A Computational Analysis of the Tournament Equilibrium Set*

Felix Brandt   **Felix Fischer**   **Paul Harrenstein**

Institut für Informatik
Universität München, Germany
`{brandtf,fischerf,harrenst}@tcs.ifi.lmu.de`

Maximilian Mair

Mathematisches Institut
Universität München, Germany
`mairm@cip.ifi.lmu.de`

Abstract

A recurring theme in AI and multiagent systems is how to select the "most desirable" elements given a binary dominance relation on a set of alternatives. Schwartz's *tournament equilibrium set* (TEQ) ranks among the most intriguing, but also among the most enigmatic, tournament solutions proposed so far in this context. Due to its unwieldy recursive definition, little is known about TEQ. In particular, its monotonicity remains an open problem to date. Yet, if TEQ were to satisfy monotonicity, it would be a very attractive solution concept refining both the Banks set and Dutta's minimal covering set. We show that the problem of deciding whether a given alternative is contained in TEQ is NP-hard. Furthermore, we propose a heuristic that significantly outperforms the naive algorithm for computing TEQ. Early experimental results support the conjecture that TEQ is indeed monotonic.

Introduction

A recurring theme in AI and multiagent systems is how to select the "most desirable" elements given a binary dominance relation on a set of alternatives. Examples are diverse and include selecting socially preferred candidates in social choice settings (*e.g.*, Fishburn, 1977; Laslier, 1997), finding valid arguments in argumentation theory (*e.g.*, Dung, 1995; Dunne, 2007), determining the winners of a sports tournament (*e.g.*, Dutta & Laslier, 1999), making decisions based on multiple criteria (*e.g.*, Bouyssou *et al.*, 2006), choosing the optimal strategy in a symmetric two-player zero-sum game (*e.g.*, Duggan & Le Breton, 1996), and singling out acceptable payoff profiles in cooperative game theory (Gillies, 1959; Brandt & Harrenstein, 2007). In social choice theory, where dominance-based solutions are most prevalent, the dominance relation can simply be defined as the pairwise majority relation, *i.e.*, an alternative a is said to dominate another alternative b if the number of individuals preferring a to b exceeds the number of individuals preferring b to a. As is well known from Condorcet's paradox (de Condorcet, 1785), the dominance relation may contain cycles and thus need not have a maximum, even if each of the underlying individual preferences does. As a consequence, the concept of maximality is rendered untenable in most cases, and a variety of so-called *solution concepts* that take over the role of maximality in non-transitive relations have been suggested (see, *e.g.*, Laslier, 1997).

The *tournament equilibrium set* (TEQ) introduced by Schwartz (1990) ranks among the most intriguing, but also among the most enigmatic, solution concepts that has been proposed for tournaments, *i.e.*, asymmetric and complete dominance relations. Due to its unwieldy recursive definition, however, preciously little is known about TEQ (Dutta, 1990; Laffond, Laslier, & Breton, 1993). In particular, whether TEQ satisfies the important property of monotonicity remains an open question to date. If it does, TEQ constitutes a most attractive tournament solution, refining both the minimal covering set and the Banks set (Laslier, 1997; Laffond, Laslier, & Breton, 1993).

Recent work in artificial intelligence and related areas has addressed the computational complexity of almost all common solution concepts (see, *e.g.*, Woeginger, 2003; Conitzer, 2006; Conitzer, Davenport, & Kalagnanam, 2006; Brandt, Fischer, & Harrenstein, 2007). The minimal covering set and the tournament equilibrium set, however, have remained notable exceptions. Laslier writes that "Unfortunately, no algorithm has yet been published for finding the minimal covering set or the tournament equilibrium set of large tournaments. For tournaments of order 10 or more, it is almost impossible to find (in the general case) these sets at hand" (Laslier, 1997, p.8). The minimal covering set has recently been shown to be computable in polynomial time (Brandt & Fischer, 2007). In this paper we prove that the same is not true for TEQ, unless P equals NP. We first give an arguably simpler alternative to Woeginger's (2003) NP-hardness proof for membership in the Banks set. Then the construction used in that proof is modified so as to obtain the analogous result for TEQ. In contrast to the Banks set, there is no obvious reason to suppose that the TEQ membership problem is in NP; it may very well be even harder. In the second part of the paper, we propose and evaluate a heuristic for computing TEQ that performs reasonably well on tournaments with up to 150 alternatives. Experiments further support the conjecture that TEQ is indeed monotonic.

Preliminaries

A *tournament T* is a pair $(A, >)$, where A is a finite set of *alternatives* and $>$ an irreflexive, anti-symmetric, and complete binary relation on A, also referred to as the *dominance relation*. Intuitively, $a > b$ signifies that alternative a beats b in a pairwise comparison. We write $\mathcal{T}$ for the class of all tournaments and have $\mathcal{T}(A)$ denote the set of all tournaments on a fixed set A of alternatives. If T is a tournament on A, then every subset X of A induces a tournament $T|_X = (X, >|_X)$, where $>|_X = \{(x, y) \in X \times X : x > y\}$.

As the dominance relation is not assumed to be transitive in general, there need not be a so-called *Condorcet winner*, *i.e.*, an alternative that dominates all other alternatives. A *tournament solution S* is defined as a function that associates with each tournament T on A a subset $S(T)$ of A. The definition of a tournament solution commonly includes the requirement that $S(T)$ be non-empty if T is defined on a non-empty set of alternatives and that it select the Condorcet winner if there is one (Laslier, 1997, p.37). For X a subset of A, we also write $S(X)$ for the more cumbersome $S(T|_X)$, provided that the tournament T is known from the context. A tournament solution S is said to be *monotonic* if for any two tournaments $T, T' \in \mathcal{T}(A)$ which only differ in that $b > a$ in T and $a > b$ in T', $a \in S(T)$ implies that also $a \in S(T')$, *i.e.*, reinforcing an alternative cannot cause it to be excluded from the solution set. Monotonicity is a vital property that all reasonable tournament solutions satisfy. In this paper, we will be concerned with two particular tournament solutions, the Banks set and Schwartz's tournament equilibrium set (TEQ). For a proper formal definition, however, we need some auxiliary notions and notations.

Let R be a binary relation on a set A. We write R^* for the transitive reflexive closure of R. By the *top cycle $TC_A(R)$* we understand the maximal elements of the asymmetric part of R^*. A subset X of A is said to be *transitive* if R is transitive on X. For $X \subseteq Y \subseteq A$, X is called *maximal transitive in Y* if X is transitive and no proper superset of X in Y is. Clearly, every transitive set is contained in a maximal transitive set. Given a set $Z = \{Z_i\}_{i \in I}$ of pairwise disjoint subsets of A, a subset X of A will be called a *choice set for Z* if it contains precisely one element from each subset $Z_i \in Z$.

In tournaments, maximal transitive sets are also referred to as *Banks trajectories*. The *Banks set $BA(T)$* of a tournament T then collects the maximal elements of the Banks trajectories.

Definition 1 (Banks set) *Let T be a tournament on A. An alternative $a \in A$ is in the* Banks set $BA(T)$ of T *if a is a maximal element of some maximal transitive set in T.*

The *tournament equilibrium set TEQ(T)* of a tournament T on A is defined as the top cycle of a particular subrelation of the dominance relation, referred to as the TEQ relation in the following. The underlying idea is that an alternative is only "properly" dominated, *i.e.*, dominated according to the subrelation, if it is dominated by an element that is selected by some tournament solution concept S. To make this idea precise, for $X \subseteq A$, we write $\overline{D}_X(a) = \{b \in X : b > a\}$ for the *dominators* of a in X, omitting the subscript when


Figure 1: Example due to Schwartz, 1990, where $BA(T) = \{a, b, c, d\}$ and $TEQ(T) = \{a, b, c\}$. The TEQ relation $\rightarrow$ is indicated by thick edges.

$X = A$. Thus, for each alternative a one examines the set $\overline{D}(a)$ of its dominators, and solves the subtournament $T|_{\overline{D}(a)}$ by means of the solution S. In the subrelation a is then only dominated by the alternatives in $S(\overline{D}(a))$. This of course, still leaves open the question as to the choice of the solution concept S. Now, in the case of *TEQ*, S is taken to be *TEQ* itself! This recursion is well-defined because for any $X \subseteq A$ and $a \in X$, the set $\overline{D}_X(a)$ is a *proper* subset of X. Thus, in order to determine the TEQ relation in a subtournament T, one has to calculate the TEQ of smaller and smaller subtournaments of T.

Definition 2 (Tournament equilibrium set) *Let $T \in \mathcal{T}(A)$. For each subset $X \subseteq A$, define the* tournament equilibrium set *TEQ(X) for X as*

$$TEQ(X) = TC_X(\rightarrow_X),$$

where $\rightarrow_X$ is defined as the binary relation on X such that for all $x, y \in X$,

$$x \rightarrow_X y \text{ if and only if } x \in TEQ(\overline{D}_X(y)).$$

Observe that the TEQ relation $\rightarrow_X$ is a subset of the dominance relation $>$ and that, if $\overline{D}_X(x) \neq \emptyset$, then there is some $y \in \overline{D}_X(x)$ with $y \rightarrow_X x$. Furthermore, Definition 2 directly yields a recursive algorithm to compute TEQ. Some reflection reveals that this *naive algorithm* requires time exponential in $|A|$ in the worst case.

It can easily be established that the Banks set and TEQ both select the Condorcet winner of a tournament if there is one. Moreover, in a cyclic tournament on three alternatives, the Banks set and TEQ both consist of all alternatives. In more complex tournaments, however, the Banks set and TEQ may differ. Consider, for example, the tournament T depicted in Figure 1. We first calculate the TEQ relation $\rightarrow$. Thus, *e.g.*, for alternative e we find $\overline{D}(e) = \{a, c, d\}$, which constitutes a three-cycle, and so $TEQ(\overline{D}(e)) = \{a, c, d\}$. Accordingly, $a \rightarrow e$, $c \rightarrow e$, as well as $d \rightarrow e$. Doing this for all alternatives, we find $TEQ(T) = \{a, b, c\}$ as the top cycle $TC(\rightarrow)$ of the relation $\rightarrow$. By contrast, the Banks set consists of the four elements a, b, c and d. E.g., $d \in BA(T)$, because $\{d, c, e\}$ is a maximal transitive set with maximal element d. Nevertheless, TEQ is always included in the Banks set.

Proposition 1 (Schwartz, 1990) *Let $T = (A, \succ)$ be a tournament. Then, $TEQ(T) \subseteq BA(T)$.*

Otherwise, little is known and much surmised about the theoretical properties of TEQ. For example, Schwartz (1990) conjectured that the top cycle of the TEQ relation is always weakly connected, a property of TEQ we will refer to as *CTC* for *connected top cycle*. Laffond, Laslier, & Breton (1993) showed that TEQ satisfying CTC is equivalent to it having a number of useful properties. In particular, TEQ is monotonic if and only if CTC holds. Moreover, CTC implies the inclusion of TEQ in the minimal covering set (see, *e.g.*, Laslier, 1997), another appealing tournament solution. Thus, if TEQ satisfies CTC it might be considered a very strong solution concept. Otherwise, TEQ lacks the vital property of monotonicity and as such it would be severely flawed as a tournament solution.

An Alternative NP-Hardness Proof for Membership in the Banks Set

We begin our investigation of the computational complexity of the TEQ membership problem by giving an alternative proof for NP-hardness of the analogous problem for the Banks set. The latter was first demonstrated by Woeginger (2003) using a reduction from graph three-colorability. Our proof works by a reduction from *3SAT*, the NP-complete satisfiability problem for Boolean formulas in conjunctive normal form with exactly three literals per clause (see, *e.g.*, Papadimitriou, 1994). It is arguably simpler than Woeginger's, and a much similar construction will be used in the next section to prove NP-hardness of membership in TEQ. The tournaments used in both reductions will be taken from a special class $\mathcal{T}^*$.

Definition 3 (The class $\mathcal{T}^*$) *A tournament $(A, \succ)$ is in the class $\mathcal{T}^*$ if it satisfies the following properties. There is some odd integer $n \geq 1$, the number of* layers *in the tournament, such that $A = C \cup U_1 \cup \cdots \cup U_n$, where $C, U_1, \ldots, U_n$ are pairwise disjoint and $C = \{c_0, \ldots, c_n\}$. Each U_i is a singleton if i is even, and $U_i = \{u_i^1, u_i^2, u_i^3\}$ if i is odd. The complete and asymmetric dominance relation $\succ$ is such that for all $c_i \in C_i$, $c_j \in C_j$, $u_i \in U_i$, $u_j \in U_j$ ($0 \leq i, j \leq n$):*

(i) $c_i \succ c_j$, *if $i > j$,*

(ii) $u_i \succ c_j$, *if $i = j$,*

(iii) $c_j \succ u_i$, *if $i \neq j$,*

(iv) $u_i \succ u_j$, *if $i < j$ and at least one of i and j is even,*

(v) $u_i^k \succ u_i^l$, *if i is odd and $k \equiv l - 1 \pmod 3$*

We also refer to c_0 by d, for "decision node" and to $\bigcup_{1 \leq i \leq n} U_n$ by U. For $i = 2k$, we have as a notational convention $U_i = Y_k = \{y_k\}$ and set $Y = \bigcup_{1 \leq 2k \leq n} Y_k$.

Observe that this definition fixes the dominance relation between any two alternatives except for some pairs of alternatives that are both in U.

As a next step in the argument, we associate with each instance of *3SAT* a tournament in the class $\mathcal{T}^*$. An instance of *3SAT* is given by a formula φ in 3-*conjunctive normal*


Figure 2: Tournament T_φ^{BA} for the 3-CNF formula $\varphi = (\neg p \vee s \vee q) \wedge (p \vee s \vee r) \wedge (p \vee q \vee \neg r)$. Omitted edges are assumed to point downwards.

form (3CNF), i.e., $\varphi = (x_1^1 \vee x_1^2 \vee x_1^3) \wedge \cdots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$, where each $x \in \{x_i^1, x_i^2, x_i^3 : 1 \leq i \leq m\}$ is a literal. For each clause $x_i^1 \vee x_i^2 \vee x_i^3$ we assume x_i^1, x_i^2 and x_i^3 to be distinct literals. We moreover assume the literals to be indexed and by X_i we denote the set $\{x_i^1, x_i^2, x_i^3\}$. For literals x we have $\bar{x} = \neg p$ if $x = p$, and $\bar{x} = p$ if $x = \neg p$, where p is some propositional variable. We may also assume that if x and y are literals in the same clause, then $x \neq \bar{y}$. We say a 3CNF $\varphi = (x_1^1 \vee x_1^2 \vee x_1^3) \wedge \cdots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$ is satisfiable if there is a choice set V for $\{X_i\}_{1 \leq i \leq m}$ such that $v' = \bar{v}$ for no $v, v' \in V$. Next we define for each 3SAT formula φ the tournament T_φ^{BA}.

Definition 4 (Banks construction) *Let φ be a 3CNF $(x_1^1 \vee x_1^2 \vee x_1^3) \wedge \cdots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$. Define $T_\varphi^{BA} = (C \cup U, \succ)$ as the tournament in the class $\mathcal{T}^*$ with $2m - 1$ layers such that for all $1 \leq j < 2m$,*

$$U_j = \begin{cases} X_i & \text{if } j = 2i - 1, \\ \{y_i\} & \text{if } j = 2i \end{cases}$$

and such that for all $x \in X_i$ and $x' \in X_j$ ($1 \leq i, j \leq m$),

$x \succ x'$ *if both $j < i$ and $x' = \bar{x}$ or both $i < j$ and $x' \neq \bar{x}$.*

Observe that in conjunction with the other requirements on the dominance relation of a tournament in $\mathcal{T}^*$, this completely fixes the dominance relation $\succ$ of T_φ^{BA}.

An example of a tournament T_φ^{BA} for a 3CNF φ is shown in Figure 2. We are now in a position to present our alternative proof that the Banks membership problem is NP-complete.

Theorem 1 *The problem of deciding whether a particular alternative is in the Banks set of a tournament is NP-complete.*

Proof: Membership in NP is obvious. For a fixed alternative d, we can simply guess a transitive subset of alternatives V with d as maximal element and verify that V is also maximal with respect to set inclusion.

For NP-hardness, we show that T_φ^{BA} contains a maximal transitive set with maximal element d if and only if φ is satisfiable. First observe that V is a maximal transitive subset with maximal element d in T_φ^{BA} only if both

(*i*) for all $1 \le i < 2m$ there is a $u \in U_i$ such that $u \in V$, and

(*ii*) there are no $1 \le i < j < 2m$, $u \in U_i$, $u' \in U_j$ with $u, u' \in V$ such that $u_j > u_i$.

Regarding (*i*), if there is an $1 \le i < 2m$ such that no element of U_i is contained in V, we can always add c_i to V in order to obtain a larger transitive set. If (*ii*) were not to hold, both i and j have to be odd for u_j to dominate u_i. However, in light of (*i*), there has to be k with $i < k < j$ and $u'' \in U_k$ such that $u'' \in V$. It follows that V is not transitive because u, u'', and u' form a cycle. If there is maximal transitive set V with maximal element d complying with both (*i*) and (*ii*), a satisfying assignment of φ can be obtained by letting all literals contained in $X \cap V$ be true.

For the opposite direction, assume that φ is satisfiable. Then there is a choice set W for $\{X_i\}_{1 \le i \le m}$ such that $x' = \bar{x}$ for no $x, x' \in W$. Obviously $V = W \cup \{y_1, \ldots, y_{m-1}\} \cup \{d\}$ does not contain any cycles and thus is transitive with maximal element d. In order to obtain a larger transitive set with a different maximal element, we need to add c_i for some $1 \le i \le m$ to V. However, $V \cup \{c_i\}$ always contains a cycle consisting of c_i, d, and u for some $u \in U_i$, contradicting the transitivity of $V \cup \{c_i\}$. We have thus shown that d is the maximal element of some maximal transitive set in T_φ^{BA} containing V as a subset. □

NP-hardness of Membership in TEQ

In this section we prove that the problem of deciding whether a particular alternative is in the TEQ of a tournament is NP-hard. To this end, we refine the construction that was used in the previous section to prove NP-completeness of membership in the Banks set.

Definition 5 (TEQ construction) *Let φ be a 3-CNF $(x_1^1 \lor x_1^2 \lor x_1^3) \land \cdots \land (x_m^1 \lor x_m^2 \lor x_m^3)$. Further for each $1 \le i < m$, let there be a set $Z_i = \{z_i^1, z_i^2, z_i^3\}$. Define T_φ^{TEQ} as the tournament $(A, >)$ in $\mathcal{T}^*$ with $4n - 3$ layers such that $A = C \cup U_1 \cup \cdots \cup U_{4n-3}$ and for all $1 \le i \le m$,*

$$U_j = \begin{cases} X_i & \text{if } j = 4i - 3, \\ Z_i & \text{if } j = 4i - 1, \\ \{y_i\} & \text{otherwise.} \end{cases}$$

As in the Banks construction, we let for all $x \in X_i$ and $x' \in X_j$ ($1 \le i, j \le m$)

$x > x'$ *if both $j < i$ and $x' = \bar{x}$ or both $i < j$ and $x' \ne \bar{x}$.*

Finally, for all $1 \le i, j \le m$, $x_i^k \in X_i$ and $z_j^l \in Z_j$,

$x_i^k > z_j^l$ *if and only if $i < j$ or both $i = j$ and $k = l$.*

Figure 3: Tournament T_φ^{TEQ} for the 3-CNF formula $\varphi = (\neg p \lor s \lor q) \land (p \lor s \lor r) \land (p \lor q \lor \neg r)$.

An example for such a tournament is shown in Figure 3.

We now proceed to show that a *3SAT* formula φ is satisfiable if and only if the decision node d is in the tournament equilibrium set of T_φ^{TEQ}. We make use of the following lemma.

Lemma 1 *Let $T = (C \cup U, >)$ be a tournament in $\mathcal{T}^*$ and let $B \subseteq C \cup U$ such that $d \in B$. Then, for each $u \in U \cap B$ there exists some $c \in C \cap B$ such that $c \to_B^* u$.*

Sketch of proof: The lemma follows from the stronger statement that for each $u \in U \cap B$ there is some $c \in C \cap B$ such that both $c \to_B^* u$ and $c \in TEQ(B)$. This can be proved by structural induction on supersets B of $\{d\}$. The basis is trivial. In the induction step, we first show that for all $u \in U \cap B$ there is some $c \in C \cap B$ with $c \to u$. Then make the following observation. Let $c_i \in C \cap B$ be such that $\overline{D}_B(c_i) \cap C = \emptyset$, *i.e.*, c_i is the alternative in C with the highest index among those included in B. By construction c_i is the Condorcet winner in $\overline{D}_B(c)$. Accordingly, $c_i \in TEQ(\overline{D}_B(c))$ and $c_i \to_B c$. Hence,

$$c_i \to_B c \text{ for all } c \in B \cap C \text{ with } c \ne c_i.$$

It follows that for all $u \in U \cap C$ we have $c_i \to_B^* u$. Therefore, $c_i \in TEQ(B)$. We may conclude that for each $u \in U \cap B$ there is some $c \in C \cap TEQ(B)$ with $c \to_B^* u$. □

We are now ready to state the main theorem of this paper.

Theorem 2 *Deciding whether a particular alternative is in the tournament equilibrium set of a tournament is NP-hard.*

Proof: By reduction from 3-SAT. Consider an arbitrary 3-DNF φ and construct the tournament $T_\varphi^{TEQ} = (C \cup U, \succ)$. This can be done in polynomial time. We show that

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

For the direction from left to right, observe that by an argument analogous to the proof of Theorem 1 it can be shown that φ is satisfiable if and only if $d \in BA(T_\varphi^{TEQ})$. So assuming that φ is not satisfiable yields $d \notin BA(T_\varphi^{TEQ})$. By the inclusion of TEQ in the Banks set (Proposition 1), it follows that $d \notin TEQ(T_\varphi^{TEQ})$.

For the opposite direction, assume that φ is satisfiable. Then there is a choice set W for $\{X_i\}_{1 \le i \le m}$ such that $x' = \bar{x}$ for no $x, x' \in W$. Obviously $W \cup \{y_1, \dots, y_{m-1}\} \cup \{z_i^j \in Z \colon x_i^j \in W\} = \{u_1, \dots, u_n\}$ contains no cycles and thus is transitive. Without loss of generality we may assume that $u_i \in U_i$ for all $1 \le i \le n$. For each $1 \le i \le n+1$, define a subset $\overline{D}_i$ of alternatives as follows. Set $\overline{D}_{n+1} = A$ and $\overline{D}_i = \bigcap_{i \le j \le n+1} \overline{D}(u_j)$ for each $1 \le i \le n$. Hence, $\overline{D}_1 \subsetneq \cdots \subsetneq \overline{D}_{n+1}$. In an effort to simplify notation, we write $\to_i$ and $\overline{D}_i(x)$ for $\to_{\overline{D}_i}$ and $\overline{D}_{\overline{D}_i}(x)$, respectively. It then suffices to prove that

$$d \in TEQ(\overline{D}_k), \text{ for all } 1 \le k \le n+1. \qquad (1)$$

The theorem then follows as the special case in which $k = n+1$. We first make the following observations concerning the TEQ relation $\to_i$ in each $\overline{D}_i$, for each $1 \le i, j \le n+1$:

(i) $u_j \in \overline{D}_i$ if and only if $j < i$,

(ii) $c_j \in \overline{D}_i$ if and only if $j < i$,

(iii) $c_i \to_{i+1} c_j$ if $j < i \le n$,

(iv) $u_i \to_{i+1} c_i$, if $i \le n$.

For (i), observe that if $j < i$, $u_j \in \overline{D}(u_i)$ by transitivity of the set $\{u_1, \dots, u_n\}$. Hence, $u_j \in \overline{D}_i$. If on the other hand $j \ge i$, then $u_j \notin \overline{D}(u_j)$ and thus $u_j \notin \overline{D}_i$. For (ii), observe that $c_j \in \overline{D}(u_i)$ for all $i \ne j$ and thus $c_j \in \overline{D}_i$ if $j < i$. However, $c_j \notin \overline{D}(u_j)$ and hence $c_j \notin \overline{D}_i$ if $j \ge i$. For (iii), merely observe that c_i is the Condorcet winner in $\overline{D}_{i+1}(c_j)$, if $j < i \le n$. To appreciate (iv), observe that by construction $\overline{D}_{i+1}(c_i)$ has to be either a singleton $\{u_i\}$ for some $u_i \in U_i$, or U_i itself. The former is the case if $U_i \subseteq Y$, or if $U_i \subseteq X$ and $i \ne n$. The latter holds if $U_i = U_n$ or if $U_i \subseteq Z$. In either case, $TEQ(\overline{D}_{i+1}(c_i)) = \overline{D}_{i+1}(c_i)$ and $u_i \to_{i+1} c_i$ holds. For the case in which $U_i \subseteq X$ with $i \ne n$, let $U_i = \{u_i, u_i', u_i''\}$. By construction, $U_{i+2} \subseteq Z$ and $u_i', u_i'' \notin \overline{D}(u_{i+2})$. Accordingly, $u_i', u_i'' \notin \overline{D}_{i+1}$. From $u_i \in \overline{D}_{i+1}$ it then follows that $\overline{D}_{i+1} \cap U_i = \{u_i\}$.

We are now in a position to prove (1) by induction on k. For $k = 1$, observe that d is a Condorcet winner in $\overline{D}_1$ and thus $d \in TEQ(\overline{D}_1)$. For the induction step, let $k = i+1$. With observation (i) we know that $u_i \in \overline{D}_{i+1}$ and, in virtue of the induction hypothesis, also that $d \in TEQ(\overline{D}_i)$. Hence,

---

Algorithm 1 Tournament Equilibrium Set

> **procedure** TEQ(X)
> $R \leftarrow \emptyset$
> $B \leftarrow C \leftarrow \arg\min_{a \in X} |\overline{D}(a)|$
> **loop**
> $R \leftarrow R \cup \{(b, a) \colon a \in C \land b \in TEQ(\overline{D}(a))\}$
> $D \leftarrow \bigcup_{a \in C} TEQ(\overline{D}(a))$
>   **if** $D \subseteq B$ **then return** $TC_B(R)$ **end if**
> $C \leftarrow D$
> $B \leftarrow B \cup C$
> **end loop**

$d \to_{i+1} u_i$. Moreover, by observations (iii) and (iv), $c_i \to_{i+1} d \to_{i+1} u_i \to_{i+1} c_i$, i.e., c_i, d and u_i constitute a $\to_{i+1}$-cycle. In virtue of Lemma 1 and observation (ii), we may conclude that $c_i \to_{i+1}^* a$ for all $a \in \overline{D}_{i+1}$. Accordingly, $\{c_i, d, u_i\} \subseteq TC_{\overline{D}_{i+1}}(\to_{i+1})$ and $d \in TEQ(\overline{D}_{i+1})$, which concludes the proof. $\qquad\square$

A Heuristic for Computing TEQ

Computational intractability of the TEQ *membership* problem implies that TEQ cannot be computed efficiently either. Nevertheless, the running time of the naive algorithm, which straightforwardly implements the recursive definition of TEQ, can be greatly reduced when assuming that TEQ satisfies CTC. This assumption can fairly be made. For if CTC were not to hold, TEQ would be non-monotonic and thus compromised as a solution concept, the issue of computing it moot.

Algorithm 1 computes TEQ by starting with the set B of all alternatives that have dominator sets of minimal size (*i.e.*, the so-called Copeland winners). These alternatives are good candidates to be included in TEQ and the small size of their dominator sets speeds up the computation of their TEQ-dominators. Then, all alternatives that TEQ-dominate any alternative in B are iteratively added to B until no more such alternatives can be found, in which case the algorithm returns the top cycle of $\to_B$. Of course, the *worst-case* running time of this algorithm is still exponential, but experimental results suggest that it outperforms the naive algorithm by a factor of about five in uniform random tournaments with up to 150 vertices (see Table 1). We implemented two versions of the naive algorithm, which differ in the subroutine that determines the top cycle. The first one uses the Floyd-Warshall algorithm with an asymptotic complexity of $O(n^3)$, whereas the second one employs Kosaraju's algorithm with a complexity of $O(n^2)$ (see, *e.g.*, Cormen *et al.*, 2001). Surprisingly, the variant relying on Floyd-Warshall performs slightly better on moderately sized instances due to factors hidden in the asymptotic notation that are amplified as a consequence of TEQ's recursive definition.

While choosing tournaments uniformly at random might be useful for benchmarking algorithms, it raises a number of conceptual problems. First, in voting and most other applications uniform random tournaments do not represent a

|A|	Floyd-Warshall	Kosaraju	Algorithm 1
Uniform random tournaments ($p = 0.5$)			
50	0.48 s	0.59 s	0.09 s
100	53.33 s	65.73 s	9.57 s
150	1 166 s	1 429 s	210 s
Structured random tournaments ($p = 0.8$)			
50	13.87 s	16.56 s	0.01 s
100	18 416 s	21 382 s	8.46 s
150	—	—	1273 s

Table 1: Experimental evaluation of algorithms that compute TEQ. Average running time for ten instances on a 3.2GHz Core2Duo machine. Both versions of the naive algorithm did not terminate within 24 hours when run on structured random tournaments with 150 vertices.

reasonably realistic model of social preferences. Secondly, these tournaments are "almost" regular and tournament solutions almost always select all alternatives in regular tournaments. One model of random tournaments that have more structure can be obtained by defining an arbitrary linear order on the alternatives $a_1, \ldots, a_m$ and letting $a_i > a_j$ for $i < j$ with probability $p > 0.5$. Letting $p = 1$ yields a "completely structured" transitive tournament. The more structure a tournament possesses, the more Algorithm 1 outperforms the naive algorithm, due to the increasing number of large dominator sets that have to be analyzed by the latter at every level of the recursion. In large structured tournaments, the performance gap becomes rather impressive (see Table 1). For example, the naive algorithm requires more than five hours to compute the TEQ of a structured random tournament with 100 vertices whereas it takes Algorithm 1 about eight seconds.[1]

We have further used the naive algorithm to try to disprove CTC (and thus TEQ's monotonicity), but failed to find a counterexample by an exhaustive search in all tournaments with up to eight vertices (roughly $27 \cdot 10^7$ tournaments when disregarding isomorphisms). We also investigated a fairly large number of uniform and structured random tournaments, again to no avail. This can be considered mild evidence that TEQ is indeed monotonic.

References

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., and Fischer, F. 2007. Computational aspects of covering in dominance graphs. In Holte, R. C., and Howe, A., eds., *Proceedings of the 22nd Conference on Artificial Intelligence (AAAI)*, 694–699. AAAI Press.

Brandt, F., and Harrenstein, P. 2007. Dominance in social choice and coalitional game theory. Presented at the 5th International Conference on Logic, Game Theory and Social Choice (LGS).

Brandt, F.; Fischer, F.; and Harrenstein, P. 2007. The computational complexity of choice sets. In Samet, D., ed., *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 82–91. ACM Press.

Conitzer, V.; Davenport, A.; and Kalagnanam, J. 2006. Improved bounds for computing Kemeny rankings. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, 620–627. AAAI Press.

Conitzer, V. 2006. Computing Slater rankings using similarities among candidates. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, 613–619. AAAI Press.

Cormen, T. H.; Leiserson, C. E.; Rivest, R. L.; and Stein, C. 2001. *Introduction to Algorithms*. MIT Press, 2nd edition.

de Condorcet, M. 1785. *Essai sur l'application de l'analyse à la probabilité de décisions rendues à la pluralité de voix*. Imprimerie Royal. Facsimile published in 1972 by Chelsea Publishing Company, New York.

Duggan, J., and Le Breton, M. 1996. Dutta's minimal covering set and Shapley's saddles. *Journal of Economic Theory* 70:257–265.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77:321–357.

Dunne, P. E. 2007. Computational properties of argumentation systems satisfying graph-theoretic constraints. *Artificial Intelligence* 171(10-15):701–729.

Dutta, B., and Laslier, J.-F. 1999. Comparison functions and choice correspondences. *Social Choice and Welfare* 16(4):513–532.

Dutta, B. 1990. On the tournament equilibrium set. *Social Choice and Welfare* 7(4):381–383.

Fishburn, P. C. 1977. Condorcet social choice functions. *SIAM Journal on Applied Mathematics* 33(3):469–489.

Gillies, D. B. 1959. Solutions to general non-zero-sum games. In Tucker, A. W., and Luce, R. D., eds., *Contributions to the Theory of Games IV*, volume 40 of *Annals of Mathematics Studies*. Princeton University Press. 47–85.

Laffond, G.; Laslier, J.-F.; and Breton, M. L. 1993. More on the tournament equilibrium set. *Mathématiques et sciences humaines* 31(123):37–44.

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

Papadimitriou, C. H. 1994. *Computational Complexity*. Addison-Wesley.

Schwartz, T. 1990. Cyclic tournaments and cooperative majority voting: A solution. *Social Choice and Welfare* 7:19–29.

Woeginger, G. J. 2003. Banks winners in tournaments are difficult to recognize. *Social Choice and Welfare* 20:523–528.

---

[1]We also had some limited success with algorithms that make use of the easy-to-prove fact that $TEQ((A, >)) = TEQ(TC(>))$. Assuming CTC, a similar preprocessing step that first computes the minimal covering set of the tournament at hand is possible.