$, moreover, we have $a > b$ for at most $\frac{1}{2}m(m-1)$ pairs (a, b) and thus $n = |N| \leq m(m-1)$. $\square$

The minimal number of voters required to obtain any majority relation has subsequently been improved by Stearns (1959) and Erdős and Moser (1964), who have eventually shown that this number is of order $\Theta(\frac{m}{\log m})$. This implies that for any fixed number of voters, there are tournaments which are not induced by any preference profile. Only little is known about the classes of majority relations that can be induced by preference profiles with small fixed numbers of voters (see Bachmeier et al., 2014).

3.2.1 Tournaments

If the number of voters is odd, there can be no majority ties and the majority relation is antisymmetric. In this case, the asymmetric part $\succ$ of the majority relation $\succsim$ is connex and irreflexive and will be referred to as the *dominance relation*.[1] A dominance relation can be conveniently represented by an oriented complete graph, a tournament (see Figure 3.1).

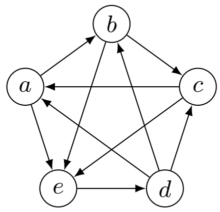

$$\begin{array}{lllllllll} 1: & a & \succ_1 & b & \succ_1 & c & \succ_1 & e & \succ_1 \; d \\ 2: & d & \succ_2 & c & \succ_2 & a & \succ_2 & b & \succ_2 \; e \\ 3: & e & \succ_3 & d & \succ_3 & b & \succ_3 & c & \succ_3 \; a \end{array}$$

Figure 3.1 A tournament $T = (\{a, b, c, d, e\}, \succ)$, which depicts the asymmetric part of the majority relation of the 3-voter preference profile on the right.

Formally, a *tournament* T is a pair $(A, \succ)$ where A is a set of vertices and $\succ$ is an asymmetric and connex relation over the vertices. Tournaments have a rich mathematical theory and many results for $C1$ SCFs have a particularly nice form if the dominance relation constitutes a tournament. Moreover, many $C1$ functions have only been defined for tournaments and possess a variety of possible generalizations to majority graphs that are not tournaments. None of these generalizations can be seen as *the* unequivocal extension of the original function. We therefore assume the dominance relation to be antisymmetric and discuss generalizations of functions in Section 3.5.[2]

The dominance relation can be raised to sets of alternatives and we write $A \succ B$ to signify that $a \succ b$ for all $a \in A$ and all $b \in B$. Using this notation, a Condorcet

[1] A relation $\succ$ is *connex* if $a \succ b$ or $b \succ a$ for all *distinct* alternatives a and b. In the absence of majority ties, $\succ$ and $\succsim$ are identical except that $\succsim$ is reflexive while $\succ$ is not.

[2] The preference profile constructed in the proof of Theorem 3.1 involves an even number of voters. It is easily seen, however, that no single additional voter, no matter what his preferences are, will affect the dominance relation $\succ$ and we may assume that every tournament is also induced by a preference profile with an *odd* number of voters. Likewise, the result by Erdős and Moser (1964) also holds for tournaments (Moon, 1968, Ch. 19, Ex. 1 (d)).

winner can be defined as an alternative a such that $\{a\} \succ A \setminus \{a\}$. For a subset of alternatives $B \subseteq A$, we will sometimes consider the restriction $\succ_B = \{(a,b) \in B \times B : a \succ b\}$ of the dominance relation $\succ$ to B. $(B, \succ_B)$ is then called a *subtournament* of $(A, \succ)$.

For a tournament $(A, \succ)$ and an alternative $a \in A$, we denote by $D(a)$ the *dominion* of a, i.e.,

$$D(a) = \{\, b \in A : a \succ b \,\},$$

and by $\overline{D}(a)$ the *dominators* of a, i.e.,

$$\overline{D}(a) = \{\, b \in A : b \succ a \,\}.$$

The *order* $|T|$ of a tournament $T = (A, \succ)$ refers to the cardinality of A.

The elements of the *adjacency matrix* $M(T) = (m_{ab})_{a,b \in A}$ of a tournament T are 1 whenever $a \succ b$ and 0 otherwise. The *skew-adjacency matrix* $G(T)$ of the corresponding tournament graph is skew-symmetric and defined as the difference of the adjacency matrix and its transpose, i.e., $G(T) = M(T) - M(T)^t$ (see Figure 3.2).

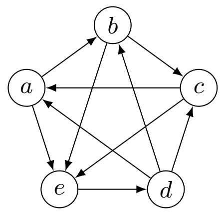

T

$$\begin{array}{c} \\ a \\ b \\ c \\ d \\ e \end{array} \begin{array}{c} \begin{array}{ccccc} a & b & c & d & e \end{array} \\ \begin{pmatrix} 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix} \end{array}$$

$M(T)$

$$\begin{array}{c} \\ a \\ b \\ c \\ d \\ e \end{array} \begin{array}{c} \begin{array}{ccccc} a & b & c & d & e \end{array} \\ \begin{pmatrix} 0 & 1 & -1 & -1 & 1 \\ -1 & 0 & 1 & -1 & 1 \\ 1 & -1 & 0 & -1 & 1 \\ 1 & 1 & 1 & 0 & -1 \\ -1 & -1 & -1 & 1 & 0 \end{pmatrix} \end{array}$$

$G(T)$

Figure 3.2 The tournament T from Figure 3.1 with its adjacency matrix $M(T)$ and its skew-adjacency matrix $G(T)$. Here, for instance, $D(a) = \{b, e\}$ and $\overline{D}(b) = \{a, d\}$.

An important structural notion in the context of tournaments is that of a component. A *component* is a nonempty subset of alternatives $B \subseteq A$ that bear the same relationship to any alternative not in the set, i.e., for all $a \in A \setminus B$, either $B \succ \{a\}$ or $\{a\} \succ B$. A *decomposition* of T is a partition of A into components.

For a given tournament $\tilde{T}$, a new tournament T can be constructed by replacing each alternative with a component. Let $B_1, \dots, B_k$ be pairwise disjoint sets of alternatives and consider tournaments $T_1 = (B_1, \succ_1), \dots, T_k = (B_k, \succ_k)$, and $\tilde{T} = (\{1, \dots, k\}, \tilde{\succ})$. The *product* of $T_1, \dots, T_k$ with respect to $\tilde{T}$, denoted by $\Pi(\tilde{T}, T_1, \dots, T_k)$, is the tournament $(A, \succ)$ such that $A = \bigcup_{i=1}^{k} B_i$ and for all $b_1 \in B_i, b_2 \in B_j$,

$$b_1 \succ b_2 \quad \text{if and only if} \quad i = j \text{ and } b_1 \succ_i b_2, \text{ or } i \neq j \text{ and } i \mathrel{\tilde{\succ}} j.$$

Here, $\tilde{T}$ is called the *summary* of T with respect to the above decomposition. In the tournament depicted in Figure 3.2, for example, $\{a, b, c\}$, $\{d\}$, and $\{e\}$ are components and $\{\{a, b, c\}, \{d\}, \{e\}\}$ is a decomposition. The tournament can therefore be seen as the product of a 3-cycle and two singleton tournaments with respect to a 3-cycle summary. Importantly, every tournament admits a unique decomposition that is minimal in a well-defined sense (Laslier, 1997, pp. 15–23).

3.2.2 Tournament Solutions

A *tournament solution* is a function S that maps each tournament $T = (A, \succ)$ to a nonempty subset $S(T)$ of its alternatives A called the *choice set*. The formal definition further requires that a tournament solution does not distinguish between isomorphic tournaments, i.e., if $h\colon A \to A'$ is an isomorphism between two tournaments $(A, \succ)$ and $(A', \succ')$, then

$$S(A', \succ') = \{h(a) : a \in S(A, \succ)\}.$$

As defined in Chapter 2 (Zwicker, 2016), an SCF is a $C1$ function if its output only depends on the dominance relation. Since the dominance relation is invariant under renaming voters, $C1$ SCFs are anonymous by definition. Moreover, due to the invariance of tournament solutions under isomorphisms, tournament solutions are equivalent to *neutral $C1$ functions*. In contrast to Laslier (1997), we do not require tournament solutions to be *Condorcet-consistent*, i.e., to uniquely select a Condorcet winner whenever one exists.

For a tournament $T = (A, \succ)$ and a subset $B \subseteq A$, we write $S(B)$ for the more cumbersome $S(B, \succ_B)$. For two tournament solutions S and S', we write $S' \subseteq S$, and say that S' is a *refinement* of S and S a *coarsening* of S', if $S'(T) \subseteq S(T)$ for all tournaments T.

The literature on rational choice theory and social choice theory has identified a number of desirable properties for (social) choice functions, also referred to as *axioms*, which can be readily applied to tournament solutions. In this section, we review three of the most important properties in this context—*monotonicity*, *stability*, and *composition-consistency*. As we will see in Section 3.3.2, another important property of SCFs—*Pareto-optimality*—is intimately connected to a particular tournament solution, the uncovered set.

A tournament solution is monotonic if a chosen alternative remains in the choice set when its dominion is enlarged, while leaving everything else unchanged.

Definition 3.2 A tournament solution S is *monotonic* if for all $T = (A, \succ)$, $T' = (A, \succ')$, $a \in A$ such that $\succ_{A\setminus\{a\}} = \succ'_{A\setminus\{a\}}$ and for all $b \in A \setminus \{a\}$, $a \succ' b$ whenever $a \succ b$,

$$a \in S(T) \quad \text{implies} \quad a \in S(T').$$

Monotonicity of a tournament solution immediately implies monotonicity of the corresponding $C1$ SCF. Note that this notion of monotonicity for irresolute SCFs is one of the weakest one could think of.

While monotonicity relates choices from tournaments of the same order to each other, the next property relates choices from different subtournaments of the same tournament to each other. Informally, stability (or self-stability) requires that a set is chosen from two different sets of alternatives if and only if it is chosen from the union of these sets.

Definition 3.3 A tournament solution S is *stable* if for all tournaments $T = (A, \succ)$ and for all nonempty subsets $B, C, X \subseteq A$ with $X \subseteq B \cap C$,

$$X = S(B) = S(C) \quad \text{if and only if} \quad X = S(B \cup C).$$

In comparison to monotonicity, stability appears to be much more demanding. It can be factorized into two conditions, $\widehat{\alpha}$ and $\widehat{\gamma}$. Condition $\widehat{\gamma}$ corresponds to the implication from left to right whereas $\widehat{\alpha}$ is the implication from right to left (Brandt and Harrenstein, 2011). $\widehat{\alpha}$ is also known as Chernoff's *postulate* 5^* (Chernoff, 1954), the *strong superset property* (Bordes, 1979), *outcast* (Aizerman and Aleskerov, 1995), and the *attention filter axiom* (Masatlioglu et al., 2012).[3] $\widehat{\alpha}$ implies *idempotency*,[4] i.e.,

$$S(S(T)) = S(T) \text{ for all } T.$$

Finally, we consider a structural invariance property that is based on components and strengthens common cloning-consistency conditions. A tournament solution is composition-consistent if it chooses the "best" alternatives from the "best" components.

Definition 3.4 A tournament solution S is *composition-consistent* if for all tournaments T, $T_1, \dots, T_k$, and $\tilde{T}$ such that $T = \prod(\tilde{T}, T_1, \dots, T_k)$,

$$S(T) \quad = \quad \bigcup_{i \in S(\tilde{T})} S(T_i).$$

Consider again the tournament given in Figure 3.2. Non-emptiness and neutrality imply that every tournament solution has to select all alternatives in a 3-cycle. It follows that every composition-consistent tournament solution has to select all five alternatives in this tournament.

Besides its normative appeal, composition-consistency can be exploited to speed up the computation of tournament solutions. Brandt et al. (2011) introduced the *decomposition degree* of a tournament as a parameter that reflects its decomposability and showed that computing any composition-consistent tournament solution is

[3] We refer to Monjardet (2008) for a more thorough discussion of the origins of this condition.

[4] Tournament solutions that fail to satisfy idempotency (such as the uncovered set) can be made idempotent by iteratively applying the tournament solution to the resulting choice sets until no further refinement is possible. The corresponding tournament solutions, however, often violate monotonicity.

fixed-parameter tractable with respect to the decomposition degree. Since computing the minimal decomposition requires only linear time, decomposing a tournament never hurts, and often helps.[5]

A weaker notion of composition-consistency, called *weak composition-consistency*, requires that for every pair of tournaments $T = (A, \succ)$ and $T' = (A, \succ')$ that only differ with respect to the dominance relation on some component Y of T, both (i) $S(T) \setminus Y = S(T') \setminus Y$, and (ii) $S(T) \cap Y \neq \emptyset$ if and only if $S(T') \cap Y \neq \emptyset$.

3.3 Common Tournament Solutions

In this section we review some of the most common tournament solutions. On top of the axiomatic properties defined in the previous section, particular attention will be paid to whether and how a tournament solution can be computed efficiently. Whenever a tournament solution is computationally intractable, we state NP-hardness of the decision problem of whether a given alternative belongs to the choice set of a given tournament. This implies hardness of computing the choice set. By virtue of the construction in the proof of Theorem 3.1, it is irrelevant whether the input for this problem is a tournament or a preference profile.

Let us start with two extremely simple tournament solutions. The trivial tournament solution *TRIV* always selects all alternatives from any given tournament. While *TRIV* does not discriminate between alternatives at all and as such is unsuitable as a tournament solution, it is easily verified that it satisfies monotonicity, stability, and composition-consistency, and, of course, can be "computed" efficiently.[6] One of the largest non-trivial tournament solutions is the set of *Condorcet non-losers (CNL)*. A *Condorcet loser* is a (unique) alternative a such that $A \setminus \{a\} \succ \{a\}$. In tournaments of order two or more, *CNL* selects all alternatives except Condorcet losers. *CNL* is barely more discriminating than *TRIV*, yet already fails to satisfy stability and composition-consistency (monotonicity is satisfied).

All tournament solutions defined in the following generalize the concept of a Condorcet winner in one way or another.

3.3.1 Solutions Based on Scores

In this section, we introduce four tournament solutions that are defined via various methods of assigning scores to alternatives: the Copeland set, the Slater set, the Markov set, and the bipartisan set.

[5] Since the representation of a tournament of order m has size $\Theta(m^2)$, the asymptotic running time of a linear time algorithm is in $O(m^2)$.

[6] Many axiomatizations of tournament solutions only require inclusive properties (i.e., properties which demand that alternatives ought to be included in the choice set under certain circumstances) and inclusion-minimality (see, e.g., Brandt et al., 2013a, pp. 224–226).

Copeland Set

The Copeland set is perhaps the first idea that comes to mind when thinking about tournament solutions. While a Condorcet winner is an alternative that dominates *all* other alternatives, Copeland's rule selects those alternatives that dominate the *most* alternatives (see, e.g., Copeland, 1951). Formally, the *Copeland set* $CO(T)$ of a tournament T consists of all alternatives whose dominion is of maximal size, i.e.,

$$CO(T) = \arg\max_{a \in A} |D(a)|.$$

$|D(a)|$ is also called the *Copeland score* of a. In graph-theoretic terms, $|D(a)|$ is the outdegree of vertex a.

In the example tournament given in Figure 3.3, $CO(T) = \{a, b\}$, since both a and b have a Copeland score of 2, whereas the Copeland score of both c and d is 1.

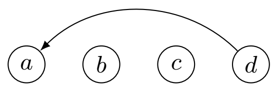

Figure 3.3 Tournament T with $MA(T) = SL(T) = \{a\}$, $CO(T) = \{a, b\}$, $UC(T) = \{a, b, d\}$, and $TRIV(T) = CNL(T) = TC(T) = \{a, b, c, d\}$. All other tournament solutions considered in this chapter coincide with UC. All omitted edges are assumed to point rightwards, i.e., $a \succ b$, $a \succ c$, $b \succ c$, $b \succ d$, and $c \succ d$.

It is straightforward to check that CO satisfies monotonicity. On the other hand, stability and composition-consistency do not hold. This can be seen by again examining the tournament in Figure 3.3. Since $CO(CO(T)) = \{a\} \neq \{a, b\} = CO(T)$, CO violates idempotency and thus stability. Moreover, as $\{\{a\}, \{b, c\}, \{d\}\}$ is a decomposition of T, composition-consistency would require that $d \in CO(T)$, which is not the case. A similar example shows that CO even violates weak composition-consistency. An axiomatic characterization of CO was provided by Henriet (1985).

CO can be easily computed in linear time by determining all Copeland scores and choosing the alternatives with maximum Copeland score.[7]

Theorem 3.5 *The Copeland set can be computed in linear time.*

It is possible to define "second-order" Copeland scores by adding the Copeland scores of all alternatives within the dominion of a given alternative. The process of iteratively computing these scores is guaranteed to converge (due to the Perron-Frobenius Theorem) and leads to a tournament solution, which is sometimes referred to as the *Kendall-Wei method* (see, e.g., Moon, 1968, Ch. 15; Laslier, 1997, pp. 54–56). Kendall-Wei scores can be computed in polynomial time by finding the eigenvector associated with the largest positive eigenvalue of the adjacency matrix.

[7] Brandt et al. (2009) have shown that deciding whether an alternative is contained in $CO(T)$ is TC^0-complete and therefore not expressible in first-order logic.

Slater Set

Although the dominance relation $\succ$ of a tournament may fail to be a strict linear order, it can be linearized by inverting edges in the tournament graph. The intuition behind Slater's rule is to select from a tournament $(A, \succ)$ those alternatives that are maximal elements (i.e., Condorcet winners) in those strict linear orders that can be obtained from $\succ$ by inverting as few edges as possible, i.e., in those strict linear orders that have as many edges in common with $\succ$ as possible (Slater, 1961).[8] Thus, Slater's rule can be seen as the unweighted analogue of Kemeny's social preference function (see Chapter 2 (Zwicker, 2016) and Chapter 4 (Fischer et al., 2016)).

Denote the maximal element of A according to a strict linear order $>$ by $\max(>)$. The *Slater score* of a strict linear order $>$ over the alternatives in A with respect to tournament $T = (A, \succ)$ is $|> \cap \succ|$. A strict linear order is a *Slater order* if it has maximal Slater score. Then, the Slater set SL is defined as

$$SL(T) = \{\max(>) : > \text{ is a Slater order for } T\}.$$

In the example in Figure 3.3, $SL(T) = \{a\}$ because $a > b > c > d$ is the only Slater order. SL satisfies monotonicity, but violates stability and composition-consistency.

Finding Slater orders is equivalent to solving an instance of the *minimum feedback arc set problem*, which is known to be NP-hard, even in tournaments.[9] Therefore, checking membership in SL is NP-hard as well.

Theorem 3.6 (Alon, 2006; Charbit et al., 2007; Conitzer, 2006) *Deciding whether an alternative is contained in the Slater set is NP-hard.*

It is unknown whether the membership problem is contained in NP. The best known upper bound for this problem is the complexity class Θ_2^p, and Hudry (2010) conjectured that the problem is complete for this class. For a more detailed discussion of the computational complexity of Slater's solution, see Hudry (2010) and Charon and Hudry (2006, 2010). Bachmeier et al. (2014) have shown that deciding membership in the Slater set remains NP-hard even when there are only 13 voters.

Although SL is *not* composition-consistent, it satisfies weak composition-consistency. Interestingly, decompositions of the tournament can be exploited to identify a *subset* of the Slater orders (see Laslier (1997, p. 66) and Conitzer (2006)).

Markov Set

Based on ideas that date back at least to Daniels (1969) and Moon and Pullman (1970), Laslier (1997) defines a tournament solution via a Markov chain. The intu-

[8] When inverting as few edges as possible in order to obtain a *Condorcet winner* (rather than a strict linear order), we get the Copeland set.

[9] Whether the minimum feedback arc set problem is NP-hard in tournaments was a long-standing open problem that was solved independently by Alon (2006), Charbit et al. (2007), and Conitzer (2006). The minimum feedback arc set problem is APX-hard (Kann, 1992) and thus does not admit a polynomial-time approximation scheme (PTAS) unless P = NP. For tournaments, however, there exists a PTAS (Kenyon-Mathieu and Schudy, 2007).

ition given by Laslier is that of a table tennis tournament in which the alternatives are players who compete in a series of pairwise comparisons. If a player wins, he will stay at the table and compete in the next match. If he loses, he will be replaced with a new random player. The goal is to identify those players who, in expectation, will win most matches.

The states of the Markov chain are the alternatives and the transition probabilities are determined by the dominance relation: in every step, stay in the current state a with probability $\frac{|D(a)|}{|T|-1}$, and move to state b with probability $\frac{1}{|T|-1}$ for all $b \in \overline{D}(a)$. The *Markov set* consists of those alternatives that have maximum probability in the chain's unique stationary distribution. Formally, the transition matrix of the Markov chain is defined as

$$Q = \frac{1}{|T|-1} \cdot (M(T) + diag(CO)),$$

where $M(T)$ is the adjacency matrix and $diag(CO)$ is the diagonal matrix of the Copeland scores. Let $\Delta(A)$ be the set of all probability distributions over A. The *Markov set* $MA(T)$ of a tournament T is then given by

$$MA(T) = \arg\max_{a \in A} \{p(a) : p \in \Delta(A) \text{ and } Qp = p\}.$$

MA tends to select significantly smaller choice sets than most other tournament solutions. In the example in Figure 3.3, $MA(T) = \{a\}$ because the stationary distribution is $\frac{4}{10}a + \frac{3}{10}b + \frac{1}{10}c + \frac{2}{10}d$. The Markov solution is also closely related to Google's PageRank algorithm for ranking websites (see Brandt and Fischer, 2007). It satisfies monotonicity, but violates stability and weak composition-consistency.

Computing p as the eigenvector of Q associated with the eigenvalue 1 is straightforward. Accordingly, deciding whether an alternative is in MA can be achieved in polynomial time.

Theorem 3.7 *The Markov set can be computed in polynomial time.*

Moreover, Hudry (2009) has pointed out that computing MA has the same asymptotic complexity as matrix multiplication, for which the fastest known algorithm to date runs in $O(m^{2.38})$.

Bipartisan Set

The last tournament solution considered in this section generalizes the notion of a Condorcet winner to lotteries over alternatives. Laffond et al. (1993a) and Fisher and Ryan (1995) have shown independently that every tournament T admits a unique *maximal lottery*,[10] i.e., a probability distribution $p \in \Delta(A)$ such that for

[10] Maximal lotteries were first considered by Kreweras (1965) and studied in detail by Fishburn (1984). The existence of maximal lotteries follows from the Minimax Theorem.

$G(T) = (g_{ab})_{a,b\in A}$,

$$\sum_{a,b\in A} p(a)q(b)g_{ab} \geq 0 \quad \text{for all } q \in \Delta(A).$$

Let p_T denote the unique maximal lottery for a tournament T. Laffond et al. (1993a) define the *bipartisan set* $BP(T)$ of T as the support of p_T, i.e.,

$$BP(T) = \{a \in A : p_T(a) > 0\}.$$

For the tournament in Figure 3.4, we have $p_T = \frac{1}{3}a + \frac{1}{3}b + \frac{1}{3}d$ and thus $BP(T) = \{a, b, d\}$. It is important to realize that the probabilities do not necessarily represent the strengths of alternatives and, that, in contrast to other score-based tournament solutions, just selecting those alternatives with *maximal* probabilities results in a tournament solution that violates monotonicity (see Laslier, 1997, pp. 145–146).

To appreciate this definition, it might be illustrative to interpret the skew-adjacency matrix $G(T)$ of T as a symmetric zero-sum game in which there are two players, one choosing rows and the other choosing columns, and in which the matrix entries are the payoffs of the row player. Then, if the players respectively randomize over rows and columns according to p_T this corresponds to the unique mixed Nash equilibrium of this game. An axiomatization of BP and an interpretation of mixed strategies in the context of electoral competition were provided by Laslier (1997, pp. 151–153) and Laslier (2000), respectively.

BP satisfies monotonicity, stability, and composition-consistency. Moreover, BP can be computed in polynomial time by solving a linear feasibility problem (Brandt and Fischer, 2008).

Theorem 3.8 *The bipartisan set can be computed in polynomial time.*

In weak tournaments—i.e., generalizations of tournaments where the dominance relation is not required to be antisymmetric (see Section 3.5)—deciding whether an alternative is contained in the bipartisan set is P-*complete* (Brandt and Fischer, 2008). Whether P-hardness also holds for tournaments is open.

3.3.2 Uncovered Set and Banks Set

If dominance relations were transitive in general, every tournament (and all of its subtournaments) would admit a Condorcet winner. The *uncovered set* and the *Banks set* address the lack of transitivity in two different but equally natural ways.

The uncovered set takes into account a particular transitive subrelation of the dominance relation, called the covering relation, and selects the maximal alternatives thereof, whereas the Banks set consists of maximal alternatives of inclusion-maximal transitive subtournaments.[11]

[11] As Brandt (2011) notes, the uncovered set contains exactly those alternatives that are Condorcet winners in inclusion-maximal subtournaments that admit a Condorcet winner.

Uncovered Set

An alternative a is said to *cover* alternative b whenever every alternative dominated by b is also dominated by a. Formally, given a tournament $T = (A, \succ)$, the covering relation C is defined as a binary relation over A such that for all distinct $a, b \in A$,

$$a \; C \; b \quad \text{if and only if} \quad D(b) \subseteq D(a).$$

Observe that $a \; C \; b$ implies that $a \succ b$ and is equivalent to $\overline{D}(a) \subseteq \overline{D}(b)$. It is easily verified that the *covering relation* C is transitive and irreflexive, but not necessarily connex. The *uncovered set* $UC(T)$ of a tournament $T = (A, \succ)$ is then given by the set of maximal elements of the covering relation, i.e.,

$$UC(T) = \{a \in A : b \; C \; a \text{ for } \textit{no } b \in A\}.$$

UC was independently proposed by Fishburn (1977) and Miller (1980) and goes back to a game-theoretic notion used by Gillies (1959).

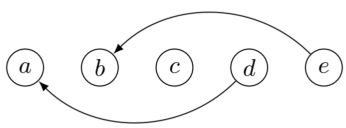

$$\begin{array}{c c} & \begin{array}{ccccc} a & b & c & d & e \end{array} \\ \begin{array}{c} a \\ b \\ c \\ d \\ e \end{array} & \begin{pmatrix} 0 & 1 & 1 & -1 & 1 \\ -1 & 0 & 1 & 1 & -1 \\ -1 & -1 & 0 & 1 & 1 \\ 1 & -1 & -1 & 0 & 1 \\ -1 & 1 & -1 & -1 & 0 \end{pmatrix} \end{array}$$

Figure 3.4 Tournament T and its skew-adjacency matrix $G(T)$. $CO(T) = SL(T) = MA(T) = \{a\}$, $BP = \{a, b, d\}$, $UC(T) = BA(T) = \{a, b, c, d\}$, and $TRIV(T) = CNL(T) = TC(T) = \{a, b, c, d, e\}$. All other tournament solutions considered in this chapter coincide with BP. Omitted edges point rightwards.

In the example in Figure 3.4, a covers e, as $D(e) = \{b\}$ and $D(a) = \{b, c, e\}$. As this is not the case for any other two alternatives, $UC(T) = \{a, b, c, d\}$. UC satisfies monotonicity and composition-consistency, but violates stability. In fact, it does not even satisfy idempotency. An appealing axiomatic characterization of UC was given by Moulin (1986).

Interestingly, UC consists precisely of those alternatives that reach every other alternative on a domination path of length at most two (Shepsle and Weingast, 1984).[12] This equivalence can be easily seen by realizing that

$$\begin{aligned} a \in UC(T) \quad & \text{if and only if there is no } b \in A \text{ such that } b \; C \; a \\ & \text{if and only if for all } b \in \overline{D}(a) \text{ there is some } c \in D(a) \text{ such that } c \succ b \\ & \text{if and only if } a \text{ reaches all } b \in A \setminus \{a\} \text{ in at most two steps.} \end{aligned}$$

This characterization can be leveraged to compute UC via matrix multiplication

[12] In graph theory, these alternatives are called the *kings* of a tournament, and they constitute the *center* of the tournament graph.

because

$$a \in UC(T) \quad \text{if and only if} \quad (M(T)^2 + M(T) + I)_{ab} \neq 0 \text{ for all } b \in A,$$

where I is the $n \times n$ identity matrix (Hudry, 2009). Hence, the asymptotic running time is $O(n^{2.38})$.[13]

Theorem 3.9 *The uncovered set can be computed in polynomial time.*

As mentioned in Chapter 2 (Zwicker, 2016), an alternative is Pareto-optimal if there exists no other alternative such that all voters prefer the latter to the former. A tournament solution is *Pareto-optimal* if its associated SCF only returns Pareto-optimal alternatives. Brandt and Geist (2014) have shown that UC is the coarsest Pareto-optimal tournament solution (see also Brandt et al., 2016). As a consequence, a tournament solution is Pareto-optimal if and only if it is a refinement of UC.

Banks set

The Banks set selects the maximal elements of all maximal transitive subtournaments. Formally, a transitive subtournament $(B, \succ_B)$ of tournament T is said to be *maximal* if there is no other transitive subtournament $(C, \succ_C)$ of T with $B \subset C$. The Banks set $BA(T)$ of a tournament is then defined as

$$BA(T) = \{\max(\succ_B) : (B, \succ_B) \text{ is a maximal transitive subtournament of } T\}.$$

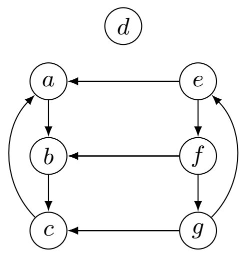

x	$\overline{D}(x)$	$TC(\overline{D}(x))$
a	$\{c, e\}$	$\{c\}$
b	$\{a, f\}$	$\{a\}$
c	$\{b, g\}$	$\{b\}$
d	$\{a, b, c\}$	$\{a, b, c\}$
e	$\{b, c, d, g\}$	$\{b\}$
f	$\{a, c, d, e\}$	$\{c\}$
g	$\{a, b, d, f\}$	$\{a\}$

Figure 3.5 Tournament T and its dominator sets. $BA(T) = \{a, b, c\}$, $UC(T) = \{a, b, c, d\}$, and $TRIV(T) = CNL(T) = TC(T) = \{a, b, c, d, e, f, g\}$. All other tournament solutions considered in this chapter coincide with BA. Omitted edges point rightwards.

The tournament in Figure 3.5 has six maximal transitive subtournaments, induced by the following subsets of A: $\{a, b, d, g\}$, $\{a, d, f, g\}$, $\{b, c, d, e\}$, $\{b, d, g, e\}$,

[13] Brandt and Fischer (2008) proved that the problem of computing UC is contained in the complexity class AC^0 by exploiting that computing the covering relation can be highly parallelized. This is interesting insofar as deciding whether an alternative lies within UC is computationally easier (in AC^0) than checking whether it is contained in CO (TC^0-complete), despite the fact that the fastest known algorithm for computing UC is asymptotically slower than the fastest algorithm for CO.

$\{c, a, d, f\}$, and $\{c, d, e, f\}$. Hence, $BA(T) = \{a, b, c\}$. Like UC, BA satisfies monotonicity and composition-consistency, but violates stability. BA was originally defined as the set of sophisticated outcomes under the amendment agenda (Banks, 1985). For more details see Section 3.4. An alternative axiomatization of the Banks set was given by Brandt (2011).

BA cannot be computed in polynomial time unless P equals NP.

Theorem 3.10 (Woeginger, 2003) *Deciding whether an alternative is contained in the Banks set is NP-complete.*

Proof Membership in NP is straightforward. Given a tournament $T = (A, \succ)$ and an alternative $a \in A$, simply guess a subset B of A and verify that $(B, \succ_B)$ is a transitive subtournament of T with $a = \max(\succ_B)$. Then, check $(B, \succ_B)$ for maximality.

For NP-hardness, we give the reduction from 3SAT by Brandt et al. (2010). Let $\varphi = (x_1^1 \vee x_1^2 \vee x_1^3) \wedge \cdots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$ be a propositional formula in 3-conjunctive normal form (3CNF). For literals x we have $\bar{x} = \neg p$ if $x = p$, and $\bar{x} = p$ if $x = \neg p$, where p is a propositional variable. We may assume that x and $\bar{x}$ do not occur in the same clause.

We now construct a tournament $T_\varphi = (A, \succ)$ with

$$A = \{c_1, \ldots, c_{2m-1}\} \cup \{d\} \cup U_1 \cup \cdots \cup U_{2m-1},$$

where for $1 \leq k \leq 2m - 1$, the set U_k is defined as follows. If k is odd, let $i = \frac{k+1}{2}$ and define $U_k = \{x_i^1, x_i^2, x_i^3\}$. If k is even, let $U_k = \{u_k\}$.

The dominance relation is defined such that $x_i^1 \succ x_i^2 \succ x_i^3 \succ x_i^1$. Moreover, for literals x_i^ℓ and $x_j^{\ell'}$ $(1 \leq \ell, \ell' \leq 3)$ with $i < j$ we have $x_i^\ell \succ x_j^{\ell'}$, unless $x_i^\ell = \bar{x}_j^{\ell'}$, in which case $x_j^{\ell'} \succ x_i^\ell$. For the dominance relation on the remaining alternatives the reader is referred to Figure 3.6.

Observe that for every maximal transitive subtournament $(B, \succ_B)$ of T_φ with $\max(\succ_B) = d$ it holds that:

(i) B contains an alternative from each U_k with $1 \leq k \leq 2m - 1$, and
(ii) for no literal x, the set B contains both x and $\bar{x}$.

For (i), assume that $B \cap U_k = \emptyset$. Since $\max(\succ_B) = d$ and $c_j \succ d$ for all $1 \leq j \leq 2m - 1$, we have $B \cap \{c_1, \ldots, c_{2m-1}\} = \emptyset$. It follows that $(B \cup \{c_k\}, \succ_{B \cup \{c_k\}})$ is transitive ($c_k \succ b$ for all $b \in B$), contradicting maximality of $(B, \succ_B)$. For (ii), assume both $x, \bar{x} \in B$. By a previous assumption then $x \in U_k$ and $\bar{x} \in U_{k'}$ for odd k and k' with $k \neq k'$. Without loss of generality assume that $k < k'$. By (i), $u_{k+1} \in B$. Then, however, $x \succ u_{k+1} \succ \bar{x} \succ x$, contradicting transitivity of $(B, \succ_B)$.

We now prove that

$$\varphi \text{ is satisfiable} \quad \text{if and only if} \quad d \in BA(T_\varphi).$$

First assume that $d \in BA(T_\varphi)$, i.e., $d = \max(\succ_B)$ for some maximal transitive

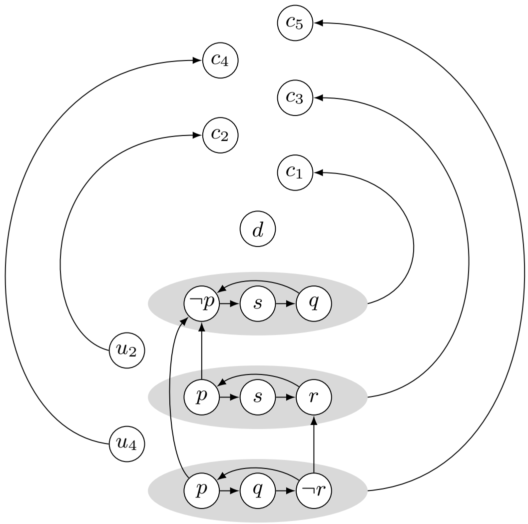

Figure 3.6 Tournament T_φ for the *3CNF* formula $\varphi = (\neg p \vee s \vee q) \wedge (p \vee s \vee r) \wedge (p \vee q \vee \neg r)$. Omitted edges point downwards.

subtournament $(B, \succ_B)$ of T_φ. Define assignment v such that it sets propositional variable p to true if $p \in B$ and to false if $\neg p \in B$. By virtue of (ii), assignment v is well-defined and with (i) it follows that v satisfies φ.

For the opposite direction, assume that φ is satisfiable. Then, there are an assignment v and literals $x_1, \ldots, x_m$ from the clauses $(x_1^1 \vee x_1^2 \vee x_1^3), \ldots, (x_m^1 \vee x_m^2 \vee x_m^3)$, respectively, such that v satisfies each of $x_1, \ldots, x_m$. Define

$$B = \{d\} \cup \{x_1, \ldots, x_m\} \cup \{u_2, u_4, \ldots, u_{2m-2}\}.$$

It is easily seen that $(B, \succ_B)$ is transitive and that $\max(\succ_B) = d$. Observe that B contains an alternative u_k from each U_k with $1 \leq k \leq 2m-1$. Hence, for each $c_k \in C$, we have $c_k \succ d \succ u_k \succ c_k$ and, thus, $(B \cup \{c_k\}, \succ_{B \cup \{c_k\}})$ is not transitive. It follows that $d = \max(\succ_{B'})$ for some maximal transitive subtournament $(B', \succ_{B'})$ with $B \subseteq B'$, i.e., $d \in BA(T_\varphi)$. $\square$

By modifying the construction only slightly and using a variant of 3SAT, Bachmeier et al. (2014) have shown that this problem remains NP-complete even when there are only 5 voters. Interestingly, finding *some* alternative in $BA(A, \succ)$ can be achieved in linear time using the following simple procedure (Hudry, 2004). Label the alternatives in A as $a_1, \ldots, a_m$ and initialize X as the empty set. Then, starting with $k = 1$, successively add alternative a_k to X if and only if a_k dominates all alternatives in X. After m steps, this process terminates and the last alternative added to X can easily be seen to be a member of the Banks set. The difficulty of computing the whole Banks set is rooted in the potentially exponential number of maximal transitive subtournaments.

3.3.3 Solutions Based on Stability

Generalizing an idea by Dutta (1988), Brandt (2011) proposed a method for refining any tournament solution S by defining minimal sets that satisfy a natural stability criterion with respect to S. Given a tournament solution S and a tournament T, a subset of alternatives $B \subseteq A$ is called S-*stable* in T if, for all $a \in A \setminus B$,

$$a \notin S(B \cup \{a\}).$$

An S-stable set B is said to be *minimal* if there is no other S-stable set C in T such that $C \subset B$. Since the set of all alternatives is finite and trivially S-stable, minimal S-stable sets are guaranteed to exist. Now for each tournament solution S, there is a new tournament solution $\widehat{S}$, which returns the union of all minimal S-stable sets in a tournament $T = (A, \succ)$, i.e.,

$$\widehat{S}(T) = \bigcup\{B \subseteq A : B \text{ is a minimal } S\text{-stable set in } T\}.$$

A crucial issue in this context is whether S admits a *unique* minimal stable set in every tournament because this is necessary for $\widehat{S}$ to satisfy stability (Brandt et al., 2014d).

In the following, we will define three tournament solutions using the notion of stable sets: the top cycle, the minimal covering set, and the minimal extending set.

Top Cycle

The top cycle TC can be defined as the unique minimal stable set with respect to CNL, the set of Condorcet non-losers, i.e.,

$$TC = \widehat{CNL}.$$

Alternatively, TC can be defined via the notion of a dominant set. A nonempty subset of alternatives $B \subseteq A$ is called *dominant* in tournament $T = (A, \succ)$ if $B \succ A \setminus B$, i.e., if each alternative in B dominates all alternatives not in B. Dominant sets are linearly ordered via set inclusion and TC returns the unique smallest dominant set. In yet another equivalent definition, TC is defined as the set of maximal elements of the transitive and reflexive closure of the dominance relation $\succ$. TC is a very elementary tournament solution and, in a slightly more general context (see Section 3.5), is also known as *weak closure maximality*, *GETCHA*, or the *Smith set* (Good, 1971; Smith, 1973; Schwartz, 1986). An appealing axiomatic characterization of the top cycle was given by Bordes (1976).

TC tends to select rather large choice sets and may even contain Pareto-dominated alternatives. In the example tournaments given in Figures 3.3, 3.4, and 3.5, TC selects the set of all alternatives because it is the only dominant set. TC satisfies monotonicity, stability, and weak composition-consistency, but violates the stronger notion of composition-consistency (see, e.g., Figure 3.3).

Since each alternative outside TC only dominates alternatives that are also outside TC and every alternative in TC dominates all alternatives outside TC, it can easily be appreciated that each alternative in TC has a strictly greater Copeland score than each alternative outside TC. Hence, $CO \subseteq TC$.

Exploiting this insight, $TC(T)$ can be computed in linear time by starting with $CO(T)$ and then iteratively adding alternatives that are not dominated by the current set. Alternatively, one can employ an algorithm, e.g., the Kosaraju-Sharir algorithm or Tarjan's algorithm, for finding the *strongly connected components* of T and then output the unique strongly connected component that dominates all other strongly connected components.[14]

Theorem 3.11 *The top cycle can be computed in linear time.*

Minimal Covering Set

A subset B of alternatives is called a *covering set* if it is UC-stable, i.e., if every $a \in A \setminus B$ is covered in the subtournament $(B \cup \{a\}, \succ_{B\cup\{a\}})$. The *minimal covering set MC* is defined as

$$MC = \widehat{UC}.$$

Dutta (1988) has shown that every tournament admits a *unique* minimal UC-stable set and that $MC \subseteq UC$. In the example in Figure 3.4, $MC(T) = \{a, b, d\}$, and hence MC is a strict refinement of UC. Observe that, for instance, $\{a, b, c\}$ is not UC-stable, as $d \in UC(\{a, b, c, d\})$. MC satisfies monotonicity, stability, and composition-consistency. Dutta also provided an axiomatic characterization of MC, which was later improved by Laslier (1997, pp. 117–120).

Laffond et al. (1993a) have shown that $BP \subseteq MC$. By virtue of Theorem 3.8, we can therefore efficiently compute a *nonempty subset* of MC. This fact can be used to compute MC by leveraging the following lemma.

Lemma 3.12 *Let $T = (A, \succ)$ be a tournament and $B \subseteq MC(A)$. Define $C = \{a \in A \setminus B : a \in UC(B \cup \{a\})\}$. Then, $MC(C) \subseteq MC(A)$.*

$MC(T)$ can then be computed by first computing the bipartisan set $BP(T)$ and then iteratively adding a specific subset of alternatives that lie outside the current set but do belong to $MC(T)$. Lemma 3.12 tells us how this subset can be found at each stage (see Algorithm 1).[15]

[14] Brandt et al. (2009) have shown that the problem of deciding whether an alternative is contained in the top cycle of a tournament is in the complexity class AC^0.

[15] Lemma 3.12 can also be used to construct a *recursive* algorithm for computing MC without making reference to BP. However, such an algorithm has exponential worst-case running time.

Algorithm 1 Minimal covering set

procedure $MC(A, \succ)$
  $B \leftarrow BP(A)$
  loop
    $C \leftarrow \{a \in A \setminus B : a \in UC(B \cup \{a\}\}$
    if $C = \emptyset$ **then return** B **end if**
    $B \leftarrow B \cup BP(C)$
  end loop

Theorem 3.13 (Brandt and Fischer, 2008) *The minimal covering set can be computed in polynomial time.*

Minimal Extending Set

A subset of alternatives is called an *extending set* if it is BA-stable. Brandt (2011) defined the *minimal extending set* $ME(T)$ as the union of all minimal extending sets of a tournament T, i.e.,

$$ME = \widehat{BA}.$$

In the tournament in Figure 3.4, we find that $ME(T) = \{a, b, d\}$. Brandt et al. (2014b) showed that $ME \subseteq BA$ and that computing ME is computationally intractable by using a construction similar to that of the proof of Theorem 3.10.

Theorem 3.14 (Brandt et al., 2014b) *Deciding whether an alternative is contained in a minimal extending set is NP-hard.*

The best known upper bound for this decision problem is the complexity class Σ_3^p. Bachmeier et al. (2014) have shown that the problem remains NP-hard even when there are only 7 voters. A relation-algebraic specification of minimal extending sets, which can be used to compute ME on small instances, was proposed by Berghammer (2014).

Brandt (2011) proved that ME satisfies composition-consistency, and conjectured that every tournament contains a unique minimal extending set. Even though this conjecture was later disproved, which implies that ME violates monotonicity and stability, it is unclear whether this seriously impairs the usefulness of ME (Brandt et al., 2013b, 2014b). The counterexample found by Brandt et al. consists of about 10^{136} alternatives and concrete tournaments for which ME violates any of these properties have never been encountered (even when resorting to extensive computer experiments).

3.3.4 Solutions Based on Retentiveness

Finally, we consider an operator on tournament solutions which bears some resemblance to the notion of minimal stable sets as introduced in the previous section. The underlying idea of *retentiveness* was first proposed by Schwartz (1990) and studied more generally by Brandt et al. (2014c).

For a given tournament solution S, we say that an alternative a is S-dominated by alternative b if b is chosen among a's dominators by S. Similarly, a nonempty set of alternatives is called S-retentive if none of its elements is S-dominated by some alternative outside the set. Formally, for a tournament solution S and a tournament $T = (A, \succ)$, a nonempty subset $B \subseteq A$ is S*-retentive* in T if for all $b \in B$ such that $\overline{D}(b) \neq \emptyset$,

$$S(\overline{D}(b)) \subseteq B.$$

An S-retentive set B in T is said to be *minimal* if there is no other S-retentive set C in T with $C \subset B$. As in the case of S-stable sets, minimal S-retentive sets are guaranteed to exist because the set of all alternatives is trivially S-retentive. Thus we can define $\mathring{S}$ as the tournament solution yielding the union of minimal S-retentive sets, i.e., for all tournaments $T = (A, \succ)$,

$$\mathring{S}(T) = \bigcup\{B \subseteq A : B \text{ is a minimal } S\text{-retentive set in } T\}.$$

As with minimal stable sets, it is important for the axiomatic properties of $\mathring{S}$ whether S admits a *unique* minimal retentive set in every tournament. It is easily verified that there always exists a unique minimal $TRIV$-retentive set, and that in fact $\mathring{TRIV} = TC$.

The Minimal TC-Retentive Set

Brandt et al. (2014c) have shown that $\mathring{S}$ inherits several desirable properties from S—including monotonicity and stability—whenever a unique minimal S-retentive set is guaranteed to exist. They went on to show that every tournament admits a unique TC-retentive set. As a consequence, the tournament solution $\mathring{TC}$—which can also be written as $\mathring{\mathring{TRIV}}$—is monotonic and stable. Also, $\mathring{TC}$ inherits efficient computability from TC and satisfies weak composition-consistency.

Theorem 3.15 (Brandt et al., 2014c) *The minimal TC-retentive set can be computed in polynomial time.*

In the tournament in Figure 3.5, the set $\{a, b, c\}$ and each of its supersets is TC-retentive. Therefore, $\mathring{TC}(T) = \{a, b, c\}$.

Tournament Equilibrium Set

Schwartz (1990) defined the *tournament equilibrium set* (TEQ) recursively as the union of all minimal TEQ-retentive sets,

$$TEQ = \mathring{TEQ}.$$

This recursion is well-defined because the order of the dominator set of any alternative is strictly smaller than the order of the original tournament. In the example in Figure 3.5, $TEQ(T) = \mathring{TC}(T) = \{a, b, c\}$, because TEQ and TC coincide on all dominator sets.

TEQ is the only tournament solution defined via retentiveness that satisfies composition-consistency. Schwartz conjectured that every tournament contains a *unique* minimal TEQ-retentive set. As was shown by Laffond et al. (1993b) and Houy (2009b,a), TEQ satisfies any one of a number of important properties including monotonicity and stability *if and only if* Schwartz's conjecture holds. Brandt et al. (2013b) showed that Schwartz's conjecture does not hold by nonconstructively disproving a related weaker conjecture surrounding ME.[16] As a consequence, TEQ violates monotonicity and stability. However, counterexamples to Schwartz's conjecture appear to be extremely rare and it may be argued that TEQ satisfies the properties for all practical purposes.

Using a construction similar to that of the proof of Theorem 3.10, it can be shown that computing TEQ is intractable.[17]

Theorem 3.16 (Brandt et al., 2010) *Deciding whether an alternative is contained in the tournament equilibrium set is NP-hard.*

There is no obvious reason why checking membership in TEQ should be in NP. The best known upper bound for this problem is the complexity class PSPACE. Bachmeier et al. (2014) have shown that this problem remains NP-hard even when there are only 7 voters. Brandt et al. (2010, 2011) devised practical algorithms for TEQ that run reasonably well on moderately-sized instances, even though their worst-case complexity is, of course, still exponential.

3.3.5 Summary

Table 3.1 summarizes the axiomatic as well as computational properties of the considered tournament solutions. There are linear-time algorithms for CO and TC. Moreover, a single element of BA can be found in linear time. Computing BA, TEQ, and SL is intractable unless P equals NP. Apparently, MC and BP fare particularly well in terms of axiomatic properties as well as efficient computability.[18]

[16] A significantly smaller counterexample for Schwartz's conjecture, consisting of only 24 alternatives, was found by Brandt and Seedig (2013). However, this smaller counterexample does not disprove the corresponding conjecture for ME.

[17] The proof of Theorem 3.16 actually shows that the membership decision problem for *any* tournament solution that is sandwiched between BA and TEQ, i.e., computing any tournament solution S with $TEQ \subseteq S \subseteq BA$, is NP-hard.

[18] Berghammer et al. (2013) have formalized the definitions of most of the considered tournament solutions using a computer algebra system, which can then be used to compute and visualize these

	Monotonicity	Stability	Composition-consistency	Computational Complexity
CO	+	–	–	in P
SL	+	–	weak	NP-hard, in Θ_2^p
MA	+	–	–	in P
BP	+	+	+	in P
UC	+	–	+	in P
BA	+	–	+	NP-complete
TC	+	+	weak	in P
MC	+	+	+	in P
ME	–	–	+	NP-hard, in Σ_3^p
$\mathring{TC}$	+	+	weak	in P
TEQ	–	–	+	NP-hard, in PSPACE

Table 3.1 *Axiomatic and computational properties of tournament solutions. All hardness results hold even for a constant number of voters. Computing UC and TC has been shown to be in AC^0 while computing CO is TC^0-complete.*

Figure 3.7 provides a graphical overview of the set-theoretic relationships between tournament solutions. It is known that *BA* and *MC* (and by the known inclusions also *UC* and *TC*) almost always select all alternatives when tournaments are drawn uniformly at random (Fey, 2008; Scott and Fey, 2012). Experimental results suggest that the same is true for *TEQ*. Interestingly, despite satisfying strong inclusive axiomatic properties such as stability and composition-consistency, *BP* is much more discriminative: For every integer $m > 1$, the average number of alternatives that *BP* selects in a labelled tournament of order m is $\frac{m}{2}$ (Fisher and Reeves, 1995; Scott and Fey, 2012).[19] Analytic results concerning the uniform distribution stand in sharp contrast to empirical observations that Condorcet winners are likely to exist in real-world settings, which implies that tournament solutions are much more discriminative than these analytical results suggest (Brandt and Seedig, 2015).

3.4 Strategyproofness and Agenda Implementation

It is well-known from the Gibbard-Satterthwaite Theorem (see Chapter 2 (Zwicker, 2016)) that only trivial resolute SCFs are strategyproof, i.e., immune against the strategic misrepresentation of preferences. Tournament solutions are irresolute by

functions. These general-purpose algorithms are however outperformed by tailor-made algorithms using matrix multiplication, linear programming, or eigenvalue decomposition (see, e.g., Seedig, 2015).

19 Brandt et al. (2014d) have shown that there is no more discriminative stable tournament solution than *BP*. In particular, there is no stable refinement of *BP*.

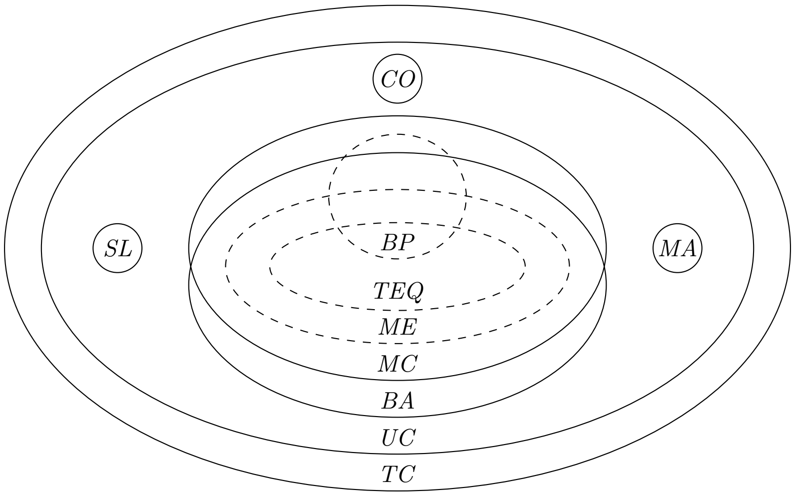

Figure 3.7 The set-theoretic relationships between tournament solutions are depicted in this Venn-like diagram. If the ellipses of two tournament solutions S and S' intersect, then $S(T) \cap S'(T) \neq \emptyset$ for all tournaments T. If the ellipses for S an S' are disjoint, however, this signifies that $S(T) \cap S'(T) = \emptyset$ for some tournament T. Thus, BA and MC are not included in each other, but they always have a nonempty intersection (see, e.g., Laslier, 1997). CO, MA, and SL are contained in UC but may be disjoint from MC and BA. The exact location of BP in this diagram is unknown but it intersects with TEQ in all known instances and is contained in MC. TEQ and ME are contained in BA, but their inclusion in MC is uncertain. Hence, the ellipses for TEQ, ME, and BP are dashed. $\mathring{TC}$ is omitted in this figure because very little is known apart from the inclusion in TC (see Brandt et al., 2015, for more details).

definition (think of a 3-cycle) and therefore the Gibbard-Satterthwaite Theorem does not apply directly.[20]

There are two ways to obtain weak forms of strategyproofness that are particularly well-suited for tournament solutions. The first one concerns the traditional notion of strategyproofness with respect to weakly dominant strategies, but incomplete preference relations over sets of alternatives, and the second one deals with the implementation of tournament solutions by means of sequential binary agendas and subgame-perfect Nash equilibrium. Each of these methods allows for rather positive results, but also comes at a cost: the first one requires a high degree of uncertainty among the voters as to how ties are broken, whereas the second one requires common knowledge of all preferences and may result in impractical voting procedures.

[20] However, the Gibbard-Satterthwaite Theorem does imply that no *resolute refinement* of any of the tournament solutions discussed in this chapter—except $TRIV$—is strategyproof. There are important extensions of the Gibbard-Satterthwaite Theorem to irresolute SCFs such as the Duggan-Schwartz Theorem (see Chapter 2 (Zwicker, 2016)). We will focus on more positive results for tournament solutions in this chapter.

3.4.1 Strategyproofness

A proper definition of strategyproofness for *irresolute* SCFs requires the specification of preferences over *sets* of alternatives. One way to obtain such preferences is to extend the preferences that voters have over individual alternatives to (not necessarily complete) preference relations over sets. A function that yields a preference relation over subsets of alternatives when given a preference relation over single alternatives is called a *set extension*. Of course, there are various set extensions, each of which leads to a different class of strategyproof SCFs (see, e.g., Gärdenfors, 1979; Barberà et al., 2004; Taylor, 2005; Brandt, 2015; Brandt and Brill, 2011).

Here, we will concentrate on two natural and well-studied set extensions due to Kelly (1977) and Fishburn (1972), respectively.[21] Let $\succsim_i$ be the preference relation of voter i and let B and C be two nonempty sets of alternatives. Then, Kelly's extension is defined by letting

$$B \succsim_i^K C \quad \text{if and only if} \quad b \succsim_i c \text{ for all } b \in B \text{ and } c \in C.$$

One interpretation of this extension is that voters are completely unaware of the tiebreaking mechanism (for example, a lottery) that will be used to pick the winning alternative.

Fishburn's extension is defined by letting

$$B \succsim_i^F C \quad \text{if and only if} \quad \begin{array}{l} b \succsim_i c \text{ for all } b \in B \text{ and } c \in C \setminus B \text{ and} \\ b \succsim_i c \text{ for all } b \in B \setminus C \text{ and } c \in C. \end{array}$$

One interpretation of this extension is that ties are broken according to some unknown linear order (e.g., the preferences of a chairman). It is easily seen that $B \succsim_i^K C$ implies $B \succsim_i^F C$.

Each set extension induces a corresponding notion of strategyproofness. An SCF f is *Kelly-strategyproof* if there is no voter i and no pair of preference profiles R and R' with $\succsim_j = \succsim_j'$ for all $j \neq i$ such that $f(R') \succ_i^K f(R)$. If such profiles exist, we say that voter i can *manipulate* f. *Fishburn-strategyproofness* is defined analogously. Note that in this definition of strategyproofness, set extensions are interpreted as fully specified preference relations according to which many choice sets are incomparable (and changing the outcome to an incomparable choice set does not constitute a manipulation). Clearly, since $B \succsim_i^K C$ implies $B \succsim_i^F C$, Fishburn-strategyproofness is stronger than Kelly-strategyproofness.

Kelly-strategyproofness may seem like an extremely weak notion of strategyproofness as only few pairs of sets can actually be compared. Nevertheless, almost all common SCFs fail to satisfy Kelly-strategyproofness because they can already be manipulated on profiles where these functions are resolute (Taylor, 2005, pp. 44–51). Brandt (2015) has shown that stability and monotonicity are sufficient for

[21] Gärdenfors (1979) attributed the second extension to Fishburn because it is the coarsest extension that satisfies a certain set of axioms proposed by Fishburn (1972).

Kelly-strategyproofness. Virtually all SCFs of interest that satisfy these conditions are tournament solutions (or weighted tournament solutions). We therefore only state the result for tournament solutions rather than for SCFs.

Theorem 3.17 (Brandt, 2015) *Every monotonic and stable tournament solution is Kelly-strategyproof. Moreover, every Condorcet-consistent coarsening of a Kelly-strategyproof tournament solution is Kelly-strategyproof.*

As a consequence, BP, each of its Condorcet-consistent coarsenings (such as MC, UC, and TC), and $\mathring{TC}$ are Kelly-strategyproof.[22] On the other hand, it can be shown that every Condorcet-consistent tournament solution that may return a single alternative in the absence of a Condorcet winner is Kelly-manipulable. It follows that CO, SL, and MA fail to be Kelly-strategyproof. More involved arguments can be used to show that ME and TEQ are not Kelly-strategyproof.

The results for Fishburn-strategyproofness are less encouraging. While it is known that TC is Fishburn-strategyproof (Brandt and Brill, 2011; Sanver and Zwicker, 2012), a computer-aided proof has shown that no refinement of UC is Fishburn-strategyproof. Since UC is the coarsest Pareto-optimal tournament solution, we have the following theorem.

Theorem 3.18 *(Brandt and Geist, 2014) There is no Pareto-optimal Fishburn-strategyproof tournament solution.*

As a consequence of this theorem, the set-theoretic relationships depicted in Figure 3.7, and other observations (Brandt and Brill, 2011), TC is the finest Fishburn-strategyproof tournament solution considered in this chapter.

3.4.2 Agenda Implementation

An important question—which has enjoyed considerable attention from social choice theorists and political scientists since the work of Black (1958) and Farquharson (1969)—is whether simple procedures exists that *implement* a particular tournament solution. This in particular concerns procedures that are based on a series of binary choices and eventually lead to the election of a single alternative. The binary choices may depend on one another and need not exclusively be between two alternatives. Such procedures are in wide use by actual committees and institutions at various levels of democratic decision-making. The most prominent among these are the *simple agenda* (or *successive procedure*) and the *amendment procedure*, both of which were initially studied in their own right by political scientists. The former is prevalent in civil law or Euro-Latin legal systems, whereas the latter is more firmly entrenched in the common law or Anglo-American legal tradition (see, e.g., Apesteguia et al., 2014).

[22] In fact, the proof even shows that these functions are *group*-strategyproof with respect to Kelly's extension.

With the simple agenda, the alternatives are ordered in a sequence $a_1, \ldots, a_m$ and subsequently successively being voted up or down by majority voting: First alternative a_1 is brought up for consideration; if a_1 is carried by a majority, it is accepted as the final decision; otherwise, a_1 is rejected and a_2 is brought up for consideration, etc.

With the amendment agenda, the alternatives are again ordered in a sequence $a_1, \ldots, a_m$ and voting then takes place in $m-1$ rounds. In the first round, a majority comparison is made between a_1, the status quo, and a_2, the amendment. The winner then goes through to the next round as the new status quo and is put in a majority contest with a_3, and so on. Figure 3.8 illustrates how these procedures can be depicted as binary trees, the leaves of which are associated with alternatives.

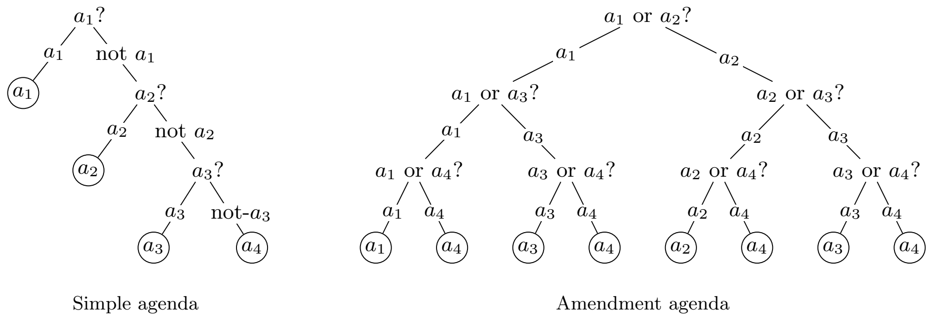

Figure 3.8 The simple agenda and the amendment agenda for four alternatives ordered as a_1, a_2, a_3, a_4.

More generally, every binary tree with alternatives at its leaves could be seen as defining a multi-stage voting procedure. Formally, an *agenda* of order m is defined as a binary tree whose leaves are labelled by an index set I. A *seeding* of a set of alternatives A of size $|I|$ is a bijection from A to I.

For the analysis of voting procedures defined by such agendas and seedings, voters can either be *sincere* or *sophisticated*. Sincere voters myopically and non-strategically vote "directly according to their preferences" whenever the agenda calls for a binary decision. If these choices are invariably between two alternatives, as in the amendment procedure, sincere voting simply comes down to voting for the more preferred alternative at each stage. We refer to Chapter 19 (Vassilevska-Williams, 2016) on knockout tournaments for this setting.

By contrast, sophisticated voters are forward looking and vote strategically. Hence, a more game-theoretic approach and 'backward inductive' reasoning is appropriate. For the remainder of this section, we assume voters to adopt *sophisticated voting strategies*, meaning that the binary tree can be "solved" by successively propagating the majority winner among two siblings to their parent, starting at the leaves and going upwards. Multi-stage sophisticated voting yields the same outcome as the one obtained by solving the *extensive-form game* as defined by the agenda using backwards induction (McKelvey and Niemi, 1978), in an important sense leveraging the strategyproofness of majority rule in settings with more than two alternatives. Similarly, the sophisticated outcome is the alternative that survives iterated elimination of weakly dominated strategies in the *strategic form game* induced by the agenda (Farquharson, 1969; Moulin, 1979).

In order to define agenda-implementability, one defines a class of agendas (one for each order m) and considers all possible seedings for each agenda. A tournament solution S is then said to be *agenda-implementable* if there exists a class of agendas such that for every tournament T, $a \in S(T)$ if and only if there is a seeding for the agenda of size $|T|$ such that its sophisticated outcome is a.

Early results on agenda implementation demonstrated that the class of simple agendas implements TC and the class of amendment agendas implements BA (Miller, 1977, 1980; Banks, 1985; Moulin, 1986). Moulin (1986), moreover, showed that agenda-implementable tournament solutions have to be weakly composition-consistent refinements of TC. As a consequence, CO and MA are not agenda-implementable. A complete characterization of agenda-implementable tournament solutions, however, had long remained elusive before Horan (2013) obtained sufficient conditions for agenda-implementability that cover a wide range of tournament solutions and almost match Moulin's necessary conditions.[23]

Theorem 3.19 (Horan, 2013) *Every weakly composition-consistent tournament solutions that chooses from among the top cycle of every component is agenda-implementable.*

As a corollary to this result it follows that—besides TC and BA—also SL, UC, MC, ME, BP, and TEQ are agenda-implementable. It should be observed, however, that the agendas actually implementing these tournament solutions may be extremely large. The size of the amendment agenda, for instance, is already exponential in the number of alternatives.[24] Moreover, Horan's proof is non-constructive and no concrete classes of agendas that implement any of the tournament solutions considered in this chapter—except the simple agenda and the amendment agenda—are known.

The fact that CO fails to be agenda-implementable has sparked some research on *approximating* Copeland winners via binary agendas. Fischer et al. (2011) showed that agenda-implementability is unachievable for any tournament solution that, from tournaments of order m, only chooses alternatives with a Copeland score at least as high as $\frac{3}{4} + O(\frac{1}{m})$ of the maximum Copeland score. Horan (2013) demonstrated the existence of agenda-implementable tournament solutions that only select

[23] A weaker version of Theorem 3.19 simply states that every composition-consistent refinement of TC is agenda-implementable.

[24] As an extreme case consider the agendas that Coughlan and Le Breton (1999) introduced to implement a refinement of the iterated Banks set (see also Laslier, 1997). The corresponding agenda of order 6 has already $2^{720!} - 1$ nodes!

	Kelly-strategyproofness	Fishburn-strategyproofness	Agenda-implementability
CO	–	–	–
SL	–	–	+
MA	–	–	–
BP	+	–	+
UC	+	–	+
BA	?	–	+ (amendment)
TC	+	+	+ (simple)
MC	+	–	+
ME	–	–	+
$\mathring{TC}$	+	–	?
TEQ	–	–	+

Table 3.2 *Strategic properties of tournament solutions. It is unknown whether BA is Kelly-strategyproof and whether* $\mathring{TC}$ *is agenda-implementable. Interestingly,* $\mathring{TC}$ *falls exactly between the necessary and sufficient conditions given by Moulin (1986) and Horan (2013).*

alternatives whose Copeland score is at least $\frac{2}{3}$ of the maximum Copeland score, improving previous results by Fischer et al. (2011).

3.4.3 ***Summary***

Table 3.2 summarizes which of the considered tournament solutions are Kelly-strategyproof, Fishburn-strategyproof, and agenda-implementable, respectively. Again, it turns out that *BP* represents a decent compromise between discriminative power and attractive axiomatic properties.

3.5 Generalizations to Weak Tournaments

So far, we assumed the majority relation to be antisymmetric, which can be justified, for instance, by assuming that there is an odd number of voters. In general, however, there may be majority ties. These can be accounted for by considering *weak tournaments* $(A, \succsim)$, i.e., directed graphs that represent the complete, but not necessarily antisymmetric, majority relation.[25]

For most of the tournament solutions defined in Section 3.3, generalizations or extensions to weak tournaments have been proposed. Often, it turns out that there are

[25] Alternatively, one can consider the strict part of the majority relation $\succ$, which is asymmetric, but not necessarily connex.

several sensible ways to generalize a tournament solution and it is unclear whether there exists a unique "correct" generalization. A natural criterion for evaluating the different proposals is whether the extension satisfies (appropriate generalizations of) the axiomatic properties that the original tournament solution satisfies.

3.5.1 The Conservative Extension

A *generic* way to generalize any given tournament solution S to weak tournaments is by selecting all alternatives that are chosen by S in some orientation of the weak tournament. Formally, a tournament $T = (A, \succ)$ is an *orientation* of a weak tournament $W = (A, \succsim')$ if $a \succ b$ implies $a \succsim' b$ for all $a, b \in A$. The *conservative extension* of S, denoted $[S]$, is defined such that, for every weak tournament W,

$$[S](W) = \bigcup_{T \in [W]} S(T),$$

where $[W]$ denotes the set of all orientations of W. Brandt et al. (2014a) have shown that $[S]$ inherits several natural properties from S, including monotonicity, stability, and composition-consistency.

An alternative interpretation of weak tournaments is in terms of a partial information setting, where the symmetric and irreflexive part of the dominance relation represents *unknown comparisons* rather than actual ties (see Chapter 10 (Boutilier and Rosenschein, 2016)). In this setting, the set of winners according to the conservative extension exactly corresponds to the set of *possible winners* of the partially specified tournament. The computational complexity of possible and necessary winners of partially specified tournaments has been studied by Aziz et al. (2012), who showed that for a number of tractable tournament solutions (such as CO, UC, and TC), possible winners—and thus the conservative extension—can be computed efficiently.

3.5.2 Extensions of Common Tournament Solutions

For many tournament solutions, *ad hoc* extensions have been proposed in the literature. In this section, we give an overview of these extensions and compare them to the conservative extension.

The *Copeland set* CO gives rise to a whole class of extensions that is parameterized by a number α between 0 and 1. The solution CO^α selects all alternatives that maximize the variant of the Copeland score in which each tie contributes α points to an alternative's score (see, e.g., Faliszewski et al., 2009). Henriet (1985) axiomatically characterized $CO^{\frac{1}{2}}$, arguably the most natural variant in this class. The conservative extension $[CO]$ does not coincide with any of these solutions. Furthermore, $[CO] \not\subseteq CO^\alpha$ for all $\alpha \in [0,1]$ and $CO^\alpha \subseteq [CO]$ if and only if $\frac{1}{2} \leq \alpha \leq 1$.

When moving from tournaments to weak tournaments, maximal lotteries are no longer unique. Dutta and Laslier (1999) have shown that the appropriate generalization of the *bipartisan set BP* is the *essential set ES*, which is given by the set of all alternatives that are contained in the support of *some* maximal lottery. The essential set coincides with the support of any quasi-strict Nash equilibrium of the game defined by the skew-adjacency matrix. It is easy to construct tournaments where ES is strictly smaller than $[BP]$, and there are also weak tournaments in which $[BP]$ is strictly contained in ES.

Duggan (2013) surveyed several extensions of the covering relation to weak tournaments. Any such relation induces a generalization of the *uncovered set UC*. The so-called *deep covering* and *McKelvey covering* relations are particularly interesting extensions. Duggan showed that for all other generalizations of the covering relation he considered, the corresponding uncovered set is a refinement of the deep uncovered set UC_D. Another interesting property of UC_D is that it coincides with the conservative extension of UC. It follows that all other UC generalizations considered by Duggan are refinements of $[UC]$.

Banks and Bordes (1988) discussed four different generalizations of the *Banks set BA* to weak tournaments. Each of these generalizations is a refinement of the conservative extension $[BA]$.

For the *top cycle TC*, Schwartz (1972; 1986) defined two different generalizations (see also Sen, 1986). *GETCHA* (or the *Smith set*) contains the maximal elements of the transitive closure of $\succsim$, whereas *GOCHA* (or the *Schwartz set*) contains the maximal elements of the transitive closure of $\succ$. GOCHA is always contained in GETCHA, and the latter coincides with $[TC]$. A game-theoretical interpretation of TC gives rise to a further generalization. Duggan and Le Breton (2001) observed that the top cycle of a tournament T coincides with the unique *mixed saddle* $MS(T)$ of the game $G(T)$, and showed that the mixed saddle is still unique for games corresponding to weak tournaments. The solution MS is nested between GOCHA and GETCHA. The computational complexity of GETCHA and GOCHA was analyzed by Brandt et al. (2009), and the complexity of mixed saddles was studied by Brandt and Brill (2012).

Generalizations of the *minimal covering set MC* using the McKelvey covering relation and the deep covering relation are known to satisfy stability. There exist weak tournaments in which $[MC]$ is strictly contained in both the McKelvey minimal covering set MC_M and the deep minimal covering set MC_D. There are also weak tournaments in which MC_M is strictly contained in $[MC]$. Computational aspects of generalized minimal covering sets have been analyzed by Brandt and Fischer (2008) and Baumeister et al. (2013).

Schwartz (1990) suggested six ways to extend the *tournament equilibrium set TEQ*—and the notion of retentiveness in general—to weak tournaments. However, all of those variants can easily be shown to lead to disjoint minimal retentive sets even in very small tournaments, and none of the variants coincides with $[TEQ]$.

It is noteworthy that, in contrast to the conservative extension, some of the extensions discussed above fail to inherit properties from their corresponding tournament solutions. For instance, GOCHA violates stability.

A further generalization of tournaments (and weak tournaments) are *weighted tournaments*, which take the size of pairwise majorities into account. Weighted tournament solutions are studied in detail in Chapter 4 (Fischer et al., 2016). Dutta and Laslier (1999) have generalized several common tournament solutions to weighted tournaments.

3.6 Further Reading

The monograph by Moon (1968) provides an excellent, but slightly outdated, overview of mathematical results about tournaments, which is nicely complemented by more recent book chapters on tournament graphs (Reid and Beineke, 1978; Reid, 2004).

The formal study of tournament solutions in the context of social choice was initiated by Moulin (1986) and sparked a large number of research papers, culminating in the definitive monograph by Laslier (1997). More recent overviews of tournament solutions, which also focus on their computational properties, were given by Brandt (2009) and Hudry (2009). There are also comprehensive studies that exclusively deal with tournament solutions based on covering (Duggan, 2013), stability (Brandt, 2011; Brandt and Harrenstein, 2011; Brandt et al., 2014d), and retentiveness (Brandt et al., 2014c), respectively. For some tournament solutions, continuous generalizations to the general spatial model are available (see, e.g., Banks et al., 2006; Duggan, 2013).

For a more extensive introduction to the vast literature on agenda-implementability, the reader is referred to Moulin (1988, Chapter 9), Laslier (1997, Chapter 8), Austen-Smith and Banks (2005, Chapter 4), and Horan (2013). For an overview of the literature on and a discussion of simple and amendment procedures, see, e.g., Apesteguia et al. (2014).

This chapter focusses on *choosing* from a tournament. For the related—but different—problem of *ranking* alternatives in a tournament, finding a ranking that agrees with as many pairwise comparisons as possible (i.e., Slater's rule) has enjoyed widespread acceptance (see, e.g., Charon and Hudry, 2010). Clearly, score-based tournament solutions such as CO and MA can easily be turned into ranking functions. Bouyssou (2004) has studied ranking functions that are defined via the successive application of tournament solutions and found that monotonic and stable tournament solutions yield particularly attractive ranking functions.

Acknowledgments

We would like to thank Florian Brandl, Johannes Hofbauer, Sean Horan, Olivier Hudry, Jean-François Laslier, Michel Le Breton, Hervé Moulin, Rolf Niedermeier, Hans Georg Seedig, and the co-editors Ulle Endriss, Jérôme Lang, and Ariel Procaccia for very helpful feedback.

Felix Brandt is supported by a Heisenberg professorship of the Deutsche Forschungsgemeinschaft under grant BR 2312/7-2. Markus Brill is supported by a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation. Paul Harrenstein is supported by the European Research Council under Advanced Grant 291528 ("RACE").

References

Aizerman, M., and Aleskerov, F. 1995. *Theory of Choice*. Studies in Mathematical and Managerial Economics, vol. 38. North-Holland.

Alon, N. 2006. Ranking Tournaments. *SIAM Journal on Discrete Mathematics*, **20**(1), 137–142.

Apesteguia, J., Ballester, M. A., and Masatlioglu, Y. 2014. A Foundation for Strategic Agenda Voting. *Games and Economic Behavior*, **87**, 91–99.

Arrow, K. J. 1951. *Social Choice and Individual Values*. 1st edn. New Haven: Cowles Foundation. 2nd edition 1963.

Arrow, K. J., and Raynaud, H. 1986. *Social Choice and Multicriterion Decision-Making*. MIT Press.

Austen-Smith, D., and Banks, J. S. 2005. *Positive Political Theory II: Strategy and Structure*. University of Michigan Press.

Aziz, H., Brill, M., Fischer, F., Harrenstein, P., Lang, J., and Seedig, H. G. 2012. Possible and Necessary Winners of Partial Tournaments. Pages 585–592 of: Conitzer, V., and Winikoff, M. (eds), *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. IFAAMAS.

Bachmeier, G., Brandt, F., Geist, C., Harrenstein, P., Kardel, K., and Seedig, H. G. 2014. k-Majority Digraphs and the Hardness of Voting with a Constant Number of Voters. Working paper.

Banks, J. S. 1985. Sophisticated Voting Outcomes and Agenda Control. *Social Choice and Welfare*, **1**(4), 295–306.

Banks, J. S., and Bordes, G. A. 1988. Voting Games, Indifference, and Consistent Sequential Choice Rules. *Social Choice and Welfare*, **5**, 31–44.

Banks, J. S., Duggan, J., and Le Breton, M. 2006. Social choice and electoral competition in the general spatial model. *Journal of Economic Theory*, **126**, 194–234.

Barberà, S., Bossert, W., and Pattanaik, P. K. 2004. Ranking Sets of Objects. Chap. 17, pages 893–977 of: Barberà, S., Hammond, P. J., and Seidl, C. (eds), *Handbook of Utility Theory*, vol. II. Kluwer Academic Publishers.

Baumeister, D., Brandt, F., Fischer, F., Hoffmann, J., and Rothe, J. 2013. The

Complexity of Computing Minimal Unidirectional Covering Sets. *Theory of Computing Systems*, **53**(3), 467–502.

Berghammer, R. 2014. Computing minimal extending sets by relation-algebraic modeling and development. *Journal of Logic and Algebraic Programming*, **83**(2), 103–119.

Berghammer, R., Rusinowska, A., and de Swart, H. 2013. Computing tournament solutions using relation algebra and RelView. *European Journal of Operational Research*, **226**(3), 636–645.

Black, D. 1958. *The Theory of Committees and Elections*. Cambridge University Press.

Bordes, G. 1976. Consistency, Rationality and Collective Choice. *Review of Economic Studies*, **43**(3), 451–457.

Bordes, G. 1979. Some more results on consistency, rationality and collective choice. Chap. 10, pages 175–197 of: Laffont, J. J. (ed), *Aggregation and Revelation of Preferences*. North-Holland.

Boutilier, C., and Rosenschein, J. 2016. Incomplete Information and Communication in Voting. Chap. 10 of: Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (eds), *Handbook of Computational Social Choice*. Cambridge University Press.

Bouyssou, D. 2004. Monotonicity of 'ranking by choosing': A progress report. *Social Choice and Welfare*, **23**(2), 249–273.

Bouyssou, D., Marchant, T., Pirlot, M., Tsoukiàs, A., and Vincke, P. 2006. *Evaluation and Decision Models: Stepping Stones for the Analyst*. Springer-Verlag.

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

Brandt, F. 2011. Minimal Stable Sets in Tournaments. *Journal of Economic Theory*, **146**(4), 1481–1499.

Brandt, F. 2015. Set-Monotonicity Implies Kelly-Strategyproofness. *Social Choice and Welfare*, **45**(4), 793–804.

Brandt, F., and Brill, M. 2011. Necessary and Sufficient Conditions for the Strategyproofness of Irresolute Social Choice Functions. Pages 136–142 of: *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*. ACM Press.

Brandt, F., and Brill, M. 2012. Computing Dominance-Based Solution Concepts. Page 233 of: *Proceedings of the 13th ACM Conference on Electronic Commerce (ACM-EC)*. ACM Press.

Brandt, F., and Fischer, F. 2007. PageRank as a Weak Tournament Solution. Pages 300–305 of: *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE)*. Lecture Notes in Computer Science (LNCS), vol. 4858. Springer-Verlag.

Brandt, F., and Fischer, F. 2008. Computing the Minimal Covering Set. *Mathematical Social Sciences*, **56**(2), 254–268.

Brandt, F., and Geist, C. 2014. Finding Strategyproof Social Choice Functions via SAT Solving. Pages 1193–1200 of: *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. IFAAMAS.

Brandt, F., and Harrenstein, P. 2010. Characterization of Dominance Relations in Finite Coalitional Games. *Theory and Decision*, **69**(2), 233–256.

Brandt, F., and Harrenstein, P. 2011. Set-Rationalizable Choice and Self-Stability. *Journal of Economic Theory*, **146**(4), 1721–1731.

Brandt, F., and Seedig, H. G. 2013. *A tournament of order 24 with two disjoint TEQ-retentive sets*. Tech. rept. http://arxiv.org/abs/1302.5592.

Brandt, F., and Seedig, H. G. 2015. On the Discriminative Power of Tournament Solutions. In: *Selected papers of the International Conference on Operations Research, OR2014*. Operations Research Proceedings. Springer-Verlag. Forthcoming.

Brandt, F., Fischer, F., and Harrenstein, P. 2009. The Computational Complexity of Choice Sets. *Mathematical Logic Quarterly*, **55**(4), 444–459.

Brandt, F., Fischer, F., Harrenstein, P., and Mair, M. 2010. A Computational Analysis of the Tournament Equilibrium Set. *Social Choice and Welfare*, **34**(4), 597–609.

Brandt, F., Brill, M., and Seedig, H. G. 2011. On the Fixed-Parameter Tractability of Composition-Consistent Tournament Solutions. Pages 85–90 of: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press.

Brandt, F., Conitzer, V., and Endriss, U. 2013a. Computational Social Choice. Chap. 6, pages 213–283 of: Weiß, G. (ed), *Multiagent Systems*, 2nd edn. MIT Press.

Brandt, F., Chudnovsky, M., Kim, I., Liu, G., Norin, S., Scott, A., Seymour, P., and Thomassé, S. 2013b. A counterexample to a conjecture of Schwartz. *Social Choice and Welfare*, **40**, 739–743.

Brandt, F., Brill, M., and Harrenstein, P. 2014a. Extending Tournament Solutions. Pages 580–586 of: *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press.

Brandt, F., Harrenstein, P., and Seedig, H. G. 2014b. Minimal Extending Sets in Tournaments. Pages 1539–1540 of: *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. IFAAMAS.

Brandt, F., Brill, M., Fischer, F., and Harrenstein, P. 2014c. Minimal Retentive Sets in Tournaments. *Social Choice and Welfare*, **42**(3), 551–574.

Brandt, F., Brill, M., Seedig, H. G., and Suksompong, W. 2014d. On the structure of stable tournament solutions. Working paper.

Brandt, F., Dau, A., and Seedig, H. G. 2015. Bounds on the Disparity and Separation of Tournament Solutions. *Discrete Applied Mathematics*, **187**, 41–49.

Brandt, F., Geist, C., and Harrenstein, P. 2016. A Note on the McKelvey Uncovered Set and Pareto Optimality. *Social Choice and Welfare*. Forthcoming.

Charbit, P., Thomassé, S., and Yeo, A. 2007. The Minimum Feedback Arc Set Problem is NP-Hard for Tournaments. *Combinatorics, Probability and Computing*, **16**(1), 1–4.

Charon, I., and Hudry, O. 2006. A branch-and-bound algorithm to solve the linear ordering problem for weighted tournaments. *Discrete Applied Mathematics*, **154**(15), 2097–2116.

Charon, I., and Hudry, O. 2010. An updated survey on the linear ordering problem for weighted or unweighted tournaments. *Annals of Operations Research*, **175**(1), 107–158.

Chernoff, H. 1954. Rational Selection of Decision Functions. *Econometrica*, **22**(4), 422–443.

Conitzer, V. 2006. Computing Slater Rankings Using Similarities among Candidates. Pages 613–619 of: *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*. AAAI Press.

Copeland, A. H. 1951. *A 'reasonable' social welfare function.* Mimeographed, University of Michigan Seminar on Applications of Mathematics to the Social Sciences.

Coughlan, P. J., and Le Breton, M. 1999. A social choice function implementable via backward induction with values in the ultimate uncovered set. *Review of Economic Design*, **4**(2), 153–160.

Daniels, H. E. 1969. Round-robin tournament scores. *Biometrika*, **56**(2), 295–299.

Duggan, J. 2013. Uncovered Sets. *Social Choice and Welfare*, **41**(3), 489–535.

Duggan, J., and Le Breton, M. 1996. Dutta's Minimal Covering Set and Shapley's Saddles. *Journal of Economic Theory*, **70**(1), 257–265.

Duggan, J., and Le Breton, M. 2001. Mixed Refinements of Shapley's Saddles and Weak Tournaments. *Social Choice and Welfare*, **18**(1), 65–78.

Dutta, B. 1988. Covering Sets and A New Condorcet Choice Correspondence. *Journal of Economic Theory*, **44**(1), 63–80.

Dutta, B., and Laslier, J.-F. 1999. Comparison Functions and Choice Correspondences. *Social Choice and Welfare*, **16**(4), 513–532.

Erdős, P., and Moser, L. 1964. On the representation of directed graphs as unions of orderings. *Publications of the Mathematical Institute of the Hungarian Academy of Science*, **9**, 125–132.

Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L., and Rothe, J. 2009. Llull and Copeland Voting Computationally Resist Bribery and Constructive Control. *Journal of Artificial Intelligence Research*, **35**, 275–341.

Farquharson, R. 1969. *Theory of Voting.* Yale University Press.

Fey, M. 2008. Choosing From a Large Tournament. *Social Choice and Welfare*, **31**(2), 301–309.

Fischer, F., Procaccia, A. D., and Samorodnitsky, A. 2011. A New Perspective on Implementation by Voting Trees. *Random Structures and Algorithms*, **39**(1), 59–82.

Fischer, F., Hudry, O., and Niedermeier, R. 2016. Weighted Tournament Solutions. Chap. 4 of: Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (eds), *Handbook of Computational Social Choice.* Cambridge University Press.

Fishburn, P. C. 1972. Even-chance lotteries in social choice theory. *Theory and Decision*, **3**(1), 18–40.

Fishburn, P. C. 1977. Condorcet Social Choice Functions. *SIAM Journal on Applied Mathematics*, **33**(3), 469–489.

Fishburn, P. C. 1984. Probabilistic Social Choice Based on Simple Voting Comparisons. *Review of Economic Studies*, **51**(4), 683–692.

Fisher, D. C., and Reeves, R. B. 1995. Optimal Strategies for Random Tournament Games. *Linear Algebra and its Applications*, **217**, 83–85.

Fisher, D. C., and Ryan, J. 1995. Tournament Games and Positive Tournaments. *Journal of Graph Theory*, **19**(2), 217–236.

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

Gillies, D. B. 1959. Solutions to General Non-zero-sum Games. Pages 47–85 of: Tucker, A. W., and Luce, R. D. (eds), *Contributions to the Theory of Games IV*. Annals of Mathematics Studies, vol. 40. Princeton University Press.

Good, I. J. 1971. A Note on Condorcet Sets. *Public Choice*, **10**(1), 97–101.

Henriet, D. 1985. The Copeland choice function: an axiomatic characterization. *Social Choice and Welfare*, **2**(1), 49–63.

Horan, S. 2013. Implementation of Majority Voting Rules. Working paper.

Houy, N. 2009a. *A few new results on TEQ*. Mimeo.

Houy, N. 2009b. Still More on the Tournament Equilibrium Set. *Social Choice and Welfare*, **32**, 93–99.

Hudry, O. 2004. A Note on "Banks Winners in Tournaments are Difficult to Recognize" by G. J. Woeginger. *Social Choice and Welfare*, **23**, 113–114.

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

Hudry, O. 2010. On the complexity of Slater's problems. *European Journal of Operational Research*, **203**(1), 216–221.

Kann, V. 1992. *On the Approximability of NP-complete Optimization Problems*. Ph.D. thesis, Royal Institute of Technology, Stockholm.

Kelly, J. S. 1977. Strategy-Proofness and Social Choice Functions Without Single-Valuedness. *Econometrica*, **45**(2), 439–446.

Kenyon-Mathieu, C., and Schudy, W. 2007. How to rank with few errors: A PTAS for weighted feedback arc set on tournaments. Pages 95–103 of: *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*. ACM Press.

Kreweras, G. 1965. Aggregation of Preference Orderings. Pages 73–79 of: *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gösing, Austria (3–27 July 1962)*.

Laffond, G., Laslier, J.-F., and Le Breton, M. 1993a. The Bipartisan Set of a Tournament Game. *Games and Economic Behavior*, **5**(1), 182–201.

Laffond, G., Laslier, J.-F., and Le Breton, M. 1993b. More on the Tournament Equilibrium Set. *Mathématiques et sciences humaines*, **31**(123), 37–44.

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

Laslier, J.-F. 2000. Interpretation of electoral mixed strategies. *Social Choice and Welfare*, **17**, 283–292.

Masatlioglu, Y., Nakajima, D., and Ozbay, E. Y. 2012. Revealed Attention. *American Economic Review*, **102**(5), 2183–2205.

May, K. 1952. A Set of Independent, Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, **20**(4), 680–684.

McGarvey, D. C. 1953. A Theorem on the Construction of Voting Paradoxes. *Econometrica*, **21**(4), 608–610.

McKelvey, R. D., and Niemi, R. G. 1978. A Multistage Game Representation of Sophisticated Voting for Binary Procedures. *Journal of Economic Theory*, **18**, 1–22.

Miller, N. R. 1977. Graph-Theoretic Approaches to the Theory of Voting. *American Journal of Political Science*, **21**(4), 769–803.

Miller, N. 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.

Monjardet, B. 2008. Statement of precedence and a comment on IIA terminology. *Games and Economic Behavior*, **62**, 736–738.

Moon, J. W. 1968. *Topics on Tournaments*. Holt, Reinhard and Winston.

Moon, J. W., and Pullman, N. K. 1970. On Generalized Tournament Matrices. *SIAM Review*, **12**(3), 384–399.

Moulin, H. 1979. Dominance Solvable Voting Schemes. *Econometrica*, **47**, 1337–1351.

Moulin, H. 1986. Choosing from a Tournament. *Social Choice and Welfare*, **3**(4), 271–291.

Moulin, H. 1988. *Axioms of Cooperative Decision Making*. Cambridge University Press.

Reid, K. B. 2004. Tournaments. Chap. 3.3, pages 156–184 of: Gross, J. L., and Yellen, J. (eds), *Handbook of Graph Theory*. CRC Press.

Reid, K. B., and Beineke, L. W. 1978. Tournaments. Pages 169–204 of: Beineke, L. W., and Wilson, R. J. (eds), *Selected Topics in Graph Theory*. Academic Press.

Sanver, M. R., and Zwicker, W. S. 2012. Monotonicity properties and their adaption to irresolute social choice rules. *Social Choice and Welfare*, **39**(2–3), 371–398.

Schwartz, T. 1972. Rationality and the Myth of the Maximum. *Noûs*, **6**(2), 97–117.

Schwartz, T. 1986. *The Logic of Collective Choice*. Columbia University Press.

Schwartz, T. 1990. Cyclic Tournaments and Cooperative Majority Voting: A Solution. *Social Choice and Welfare*, **7**(1), 19–29.

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

Seedig, H. G. 2015. *Majority Relations and Tournament Solutions: A Computational Study*. Ph.D. thesis, Technische Universität München.

Sen, A. K. 1986. Social Choice Theory. Chap. 22, pages 1073–1181 of: Arrow, K. J., and Intriligator, M. D. (eds), *Handbook of Mathematical Economics*, vol. 3. Elsevier.

Shepsle, K. A., and Weingast, B. R. 1984. Uncovered Sets and Sophisticated Outcomes with Implications for Agenda Institutions. *American Journal of Political Science*, **28**(1), 49–74.

Slater, P. 1961. Inconsistencies in a Schedule of Paired Comparisons. *Biometrika*, **48**(3–4), 303–312.

Smith, J. H. 1973. Aggregation of Preferences with Variable Electorate. *Econometrica*, **41**(6), 1027–1041.

Stearns, R. 1959. The Voting Problem. *American Mathematical Monthly*, **66**(9), 761–763.

Taylor, A. D. 2005. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press.

Vassilevska-Williams, V. 2016. Knockout Tournaments. Chap. 19 of: Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (eds), *Handbook of Computational Social Choice*. Cambridge University Press.

Woeginger, G. J. 2003. Banks Winners in Tournaments are Difficult to Recognize. *Social Choice and Welfare*, **20**(3), 523–528.

Zwicker, W. S. 2016. Introduction to the Theory of Voting. Chap. 2 of: Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (eds), *Handbook of Computational Social Choice*. Cambridge University Press.