



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS



The University of
Nottingham

Discussion Paper No. 2010-04

Vincent Anesi
March 2010

A New Old Solution for Weak
Tournaments

CeDEx Discussion Paper Series

ISSN 1749 - 3293



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit <http://www.nottingham.ac.uk/economics/cedex/> for more information about the Centre or contact

Karina Terry
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0) 115 95 15620
Fax: +44 (0) 115 95 14159
karina.terry@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

<http://www.nottingham.ac.uk/economics/cedex/papers/index.html>

A New Old Solution for Weak Tournaments

Vincent Anesi
University of Nottingham

March 23, 2010

Abstract

This note uncovers new properties of the von Neumann-Morgenstern solution in weak tournaments and majoritarian games. We propose a new procedure for the construction of choice sets from weak tournaments, based on dynamic stability criteria. The idea is to analyze dynamic versions of the tournament game introduced by Lafond, Laslier and Le Breton (1993) [The bipartisan set of a tournament game. *Games and Economic Behavior* 5, 182-201]. The exploration of a specific class of Markov perfect equilibria in these “dynamic tournament games” yields a new solution concept for weak tournaments — the A-stable set. The alternatives in an A-stable set constitute persistent, long-run policy outcomes in the corresponding dynamic tournament games. We find that, in any weak tournament, the class of A-stable sets coincides with that of von Neumann-Morgenstern stable sets.

Keywords: Electoral competition, stable set, stationary Markov equilibrium, weak tournament, solution theory.

1 Introduction

A “weak tournament” is a pair of a finite set of alternatives and an asymmetric binary relation. When the latter is total as well as asymmetric, the pair is referred to as a “tournament.” These definitions are entirely general — tournaments arise in many areas (sports, preference-driven choice, biometrics, . . .) — but a specific interpretation is used by social choice theorists who think of the binary relation as the majority preference relation of a group of voters on the set of alternatives. In concert with this interpretation, majority voting is often viewed as choosing from a weak tournament, the winning alternatives (or “Condorcet winners”) being those that are maximal with respect to the binary relation. The difficulty with using this approach to predict majority-voting outcomes lies in the fact that the majority preference relation may have no maximal element; that is, there may be no alternative that defeats or ties all other alternatives.¹

Social choice theorists have devoted considerable attention to problems associated with the construction of nonempty choice sets from weak tournaments.² One approach, initiated by Laffond et al. (1993), has been to apply game-theoretic equilibrium concepts to a special class of two-player zero-sum games, called “tournament games.” Given a weak tournament $T = (X, R)$, the players in the corresponding tournament game are two office-motivated candidates who compete in an election by choosing policy platforms from X . If one candidate’s choice beats the other’s according to R then her payoff is 1 and the other’s is -1 . Otherwise, both receive payoffs of 0. If T is a tournament, the bipartisan set of T is the support of the unique mixed strategy Nash equilibrium of this game (Laffond et al., 1993). For any weak tournament T , Dutta and Laslier (1999) define the essential set of T as the support of the unique mixed strategy equilibrium with maximal support. Duggan and Le Breton (1996, 2001) also construct choice sets for weak tournaments by applying Shapley’s saddles to the corresponding tournament games.

¹Ties are possible since the definition of a weak tournament does not assume that the binary relation is total.

²This short note is not the place for a full review of the different solution sets that have been proposed for tournaments and weak tournaments. The reader is referred to Moulin (1986), Dutta (1988), Laslier (1997), Peris and Subiza (1999), and Hudry (2009) for exhaustive accounts of this literature.

Our aim in this note is to propose a solution that would account for a different form of stability in tournament games: policy persistence. The motivation for this new solution derives from the large and growing literature on dynamic models of Downsian electoral competition, dating back to the seminal work of Kramer (1977).³ To any tournament $T = (X, R)$, we thus associate dynamic adaptations of the tournament game described above. In particular, in order to study dynamic stability we assume that in any period: (i) the challenger can choose any platform from X , while the incumbent is bound to her previous choice; and (ii) a set of farsighted and rational voters/players, whose majority preference relation on X is R , decide to retain or to replace the incumbent. Evidently, one can contrive many electorates underlying the majority relation R and, therefore, many “dynamic tournament games” corresponding to T . We concentrate on stationary Markov perfect equilibria of these games when voters are sufficiently farsighted, and define a new solution concept for weak tournaments — the A-stable set. An A-stable set of T is a subset Y of X with the following property: *every* dynamic tournament game corresponding to T has an absorbing equilibrium such that the set of alternatives implemented in its absorbing states is precisely Y . Put differently, each alternative in an A-stable set satisfies the following stability condition: *every* dynamic tournament game has an equilibrium in which the first candidate offering this alternative gains office and will remain there, reenacting the same alternative in all subsequent periods.

While the A-stable set solution seeks to impose new stability criteria, it is not unrelated to existent theory and, as a matter of fact, not new at all. We indeed establish that it is equivalent to a famed cooperative solution — von Neumann-Morgenstern (vNM) stable sets; that is, the class of A-stable sets coincides with that of vNM stable sets of any weak tournament. This finding has important implications for tournament solution theory since the concept of stable set, introduced by von Neumann and Morgenstern (1944), has so

³The subsequent literature is enormous. Recent contributions include Bendor et al. (2006) who review the literature on adaptive parties, and Forand (2009) who surveys noncooperative models of dynamic electoral competition. See also Banks' and Duggan's (2008) literature review. To the best of our knowledge, this note constitutes the first attempt to characterize stationary Markov equilibrium outcomes of a dynamic game of Downsian electoral competition with incumbent policy persistence.

far received little attention from theorists who study weak tournaments. The implications of our finding, however, go further than simply raising the argument that vNM stable sets constitute a relevant solution for weak tournaments. Indeed, the skeptical attitude of social choice theorists and political scientists about vNM stable sets is not confined to weak tournaments, but applies to the entire theory of collective choice. One reason for this attitude might be that vNM stable sets may fail to exist.⁴ But the main reason is the absence of credible “stories” of individual interaction that would provide interpretations for vNM stable sets in the context of collective choice. Paraphrasing McKelvey et al. (1978), “their definition [. . .] contains no behavioral justification for supposing that players adopt outcomes in them. As such, [vNM stable] sets are mathematical inventions without a behavioral rationale.” This note offers such a rationale: Alternatives in vNM stable sets constitute persistent, long-run policy outcomes of electoral competition between office-motivated candidates.⁵

The rest of the note is organized as follows. Section 2 introduces our method of formulating A-stable sets by using stationary Markov perfect equilibria of dynamic tournament games. Section 3 then establishes the equivalence between A-stable sets and vNM stable sets. The appendix contains the proofs of our results.

2 Notation and Definitions

Weak Tournaments, Electorates and Repeated Elections

A *weak tournament* is a pair $T = (X, R)$, where X is a finite set of alternatives (or “vertices”), and R is an asymmetric binary relation on X . T is a *tournament* if R is complete as well as asymmetric. One way to think about weak tournaments is to consider the relation R as being related to the views of an electorate regarding alternatives in X : the expression “ xRy ” represents the statement “ x is majority-preferred to y .” (Although

⁴Existence and uniqueness of vNM stable sets in voting games is discussed in Muto (1984), and Le Breton and Weber (1992).

⁵Different cooperative and noncooperative interpretations of vNM stable sets in the context of committee voting are offered in Anesi (2006, 2009).

$R \subseteq X^2$, we adopt the usual convention of writing “ xRy ” instead of “ $(x, y) \in R$ ” throughout this note.)

Formally, given a weak tournament $T = (X, R)$, a T -electorate is a pair (N, u) such that

- (i) $N = \{1, \dots, n\}$ is a finite set of voters;
- (ii) $u = (u_1, \dots, u_n)$ is a utility profile, where $u_i \in \mathbb{R}^X$ is an injection for each $i \in N$ (no indifference) and, for all $(x, y) \in X^2$,

$$\{i \in N : u_i(x) > u_i(y)\} > \frac{1}{2} \text{ if and only if } xRy .$$

The weak tournament T with a T -electorate (N, u) defines a dynamic game of Downsian electoral competition Γ , which we refer to as a *dynamic T -game*. Specifically, the players in Γ are two candidates, labeled α and β , and the n voters in N , who participate in an infinite sequence of elections. Following each history, period t begins under the shadow of an ongoing state $\mathbf{s}^{t-1} = (\iota^{t-1}, x^{t-1})$, in place from the previous period : $\iota^{t-1} \in \{\alpha, \beta\}$ is the candidate elected in period $t - 1$ (and therefore the incumbent at the start of period t); $x^{t-1} \in X$ is the policy the latter implemented once elected and, by assumption, must also defend in the upcoming election.⁶ In the period- t election, the *challenger* $c \in \{\alpha, \beta\}$, $c \neq \iota^{t-1}$, announces a policy platform y freely chosen from X . Faced with a choice between alternatives y and x^{t-1} , each voter i chooses to vote for one or other of the two candidates. If the proportion of voters casting ballots for the incumbent is at least $1/2$, then the incumbent wins the election and $(\iota^t, x^t) = (\iota^{t-1}, x^{t-1})$; otherwise challenger c wins, and $(\iota^t, x^t) = (c, y)$. The pair $\mathbf{s}^t = (\iota^t, x^t)$ thus becomes the state at the start of period $t + 1$. This process continues *ad infinitum*. The initial state, \mathbf{s}^0 , is chosen by Nature according to some probability measure \wp on $\mathbf{S} \equiv (\{\alpha, \beta\} \times X)$ such that $\wp(\{\mathbf{s}\}) > 0$ for every $\mathbf{s} \in \mathbf{S}$.

In every period t , once alternative x^t has been implemented, every voter i receives an instantaneous payoff $(1 - \delta_i) u_i(x^t)$, where $\delta_i \in (0, 1)$ is her discount factor. Thus, voter i 's payoff from a sequence of states $\{(\iota^t, x^t)\}_{t=1}^\infty$ is given by $(1 - \delta_i) \sum_{t=1}^\infty \delta_i^{t-1} u_i(x^t)$. Each candidate $c \in \{\alpha, \beta\}$ is solely motivated by winning office, so that c 's instantaneous payoff

⁶Originally put forward by Downs (1957) and first formalized by Kramer (1977) and Wittman (1977), this assumption is extensively discussed in Bendor et al. (2006), who propose a different interpretation.

in period t is representable by $(1 - \delta_c) \pi_c(\iota^t)$, where $\delta_c \in (0, 1)$ is her discount factor and

$$\pi_c(\iota) \equiv \begin{cases} 1 & \text{if } \iota = c, \\ -1 & \text{otherwise.} \end{cases}$$

Hence, candidate c 's payoff from a sequence of states $\{(\iota^t, x^t)\}_{t=1}^{\infty}$ is $(1 - \delta_c) \sum_{t=1}^{\infty} \delta_c^{t-1} \pi_c(\iota^t)$.

For any weak tournament T , let $\mathcal{G}(T)$ be the set of dynamic T -games; that is, the set of games induced by T -electorates.

Strategies

Let T be a weak tournament, and let $\Gamma \in \mathcal{G}(T)$ be induced by (N, u) . A *history* at some stage of Γ describes all that has transpired in the previous periods and stages (the sequence of incumbents and challengers, their respective platforms and the associated pattern of votes). In general, a (behavior) strategy σ_l for player $l \in \{\alpha, \beta\} \cup N$ is a mapping that assigns a probability distribution over intended actions (what platform to announce, how to vote) to all conceivable histories at which l is active. Since more detailed notation is required only for *stationary Markov pure strategies* in what follows, we shed unneeded generality and provide a formal definition only of such strategies.

For each $c \in \{\alpha, \beta\}$, let $\sigma_c \in X^{\mathbf{S}}$ denote c 's strategy and, for each $(\iota, x) \in \mathbf{S}$, let $\sigma_c(\iota, x)$ be the alternative offered by candidate c in state (ι, x) , with the restriction that $\sigma_c(\iota, x) = x$ whenever $\iota = c$. A strategy for voter $i \in N$ is denoted by $\sigma_i \in \{\alpha, \beta\}^{\mathbf{S} \times X}$ where, for every $(\mathbf{s}, x) \in \mathbf{S} \times X$, $\sigma_i(\mathbf{s}, x)$ is the candidate whom i votes for when the challenger offers x in state \mathbf{s} . The full collection $\sigma = \{(\sigma_c)_{c \in \{\alpha, \beta\}}, (\sigma_i)_{i \in N}\}$ is a stationary Markov pure strategy profile.

Every stationary Markov pure strategy profile σ generates a probability distribution over infinite sequences of states $\{\mathbf{s}^t\}_{t=1}^{\infty}$.⁷ Say that σ is *absorbing* if $\{\mathbf{s}^t\}_{t=1}^{\infty}$ is convergent. The set of absorbing states of σ is then defined as $A(\sigma) \equiv \{\mathbf{s} \in \mathbf{S} : \tau^\sigma(\mathbf{s}) = \mathbf{s}\}$, where $\tau^\sigma \in \mathbf{S}^{\mathbf{S}}$ is the (deterministic) transition process engendered by σ .

⁷Owing to Nature randomizing over the initial state \mathbf{s}^0 , this distribution is not degenerate.

Equilibrium and Stability

Following the previous literature, our main focus will be on pure-strategy *stage-undominated stationary Markov perfect equilibria (SMPEs)*, i.e., pure strategy subgame perfect equilibria with the following two properties: (i) all players use stationary Markov strategies; and (ii) at any voter history no voter uses a dominated strategy (Baron and Kalai, 1993). And because our goal is to formulate choices sets from weak tournaments with dynamic noncooperative foundations, we further restrict attention to a class of “farsighted” SMPEs.

Definition 1. *Let T be a weak tournament, and let $\Gamma \in \mathcal{G}(T)$. A strategy profile σ is a farsighted equilibrium of Γ if and only if: (i) σ is absorbing; and (ii) there exists a threshold $\delta_0 \in (0, 1)$ such that σ is a pure strategy stage-undominated SMPE of Γ whenever $\delta_i \in (\delta_0, 1)$ for all $i \in N$.*

A strategy profile is absorbing if the sequence of states it engenders ultimately settles down (instead of cycling). An absorbing strategy profile is a farsighted equilibrium if it is a pure strategy stage-undominated SMPE under the presumption that voters can (and do) anticipate further changes when they decide to replace the current incumbent.

Farsighted equilibria thus provide a new method, with a noncooperative notion of farsighted stability, of formulating choices sets from every weak tournament T . We are particularly interested in sets of the form

$$S(\sigma) \equiv \{x \in X : (c, x) \in A(\sigma), \forall c \in \{\alpha, \beta\}\}$$

where σ is a farsighted equilibrium in some $\Gamma \in \mathcal{G}(T)$.

Definition 2. *Let $T = (X, R)$ be a weak tournament. We say that a nonempty set $Y \subseteq X$ is an A-stable set of T if and only if, for all $\Gamma \in \mathcal{G}(T)$, there exists a farsighted equilibrium of Γ , σ_Γ , such that $S(\sigma_\Gamma) = Y$.*

The class of A-stable sets of T is denoted by $\mathcal{A}(T)$. Every element $Y \in \mathcal{A}(T)$ exhibits strong stability properties: For any dynamic T -game, there exists a farsighted equilibrium in which each alternative in Y has a positive probability of being selected by a challenger

(or Nature), and, when it is selected, it remains as the policy in all subsequent periods.⁸ Once a policy in Y has been enacted, there is no other alternative and no majority coalition in favor of changing (with the farsighted implications factored in) the prevailing policy to that alternative. In this sense, A-stability accounts for policy persistence. The analysis of A-stable sets is the subject matter of the next section.

3 Farsighted Equilibrium Outcomes and A-Stable Sets

Our next step is to characterize A-stable sets, and in so doing connect this new solution to classical notions in cooperative game theory. It is useful to begin by recalling the definition of a vNM stable set for weak tournaments. A vNM stable set of weak tournament (X, R) is a nonempty subset V of X satisfying the following two conditions:

(IS) $\forall(x, y) \in V: \neg(xRy)$;

(ES) $\forall x \notin V, \exists y \in V : yRx$.

These two conditions are called *internal stability* and *external stability*, respectively. The class of vNM stable sets of a weak tournament T is denoted by $\mathcal{V}(T)$. Our first result states that any element of $\mathcal{V}(T)$ must be an A-stable set of T .

Proposition 1. *Let T be a weak tournament. Every vNM stable set of T is an A-stable set of T ; that is, $\mathcal{V}(T) \subseteq \mathcal{A}(T)$.*

This result immediately prompts the following question: Is the converse also true? Were the answer “yes,” the class of vNM stable sets would completely characterize the class of A-stable sets of every weak tournament. The following lemma is a first step to demonstrating that this is actually the case.

Lemma 1. *Let T be a weak tournament. The following is true in every $\Gamma \in \mathcal{G}(T)$: if σ is a stage-undominated SMPE of Γ such that $S(\sigma) \neq \emptyset$, then $S(\sigma)$ satisfies internal stability in T .*

⁸The A-stable set as a concept of dynamic stability resembles in spirit the concept of dynamically stable state which Acemoglu et al. (2008) introduce in a dynamic bargaining setting.

An immediate implication of this lemma is that, for any $T = (X, R)$, $\mathcal{A}(T)$ is a subset of the class of internally stable subsets of X (given R). Showing that it is also a subset of the class of externally stable subsets of X uses the following fact.

Lemma 2. *Let T be a weak tournament. The following is true in every $\Gamma \in \mathcal{G}(T)$: if σ is a farsighted equilibrium of Γ such that $S(\sigma) \neq \emptyset$, then $S(\sigma)$ satisfies external stability in T .*

By definition, every farsighted equilibrium path must eventually converge to some absorbing set. Lemma 2 states that this absorbing set must satisfy external stability. Combined with Definition 1, this lemma thus shows that every A-stable set satisfies external stability. By definition of a vNM stable set, we thus obtain the following

Proposition 2. *Let T be a weak tournament. Every A-stable set of T is a vNM stable set of T ; that is, $\mathcal{A}(T) \subseteq \mathcal{V}(T)$.*

Considered together, Propositions 1 and 2 establish the equivalence between A-stable and vNM stable sets. This is formally stated in the next corollary.

Corollary 1. *Let $T = (X, R)$ be a weak tournament. A set of alternatives $Y \subseteq X$ is an A-stable set of T if and only if it is a vNM stable set of T ; that is, $\mathcal{A}(T) = \mathcal{V}(T)$.*

4 Implications and Concluding Remarks

This note uncovers new properties of the vNM solution in weak tournaments using a traditional methodology based on tournament games (Laffond et al., 1993). As described in Section 2, the static tournament game needs first to be amended to account for both dynamic stability and farsighted behavior. We then formulate a solution concept satisfying the criteria of absorption and durability in all the dynamic tournament games thus obtained. Our analysis reveals that the vNM solution is the only solution that meets these criteria. Apart from benchmarking the concept of A-stable set, this result is of independent interest because it reveals new properties of vNM stable sets, which go beyond the internal and external stability criteria.

Another notable implication is the following. As explained in the introduction, it has been fairly common among political scientists to regard the vNM solution as inappropriate for predictive purposes in the context of voting (as in weak tournaments); the main reason for this being the absence of a behavioral rationale underlying its definition. Providing noncooperative foundations for vNM stable sets, this note can be seen as a response to this skeptical view: Alternatives in vNM stable sets constitute absorbing and durable policy outcomes in an important class of dynamic electoral competition games.

Although our analysis of dynamic tournament games has already helped us discover interesting properties of vNM stable sets (and because our interest has been initially limited to weak tournaments), the electoral framework we have presented is very simple. Possible extensions include infinite policy spaces, policy-motivated candidates, multicandidate elections, and government formation to name a few. Its simplicity notwithstanding, we view this note as a first step towards disentangling the relationship between vNM stable sets and electoral outcomes.

Appendix

In a dynamic tournament game, each pure strategy Markov strategy profile σ induces a transition function $\tau^\sigma \in \mathbf{S}^{\mathbf{S}}$ and, with it, a continuation payoff from each state $(\iota, x) \in \mathbf{S}$ for every player. This payoff is given by

$$V_i^\sigma(\iota, x) \equiv (1 - \delta_i) u_i(x) + \delta_i V_i^\sigma(\tau^\sigma(\iota, x))$$

for each voter $i \in N$, and by

$$V_c^\sigma(\iota, x) \equiv (1 - \delta_c) \pi_c(\iota) + \delta_c V_c^\sigma(\tau^\sigma(\iota, x))$$

for each candidate $c \in \{\alpha, \beta\}$.

Proof of Proposition 1

Let T be a weak tournament, and let $V \in \mathcal{V}(T)$. To prove the proposition, we need to show that for every $\Gamma \in \mathcal{G}(T)$ with electorate (N, u) , there exists $\delta_0 \in (0, 1)$ such that the

following statement is true whenever $\min_{i \in N} \delta_i > \delta_0$: There is an absorbing pure strategy stage-undominated SME σ^* of Γ such that $S(\sigma^*) = V$.

Consider such a game $\Gamma \in \mathcal{G}_T$. We first define σ^* and δ_0 . Let f be a function in V^X that satisfies the following conditions: if $x \in V$, then $f(x) = x$; if $x \notin V$, then $f(x) \in \{y \in V : yRx\}$. Since V satisfies external stability, such a function must exist. The stationary Markov strategy profile σ^* is then defined as follows. At the platform-announcement stage of every period starting with state (ι, x) , candidate c 's strategy is as follows:

$$\sigma_c^*(\iota, x) = \begin{cases} x & \text{if } \iota = c \\ f(x) & \text{if } \iota \neq c . \end{cases} \quad (1)$$

At the voting stage of every period with state (x, ι) , each voter i plays as follows

$$\sigma_i^*((x, \iota), y) = \begin{cases} \iota & \text{if } u_i(f(x)) > u_i(f(y)) , \\ c \neq \iota & \text{otherwise,} \end{cases} \quad (2)$$

if $y \neq f(x)$, and

$$\sigma_i^*((x, \iota), f(x)) = \begin{cases} \iota & \text{if } u_i(x) > u_i(f(x)) , \\ c \neq \iota & \text{otherwise.} \end{cases} \quad (3)$$

Inspection of the definition of σ^* reveals that that it is absorbing and $S(\sigma^*) = V$. Recall that any state in \mathbf{S} can be selected with positive probability by Nature. If the initial alternative in \mathbf{s}^0 , say x^0 , is an element of V , the period-1 challenger offers a policy, say y , which is always rejected. Inspection of (2)-(3) indeed reveals that y is accepted if and only if there is a majority coalition of voters who all prefer $f(y)$ to $f(x^0)$. But this is impossible since $f(x^0)$ and $f(y)$ both belong to V which is internally stable.

Suppose now that x^0 , is not a member of V . This implies the set of voters $N_0 \subseteq N$ of voters who prefer $f(x^0)$ to x^0 is a majority coalition. According to (1), the challenger offers $f(x^0) \in V$ and wins (condition (3)). She is then retained and reenacts $f(x^0)$ in all subsequent periods (condition (2)). This proves that σ^* is absorbing and $S(\sigma^*) = V$.

We now turn to the definition of the threshold δ_0 . For each $i \in N$, let

$$\begin{aligned}\Delta_i^{\max} &\equiv \max_{(x,y) \in X^2} [u_i(x) - u_i(y)] \text{ , and} \\ \Delta_i^{\min} &\equiv \min_{(x,y) \in X^2: u_i(x) > u_i(y)} [u_i(x) - u_i(y)] \text{ .}\end{aligned}$$

We then define δ_0 as

$$\delta_0 \equiv \frac{\Delta^{\max}}{\Delta^{\min} + \Delta^{\max}} \in (0, 1)$$

where $\Delta^{\min} \equiv \min_{i \in N} \Delta_i^{\min}$ and $\Delta^{\max} \equiv \max_{i \in N} \Delta_i^{\max}$.

Assume now that $\min_{i \in N} \delta_i > \delta_0$. To complete the proof of Proposition 1, we need to show that σ^* is a stage-undominated SMPE. We do so in two easy-to-prove steps.

Claim 1: For every $i \in N$, and all states (ι, x) and (c, y) , $\sigma_i^((\iota, x), y) \neq \iota$ if and only if $V_i^{\sigma^*}(c, y) > V_i^{\sigma^*}(\iota, x)$.*

By construction, for every $i \in N$ and every $(\iota, x) \in \mathbf{S}$,

$$V_i^{\sigma^*}(\iota, x) = (1 - \delta_i) u_i(x) + \delta_i u_i(f(x)) \quad (4)$$

(= $u_i(x)$ if $x \in V$). Therefore, $V_i^{\sigma^*}(c, y) > V_i^{\sigma^*}(\iota, x)$ if and only if

$$(1 - \delta_i) [u_i(y) - u_i(x)] + \delta_i [u_i(f(y)) - u_i(f(x))] > 0 \text{ .} \quad (5)$$

Suppose first that $y \neq f(x)$. As $\delta_i > \delta_0$, inequality (5) holds if and only if $u_i(f(y)) - u_i(f(x)) > 0$. From condition (2), this implies that $\sigma_i^*((\iota, x), y) \neq \iota$ if and only if $V_i^{\sigma^*}(c, y) > V_i^{\sigma^*}(\iota, x)$.

Suppose now that $y = f(x)$. Equality (4) then implies

$$V_i^{\sigma^*}(c, y) - V_i^{\sigma^*}(\iota, x) = (1 - \delta_i) [u_i(f(x)) - u_i(x)] \text{ .}$$

As a consequence, $u_i(f(x)) > u_i(x)$ (and therefore $\sigma_i^*((\iota, x), y) \neq \iota$) if and only if $V_i^{\sigma^*}(c, y) - V_i^{\sigma^*}(\iota, x) > 0$. This proves Claim 1.

Claim 2: At the platform-announcement stage of any period starting with state (ι, x) , no candidate $c \in \{\alpha, \beta\}$ can gain by deviating from platform $\sigma_c^(\iota, x)$ and conforming to σ_c^* thereafter.*

If $c = \iota$, then Claim 2 is trivial. Suppose therefore that $c \neq \iota$; that is, c is the challenger in state (ι, x) . If x belongs to V , internal stability implies that for every $y \in X$ there is no majority coalition whose members all strictly prefer $f(y)$ to $x = f(x)$. As a result, any platform announcement — including $\sigma_c^*(\iota, x) = f(x) = x$ — is optimal for c and, consequently, she has no profitable deviation.

If x does not belong to V , then c can either conform to σ_c^* by announcing $f(x)$, or deviate by announcing any other alternative. From conditions (1)-(3), her payoff from announcing $f(x)$ is $1/(1 - \delta_c)$ (i.e. she gets elected forever). As this is the highest payoff she can possibly obtain, she has no profitable deviation from σ_c^* .

By the one-shot deviation principle, Claims 1 and 2 establish that σ^* is a stage-undominated SMPE, thus completing the proof of the proposition.

Proof of Lemma 1

Let $T = (X, R)$ be a tournament, and let σ be a stage-undominated SME of $\Gamma \in \mathcal{G}(T)$ such that $S(\sigma) \neq \emptyset$. If $|S(\sigma)| = 1$, Lemma 1 is trivial; so assume $|S(\sigma)| \geq 2$.

Suppose, contrary to the statement of the lemma, that $S(\sigma)$ does not satisfy internal stability. This implies that there are two alternatives in $S(\sigma)$, say x and y , such that xRy . By definition of $S(\sigma)$, states (c, x) and (c, y) are fixed points of τ^σ for each c in $\{\alpha, \beta\}$. An immediate consequence of xRy is therefore that there is a majority coalition M such that

$$V_i^\sigma(\beta, x) = u_i(x) > u_i(y) = V_i^\sigma(\alpha, y)$$

for each $i \in M$. But this implies that candidate β could announce x and gain office forever in state (α, y) ; a contradiction with σ being a stage-undominated SME of Γ .

Proof of Lemma 2

Let $T = (X, R)$ be a tournament. To prove Lemma 2, we need to show that, for every $\Gamma \in \mathcal{G}(T)$, there exists $\tilde{\delta} \in [0, 1)$ such that the following statement is true whenever $\min_{i \in N} \delta_i > \tilde{\delta}$: If σ is an absorbing stage-undominated SMPE of Γ and $S(\sigma) \neq \emptyset$, then $S(\sigma)$ satisfies external stability.

Let $\Gamma \in \mathcal{G}(T)$ be the game induced by T -electorate (N, u) . We define $\tilde{\delta}$ as follows. For each $i \in N$, let

$$d_i \equiv \max_{(x,y) \in X^2: u_i(x) > u_i(y)} \left[\frac{u^{\max} - u_i(x)}{u^{\max} - u_i(y)} \right]^{\frac{1}{|X|-1}},$$

where $u^{\max} \equiv \max_{i \in N} \max_{x \in X} u_i(x)$. By construction, $d_i \in [0, 1)$ for every $i \in N$, and therefore $\max_{i \in N} d_i \equiv \tilde{\delta} \in [0, 1)$.

Suppose now that there is an absorbing pure-strategy SMPE, say σ such that $S(\sigma) \neq \emptyset$. Suppose further that, contrary to the statement of the lemma, $S(\sigma)$ does not satisfy external stability. This implies that there exists $\hat{x} \notin S(\sigma)$ such that $\neg[yR\hat{x}]$ for all $y \in S(\sigma)$. As σ is absorbing, this in turn implies that there exists a finite sequence of states $\{\mathbf{s}^m\} = \{(\iota^m, y^m)\}_{m=0, \dots, \bar{m}}$, $1 \leq \bar{m} \leq |X|$, such that: $y^0 = \hat{x}$, $y^{\bar{m}} \in S(\sigma)$, $\mathbf{s}^m \neq \mathbf{s}^{m+1}$ and $\tau^\sigma(\mathbf{s}^m) = \mathbf{s}^{m+1}$ for each $m = 0, \dots, \bar{m} - 1$. To prove Lemma 2, we must show that this is impossible when $\min_{i \in N} \delta_i > \tilde{\delta}$.

Since $y^{\bar{m}} \in S(\sigma)$, $\neg[y^{\bar{m}}R\hat{x}]$. In every majority coalition $M \subseteq N$ ($|M| > n/2$), there is consequently a voter i_M who prefers \hat{x} to $y^{\bar{m}}$: $u_{i_M}(\hat{x}) > u_{i_M}(y^{\bar{m}})$. Hence,

$$\delta_{i_M}^{\bar{m}-1} > \tilde{\delta}^{|X|-1} \geq \frac{u^{\max} - u_{i_M}(\hat{x})}{u^{\max} - u_{i_M}(y^{\bar{m}})},$$

and therefore

$$\begin{aligned} \frac{V_{i_M}^\sigma(\iota^0, \hat{x}) - V_{i_M}^\sigma(\mathbf{s}^1)}{1 - \delta_{i_M}} &= u_{i_M}(\hat{x}) - V_{i_M}^\sigma(\mathbf{s}^1) \\ &= u_{i_M}(\hat{x}) - (1 - \delta_{i_M}) \sum_{m=1}^{\bar{m}-1} \delta_{i_M}^{m-1} u_{i_M}(y^m) - \delta_{i_M}^{\bar{m}-1} u_{i_M}(y^{\bar{m}}) \\ &\geq \delta_{i_M}^{\bar{m}-1} [u^{\max} - u_{i_M}(y^{\bar{m}})] - [u^{\max} - u_{i_M}(\hat{x})] > 0. \end{aligned}$$

As a consequence, every majority coalition M includes a voter i_M who, in state (ι^0, \hat{x}) , is strictly better-off retaining the incumbent when the challenger announces y^1 . As σ is a stage-undominated SMPE, $\sigma_{i_M}((\iota^0, \hat{x}), y^1) = \iota^0$. This is a contradiction with $\tau^\sigma(\iota^0, \hat{x}) = \mathbf{s}^1$.

References

- Acemoglu, D., Egorov, G., Sonin, K., 2008. Dynamics and Stability of Constitutions, Coalitions, and Clubs. MIT Department of Economics Working Paper No. 08-15.
- Anesi, V., 2006. Committees with Farsighted Voters: A New Interpretation of Stable Sets. *Social Choice and Welfare* 27, 595-610.
- Anesi, V., 2010. Noncooperative Foundations of Stable Sets in Voting Games. Forthcoming in *Games and Economic Behavior*.
- Banks, J.S., Duggan, J., 2008. A Dynamic Model of Democratic Elections in Multidimensional Policy Spaces. *Quarterly Journal of Political Science* 3, 269-299.
- Baron, D.P., Kalai, E., 1993. The simplest equilibrium of a majority rule game. *Journal of Economic Theory* 61, 290-301.
- Bendor, J., Mookherjee, D., Ray, D., 2006. Satisficing and Selection in Electoral Competition. *Quarterly Journal of Political Science* 1, 171-200.
- Downs, A., 1957. *An Economic Theory of Democracy*. New York, Harper.
- Duggan, J., Fey, M., 2006. Repeated Downsian Electoral Competition. *International Journal of Game Theory* 35, 39-69.
- Duggan, J., Le Breton, M., 1996. Dutta's Minimal Covering Set and Shapley's Saddles. *Journal of Economic Theory* 70, 257-265.
- Duggan, J., Le Breton, M., 2001. Mixed refinements of Shapley's saddles and weak tournaments. *Social Choice and Welfare* 18, 65-78.
- Dutta, B., 1988. Covering Sets and a New Condorcet Choice Correspondence. *Journal of Economic Theory* 44, 63-80.
- Dutta, B., Laslier, J.-F., 1999. Comparison Functions and Choice Correspondences. *Social Choice and Welfare* 16, 513-532.

- Forand, J. G., 2009. Two-Party Competition with Persistent Policies. Unpublished.
- Hudry, O., 2009. A Survey on the Complexity of Tournament Solutions. *Mathematical Social Sciences* 57, 292-303.
- Kramer, G.H., 1977. A dynamical model of political equilibrium. *Journal of Economic Theory* 19, 565-567.
- Laffond, G., Laslier, J.-F., Le Breton, M., 1993. The Bipartisan Set of a Tournament Game. *Games and Economic Behavior* 5, 182-201.
- Laslier, J.-F., 1997. *Tournament solutions and majority voting*. Springer, Berlin.
- Le Breton, M., Weber, S. 1992. A note on the core and Von Neumann-Morgenstern solutions of simple games. *Social Choice and Welfare* 9, 57-61.
- McKelvey, R.D., Ordeshook P.C., Winer, M.D., 1978. The Competitive Solution for N -Person Games Without Transferable Utility, With an Application to Committee Games. *American Political Science Review* 72, 599-615.
- Moulin, H., 1986. Choosing from a Tournament. *Social Choice and Welfare* 3, 271-291.
- Muto, S., 1984. Stable Sets for Simple Games with Ordinal Preferences. *Journal of Operations Research Society Japan* 27, 250-258.
- Peris, J.E., Subiza, B., 1999. Condorcet Choice Correspondences for Weak Tournaments. *Social Choice and Welfare* 16, 217-231.
- von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Wittman, D., 1977. Candidates with policy preferences: A dynamic model. *Journal of Economic Theory* 14, 180-189.