

Equilibrium Refinement in Dynamic Voting Games*

Daron Acemoglu
MIT

Georgy Egorov
Northwestern University

Konstantin Sonin
New Economic School

October 2009.

Abstract

We propose two related equilibrium refinements for voting and agenda-setting games, Sequentially Weakly Undominated Equilibrium (SWUE) and Markov Trembling Hand Perfect Equilibrium (MTHPE), and show how these equilibrium concepts eliminate non-intuitive equilibria that arise naturally in dynamic voting games and games in which random or deterministic sequences of agenda-setters make offers to several players. We establish existence of these equilibria in finite and infinite (for MTHPE) games, provide a characterization of the structure of equilibria, and clarify the relationship between the two concepts. Finally, we show how these concepts can be applied in a dynamic model of endogenous club formation.

Keywords: voting, agenda-setting games, Markov trembling-hand perfect equilibrium.

JEL Classification: D72, C73.

*Daron Acemoglu gratefully acknowledges financial support from the National Science Foundation.

1 Introduction

In many political economy problems, including models of legislative bargaining (e.g., Baron, Diermeier, and Fong, 2008, Battaglini and Coate, 2007, 2008, Diermeier and Fong, 2009, Duggan and Kalandrakis, 2007, Romer and Rosenthal, 1978) and models of political coalition formation (e.g., Acemoglu, Egorov, and Sonin, 2008), agents participate in multiple rounds of proposal making and strategic voting. Noncooperative game-theoretic formulations of these models typically specify the extensive form of the game, in particular, the procedures through which proposals are made and the voting protocols (e.g., whether voting is sequential or simultaneous). These “details” as well as the exact notion of equilibrium often have a major impact on the set of equilibria and may introduce “non-intuitive” equilibria in many of these games. This is a significant challenge for applied analysis in this area. It would be desirable for features of the game that appear unimportant and difficult to map to reality (such as the exact order of proposals and the details of the voting protocol) not to have a major impact on predictions, and for non-intuitive equilibria not to emerge as predictions. Therefore, it appears important for political economy theory to have equilibrium concepts that (i) can be applied to a wide variety of political situations; (ii) rule out non-intuitive equilibria; and (iii) make broad predictions that are independent of the procedural details specified in the extensive form of the political game. The next example illustrates how non-intuitive equilibria can emerge in even the simplest political games.

Example 1 Consider three individuals, each strictly preferring option a to b . One of these two options will be implemented by voting. Voting is simultaneous; the option that receives the majority will be implemented. It is straightforward to verify that the natural equilibrium, in which all three individuals vote for option a , is a Nash equilibrium. However, the non-intuitive possibility that all three individuals vote for option b is also a Nash equilibrium. When any two players vote for b , it is a weak best response for the third one to do so as well. Naturally, we do not think that all three individuals supporting option b is a reasonable prediction.

The non-intuitive prediction implied by the notion of Nash equilibrium in the previous example is generally ignored, either by using one of two refinements, Nash equilibria in (weakly) undominated strategies or trembling-hand perfect equilibria (THPE), or by changing the game so that instead of simultaneous voting there is sequential voting. It is evident why the first approach would work as voting for b is a weakly dominated strategy for each player. The second approach also works for relatively intuitive reasons: with sequential voting, the game becomes

dynamic, and it is easy to show that in any subgame-perfect equilibrium (SPE) at least two players vote for a , so option a is chosen. Moreover, this result applies regardless of the order in which the three individuals vote. Nevertheless, sequential voting has the potential disadvantage that in practice most votes are not cast sequentially, but simultaneously (and with secret ballots). Therefore, this approach “works,” though it does so by changing the environment to one that is a worse approximation to the situation being modeled. Moreover, this discussion highlights that the set of predictions might crucially depend on seemingly unimportant procedural details such as whether voting is sequential or simultaneous.

In addition, while these two refinements eliminate non-intuitive equilibria in static voting games, many interesting political economy problems, such as those mentioned at the beginning, involve dynamic proposal making and voting. In this case, these refinements do not necessarily rule out non-intuitive equilibria. This is illustrated in the next example.

Example 2 Consider the same environment as in Example 1, except that voting takes place at two dates, 1 and 2. If a unanimity of voters prefers one of the two options in period 1, then this option is implemented. Otherwise, there is a majority vote at date 2, identical to that in Example 1. In addition, we assume that delay is costly. Now because the game is multi-stage, the relevant equilibrium concept is SPE. It can be verified that all three individuals voting for option b in both dates is an SPE. This is because all three players voting for b is an equilibrium in the subgame starting from date 2. Then, expecting option b being implemented at the date 2, all players would be best responding by voting for b at date 1. It can also be verified that eliminating weakly dominated strategies does not change this outcome, since voting for option b at date 1 is not weakly dominated because of costly delay.

Example 2 shows that simply eliminating weakly dominated strategies is not sufficient to resolve the problems that arise in dynamic voting games. Nevertheless, it can also be used to illustrate that the sequential elimination of weakly dominated strategies is typically more powerful. For example, the strategy “vote for b in both periods” is weakly dominated by the strategy “vote for b in period 1, vote for a in period 2”, and the strategy “vote for a in period 1, for b in period 2” is weakly dominated by “vote for a in both periods”. After these weakly dominated strategies are eliminated, the strategy “vote for b in period 1 and a in period 2” becomes weakly dominated by “vote for a in both periods” in the reduced game. Note that this process of sequential elimination requires that all weakly dominated strategies are eliminated simultaneously.¹ We can strengthen the point even further and show that even allowing for

¹See Duggan (2003) for a discussion of problems with this equilibrium refinement.

sequential elimination of weakly dominated strategies may not always rule out non-intuitive equilibria (see Example 5 in the Appendix).

As already mentioned above, a powerful equilibrium concept that often eliminates non-intuitive predictions is trembling-hand perfect equilibria (THPE). However, the use of trembling-hand perfection concept in dynamic games requires some caution (e.g., Osborne and Rubinstein, 1994, pp. 246-252). One way to define the THPE in dynamic games is to allow for correlated trembles if a player has a move at different stages of the game; a more restrictive generalization is to focus on uncorrelated trembles. Example 6 in the Appendix shows that neither of these generalizations is sufficient to eliminate non-intuitive equilibria in certain dynamic voting games.

Motivated by these problems, this paper introduces two simple, easy-to-use and intuitive equilibrium concepts for agenda-setting games: the *Sequentially Weakly Undominated Equilibrium (SWUE)* and *Markov Trembling Hand Perfect Equilibrium (MTHPE)*. We refer to a perfect-information game as *an agenda-setting game* if at each stage of this game either one of the players (or, possibly, Nature) moves, or there is a voting among a subset of players (a formal definition is provided below). Dynamic voting games, in which there are several stages of voting over proposals, are a special case of agenda-setting games. In addition, most dynamic multilateral bargaining games can be cast as agenda-setting games. In this context, SWUE is an intuitive generalization of the idea of eliminating weakly dominated strategies, and uses this procedure iteratively starting from the terminal subgames. We prove existence of SWUE in finite games. One disadvantage of SWUE is that it is difficult to generalize or apply to infinite games. Another, and a more serious one, is that it does not necessarily rule out all non-intuitive equilibria, as Example 5 in the Appendix shows. Our alternative notion, MTHPE, which focuses on Markovian trembles, has three desirable features. First, it eliminates all non-intuitive equilibria in agenda-setting games. Second, it can be applied easily and straightforwardly for a larger class of games, including infinite games with finite number of non-equivalent subgames. Third, we will provide a tight characterization of the structure of MTHPE in agenda-setting games and show that several “details,” such as whether voting is simultaneous or sequential, do not affect the predictions in these games. In particular, we will show that the set of MTHPE of agenda-setting games with simultaneous voting is identical to the set of MTHPE (and of SPE) of the equivalent sequential voting games. This result is also useful, since games with sequential voting, though admittedly less realistic, allow a straightforward characterization of SPEs (and we will show of MTHPEs) by backward induction.²

²We are not aware of any other equilibrium concept or refinement that has these three desirable features.

We will first establish the existence of MTHPE for a broad class of games, including all agenda-setting games, and then provide a characterization of the set of MTHPE for agenda-setting games. While MTHPE and SWUE are not subsets of one another in general, we show that for agenda-setting games (and, more broadly, for finite games) an MTHPE is always a SWUE (so within this class of games MTHPE is a stronger concept).

Our main characterization result as well as our result on the equivalence of the set of equilibria in simultaneous and sequential voting games rely on the notion of *regularization* which we introduce. A regularization replaces each stage of a dynamic game where multiple players vote simultaneously by a sequence of actions (stages) in which one player moves at a time. This implies that whenever several players move simultaneously, there are many regularizations. However, all the resulting games have the same set of equilibrium payoffs. A key part of our characterization result is that the set of MTHPE of the initial (simultaneous voting) game is equivalent to the set of SPE of *any* regularization. This result both implies that which regularization is chosen is immaterial and shows that computing the set of MTHPE is generally quite straightforward and can be done by choosing an arbitrary regularization (an arbitrary sequence of voting stages corresponding to the original game). Note, however, that the Markovian trembles restriction in MTHPE is critical for these results: Example 6 in the Appendix shows that the above correspondence between the equilibria of simultaneous voting game and its regularization breaks down when we use the standard THPE concept instead of MTHPE.

Finally, we show how either or both concepts can be applied to a dynamic voting model proposed and analyzed in Acemoglu, Egorov and Sonin (2008) under sequential voting. This game can be interpreted as one of elimination (as in Soviet Politburo) or as a game of dynamic coalition formation, similar to previous pioneering work in this area, Roberts (1999) and Barbera, Maschler, and Shalev (2001). In Roberts (1999), the club membership changes subject to approval by a majority of the current members; in Barbera, Maschler, and Shalev (2001), any member of the club might unilaterally admit a new member to the club.³ Although Section 4 applies our main results to a game of dynamic coalition formation, our equilibrium concept can also be straightforwardly applied to dynamic legislative bargaining games. For example, recent interesting work by Diermeier and Fong (2009) studies legislative bargaining with persistent agenda-setting power and uses a solution concept that is a special case of MTHPE. Applying MTHPE in their model would give identical results.

³See also Jack and Lagunoff (2006) and Granot, Maschler, and Shalev (2002), where unanimity is required for admission to a club.

Although the literature on voting games is vast (see Austen-Smith and Banks, 1999, for a modern treatment), the notions of MTHPE and SWUE are new. Two previous contributions are particularly noteworthy. Moulin (1979) pioneered application of dominance solvability to political games. Duggan (2003) discusses the possibility of non-intuitive equilibria in political games, suggests several solutions and also contains a list of related open problems.

The rest of the paper is organized as follows. Section 2 introduces the concepts of Markov Trembling-Hand Perfect Equilibrium and Sequentially Weakly Undominated Equilibrium. Section 3 contains the main results, while Section 4 provides an application of MTHPE. Section 5 concludes and the Appendix contains two additional examples.

2 Setup

2.1 Setup and Notation

Consider a general n -person T -stage game (where T is a natural number or ∞), where each individual can take an action at every stage. Let the action profile of each individual i be

$$a^i = (a_1^i, \dots, a_T^i) \text{ for } i = 1, \dots, n,$$

with $a_t^i \in A_t^i$ and $a^i \in A^i = \prod_{t=1}^T A_t^i$. Let $h_t = (a_1, \dots, a_t)$ be the history of play up to stage t (not including stage t), where $a_s = (a_s^1, \dots, a_s^n)$, so h_0 is the history at the beginning of the game, and let H_t be the set of histories h_t for $t : 0 \leq t \leq T - 1$. We denote the set of all potential histories up to date t by $H^t = \bigcup_{s=0}^t H_s$. Let t -continuation action profiles be

$$a^{i,t} = (a_t^i, a_{t+1}^i, \dots, a_T^i) \text{ for } i = 1, \dots, n,$$

with the set of continuation action profiles for player i denoted by $A^{i,t}$. Symmetrically, define t -truncated action profiles as

$$a^{i,-t} = (a_1^i, a_2^i, \dots, a_{t-1}^i) \text{ for } i = 1, \dots, n,$$

with the set of t -truncated action profiles for player i denoted by $A^{i,-t}$. We also use the standard notation a^i and a^{-i} to denote the action profiles for player i and the action profiles of all other players (similarly, A^i and A^{-i} are used for the sets of actions). The payoff functions for the players depend only on actions, i.e., player i 's payoff is given by

$$u^i(a^1, \dots, a^n).$$

We assume that even if T is infinite, the game admits only a finite set of payoff-relevant classes of histories. We also define the restriction of the payoff function u^i to a continuation play

$(a^{1,t}, \dots, a^{n,t})$ as $u^i \left(a^{1,-t}, \dots, a^{n,-t} : a_t^1, \dots, a_t^n : a^{1,t+1}, \dots, a^{n,t+1} \right)$. This function specifies the utility of player i when the action profile is $(a^{1,-t}, \dots, a^{n,-t})$ up to and including time $t-1$; then, it is (a_t^1, \dots, a_t^n) at time t , and then finally, it is restricted to be $a^{1,t}, \dots, a^{n,t}$ from t onwards. Symmetrically, this payoff function can also be read as the utility from continuation action profile $(a^{1,t+1}, \dots, a^{n,t+1})$ given that up to time t , the play has consisted of the action profile $(a^{1,-t}, \dots, a^{n,-t} : a_t^1, \dots, a_t^n)$.

A (possibly mixed) strategy for player i is

$$\sigma^i : H^{T-1} \rightarrow \Delta(A^i),$$

where $\Delta(X)$ denotes the set of probability distributions defined over the set X .

Denote the set of strategies for player i by Σ^i . A t -truncated strategy for player i (corresponding to strategy σ^i) specifies plays only until time t (including time t), i.e.,

$$\sigma^{i,-t} : H^{t-1} \rightarrow \Delta(A^{i,-t}).$$

The set of truncated strategies is denoted by $\Sigma^{i,-t}$. A t -continuation strategy for player i (corresponding to strategy σ^i) specifies plays only after time t (including time t), i.e.,

$$\sigma^{i,t} : H^{t-1} \setminus H^{t-2} \rightarrow \Delta(A^{i,t}),$$

where $H^{t-1} \setminus H^{t-2}$ is the set of histories until time t or longer.

With a slight abuse of notation, we will also use the same utility function defined over strategies (as actions) and write

$$u^i(\sigma^{i,t}, \sigma^{-i,t} | h^{t-1})$$

to denote the continuation payoff to player i after history h^{t-1} when he uses the continuation strategy $\sigma^{i,t}$ and other players use $\sigma^{-i,t}$. We also use the notation $u^i \left(\sigma^{1,t}, \dots, \sigma^{n,t} : \sigma^{1,t+1}, \dots, \sigma^{n,t+1} | h^{t-1} \right)$ as the payoff from strategy profile $(\sigma^{1,t}, \dots, \sigma^{n,t})$ at time t restricted to the continuation strategy profile $(\sigma^{1,t+1}, \dots, \sigma^{n,t+1})$ from $t+1$ onwards, given history h^{t-1} . Similarly, we use the notation

$$u^i(a^{i,t}, a^{-i,t} | h^{t-1})$$

for the payoff to player i when he chooses the continuation action profile $a^{i,t}$ and others choose $a^{-i,t}$ given history h^{t-1} .

We start by providing the standard definitions of Nash equilibria and subgame perfect Nash equilibria.

Definition 1 A strategy profile $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ is a *Nash Equilibrium* if

$$u^i(\hat{\sigma}^i, \hat{\sigma}^{-i}) \geq u^i(\sigma^i, \hat{\sigma}^{-i}) \text{ for all } \sigma^i \in \Sigma^i \text{ and for all } i = 1, \dots, n.$$

Definition 2 A strategy profile $(\hat{\sigma}^1, \dots, \hat{\sigma}^N)$ is a *Subgame Perfect Nash Equilibrium* if

$$u^i(\hat{\sigma}^{i,t}, \hat{\sigma}^{-i,t} \mid h^{t-1}) \geq u^i(\sigma^{i,t}, \hat{\sigma}^{-i,t} \mid h^{t-1}) \text{ for all } h^{t-1} \in H^{t-1},$$

for all t , for all $\sigma^{-i} \in \Sigma^{-i}$ and for all $i = 1, \dots, n$.

Sequentially Weakly Undominated Equilibrium

Towards introducing weakly dominant strategies, let us take a small digression and consider a one stage game with actions (a^1, \dots, a^n) .

Definition 3 We say that $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ is a *weakly undominated equilibrium* if for each $i = 1, \dots, n$, for any $\sigma^i \in \Sigma^i$ we have

$$u^i(\hat{\sigma}^i, \hat{\sigma}^{-i}) \geq u^i(\sigma^i, \hat{\sigma}^{-i})$$

(i.e., it is a Nash equilibrium) and for each $i = 1, \dots, n$, there does not exist $\tilde{\sigma}^i \in \Sigma^i$ such that

$$u^i(\tilde{\sigma}^i, \sigma^{-i}) \geq u^i(\hat{\sigma}^i, \sigma^{-i})$$

for all $\sigma^{-i} \in \Sigma^{-i}$ with at least one strict inequality.

Naturally, such equilibrium always exists in static finite games. Let us next extend it to the general T -stage game for T finite. A weakly undominated strategy equilibrium in this last stage of the game is defined analogously to Definition 3.

Definition 4 Take any $t : 1 \leq t \leq T + 1$. Strategy profile $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ is a *h^{t-1} -sequentially weakly undominated equilibrium* if $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ is a h^t -sequentially weakly undominated equilibrium for any h^t that may occur after h^{t-1} and for each $i = 1, \dots, n$, for any $\sigma^{i,t} \in \Sigma^{i,t}$,

$$u^i\left(\hat{\sigma}^{i,t}, \hat{\sigma}^{-i,t} : \hat{\sigma}^{1,t+1}, \dots, \hat{\sigma}^{N,t+1} \mid h^{t-1}\right) \geq u^i\left(\sigma^{i,t}, \hat{\sigma}^{-i,t} : \hat{\sigma}^{1,t+1}, \dots, \hat{\sigma}^{N,t+1} \mid h^{t-1}\right) \quad (1)$$

and there does not exist $\tilde{\sigma}^{i,t} \in \Sigma^{i,t}$ such that

$$u^i\left(\tilde{\sigma}^{i,t}, \sigma^{-i,t} : \hat{\sigma}^{1,t+1}, \dots, \hat{\sigma}^{N,t+1} \mid h^{t-1}\right) \geq u^i\left(\hat{\sigma}^{i,t}, \sigma^{-i,t} : \hat{\sigma}^{1,t+1}, \dots, \hat{\sigma}^{N,t+1} \mid h^{t-1}\right) \quad (2)$$

for all $\sigma^{-i,t} \in \Sigma^{-i,t}$ with at least one strict inequality.

Definition 4 is inductive, but this induction is finite. Indeed, if history h^{t-1} leads to a terminal node, then the first condition is satisfied automatically, because the history will not be recorded any further. In words, we first hypothesize that there exists a h^t -sequentially weakly undominated equilibrium, impose that it will be played from time $t + 1$ onwards and then look for a weakly dominant strategy profile at stage t of the game.

Definition 5 Strategy profile $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ of a finite extensive-form game is a *Sequentially Weakly Undominated Equilibrium (SWUE)* if it is a h^0 -sequentially weakly undominated equilibrium.

Markov Trembling-Hand Perfect Equilibria

We now introduce the notion of Markov Trembling-Hand Perfect Equilibria (MTHPE). For this, we need to define a Markovian strategy (Maskin and Tirole, 2001).

Definition 6 A continuation strategy $\sigma^{i,t}$ is *Markovian* if

$$\sigma^{i,t}(h^{t-1}) = \sigma^{i,t}(\tilde{h}^{t-1})$$

for all $h^{t-1}, \tilde{h}^{t-1} \in H^{t-1}$ such that for any $a^{i,t}, \tilde{a}^{i,t} \in A^{i,t}$ and any $a^{-i,t} \in A^{-i,t}$

$$u^i(a^{i,t}, a^{-i,t} | h^{t-1}) \geq u^i(\tilde{a}^{i,t}, a^{-i,t} | h^{t-1})$$

implies that

$$u^i(a^{i,t}, a^{-i,t} | \tilde{h}^{t-1}) \geq u^i(\tilde{a}^{i,t}, a^{-i,t} | \tilde{h}^{t-1}).$$

Definition 7 We say that a strategy profile $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ of an extensive-form game in agent-strategic form is *Markov Trembling-Hand Perfect Equilibrium (MTHPE)* if there exists a sequence of totally mixed Markovian strategy profiles $\{(\hat{\sigma}^1(m), \dots, \hat{\sigma}^n(m))\}_{m \in \mathbb{N}}$ (meaning that continuation strategy $\sigma^{i,0}(m)$ is Markovian for all $i = 1, \dots, n$ and all $m \in \mathbb{N}$) such that $(\hat{\sigma}^1(m), \dots, \hat{\sigma}^n(m)) \rightarrow (\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ as $m \rightarrow \infty$ and

$$u^i(\hat{\sigma}^i, \hat{\sigma}^{-i}(m)) \geq u^i(\sigma^i, \hat{\sigma}^{-i}(m)) \text{ for all } \sigma^i \in \Sigma^i, \text{ for all } m \in \mathbb{N} \text{ and for all } i = 1, \dots, n.$$

We say that a strategy profile $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ of an extensive-form game in strategic form is *Markov Trembling-Hand Perfect Equilibrium (MTHPE)* if it is MTHPE in the corresponding agent-strategic form game.

Note that MTHPE is defined directly in the agent-strategic form in order to avoid standard problems that arise when trembling hand perfection is defined on the strategic form (e.g., Selten,

1975, Osborne and Rubinstein, 1994, pp. 246-252). Naturally, MTHPE is a stronger concept than THPE.⁴ Nevertheless, we establish its existence under general assumptions (Theorem 1). In a class of political games, which we refer to as *agenda-setting games*, we prove that a pure-strategy MTHPE exists, and moreover every MTHPE is a SWUE (Theorem 2). Nevertheless, these two equilibrium concepts do not always coincide. First, a MTHPE always exists, while SWUE may not. Second, in an infinite game, there may exist SWUEs that are not MTHPE.

2.2 Agenda-Setting Games

Let us first define general agenda-setting games, a class of extensive form games with perfect information that includes most voting games as a special case.⁵ In Section 3, we establish the existence of MTHPE and SWUE for these games.

To define a general agenda-setting game, we first need a definition of a stage. Inside one stage, information sets are possible as the definition should be general enough to nest simultaneous voting.

Definition 8 A game Γ in extensive form with a set of players N is called an *agenda-setting game* if it satisfies the following conditions.

1. Game Γ consists of a (possibly infinite) number of stages, where each stage Θ is either:
 - (a) a “proposing stage”, i.e., a decision node and a number of possible actions of a single player (possibly Nature), the agenda-setter at stage Θ .
 - (b) a “voting stage”, i.e., a connected subgraph of Γ , in which each player $i \in N$ has at most one decision node and two actions at this node, say $a_i^y(\Theta)$ and $a_i^n(\Theta)$.
2. For each voting stage Θ , there are only two equivalence classes of continuation-payoff-identical subgames following terminal nodes of Θ in game Γ , say $y(\Theta)$ and $n(\Theta)$.
3. In each voting stage Θ , for each player $i \in N$ and for any other players’ actions held fixed, action $a_i^y(\Theta)$ does not decrease the probability of moving into a subgame that belongs to the equivalence class $y(\Theta)$.

⁴Another trembling hand refinement used in the literature, *truly perfect equilibrium*, is stronger than our notion of MTHPE. A truly perfect equilibrium requires strategies from σ to be best responses to all fully mixed profiles in some neighborhood of σ rather than to any one sequence of profiles in the standard notion of trembling-hand perfection and to one sequence of Markovian profiles in the case of MTHPE. However, this equilibrium concept fails to exist in many games, including our dynamic game of coalition formation (except in some special cases).

⁵The study of agenda-setting games was pioneered by Romer and Rosenthal (1978) and Baron and Ferejohn (1989).

4. For any two nodes ξ and ξ' of game Γ , if ξ and ξ' belong to one information set, then there exists a voting stage Θ' of game Γ that contains both ξ and ξ' .

This definition states that any game in which one of the agents makes a proposal and others vote in favor or against this proposal is an agenda-setting game (and imposes the intuitive restriction that more “yes” votes do not make a proposal less likely to succeed). Agenda-setting games naturally cover all dynamic voting games in which at some stages proposals are made, and then subsequently, a subset of, or all, players vote in favor or against these proposals. They also include several dynamic bargaining games, where a division of a pie is offered at a certain stage, and then this needs to be accepted by the other participants acting simultaneously or in sequence, as special cases.

3 Characterization

3.1 Existence Results

We start by observing that MTHPE exists in a very broad class of games. We later prove that it is, nonetheless, a sufficiently strong equilibrium refinement to rule out all non-intuitive equilibria in agenda-setting games.

Theorem 1 *Any finite or infinite extensive-form game with a finite number of stages has a MTHPE (possibly in mixed strategies).*

Proof. Suppose first that T is finite. Consider a perturbation of the original game where each player $i \in N$ is restricted to play each action at each stage with probability $\eta_{k,t}^i \geq \underline{\eta} > 0$, where $\underline{\eta}$ is sufficiently small. By the standard fixed point theorem argument, this perturbed game has a Nash equilibrium; moreover, the fixed point theorem applies if we restrict our attention to Markovian strategies only. Therefore, the perturbed game has a Markov Perfect Equilibria $(a^1(\underline{\eta}), \dots, a^n(\underline{\eta}))$. Because the action space has finite dimensions and is thus compact, we can choose a sequence $\underline{\eta}^1, \underline{\eta}^2, \dots$ which converges to 0 such that $(a^1(\underline{\eta}^k), \dots, a^n(\underline{\eta}^k))$ has a limit. This limit would be a trembling-hand perfect equilibrium in Markovian strategies, i.e., an MTHPE. ■

If T is infinite, the previous reasoning applies straightforwardly if instead of “stages” we use “classes of stages with payoff-equivalent histories”.

Theorem 2 *1. Any (finite) agenda-setting game has a MTHPE in pure strategies.*

2. In any finite game, a MTHPE is a SWUE.

3. Any finite agenda-setting game has a SWUE in pure strategies.

Proof. (1) We proceed by induction on the number of stages. The base is trivial; indeed, consider a one-shot agenda-setting game. If this stage is one-player move, then, evidently, the action which maximizes his utility constitutes a pure-strategy MTHPE, since it is Markovian and trembling-hand perfect. If the single stage is voting, then each player i weakly prefers one of the outcomes to another. It is straightforward to verify that voting for a weakly preferred outcome is an MTHPE.

We next proceed with the induction step. Suppose that we have proved the existence of pure-strategy MTHPE in all agenda-setting games with number of stages less than T ; take an agenda-setting game with T stages. Consider its first stage. We start with the case where there is one player i making a choice between k actions a_1, \dots, a_k . In each of k corresponding subgames there exists (by the induction hypothesis) a pure strategy MTHPE; we can choose the same MTHPE for isomorphic subgames. Therefore, there exist sequences of strategy profiles $\sigma_j^1, \sigma_j^2, \dots$ for each $j \in \{1, \dots, k\}$ which converge to a pure strategy MTHPE for each such j and which are Nash equilibria in the corresponding $T - 1$ -stage games. Moreover, we can require that if two subgames (for $j = j_1, j_2$) are isomorphic, then $\sigma_{j_1}^n = \sigma_{j_2}^n$ for any n . Now, consider player i at first stage. For any n , there is j_n such that a_{j_n} is weakly better than other actions. Since there are a finite number of actions, there is action a_j that is weakly better than other actions for infinitely many values of n . It is now straightforward to prove that a_j , along with the chosen MTHPE, forms a pure strategy MTHPE of the whole game.

Now consider the case where the initial stage is a voting stage. Then there are, essentially, (at most) two subgames; take a pure-strategy MTHPE in each (if they are isomorphic, take the same MTHPE), and we can similarly construct two sequences of strategy profiles $\sigma_y^1, \sigma_y^2, \dots$ and $\sigma_n^1, \sigma_n^2, \dots$. Each player i has two actions, y and n , in stage 1. Consider one of the players (i) and take any m . If in the subgames strategies σ_y^m and σ_n^m are played, respectively, then player i weakly prefers to choose one of the actions, y or n , to another one. There is an action $a(i)$ which is weakly preferred for infinite number of m 's. Now consider two subsequences of sequences $\sigma_y^1, \sigma_y^2, \dots$ and $\sigma_n^1, \sigma_n^2, \dots$ which are formed by values of m for which $a(i)$ is weakly better for player i . Take another player, j , and repeat the procedure; then we will get action $a(j)$ and two subsequences of the previous subsequences. Since the number of players is finite, we can proceed in this way, and for any player i find action $a(i)$. It is now evident that the MTHPEs chosen for the stages starting from the second one and actions $a(i)$ for each player i at the first stage form a pure strategy MTHPE.

(2) Take any strategy profile σ that forms a MTHPE. This is proved by backward induction on the number of stages in the game. Suppose that the Theorem has been proved for games with $q' < q$ stages. Consider an agenda-setting game with q stages and take any MTHPE in it. By induction, this MTHPE, when truncated to any of the game's proper subgames, forms a SWUE. Consider its first stage.

Suppose that only one player i moves at this stage; denote his expected utility (in this MTHPE) from taking action a at first stage by u_i^a . If action a^* is an action played with a non-zero probability in equilibrium then $u_i^{a^*} \geq u_i^a$ for any other feasible action a (otherwise there would exist a payoff-improving deviation). Hence, all actions played in a MTHPE with a non-zero probability yield the same expected utility for player i , and this utility is the maximum possible over the set of feasible actions. Hence, this MTHPE is a SWUE.

Now consider the other situation where the first stage is a voting stage. Consider a profile σ' consisting of fully mixed strategies and suppose that it is η -close to σ for a small η . Depending on how other players vote, three mutually exclusive situations are possible: proposal is accepted regardless of how player i votes, it is rejected regardless of how he votes, and player i is pivotal; let μ^+ , μ^- , and μ^p be the respective probabilities of these events. By definition, $\mu^+ + \mu^- + \mu^p = 1$, and by assumption $\mu^p > 0$. Voting for the proposal yields $(\mu^+ + \mu^p)u_i^{+'} + \mu^-u_i^{-'}$ in expectation, voting against it yields $\mu^+u_i^{+'} + (\mu^- + \mu^p)u_i^{-'}$ where $u_i^{+'}$ and $u_i^{-'}$ are i 's utilities from acceptance and rejection of the proposal if profile σ' is played. Thus, if $u_i^{+'} > u_i^{-'}$, then player i 's sole best response is voting for the proposal, and if $u_i^{+'} < u_i^{-'}$ it is voting against it. If η is sufficiently small then $u_i^+ > u_i^-$ implies $u_i^{+'} > u_i^{-'}$, and thus by definition of MTHPE player i must support the proposal in equilibrium with probability one. Similar reasoning applies to the case $u_i^+ < u_i^-$.

Now take any player i who participates in voting. If $u_i^+ > u_i^-$, then he votes for the proposal in this MTHPE. This is a weakly dominant strategy for him (given continuation strategies of himself and other players). Similarly, if $u_i^+ < u_i^-$, then the strategy he plays in this MTHPE is weakly dominant. If, $u_i^+ = u_i^-$ or the player is never pivotal, then any strategy is weakly undominated. Therefore, for any player, the strategy he plays in this MTHPE is weakly undominated, and thus this MTHPE is a SWUE. This completes the induction step.

(3) This follows from (1) and (2) of this Theorem. ■

Existence of an MTHPE equilibrium in pure strategies is at first a surprising result and will be exploited below for a full characterization of these equilibria in agenda-setting games (existence of an SWUE in pure strategies, which can be computed by a backward-induction procedure, is obvious).




Figure 3: A Game With SWUE which is not MTHPE

3.2 Relationship Between MTHPE and SWUE

However, SWUE and MTHPE concepts are not equivalent, and neither is in general stronger than the other (though see Theorem 4 below).

Example 3 There exists an agenda-setting game where an SWUE is not an MTHPE. Consider a game of two players with extensive form depicted on Figure 3. This is an agenda-setting game, because at each stage only one player has a (non-trivial) move. It game has exactly one MTHPE (R, r) . However, there are two SWUEs: (R, r) and (L, r) . The latter is not MTHPE, because if there is a non-zero chance that player 2 will play l , player 1 is better off choosing R .

As an alternative to the concept of SWUE, one might consider Sequentially Weakly Dominant Equilibria, defined as follows:

Definition 9 Take any $t : 1 \leq t \leq T + 1$. Strategy profile $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ is a h^{t-1} -sequentially weakly dominant equilibrium if $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ is a h^t -sequentially weakly dominant equilibrium for any h^t that may occur after h^{t-1} and

$$u^i \left(\hat{\sigma}^{i,t}, \sigma^{-i,t} : \hat{\sigma}^{1,t+1}, \dots, \hat{\sigma}^{N,t+1} \mid h^{t-1} \right) \geq u^i \left(\sigma^{i,t}, \sigma^{-i,t} : \hat{\sigma}^{1,t+1}, \dots, \hat{\sigma}^{N,t+1} \mid h^{t-1} \right)$$

for all $\sigma^{i,t} \in \Sigma^{i,t}$, for all $\sigma^{-i,t} \in \Sigma^{-i,t}$, and for all $i = 1, \dots, N$.

(Here, we use the convention that any profile is h^T -sequentially weakly dominant equilibrium.) Strategy profile $(\hat{\sigma}^1, \dots, \hat{\sigma}^n)$ is a Sequentially Weakly Dominant Equilibrium (SWDE) if it is a h^0 -sequentially weakly dominant equilibrium.

Note that we do not require that the inequality holds strictly for at least one $\sigma^{i,t}, \sigma^{-i,t}$ in order to ensure existence in situations where one player is indifferent between different outcomes. Clearly, SWDE is a much stronger equilibrium concept than SWUE. Nevertheless, it can be shown that for agenda-setting games the notions of SWDE and SWUE coincide, though naturally for general finite games SWDE is a “strictly” stronger concept and may fail to exist. More precisely, we can establish the following result.

Theorem 3 1. *In any finite agenda-setting game, a SWUE is a SWDE (and thus a MTHPE is a SWDE). In particular, such games have SWDE in pure strategies.*

2. *In any finite game, a SWDE is a SWUE (so for agenda-setting games these notions coincide).*

3. *There exists a finite game where a MTHPE is not a SWDE.*

4. *There exists an agenda-setting game where a SWDE is not a MTHPE.*

Proof. (1) This is easily proven by induction. Take any SWUE σ . The base is trivial (essentially, it may be proved with exactly the same words as the induction step). Step: take any stage; if only a single player i acts, then both SWUE and SWDE allow for any action that is a best response of player i given that σ is played from that stage on. Suppose the stage involves voting; take some player i . If he is pivotal with probability zero, then any strategy would constitute a SWDE, in particular, the one he plays in profile σ . Now assume that he is pivotal with a positive probability. Consider first the case where i plays a pure strategy; then without loss of generality assume that SWUE prescribes that player i votes “for” the proposal. Denotes his continuation payoffs if the proposal is accepted and is rejected by u_i^+ and u_i^- , respectively. then in SWUE σ , $u_i^+ \geq u_i^-$ must hold (according to Definition 4; otherwise voting “for” would be a dominated strategy at this stage). But in that case, voting “for” is also a weakly dominant strategy: indeed, voting “against” is the only alternative action, and hence $u_i^+ \geq u_i^a$ for any action a . Finally, if σ prescribes a mixed strategy (and i is pivotal with a positive probability), it must be that $u_i^+ = u_i^-$. But then again, any player i ’s strategy at this stage forms a SWDE after the corresponding history. This proves the induction step.

(2) Again, this is proved by induction. Take any finite game and SWDE σ . The base is again trivial. Step: take any player i and suppose that he plays strategy $\hat{\sigma}^{i,t} \in \Sigma^{i,t}$. This implies, by definition, that

$$u_i(\hat{\sigma}^{i,t}, \sigma^{-i,t}) \geq u_i(\sigma^{i,t}, \sigma^{-i,t}) \quad (3)$$

(we drop conditioning on past and future actions for brevity) for any $\sigma^{i,t}$ and $\sigma^{-i,t}$. This immediately implies that condition (1) holds for any $\sigma^{i,t} \in \Sigma^{i,t}$. Now assume, to obtain a contradiction, that there is $\tilde{\sigma}^{i,t} \in \Sigma^{i,t}$ such that for at least one combination of other players’ strategies $\sigma^{-i,t} \in \Sigma^{-i,t}$ we have

$$u_i(\tilde{\sigma}^{i,t}, \sigma^{-i,t}) > u_i(\hat{\sigma}^{i,t}, \sigma^{-i,t})$$

(recall that Definition 4 requires that one of the inequalities (2) be strict). However, this last condition violates (3). This contradiction completes the proof.

(3) Consider a one-stage game with two players making simultaneous moves with payoff matrix

$$\begin{array}{cc} & \begin{array}{c} l \\ r \end{array} \\ \begin{array}{c} L \\ R \end{array} & \begin{pmatrix} (1,1) & (0,0) \\ (0,0) & (1,1) \end{pmatrix} . \end{array}$$

This game does not have SWDE, because it is one-stage and in that only stage neither of the players has a weakly dominant strategy. It is straightforward to check, however, that both (L, l) and (R, r) are MTHPEs of this game.

(4) This follows from Example 3. Indeed, in this example the SWUE is also a MTHPE, as any SWUE is SWDE in an agenda-setting game. Hence, there is an agenda-setting game where SWDE is not MTHPE. ■

3.3 MTHPE in Agenda-Setting Games

In this subsection, we provide a characterization of the structure of MTHPE in agenda-setting games. We define the notion of a *regularization* of an agenda-setting game. A regularization of an agenda setting game transforms it into an alternative game in which at each stage only one player takes an action, i.e., stages where there are simultaneous moves are replaced by a sequence of stages corresponding to each one of these moves.

Definition 10 *Suppose that Γ is a (finite) agenda-setting game. We call game Γ' a regularization of game Γ if it is obtained from Γ by substituting every stage where more than one player moves by a sequence of stages such that each stage only one player takes an action. If there is voting among players $X \subset N$ at stage ξ , then ξ is substituted by $|X|$ stages in which players $i \in X$ move sequentially. Player i has the same two actions that he had at stage ξ , and the game proceeds with subgame $y(\xi)$ after profile $\{a_i(\xi)\}_{i \in X}$ is played if and only if the original game Γ proceeds with $y(\xi)$ after $\{a_i(\xi)\}_{i \in X}$ is played at stage ξ .*

Clearly there are several regularizations of an agenda-setting game that has simultaneous actions by two or more players at the same stage (and the regularization of an agenda-second game is itself an agenda-setting game, with the unique regularization that is itself). Crucially, however, each of these regularizations has the same set of terminal payoffs. We next show that different regularizations will also lead to the same set of equilibrium payoffs to all players.

To establish this result, first note that any strategy profile σ of game Γ can be naturally mapped to a strategy profile σ' of a regularization Γ' , so that players play the same actions as in profile σ . With a slight abuse of notation, we say that σ is a strategy profile of game Γ' . Notice that the converse of the statement is not true, since σ' may involve a pattern of

actions that cannot be captured by σ , for example, when later actions depend on earlier actions. Nevertheless, the next theorem shows a strong equivalence between the set of MTHPE of an agenda-second game and the set of MTHPE (and SPE) of any regularization of the same game.

Theorem 4 *Suppose Γ is a finite agenda-setting game.*

1. *If σ is a MTHPE of game Γ , then it is an MTHPE (and thus an SWUE and a SPE) in any regularization Γ' .*
2. *Conversely, if σ' is a MTHPE in regularization Γ' , then there exists an MTHPE σ in Γ in which all players obtain the same payoffs as in σ' .*

Proof. (1) Take MTHPE σ of game Γ . Denote the corresponding strategy profile in game Γ' by σ' to avoid confusion. Consider a sequence of fully mixed profiles $(\sigma^1(r), \dots, \sigma^n(r))$ which converge to the equilibrium profile $(\sigma^1, \dots, \sigma^n)$ as $r \rightarrow \infty$ and such that the $\sigma^{i,t}$ is a best response to $\sigma(r)$ for every player i at every stage or node ξ (this exists by the definition MTHPE). Since the game is finite, we may, without loss of generality, assume that at every voting stage every player that participates in the voting has the same preferences over the two alternatives y and n for all r 's in the sequence (i.e., we can always take such subsequence). This implies that at voting stage, each player either strictly prefers y , or strictly prefers n , or is indifferent between the two for all r .

Consider the corresponding profiles $(\sigma'^1(r), \dots, \sigma'^n(r))$ of the regularized game Γ' ; obviously, they are fully mixed and converge to σ' as $r \rightarrow \infty$. Let us verify that σ' is a best response to $(\sigma'^1(r), \dots, \sigma'^n(r))$ for every player i at any node ξ and for any r . If node ξ corresponds to a single action by some player in game Γ , then it is obvious, since the payoffs of all actions are the same, as the continuations of profiles $(\sigma'^1(r), \dots, \sigma'^n(r))$ and $(\sigma^1(r), \dots, \sigma^n(r))$ have the same distribution over the payoffs at terminal nodes. Consider the case where ξ corresponded to a voting in game Γ . If i were indifferent about the outcomes of the voting in Γ (provided that $(\sigma^1(r), \dots, \sigma^n(r))$ is played after that voting), then he is indifferent in the game Γ' as well, for the same reason as before. If he strictly preferred one of the outcomes (say, $y(\xi)$), then he must have picked it in game Γ (since he was pivotal with a positive probability since strategies were fully mixed). In the game Γ' , if profile $(\sigma'^1(r), \dots, \sigma'^n(r))$ is played, his choice of action $a_i^y(\xi)$ cannot decrease the likelihood that the corresponding subgame is reached after this voting, because other players' strategies in this (now sequential) voting do not depend on his action. Consequently, choosing $a_i^y(\xi)$ is still a best response for him. The case where player i strictly preferred outcome $n(\xi)$ is treated similarly. This establishes that profile σ' consists of

best responses to $(\sigma'^1(r), \dots, \sigma'^n(r))$ for all players, at all stages, and for all r 's, i.e., σ' is a MTHPE.

(2) Take any MTHPE σ' in regularization Γ' . Consider a sequence of fully mixed profiles $(\sigma'^1(r), \dots, \sigma'^n(r))$ which converge to the equilibrium profile $(\sigma'^1, \dots, \sigma'^n)$ as $r \rightarrow \infty$ and such that the $\sigma'^{i,t}$ is a best response to $\sigma'(r)$ for every player i at every stage or node ξ . As in part 1, we may, without loss of generality, assume that at every voting stage every player who participates in the voting has the same preferences over the two alternatives y and n for all r 's in the sequence (note that these preferences are well-defined, since σ' must consist of Markovian strategies).

For any voting node ξ in game Γ , take the probability $p(\xi)$ that the proposal will be accepted in the corresponding voting in game Γ' under profile σ' . The probability is taken conditional on reaching this node; note that this probability is well defined in the sense that it is the same for all votings in Γ' that may correspond to node ξ in Γ since σ' consists of Markovian strategies. One can easily check the following statements about votings in Γ' . If $p(\xi) < 1$, then the set of players who strictly prefer $y(\xi)$ to $n(\xi)$ (for all r 's) is not sufficient for $y(\xi)$ to be reached even if they all voted $a_i^y(\xi)$. Similarly, if $p(\xi) > 0$, then the set of players who strictly prefer $n(\xi)$ to $y(\xi)$ is not sufficient for $n(\xi)$ to be reached. Moreover, if $p(\xi) \in (0, 1)$, then there is a player who is indifferent between $y(\xi)$ and $n(\xi)$ subgames for all r .

Let us now define strategy profile σ as follows. At a node where player i makes a single decision, let him choose the same action as in corresponding nodes of game Γ' (there, he chooses the same action as σ' is Markovian). At a voting node ξ with $p(\xi) = 0$, let all players who prefer n to y or were indifferent (for all r above) choose $a_i^n(\xi)$, and let the rest choose $a_i^y(\xi)$. If $p(\xi) = 1$, let all players who prefer y to n or are indifferent choose $a_i^y(\xi)$ and let the rest choose $a_i^n(\xi)$. Finally, if $p(\xi) \in (0, 1)$, then let all players who strictly preferred y choose $a_i^y(\xi)$, all players who strictly preferred n choose $a_i^n(\xi)$, and let those who were indifferent mix in such a way that the probability of acceptance is $p(\xi)$ (this is possible, since neither of the two groups with strict preferences is sufficient to enforce either of the decisions).

Let us verify that σ is a MTHPE. Observe that for any r we can take a fully mixed profile $(\sigma^1(r), \dots, \sigma^n(r))$ which converges to σ and that for every subgame in Γ have the same payoffs for all players as profiles $(\sigma'^1(r), \dots, \sigma'^n(r))$ chosen above. Then at each node of Γ players play best responses, due to the way we defined profile σ . This completes the proof. ■

This result establishes the equivalence of agenda-setting games with simultaneous and sequential votings once a proper equilibrium concept (MTHPE) is used. More importantly, it also highlights the importance of MTHPE as an equilibrium refinement. Example 6 shows that

part (1) of Theorem 4 would not hold if we considered THPE instead of MTHPE. Example 5 in the Appendix shows, among other things, that it would also not hold if we considered sequential elimination of weakly-dominated strategies (and in particular, with this notion certain non-intuitive equilibria still survive). Thus, the equivalence result for SWUE does not hold.

Finally, the next theorem establishes a “generic uniqueness” of MTHPE in agenda-setting games. Example 6 in the Appendix again shows that it would not hold for THPE.

Theorem 5 *Suppose agenda-setting game Γ is such that for any two terminal nodes, if one player obtains the same payoffs in these two nodes, then each player obtains the same payoffs in both nodes. Then in any two MTHPEs σ_1 and σ_2 , all players obtain the same payoffs, and these payoffs are the same as those in any SPE of any regularization of Γ' .*

Proof. We prove the following statement by induction by the number of stages: (a) for any two MTHPEs σ_1 and σ_2 all players get the same payoffs, and (b) if at some MTHPE σ players get a certain expected payoff vector, then there exists a terminal node where players get these payoffs. Let us prove the stage (the base is trivial). If the first node of the game involves player i making a decision, then of all available actions he will choose one which maximizes his utility. If there are several such actions then, by induction, the payoff vectors of the corresponding subgames are realized at some terminal nodes of the subgames. Since player i is indifferent, he gets the same payoff, and by assumption all players get the same payoff. This proves part (a) of the statement; part (b) is trivial. Consider the case where the first stage involves voting. By induction, either all players are indifferent between the alternatives y and n , or all players have a strict preference. In the first case, all MTHPEs involve the same payoffs for trivial reasons, while in the second, the MTHPE strategies of players are unique, and hence the payoffs are the same. Again, part (b) is trivial. This proves the induction step.

To complete the proof, we need to show that these payoffs coincide with the payoffs in any SPE of any regularization. By Theorem 4, all players get these payoffs in any MTHPE of the game (and thus in some SPE of the game). It remains to prove that in a game where players move one at a time and where in any two terminal nodes either all players get the same payoffs or all players get different payoffs, the payoffs of all players in any two SPEs are the same. This proof may be done by induction, similarly to the one in the beginning of this proof. ■

4 Application: the Elimination Game

To demonstrate the usefulness of the MTHPE and SWUE concepts, we now use a version of the endogenous club formation game (or “political elimination” game) studied in Acemoglu, Egorov,

and Sonin (2008). This is an infinite-horizon dynamic game with the following structure: there is an initial collection of individuals, $N = \{1, \dots, n\}$, with each individual $i \in N$ endowed with political power $\gamma_i > 0$ (for example, representing individual i 's "guns"). At each date $t \geq 1$, there exists a ruling coalition $X_t \subset N$, and members of this ruling coalition may vote to exclude some of its members, thus determining a new ruling coalition $X_{t+1} \subset X_t$. At time $t = 0$, we have $X_0 = N$. The total power of any coalition $X \subset N$ is

$$\gamma_X \equiv \sum_{i \in X} \gamma_i.$$

We assume that powers $\{\gamma_i\}_{i \in N}$ are *generic* in the sense that if $X, Y \subset N$ and $X \neq Y$, then $\gamma_X \neq \gamma_Y$.

In any period $t \geq 1$, the voting proceeds as follows. An agenda-setter, which is determined by Nature at random (or through any fixed protocol) proposes an alternative coalition $A_{t,i}$ which is a subcoalition of the current coalition. All individuals who are entitled to do so vote for or against $A_{t,i}$. Alternative $A_{t,i}$ becomes the next ruling coalition, i.e., $X_{t+1} = A_{t,i}$, if and only if it receives an absolute majority of the available "weighted votes," where votes are weighted by the power of each member of the junta, so that an individual with a greater γ_i has proportionately more votes. If the proposal $A_{t,i}$ does not receive an absolute majority, then the next agenda-setter nominates a proposal and so on. In case no proposal is accepted, $X_{t+1} = X_t$.

The preferences (for each $i \in N$) consist of two parts. The first is utility from power,

$$U_{i,t}^+ = (1 - \beta) \mathbb{E}_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \frac{\gamma_i}{\gamma_{X_\tau}} \mathbf{I}_{i \in X_\tau}.$$

Here $\beta \in (0, 1)$ is the discount factor common across all individuals, $\mathbf{I}_{i \in X_\tau}$ is the indicator function for individual j being a member of the ruling coalition at time t , and \mathbb{E}_t denotes expectations at time t . The term γ_i/γ_{X_t} represents the power of the individual relative to other ruling coalition members. It can be motivated from the division of a unit size pie among the coalition members in proportion to their power. The important implication of this functional form is that each player obtains greater utility when the power of the ruling coalition is smaller. Therefore, each prefers to be a member of a smaller ruling coalition.⁶

Formally, the extensive form of the game is as follows. We assume that an agenda-setter is picked randomly by nature; yet all the results go through for any fixed protocol of agenda-setters.

1. At each stage, $j = 0, 1, \dots$, the game starts with an intermediary coalitions by $N_j \subset N$ (with $N_0 = N$).

⁶These payoffs are a special case of the more general payoffs introduced in Acemoglu, Egorov and Sonin (2008) and are adopted to simplify the notation here (without affecting any of the results).

2. Nature randomly picks agenda setter $i_{j,q} \in N_j$ for $q = 1$ (i.e., a member of the coalition N_j).
3. Agenda setter $i_{j,q}$ proposes a coalition $X_{j,q} \in P(N_j)$.
4. All players in $X_{j,q}$ simultaneously vote over this proposal; let $v_{j,q}(i, X_{j,q}) \in \{\tilde{y}, \tilde{n}\}$ be the vote of player $i \in X_{j,q}$. Let $\text{Yes}\{X_{j,q}\}$ be the subset of $X_{j,q}$ voting in favor of this proposal (i.e. $\text{Yes}\{X_{j,q}\} \equiv \{i \in X_{j,q} \mid v_{j,q}(i, X_{j,q}) = \tilde{y}\}$). Then, if

$$\sum_{i \in \text{Yes}\{X_{j,q}\}} \gamma_i > \alpha \sum_{i \in N_j} \gamma_i,$$

i.e., if $X_{j,q}$ has the absolute majority of the weighted votes within N_j , then the game proceeds to step 5; otherwise the game proceeds to step 6.

5. If $X_{j,q} = N_j$, then we proceed to step 7 and the game ends. Otherwise players from $N_j \setminus X_{j,q}$ are eliminated, players from $X_{j,q}$ add $-\varepsilon$ to their payoff, and the game proceeds to step 1 with $N_{j+1} = X_{j,q}$ (and j increases by 1).
6. If $q < |N_j|$, then next agenda setter $i_{j,q+1} \in N_j$ is randomly picked by nature such that $i_{j,q+1} \neq i_{j,r}$ for $1 \leq r \leq q$ (i.e., it is picked among those who have not made a proposal at stage j) and the game proceeds to step 3 (with q increased by 1). Otherwise, we proceed to step 7.
7. N_j becomes the ruling coalition of this terminal node, and each player $i \in N_j$ adds $w_i(N_j)$ to his payoff.

Coalitions that emerge during the game form a sequence $N_0 \supset N_1 \supset \dots \supset N_{\bar{j}}$ where \bar{j} is the number of coalitions (except initial one) that emerges during the game. Whenever there is a non-trivial transition, each player incurs a disutility cost equal to $\varepsilon > 0$ (ε can be arbitrarily small). Summing over the payoffs at each node, the payoff of each player i is given by

$$U_i = w_i(N_{\bar{j}}) - \varepsilon \sum_{1 \leq j \leq \bar{j}} I_{N_j}(i), \quad (4)$$

where $I_X(\cdot)$ is the indicator (characteristic) function of set X . This payoff function captures the fact that individuals' overall utility in the game is related to their share w_i and to the number of rounds of elimination in which the individual is involved in (the second term in (4)). The number of players eliminated equals $|N| - |N_{\bar{j}}|$, and there are totally \bar{j} rounds of elimination. With a slight abuse of terminology, we refer to j above as “the stage of voting.”

The next two examples show that, as in Examples 1 and 2 in the Introduction, in this game there are also several “non-intuitive” SPEs, and in fact, there are several “non-intuitive” Markovian SPEs.

Example 4 Consider $N = \{A, B, C, D\}$, with $\gamma_A = 3, \gamma_B = 4, \gamma_C = 5$ and $\gamma_D = 10$, and suppose that $\alpha = 1/2$. It can be verified that there is a MTHPE where A, B , or C proposes the coalition $\{A, B, C\}$, and following this proposal, A, B and C vote for in favor of this coalition (and if D is the first to make a proposal, he proposes $\{A, B, C, D\}$ which is rejected by the rest). The fact that $\{A, B, C\}$ should be the ultimate ruling coalition is intuitive: any two-person subset of it is “unstable” and would lead to the elimination of one of its members. And clearly, A, B , and C would prefer $\{A, B, C\}$ to $\{A, B, C, D\}$. Nevertheless, there exists another SPE, where all players make proposals according to a pre-fixed protocol and everybody votes against any proposal. This is indeed an SPE, because no deviation by a single player can lead to acceptance of any proposal. Consequently, both $\{A, B, C\}$ and $\{A, B, C, D\}$ can emerge as outcomes in a SPE, even though the former is not a reasonable outcome, since A, B , and C have enough votes to eliminate D . We can also note that all players are using Markovian strategies. Yet, this equilibrium is not intuitive in the sense that it is resulting from an artificial coordination failure among A, B , and C .

The main result for this application is given in the following proposition and shows how the application of MTHPE removes the non-intuitive equilibrium leading to $\{A, B, C, D\}$.

Proposition 1 *Let the mapping ϕ be defined as*

$$\phi(X) = \arg \min_{Y \in \{Z \subset N: \gamma_Z < \gamma_X; \gamma_Z > \gamma_{X \setminus Z}; \phi(Z) = Z\} \cup \{X\}} \gamma_Y. \quad (5)$$

Then there exists a pure strategy SWUE (MTHPE) where the ultimate ruling coalition is reached after at most one stage of elimination and coincides with $\phi(N)$ with probability 1, and the payoff of each $i \in N$ is given by

$$U_i(N) = w_i(X) - \varepsilon I_{\phi(N)}(i) I_{\{\phi(N) \neq N\}}. \quad (6)$$

Vice versa, in any SWUE in pure or mixed strategies (and in any MTHPE in pure or mixed strategies) the ultimate ruling coalition is $\phi(N)$, and it is necessarily reached after at most one stage of elimination, and players’ payoffs are given by (6).

Proof. Acemoglu, Egorov, and Sonin (2008) considered a version of this game with sequential voting and established that the unique SPE is characterized by mapping (5), with payoffs (6) as

given in the proposition. Given the notion of regularization introduced here, this implies that they considered a regularization of the game here. Then Theorem 5 immediately applies and establishes that the unique MTHPE is given by (5) and (6), completing the proof. ■

This proposition provides an illustration of the usefulness of the concept of MTHPE. Instead of considering a more complicated game with sequential voting as in Acemoglu, Egorov, and Simon (2008), one can consider a simpler game and obtain exactly the same results (eliminating all non-intuitive equilibria).

5 Conclusion


In many political economy models, which involve voting or multilateral bargaining, or more generally, in what we defined as agenda-setting games, several Nash or subgame perfect (Nash) equilibria are “non-intuitive”. In addition, details concerning voting and proposal making procedures often have a major effect on the set of equilibria. In many applied models, non-intuitive equilibria are eliminated (sometimes without specifying the exact notion of equilibrium), and modelers often have to make sure that the “right assumptions” are imposed so that unimportant details do not influence conclusions.

In this paper, we proposed two related equilibrium refinements for voting and agenda-setting games, Sequentially Weakly Undominated Equilibrium (SWUE) and Markov Trembling Hand Perfect Equilibrium (MTHPE). These concepts are widely applicable, easy to use and eliminate non-intuitive equilibria that arise naturally in dynamic voting games and games in which random or deterministic sequence of agenda setters make offers to several players. We also provided a tight characterization of MTHPE in agenda-setting games and showed that the set of MTHPE does not depend on features such as whether voting is simultaneous or sequential. We believe that these concepts can be applied in a wide variety of political economy models and can both eliminate the need to devise ad hoc equilibrium refinements to remove non-intuitive equilibria and ensure that major predictions do not depend auxiliary assumptions that are difficult to map to reality (such as whether voting is simultaneous or sequential).

References

- Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin (2008) Coalition Formation in Non-Democracies, *Review of Economic Studies*, 75(4): 987-1009.
- Austen-Smith, David and Jeffrey Banks (1999) *Positive Political Theory*. Ann Arbor, U.Michigan Press.
- Barbera, Salvador, Michael Maschler, and Jonathan Shalev (2001) Voting for voters: A model of electoral evolution, *Games and Economic Behavior*, 37: 40-78.
- Baron, David, Daniel Diermeier and Pohan Fong (2008) "A Dynamic Theory of a Parliamentary Democracy," mimeo.
- Baron, David and John Ferejohn (1989) "The Power to Propose." In *Models of Strategic Choice in Politics*, ed. Peter C. Ordeshook. Ann Arbor, MI: University of Michigan Press.
- Battaglini, Marco and Stephen Coate (2007) "Inefficiency in Legislative Policy-Making: A Dynamic Analysis." *American Economic Review*, 97(1): 118-149.
- Battaglini, Marco and Stephen Coate (2008) "A Dynamic Theory of Public Spending, Taxation and Debt." *American Economic Review*, 98(1): 201-236.
- Diermeier, Daniel and Pohan Fong (2009) "Legislative Bargaining with Reconsideration", mimeo.
- Duggan, John (2003) A Note on Backward Induction, Iterative Elimination of Weakly Dominated Strategies, and Voting in Binary Agendas, mimeo.
- Duggan, John and Tasos Kalandrakis (2007) "Dynamic Legislative Policy Making." mimeo.
- Fudenberg, Drew and Jean Tirole (1991) *Game Theory*, MIT Press, Cambridge, MA.
- Granot, Daniel, Michael Maschler, and Jonathan Shalev (2002) Voting for voters: the unanimity case, *International Journal of Game Theory* 31:155–202
- Jack, William and Roger Lagunoff (2006) Dynamic Enfranchisement, *Journal of Public Economics*, 90 - 4-5, 551-572.
- Maskin, Eric and Jean Tirole (2001) "Markov Perfect Equilibrium I: Observable Actions", *Journal of Economic Theory*, 100 (2), 191-219.
- Moulin, Herve (1979) Dominance Solvable Voting Schemes, *Econometrica*, 47 - 6, 1337-1351
- Osborne, Martin and Ariel Rubinstein (1994) *A Course in Game Theory*. MIT Press, Cambridge, MA.
- Romer, Thomas and Howard Rosenthal (1978) "Political Resource Allocation, Controlled Agendas, and the Status Quo." *Public Choice*, 33(4): 27-43.
- Selten, Reinhard (1975) Reexamination of the perfectness concept for equilibrium points in

extensive games, *International Journal of Game Theory* 4:25–55



Appendix

Example 5 Consider the following extensive-form game. In this example, player 3 gets payoff 0 in all terminal nodes. However, his choice of action may affect the payoffs of players 1 and 2. First, players 1 and 2 simultaneously choose their strategies. (We can think about this as a simultaneous vote.) If they play strategies R and r , respectively, then both get the payoff of 2. Otherwise, player 3's choice determine whether they both get 1 or 0. One would expect that players 1 and 2 should play R and r , respectively. This is indeed what happens in any MTHPE. (There are multiple MHTPEs as player 3 is indifferent between the outcomes).

Now suppose that players 1 and 2 choose their actions sequentially (which corresponds to sequential process of casting votes). If they play (R, r) , both get the payoff of 2, and this is arguably the most intuitive outcome. However, the following unintuitive SPE cannot be eliminated by sequential elimination of weakly dominated strategies. Consider the strategy-profile (L, l, a_1, b_2, b_3) , which is an SPE. None of player 3's strategies is weakly dominated. Then, neither L nor l can be eliminated.

Example 6 Consider a game of three players with extensive form and payoffs as shown on Figure 4. The first two players vote, and if both vote for the 'right', all three players receive first-best; if one of them votes for the 'left' then the third player chooses between 'moderate' and 'bad'. All players receive the same in all terminal nodes, so there is no strategic conflict between them.

Equilibrium $(R, r, (a_1, a_2, a_3))$ is trembling-hand perfect, but so is $(L, l, (a_1, a_2, a_3))$ where efficiency is not achieved because of 'herding' in voting (note that neither L nor l are dominated strategies: for instance, L is best response to second player playing l and third player playing (a_1, b_2, b_3)). Indeed, take some η and consider

$$\sigma^n = ((1 - \eta^3) L + \eta^3 R, (1 - \eta^3) l + \eta^3 r, ((1 - \eta^2) a_1 + \eta^2 b_1, (1 - \eta) a_2 + \eta b_2, (1 - \eta) a_3 + \eta b_3)).$$

Evidently, player 3 (and all his agents in agent-strategic form) are better off choosing a_1 over




Figure 4: A Game With Herding in Trembling-Hand Perfect Equilibrium.

b_1 , a_2 over b_2 , and a_3 over b_3 . Now consider payoffs of player 1 choosing L or R . If he chooses L , he obtains $u_L = 5((1 - \eta^3)(1 - \eta^2) + \eta^3(1 - \eta)) = 5 - 5\eta^2 - 5\eta^4 + 5\eta^5$. If he chooses R , he obtains $u_R = 5((1 - \eta^3)(1 - \eta)) + 7\eta^3 = 5 - 5\eta + 2\eta^3 + 5\eta^4$. Hence, For small η , player 1 should put all weight to L , and a similar argument would show that player 2 should put all weight to l . This proves that $(L, l, (a_1, a_2, a_3))$ is also a trembling-hand perfect equilibrium.

The effect that Example 6 emphasizes would not be the case if fully mixed profiles σ^n were required to be Markovian, which is what our definition of MTHPE imposes. Indeed, it is a natural restriction to require that in the three subgames where player 3 moves and payoffs are identical, his mixed action profile σ^n should lead to identical place. In that case, the increase of utility of player 1 due to the possibility of player 2 playing r instead of l would be not be offset by worse development in the subgame if he still plays l .