

Endogenous agenda formation processes with the one-deviation property

HANNU VARTIAINEN

Helsinki Center of Economic Research, University of Helsinki

We study collective choice via an endogenous agenda setting process. At each stage, a status quo is implemented unless it is replaced by a majority (winning coalition) with a new status quo outcome. The process continues until the prevailing status quo is no longer challenged. We impose a one-time deviation restriction on the feasible policy processes, reflecting the farsightedness of voters. The key feature of the solution is history dependence. The existence of the solution is proven by iterating a version of the uncovered set. We show that the resulting fixed point is the largest set of outcomes that can be implemented via any policy process that meets the one-deviation restriction. Finally, we relate our solution to a concrete noncooperative model and show that it can be interpreted as a refinement of the solution of Bernheim and Slavov (2009) in the context of repeated voting, and of the solution of Konishi and Ray (2003) and Vartiainen (2011) in the context of coalition formation.

KEYWORDS. Voting, history dependence, one-deviation principle, covering.

JEL CLASSIFICATION. C71, C72.

1. INTRODUCTION

Endogenous agenda setting is an important and difficult question in political choice theory. For example, the famous chaos theorems (McKelvey 1979, Rubinstein 1979, Bell 1981, Schofield 1983) state that with very relaxed conditions concerning how voters are distributed in the policy space, an agenda can be created where it is possible to start at any status quo alternative, and with a succession of majority comparisons, end at any other specified alternative in the policy space. Hence, without a presumed institutional structure, there seems to be only little hope of reaching *any* predictions of the actual political choice.

Lack of farsightedness of agents is a well known limitation of the chaos argument, however. With the agenda amendment procedure studied by McKelvey and others, it is

Hannu Vartiainen: hannu.vartiainen@helsinki.fi

I am indebted to John Duggan for pointing out an error in an earlier version of this paper, which led to a considerable generalization of the main result. I am grateful to the co-editor Nicola Persico and two referees for comments and suggestions that significantly improved the paper. I also thank Vincent Anesi, Klaus Kultti, Roger Myerson, Hannu Nurmi, Marco Ottaviani, Ariel Rubinstein, Hannu Salonen, Daniel Seidmann, and Juuso Välimäki for insightful comments and discussions, and Lin Zhang for generous feedback at the final stage of the project. All errors are mine.

Copyright © 2014 Hannu Vartiainen. Licensed under the Creative Commons Attribution-NonCommercial License 3.0. Available at <http://econtheory.org>.

DOI: 10.3982/TE1225

the outcome in the end of the dominance chain—not the next one in line—that should be used for payoff comparisons. Hence, not all the dominance chains described by the chaos argument are really feasible.¹ But this observation raises a difficult question, known in the coalition formation literature as the *prediction problem* (see Ray 2007): profitability of a blocking depends on *which* one of the feasible dominance paths is to be played. That is, one needs to have a theory that simultaneously explains which outcomes are blocked and where the play converges via transitory blockings. Clearly, the transitory blockings should also reflect, in an appropriate sense, coalitional optimality.

A natural research strategy is to add forward looking strategic thinking into the model or to model the whole decision making procedure as a concrete noncooperative game.² But since the problem is open ended—there is no final stage from which to start the recursion—it is not clear how it should be solved in general contexts. How to find a consistent, applicable, and powerful solution that also exists in a general class of collective choice situations is an open question.

This paper develops a solution of farsighted and endogenous political choice that is based on the standard *one-deviation property*. Our solution builds on a general model of coalition formation owing to Konishi and Ray (2003) and Vartiainen (2011), and on a model of sequential political choice owing to Bernheim and Slavov (2009). The benefit of our solution is that it is simple and parsimonious, and relates in a natural way to the usual principles of dynamic optimization. It is also more robust, has more cutting power, and requires less information on the part of the modeler than the solutions on which it builds. Importantly, our solution exists under very relaxed conditions.

Bernheim and Slavov (2009) study a policy program that specifies an infinite stream of social states and is robust against one-time deviations by *any* majority coalition. The problem is that their solution is not applicable in the classic one-time decision scenario. It is critical for their results that society cannot commit to the status quo or “implement” it. In contrast, the general model of Konishi and Ray (2003) is also well suited to the analysis of one-time decision making, as demonstrated in Vartiainen (2011). However, the solution of Konishi and Ray (2003) is less demanding than that of Bernheim and Slavov (2009) in the sense that, at each instant of time, the coalitional act must be efficient only for a *single* coalition, given the continuation play. This is not an unproblematic assumption.³ Why would the other coalitions stay passive when one coalition takes the game to a direction that is undesirable to them? Ideally, an institution-free model would not be sensitive to such ad hoc restrictions on coalitional behavior.

We show that the strengths of Konishi and Ray (2003), Bernheim and Slavov (2009), and Vartiainen (2011) can be combined to obtain a robust prediction of endogenous and farsighted political choice. Specifically, we study the canonical model of McKelvey (1979) and others, where a social outcome is implemented once and for all through an endogenous amendment agenda. At each stage, there is a status quo that may be

¹ See Chwe (1994).

² For recent contributions, see Acemoglu et al. (2012), Anesi (2006, 2010, 2012), Anesi and Seidmann (forthcoming), Duggan and Banks (2008), Diermeier and Fong (2012), Dutta et al. (2001, 2002, 2004), Duggan (2006), Penn (2006, 2009), Roberts (2007), and Bernheim and Slavov (2009).

³ See Konishi and Ray (2003, Appendix A) and Xue (1998).

challenged with another social alternative. The choice between the two is made by majority voting, and the winner becomes the new status quo. The process continues until the prevailing status quo is no longer challenged. Our solution concept—the one-deviation property—is a strengthened version of the solution of Konishi and Ray (2003) and Vartiainen (2011), requiring that, for each history of blocked status quos, the prescribed political act is optimal for *all* the majority coalitions in light of the eventual consequences of the acts.⁴ In contrast to Bernheim and Slavov (2009), an outcome is implemented once and for all if a majority does not challenge it.

An attractive feature of the model studied in this paper is that it is detail-free. As robustness is required against all feasible coalitional moves, no knowledge is needed of which coalition has the chance to make a move or even what are the identities of the coalitions that have acted in the past. In this sense, the model is quite different from Vartiainen (2011), which explicitly relies on such information. As a consequence, the current model provides a more robust prediction of endogenous political choice. It also has strictly more cutting power, as we show with an example. The third aspect in which the current model is stronger is the domain assumption: we allow all compact spaces whereas Vartiainen (2011) critically hinges on the finiteness of the set of social alternatives.⁵

Moreover, our solution imposes more constraints on behavior than Bernheim and Slavov (2009) since, in our model, the value of the future play cannot be lower than that of the current status quo for any of the majority coalition. Consequently, our solution has more cutting power, as we show with an example.

Our focus is on policy programs that implement an outcome in finite time after all histories. Such terminating policy programs that satisfy the one-deviation property, or equilibrium policy programs for short, are characterized directly in terms of the outcomes they implement and the underlying social preferences. This characterization is complete in the sense that we identify a consistency property that is met by any set of social alternatives that are implementable via an equilibrium policy program. Conversely, we show that any such set of alternatives can be implemented via some equilibrium policy program.

The main contribution of this paper is a general existence result. The conditions under which the existence of the solution is proven encompass, for instance, the spatial model of electoral decision making. The result uses a novel covering technique, which we dub *quasi-covering*. We show that the iterated version of the quasi-uncovered set, the *ultimate quasi-uncovered set*, satisfies the desired consistency property.⁶ Moreover, the ultimate quasi-uncovered set contains all sets with this property.

History dependence of the policy program is crucial to the results of this paper, particularly for the existence. History dependence is also the key aspect that differentiates our model from most of the literature on endogenous political choice.⁷ Anesi and

⁴Ray and Vohra (1997) are among the first to study coalitional behavior by using equilibrium reasoning.

⁵The focus of Konishi and Ray (2003) and Vartiainen (2011) is on the general coalitional model due to Chwe (1994).

⁶The uncovered set is due to Fishburn (1977) and Miller (1980). The ultimate uncovered set is studied by Dutta (1988).

⁷Bernheim and Slavov (2009) is a notable exception.

Seidmann (forthcoming) and Anesi (2006, 2010, 2012) study agenda formation in related bargaining models where, at each date, the current status quo is voted against an outcome proposed by an agent, and the winner becomes the new status quo on the next date. Their general finding is that the absorbing states of Markov perfect equilibria are characterized by von Neumann–Morgenstern (vNM) stable sets independently of the way the proposing player is recognized (see also Diermeier and Fong 2012, Acemoglu et al. 2012). This is in line with our observation that any history independent policy program constitutes a vNM stable set (with respect to the underlying social preference relation). Thus the vNM solution captures the essential features of history independent collective choice behavior when the choice can be freely amended.⁸ This is an important observation, as the vNM stable set, which is a solution with pedigree, has been lacking a proper noncooperative foundation. The caveat is that a vNM solution may not exist. General existence requires history dependence, which is the theme of this paper.

Even though our policy programs are history dependent, the relevant part of histories concerns only the past status quos, not the identities of the coalitions that induced the status quos. This is important, as now the solution does not hinge on the specifics of the institutional structure, for example, which coalition is allowed to challenge the status quo if there are many that are willing to or what the preferences are of the agents in a coalition. This informational flexibility also permits secret voting ballots, which are often used in real applications. Moreover, the results of Anesi and Seidmann (forthcoming) suggest that our solution could be of use in analyzing history dependent equilibria of noncooperative agenda setting games more generally.⁹ Along these lines, we show that there is an equivalence between equilibrium policy processes and strong subgame perfect equilibria of a game where agenda proposals at each stage are made via plurality voting.

An important benefit of “detail-freeness” is that our solution can be expressed in reduced form by appealing only to the underlying majority relation. Because of this property, it is easy to relate our model to other solutions in the social choice literature. For example, we show that the ultimate quasi-uncovered set—the maximal set of outcomes implementable via our solution—is always a subset of the ultimate (deep) uncovered set (Duggan 2013). However, *if* the latter happens to be also externally stable, then the two concepts coincide. In fact, any covering set (Dutta 1988) is shown to satisfy the desired consistency property.

As a side product, this paper also adds to the literature on uncovered sets and their derivatives in general domains, an issue that has received attention recently (e.g., Penn 2006, Duggan 2013, Dutta et al. 2005). For example, it has been viewed problematic that the uncovered set and its iterations may lack external stability. Without external stability, how can one guarantee that the play eventually converges back to the uncovered set

⁸Anesi and Seidmann (forthcoming) also analyze comparative static properties of Markov perfect equilibria of an agenda amendment game. They find that the eventually implemented outcome may heavily depend on the details of the protocol, but the *set* of implementable outcomes is determined by the primitives, i.e., social preferences. This also applies to our model.

⁹That is, subgame perfect Nash equilibria of agenda setting games that are not sensitive to the recognition rule.

(or its iteration) once it departs from this set? The benefit of our approach is that we describe political choices also at off-equilibrium histories. Indeed our result implies that a deviation from the ultimate quasi-uncovered set leads back to the set in a way that makes the original deviation unprofitable, even if the set does not satisfy external stability.

Hence the ultimate quasi-uncovered set characterizes precisely the equilibrium outcomes of this game. As the ultimate quasi-uncovered set is a subset of the (deep) uncovered set and its iterations, our results refine the conventional wisdom that under a variety of institutional settings, it is the uncovered set that describes the outcomes that can be implemented (Miller 1980, Shepsle and Weingast 1984, Banks 1985).

Duggan (2006) provides a general existence result for a game of endogenous agenda formation in which the agenda is formed by an ex ante known finite sequence of proposers. The constructed agenda is then voted on. Dutta et al. (2002, see also 2001, 2004) consider endogenous agendas in a less structured setting, imposing only consistency conditions on the outcomes of the process. Importantly, also they assume a bounded maximum length of the resulting agenda that again permits iterating the solution backward. Penn (2008) studies an unbounded proposal process, where players stop amending the agenda only when the constructed agenda is stable against changes, given the forthcoming voting under the agenda. Also she finds that in the divide-the-dollar setup, the set of feasible outcomes is a subset of the vNM stable set associated to the problem.

The paper is organized as follows. Section 2 introduces the model and defines the solution concept. In Section 3, the solutions is characterized. Section 4 derives the existence result. Section 5 provides a noncooperative interpretation for the model and in Section 6, a precise connection to the models of Konishi and Ray (2003), Vartiainen (2011), and Bernheim and Slavov (2009) is demonstrated. Section 7 concludes with discussion.

2. THE SETUP

There is a nonempty set of social alternatives X . We assume that X is a *compact* topological space.

Social preferences over X are summarized by a binary relation $R \subseteq X \times X$. Denote the asymmetric part of R by $P \subset X \times X$. Then P is the dual of R in the sense that $x P y$ if and only if *not* $y R x$.

We assume that R is a *closed* subset of $X \times X$ and is *complete*, i.e., $x R y$ and/or $y R x$ for all $x, y \in X$.¹⁰ Equivalently, P is an open subset of $X \times X$ such that *not* $x P y$ and $y P x$ for all $x, y \in X$.

The above assumptions hold in the remainder of the paper without further notice. They are the only assumptions we impose on the domain or on the social preferences.

Denote the *upper section* and the *strict upper section* of R at x , respectively, by $R(x) = \{y \in X : y R x\}$ and $P(x) = \{y \in X : y P x\}$. The *lower section* of R at x is then denoted by $R^{-1}(x) = \{y \in X : x R y\}$. By the closed graph theorem, $R(\cdot)$ and $R^{-1}(\cdot)$

¹⁰We use the standard notation $x R y$ for $(x, y) \in R$.

are upper-hemicontinuous correspondences, and by the open graph theorem, $P(\cdot)$ is a lower-hemicontinuous correspondence. These continuity notions are needed in Section 4, when we prove the existence of the solution.

Policy program

A *path* is a *finite* sequence $\bar{x} = (x_0, \dots, x_K)$ of outcomes. Denote the final element of a path (x_0, \dots, x_K) by

$$\mu[(x_0, \dots, x_K)] = x_K.$$

Denote the set of paths, i.e., *histories*, by $H = \bigcup_{k=0}^{\infty} X^k$, with \emptyset as the initial history. Using the terminology of Bernheim and Slavov (2009), a *policy program* σ specifies a social action given the history of past actions $\sigma: H \rightarrow X \cup \{\text{STOP}\}$. Since σ is conditioned on all the histories, we may typically use a shorthand notation $\sigma(h, x_t)$ for $\sigma(x_0, \dots, x_t)$, where $h = (x_0, \dots, x_{t-1})$. The interpretation of a policy program is that if $\sigma(h, x) = y \in X$, then after a history h of status quos, the current status quo x is successfully challenged by a winning coalition (e.g., majority) with outcome y , which then becomes the new status quo, and if $\sigma(h, x) = \text{STOP}$, then all the winning coalitions agree on implementing x and this action is put in force. Thus, a policy program specifies how the sequence of status quos evolves and which outcome—if any—eventually is implemented.

Denote, in the usual way, by $\sigma^t(h)$ the t th iteration of σ starting from h , i.e., $\sigma^0(h) = \sigma(h)$ and $\sigma^t(h) = \sigma(h, \sigma^0(h), \dots, \sigma^{t-1}(h))$ for all $t = 1, \dots$. A policy program σ is *terminating* if, for any $h \in H$, there is $T < \infty$ such that $\sigma^{T+1}(h) = \text{STOP}$ (T may depend on h). That is, after history h , the policy program will eventually implement the outcome $\sigma^T(h)$.

Our focus is on terminating policy programs. That is, at the outset, we preclude complex dynamics such as infinite cycling. Terminating programs are easy to interpret if the political process concerns a one-shot policy decision. With a terminating program, political actions could reflect negotiation prior to a binding one-shot agreement. Terminating programs are also easier to describe and compute.

We note that requiring that an agreement be achieved in finite time reduces the flexibility of the political process. In general, this makes it harder—not easier—to find a solution that meets the desired stability properties.¹¹

Let $\bar{\sigma}(h)$ denote the sequence of status quos in X that is induced by the program σ from the history h onward:

$$\bar{\sigma}(h) = (\sigma^0(h), \sigma^1(h), \dots).$$

If σ is terminating, then $\bar{\sigma}(h)$ is finitely long and $\mu[\bar{\sigma}(h)]$ is well defined for all h . Specifically, for a terminating policy program σ , if a policy action $a \in X \cup \{\text{STOP}\}$ is chosen at

¹¹Any terminating policy program that is robust against one-time deviations after all histories is, in general, robust against such deviations, even if infinite cycling is allowed. However, then there may also be *nonterminating* programs that meet the one-deviation criterion. Thus, finding a solution might become easier, but, at the same time, characterizing them could become more difficult.

history $(h, x) \in H$, then

$$\mu[\bar{\sigma}(h, x, a)] = \begin{cases} \mu[\bar{\sigma}(h, x, y)] & \text{if } a = y \in X \\ x & \text{if } a = \text{STOP}. \end{cases}$$

In particular, $\mu[\bar{\sigma}(h, \sigma(h))] = \mu[\bar{\sigma}(h)]$.

The solution

Our equilibrium condition, which is just a version of the standard one-deviation principle, is defined next.

DEFINITION 1 (One-deviation property). A history dependent terminating policy program σ satisfies the *one-deviation property* if

$$\mu[\bar{\sigma}(h)] R \mu[\bar{\sigma}(h, a)] \quad \text{for all } a \in X \cup \{\text{STOP}\}, \text{ for all } h \in H. \quad (1)$$

That is, after each history, a majority will not want to change the prescribed action given the eventual consequences. Since the program is terminating, the consequences are always well defined. It is important that the no profitable one-time deviations requirement is also imposed on off-equilibrium histories. Whereas the final deviation of any *finitely long* deviation sequence also violates the one-deviation property, the property also implies robustness against all finitely long deviations.

This solution is connected to the dynamic Condorcet winner solution of Bernheim and Slavov (2009) and the dynamic equilibrium processes of coalition formation due to Konishi and Ray (2003) and Vartiainen (2011). A closer comparison with these concepts is relegated to Section 6.

Implementable outcomes and Condorcet consistency

Note that a majority coalition can always guarantee the status quo x by choosing **stop**. Therefore, the one-deviation property implies that

$$\mu[\bar{\sigma}(h, x)] R x \quad \text{for all } (h, x) \in H. \quad (2)$$

That is, the outcome that is implemented if the equilibrium path is followed must not be majority dominated by any element along the path.

We say that the set Y of alternatives is *implementable* via a dynamic policy program σ if

$$Y = \{x \in X : \sigma(h, x) = \text{STOP for some } h \in H\}.$$

That is, for each element x of Y , there is a history (h, x) such that x is implemented. What the initial status quo is may affect the alternative that will be implemented in Y , but not the set Y itself. The sets of implementable outcomes are the main object of our study.

Before going to the main results of the paper, we observe that our solution passes the test of being Condorcet consistent. An outcome x is a *Condorcet winner* if $x R y$ for

all outcomes y . It is a *strong Condorcet winner* if $x P y$ for all outcomes $y \neq x$. It is not difficult to see the following extensions:

- (a) If z is a Condorcet winner, then there is a terminating policy program σ that meets the one-deviation property such that z is implementable via σ .
- (b) If z is a strong Condorcet winner, then z is the only outcome that is implementable via any terminating policy program σ that meets the one-deviation property.

3. CHARACTERIZATION

In this section, we characterize terminating policy programs that meet the one-deviation property. The characterization is given directly in terms of outcomes that are implementable via such policy programs. With this purpose in mind, we define the following solution concept for social choice problems.

DEFINITION 2 (Consistent choice set). A nonempty set $C \subseteq X$ is a *consistent choice set* if, for any $x \in C$ and for any $y \in X$, there is $z \in C$ such that $z \in R(y) \setminus P(x)$.

The definition of a consistent choice set can be rephrased in the form of the familiar two-step principle. If x is in C , then for any other outcome y , there is another outcome z —possibly x itself—in C such that $x R z R y$. Hence, any element x in the set is reachable from any other element y with at most two dominance steps such that the intermediate step, z , is *also* in the set. The difference between this concept and the notion of deep uncovered set (see Duggan 2013) is that the latter allows the intermediate step z to be *any* element in X . Hence the notion of a consistent choice set is more stringent and its existence is not implied by the existence of the deep uncovered set. Moreover, there may be many consistent choice sets.

Now we characterize terminating policy programs that meet the one-deviation property through the concept of a consistent choice set.

LEMMA 3. *Let a terminating policy program σ satisfy the one-deviation property. Then the set Y of outcomes that are implementable via σ is a consistent choice set.*

PROOF. We show that Y satisfies Definition 2. Take any $(h, x) \in H$ such that $\sigma(h, x) = \text{STOP}$. Then $\mu[\bar{\sigma}(h, x)] = x \in Y$. Take any $y \in X$ and let $z = \mu[\bar{\sigma}(h, x, y)] \in Y$. By (2), $z R y$, or $z \in R(y)$. By Definition 1, $x R z$ or $z \notin P(x)$, as desired. \square

We now show that the converse of this result holds too. Let C be a consistent choice set. Our task is to construct a terminating policy program $\sigma^C : H \rightarrow X$ in such a way that σ^C meets the one-deviation property and implements the outcomes in C . To this end, construct a function $z : X \times X \rightarrow X$ such that, for any $x \in C$ and for any $y \in X$,

$$z(x, y) \in C \cap R(y) \setminus P(x). \quad (3)$$

By Definition 2, $C \cap R(y) \setminus P(x)$ is nonempty and hence $z(x, y)$ is well defined.

Interpret C as an index set and construct recursively a partitioning $\{H_x\}_{x \in C}$ of H as follows. Initial step: Choose $x_0 \in C$ and $\emptyset \in H_{x_0}$. Inductive step: If $h \in H_x$, then, for any $y \in X$, let

$$(h, y) \in \begin{cases} H_y & \text{if } y \in C \setminus P(x) \\ H_{z(x, y)} & \text{if } y \notin C \setminus P(x). \end{cases} \quad (4)$$

By induction, each history h becomes allocated to exactly one element of $\{H_x\}_{x \in C}$.

An element of the partition $\{H_x\}_{x \in C}$ is called a *phase*. The role of a phase is to indicate how the policy program reacts to the status quo. Given the function z and the collection phases $\{H_x\}_{x \in C}$, let the policy program σ^C satisfy, for all $h \in H_x$,

$$\sigma^C(h, y) = \begin{cases} \text{STOP} & \text{if } y \in C \setminus P(x) \\ z(x, y) & \text{if } y \notin C \setminus P(x). \end{cases} \quad (5)$$

We now give a verbal interpretation of the constructed policy program σ^C . Think of R as the majority relation. The policy program is constructed so that it implements any element in C and such that any deviating majority coalition will be punished. The punishment is achieved by implementing an outcome in C that the deviating coalition does not prefer relative to the outcome that was originally to be implemented. The role of a phase in the construction is to store in memory which majority is to be punished. The z -function specifies the majority whose job it is to implement the punishment (by stopping the program). Transition between phases (4) determines when and how the majority that is to be punished should be changed. The circularity in punishments makes the program robust against profitable majority deviations in all phases, i.e., after all histories.

Of course, the construction of σ^C is feasible only due to the assumed characteristics of the consistent choice set C . The existence of such a set is a separate issue, which the theme of the next section.

We now prove formally that the policy program σ^C satisfies the one-deviation property.

LEMMA 4. *Policy program σ^C as constructed in (3)–(5) is terminating and satisfies the one-deviation property.*

PROOF. To demonstrate that σ^C is terminating, it suffices to show that, for any $x \in C$, for any $h \in H_x$, and for any $y \in X$,

$$\mu[\bar{\sigma}^C(h, y)] = \begin{cases} y & \text{if } y \in C \setminus P(x) \\ z(x, y) & \text{if } y \notin C \setminus P(x). \end{cases} \quad (6)$$

Consider the two cases separately.

- (i) Let $y \in C \setminus P(x)$. By (5), $\sigma^C(h, y) = \text{STOP}$ and, hence, $\mu[\bar{\sigma}^C(h, y)] = y$, as desired.
- (ii) Let $y \notin C \setminus P(x)$. By (5), $\sigma^C(h, y) = z(x, y)$. Since, by (4), $(h, y) \in H_{z(x, y)}$ and, by (3), $z(x, y) \in C \setminus P(z(x, y))$, it follows by (5) that $\sigma^C(h, y, \sigma^C(h, y)) = \sigma^C(h, y, z(x, y)) = \text{STOP}$. Thus, $\mu[\bar{\sigma}^C(h, y)] = z(x, y)$, as desired.

It follows from (6) and (3) that either $y \in C \setminus P(x)$ and $\mu[\bar{\sigma}^C(h, y)] = y$ or $y \notin C \setminus P(x)$ and $\mu[\bar{\sigma}^C(h, y)] = z(x, y) \in C \setminus P(x)$. Therefore,

$$\mu[\bar{\sigma}^C(h, y)] \in C \setminus P(x) \quad \text{for all } h \in H_x, \text{ for all } y \in X, \text{ and for all } x \in C. \quad (7)$$

We now show that σ^C satisfies the one-deviation property. Take any $(h, y) \in H$ such that $h \in H_x$. It suffices that a one-time deviation from $\sigma^C(h, y)$ is not profitable. There are again two cases to consider.

- (i) Let $y \in C \setminus P(x)$. By (5), $\sigma^C(h, y) = \text{STOP}$ and, by (6), $\mu[\bar{\sigma}^C(h, y)] = y$. A deviation to $w \in X$ leads to the history (h, y, w) . Since, by (4), $(h, y) \in H_y$, applying (7) yields

$$\mu[\bar{\sigma}^C(h, y, w)] \in C \setminus P(y).$$

Hence the deviation is not profitable.

- (ii) Let $y \notin C \setminus P(x)$. By (5), $\sigma^C(h, y) = z(x, y)$ and, by (6), $\mu[\bar{\sigma}^C(h, y)] = z(x, y)$. There are two kinds of deviations. (a) A deviation to STOP implements y . By (3), $z(x, y) \in R(y)$; thus, the deviation is not profitable. (b) A deviation to $w \in X \setminus \{z(x, y)\}$ leads to the history (h, y, w) . Since, by (4), $(h, y) \in H_{z(x, y)}$, applying (7) yields

$$\mu[\bar{\sigma}^C(h, y, w)] \in C \setminus P(z(x, y)).$$

Hence the deviation is not profitable. □

By Lemma 3, a set Y of alternatives is implementable via a terminating policy program that meets the one-deviation property only if Y is a consistent choice set. Conversely, by Lemma 4, outcomes of any consistent choice can be implemented via a terminating policy program that meets the one-deviation property. We compound these observations into the following characterization.

THEOREM 5. *Set Y of alternatives is implementable via a terminating policy program that satisfies the one-deviation property if and only if Y is a consistent choice set.*

The application of the characterization result is illustrated in the following two examples. The first concerns the classic Condorcet cycle in which rational decision making is often deemed pathologically difficult.

EXAMPLE 6. Let $X = \{x, y, z\}$, and let $x P y$, $y P z$, and $z P x$. Then X itself is a consistent choice set. Hence there is a terminating policy program that meets the one-deviation property that can commit to implement any element of X . ◇

It is instructive to construct a concrete policy program that meets the one-deviation property in the case of a Condorcet cycle. Identify a partitioning $\{H_x, H_y, H_z\}$ of the set of histories H recursively as follows. Initial step: $\emptyset \in H_w$ for some $w \in X$. Inductive step: for any $h \in H$, the following statements hold:




FIGURE 1.

- If $h \in H_x$, then $(h, x) \in H_x$ and $(h, y), (h, z) \in H_y$.
- If $h \in H_y$, then $(h, y) \in H_y$ and $(h, z), (h, x) \in H_z$.
- If $h \in H_z$, then $(h, z) \in H_z$ and $(h, x), (h, y) \in H_x$.

Choose a policy program that is conditioned on the past history and the current status quo as follows:

- If $h \in H_x$, then $\sigma(h, x) = \sigma(h, y) = \text{STOP}$ and $\sigma(h, z) = y$.
- If $h \in H_y$, then $\sigma(h, y) = \sigma(h, z) = \text{STOP}$ and $\sigma(h, x) = z$.
- If $h \in H_z$, then $\sigma(h, z) = \sigma(h, x) = \text{STOP}$ and $\sigma(h, y) = x$.

To see that the constructed σ satisfies the one-deviation property, let, say, $h \in H_x$. Obediently following σ implements x in one step under (h, x) , y in one step under (h, y) , and y in two steps under (h, z) . Note that $x, y \in R^{-1}(x)$. A one-time deviation to $\sigma(h, x)$ implements, in two steps, an outcome in $R^{-1}(x)$, a one-time deviation to $\sigma(h, y)$ implements, in two steps, an outcome in $R^{-1}(y)$, and a one-time deviation to $\sigma(h, z)$ implements either z immediately or, in two steps, an outcome in $R^{-1}(y)$. Hence no one-time deviation is profitable.

The next example demonstrates how the one-deviation property restricts outcomes from being implementable.¹²

EXAMPLE 7. Let $X = \{x, y, z, w\}$, and let $y P x$, $z P y$, $w P z$, $x P w$, $x P z$, and $y P w$ (see Figure 1, where $x \rightarrow y$ means $x P y$, etc.). Then the unique consistent choice set is $\{x, y, z\}$ and, hence, w cannot be implemented via a terminating policy program that meets the one-deviation property. \diamond

To see why w cannot be implemented in Example 7, suppose that there is a history after which the status quo w is implemented. Consider a one-time deviation where the status quo w is replaced with x . Since the deviation is not profitable, the outcome that becomes implemented cannot be x . Since the only alternative that is (at least weakly) preferred to the status quo x is y , it follows that y must be implemented. But since y is also preferred to w , the original deviation from implementing w is profitable. Hence, w cannot be implemented by a policy program that meets the one-deviation property.

¹²See Vartiainen (2011, Figure 1) and Konishi and Ray (2003, Example 10) for other treatments of this case.

History dependence is crucial for Theorem 5. To see this, consider a *simple* (or memoryless or Markovian) policy program, where the policy action depends only on the status quo alternative, not on the history. A simple policy program can be described by a function $\sigma: X \rightarrow X \cup \{\text{STOP}\}$, where $\sigma(x)$ is the action when x is the status quo.

Simplicity is a harsh restriction on a policy program. To understand the nature of the constraint, let us define a version of another well known solution. A set $V \subseteq X$ is a (von Neumann–Morgenstern) *weakly stable set* if (i) $x \in V$ implies that there is *no* $y \in V$ such that $y P x$ and (ii) $x \notin V$ implies that there is $y \in V$ such that $y R x$.¹³

A weakly stable set is a consistent choice set, but a consistent choice set need not be a weakly stable set. The proof of the following proposition appears in Section A.1.

PROPOSITION 8. *The set V of alternatives is implementable by a simple policy program that meets the one-deviation property if and only if V is a weakly stable set.*

It is well known that a weakly stable set—and, hence, a simple policy program that satisfies the one-deviation property—often does not exist. In fact, if R is a tournament relation (i.e., $x R y$ and $y R x$ imply $x = y$), then V is a weakly stable set if and only if $V = \{x\}$, where x is the strong Condorcet winner.¹⁴ In the absence of such an element, any terminating policy program that satisfies the one-deviation property *must* be history dependent. Our next task is to show that such a policy program always exists.

4. EXISTENCE

In this section, we prove the main result of the paper: a consistent choice set and, hence, a terminating policy program, that satisfies the one-deviation property exists. We develop an iterative procedure that identifies the maximal consistent choice set. Our construction is related to the notion of deep covering and its iterations (see Duggan 2013).

Given a subset B of X , we say that x *quasi-covers* y under B if

$$R(x) \cap B \subseteq P(y) \cap B \quad \text{and} \quad x, y \in X. \quad (8)$$

The prefix “quasi” refers to the property of (8) that covering between two elements in X is restricted to B . When $B = X$, quasi-covering coincides with deep covering, but when $B \neq X$, the two notions differ. A more detailed comparison between quasi-covering and deep covering appears at the end of this section.

Denote the maximal elements of the quasi-covering relation under B by $QU(B)$, the *quasi-uncovered* set under B :

$$QU(B) := \{x \in X : x \text{ is not quasi-covered by } y \in X \text{ under } B\}.$$

¹³The qualifier “weak” refers to the fact that external stability is defined with respect to R rather than P (see, e.g., Anesi and Seidmann forthcoming). When R is a tournament relation, this difference becomes vacuous.

¹⁴Existence conditions can be relaxed when the social preference relation is inherited from a richer voting structure than majority voting or when the underlying choice domain is restricted (see Penn 2008, Anesi 2010, 2012, Anesi and Seidmann forthcoming).

That is, $QU(B)$ comprises all alternatives in X that are not quasi-covered under B by any element in X . Note that $QU(B)$ need *not* be a subset of B or even have common elements with B . Even though the quasi-covering relation is transitive, the existence of a maximal element could be an issue. The next result verifies that these concerns are not warranted for *any* B .¹⁵

LEMMA 9. *The set $QU(B)$ is nonempty for any $B \subseteq X$.*

PROOF. By the Hausdorff maximal principle, there is a maximal subset $\{x_\alpha\}_{\alpha \in A}$ of X that is totally ordered by the quasi-covering relation under B such that x_β quasi-covers x_α under B if $\beta > \alpha$. Since $P(x_\beta) \subseteq R(x_\beta)$,

$$P(x_\beta) \cap B \subseteq P(x_\alpha) \cap B \quad \text{if } \beta > \alpha. \quad (9)$$

Since $P(x)$ is an open set for all x , $P(\cdot) \cap B$ is lower hemicontinuous as a correspondence. Let $\bar{x} \in X$ be an accumulation point of $\{x_\alpha\}$. Such a point exists since X is a compact space. By (9), and since $P(\cdot) \cap B$ lower hemicontinuous, $P(\bar{x}) \cap B \subseteq \bigcap_\alpha P(x_\alpha) \cap B$. We claim that \bar{x} is not quasi-covered under B . If it were, then there is $y \in X$ such that $R(y) \cap B \subseteq P(\bar{x}) \cap B$. But then also $R(y) \cap B \subseteq P(x_\alpha) \cap B$ for all $\alpha \in A$, which contradicts the assumption that $\{x_\alpha\}$ is a maximal subset of X that is totally ordered by the quasi-covering relation. \square

Duggan (2013) shows that the deep uncovered set of a closed subset B of X is closed. We complement this result by showing that the quasi-uncovered set under a closed set B is also closed.

LEMMA 10. *Let B be a closed subset of X . Then $QU(B)$ is closed.*

PROOF. Suppose that $QU(B)$ is not a closed subset of a compact space, i.e., compact. Then there is a net $\{x_\alpha\} \subseteq QU(B)$ with a limit point $x \in X$ such that $x \notin QU(B)$. Since x is quasi-covered under B , there is $y \in B$ such that $R(y) \cap B \subseteq P(x) \cap B$. Since $x_\alpha \in QU(B)$ for all α , also $R(y) \cap B \not\subseteq P(x_\alpha) \cap B$ for all α . That is, there is z_α such that $z_\alpha \in R(y) \cap R^{-1}(x_\alpha) \cap B$ for all α . Since B is a closed subset of a compact space, $\{z_\alpha\}$ has a subnet $\{z_\beta\}$ with a limit $z \in B$. Then $\{x_\beta\}$ also has the limit x . Since $R(y)$ is a closed set and $R^{-1}(\cdot)$ is upper hemicontinuous, $z \in R(y) \cap R^{-1}(x) \cap B$. But then y does not quasi-cover x under B , a contradiction. \square

While $QU(B)$ may, in general, contain elements outside B , we now show that this is not the case when B is obtained by iterating the quasi-uncovered set from the grand set X . The iterations are defined by transfinite recursion as by transfinite recursion as follows:¹⁶

- $QU^0 = X$
- $QU^\alpha = QU(\bigcap_{\beta < \alpha} QU^\beta)$ for any ordinal $\alpha > 0$.

¹⁵Duggan (2013) shows the nonemptiness of the deep uncovered set under compact B .

¹⁶Finitely many inductive steps suffice when X is a finite set.

LEMMA 11. *For any ordinal $\alpha > 0$, $QU^\alpha \subseteq \bigcap_{\beta < \alpha} QU^\beta$.*

PROOF. The claim holds for $\alpha = 1$. Assume that the claim holds for all ordinals $\beta < \alpha$ for a given ordinal $\alpha > 1$. By transfinite induction, it suffices to show that the claim holds also for α . Let $x \notin \bigcap_{\beta < \alpha} QU^\beta$. Then there is $\beta^* < \alpha$ such that $x \notin QU^{\beta^*}$. By construction, $QU^{\beta^*} = QU(\bigcap_{\gamma < \beta^*} QU^\gamma)$. Thus there is $y \in X$ such that

$$R(y) \cap \left(\bigcap_{\gamma < \beta^*} QU^\gamma \right) \subseteq P(x) \cap \left(\bigcap_{\gamma < \beta^*} QU^\gamma \right). \quad (10)$$

Now $\bigcap_{\beta < \alpha} QU^\beta \subseteq QU^{\beta^*} \subseteq \bigcap_{\gamma < \beta^*} QU^\gamma$, where the first inclusion follows from $\beta^* < \alpha$ and the second follows from the supposition. Thus, by (10),

$$R(y) \cap \left(\bigcap_{\beta < \alpha} QU^\beta \right) \subseteq P(x) \cap \left(\bigcap_{\beta < \alpha} QU^\beta \right),$$

which implies $x \notin QU(\bigcap_{\beta < \alpha} QU^\beta) = QU^\alpha$. Hence $QU^\alpha \subseteq \bigcap_{\beta < \alpha} QU^\beta$. \square

The following corollary is an immediate consequence of Lemma 11.

COROLLARY 12. *For any two ordinals $\beta < \alpha$, $QU^\alpha \subseteq QU^\beta$.*

This corollary implies that $QU(QU^{\alpha-1}) = QU^\alpha$ for any ordinal α that has a predecessor. When X is a *finite* set, this holds for every ordinal. Thus, in this case, the iterative procedure has a simple form: $QU^0 = X$ and $QU^{k+1} = QU(QU^k)$ for all $k = 0, 1, \dots$. The procedure also ends in finite time, i.e., there is a K such that $QU^K = QU(QU^K)$. However, in a general case, such final stage K may not exist. Instead, to reach a fixed point of the $QU(\cdot)$ correspondence, a more elaborate argument is needed.

LEMMA 13. *The set QU^α is closed and nonempty for any ordinal α .*

PROOF. The proof is by transfinite induction. Since X is compact and nonempty, QU^0 is closed and nonempty. Suppose that QU^β is closed and nonempty for all ordinals $\beta < \alpha$. We show that it holds also for the ordinal α . Since, by Corollary 12 and the assumption, $\bigcap_{\beta < \alpha} QU^\beta$ is an intersection of nested closed and nonempty sets, it is closed and nonempty. By Lemmata 9 and 10, QU^α is closed and nonempty. \square

Our aim is to show that the *ultimate quasi-uncovered set* UQU , defined constructively by

$$UQU := \bigcap_{\alpha} QU^\alpha,$$

constitutes a consistent choice set. Note that the *existence* of UQU is not an issue: it is a complement of a union of a collection of well defined sets.

We first argue that UQU is a fixed point of the $QU(\cdot)$ correspondence.

LEMMA 14. *We have $UQU = QU(UQU)$.*

PROOF. \supseteq : Suppose that $x \in QU(\bigcap_{\alpha} QU^{\alpha}) \setminus \bigcap_{\alpha} QU^{\alpha}$. Then there is α such that $x \notin QU^{\alpha} = QU(\bigcap_{\beta < \alpha} QU^{\beta})$. Hence, there is $y \in X$ such that y quasi-covers x under $\bigcap_{\beta < \alpha} QU^{\beta}$. But $\bigcap_{\alpha} QU^{\alpha} \subseteq \bigcap_{\beta < \alpha} QU^{\beta}$ implies that y quasi-covers x also under $\bigcap_{\alpha} QU^{\alpha}$, contradicting $x \in QU(\bigcap_{\alpha} QU^{\alpha})$.

\subseteq : Suppose that $x \in \bigcap_{\alpha} QU^{\alpha} \setminus QU(\bigcap_{\alpha} QU^{\alpha})$. Then there is $y \in X$ such that y quasi-covers x under $\bigcap_{\alpha} QU^{\alpha}$. Since $QU^{\alpha} = QU(\bigcap_{\beta < \alpha} QU^{\beta})$, y does not quasi-cover x under $\bigcap_{\beta < \alpha} QU^{\beta}$ for any α . Thus, for any α , there is a z_{α} such that

$$z_{\alpha} \in R(y) \cap R^{-1}(x) \cap \left(\bigcap_{\beta < \alpha} QU^{\beta} \right).$$

Since X is compact, we can assume that $\{z_{\alpha}\}$ is a convergent net with a limit z . By Lemma 11, the sequence of sets $\{\bigcap_{\beta < \gamma} QU^{\beta}\}_{\gamma \geq \alpha}$ is nested and, hence, $\{z_{\gamma}\}_{\gamma \geq \alpha} \subseteq \bigcap_{\beta < \alpha} QU^{\beta}$ for all α . Since, by Lemma 13, $\bigcap_{\beta < \alpha} QU^{\beta}$ is a closed set, $z \in \bigcap_{\beta < \alpha} QU^{\beta}$. Since this holds for any α , $z \in \bigcap_{\alpha} \bigcap_{\beta < \alpha} QU^{\beta} = \bigcap_{\alpha} QU^{\alpha}$. Since $R(\cdot)$ and $R^{-1}(\cdot)$ are upper hemicontinuous correspondences, also $z \in R(y) \cap R^{-1}(x)$. Thus

$$z \in R(y) \cap R^{-1}(x) \cap \left(\bigcap_{\alpha} QU^{\alpha} \right).$$

But this contradicts the assumption that y quasi-covers x under $\bigcap_{\alpha} QU^{\alpha}$. \square

By Lemma 14, no element in UQU is quasi-covered in UQU by *any* element in X . The next theorem, which is the main result of this paper, uses this property to establish that UQU is a consistent choice set. It is also argued that UQU is the maximal consistent choice set (in the sense of set inclusion).

THEOREM 15. *The set UQU is a consistent choice set. Moreover, UQU contains as a subset any consistent choice set.*

PROOF. We first argue that UQU is a consistent choice set. By Lemma 13 and Corollary 12, $\{QU^{\alpha}\}$ is a collection of nested, closed, and nonempty sets in a compact space. Thus, their intersection, UQU , is closed and nonempty. Take any $x \in UQU$ and any $y \in X \setminus \{x\}$. Since, by Lemma 14, x is not quasi-covered by $y \in X$ under UQU , there is z such that $z \in R(y) \cap UQU \setminus P(x)$, as required.

We now demonstrate that UQU contains as a subset any consistent choice set. Let C be a consistent choice set. We show that $C \subseteq UQU$. By the definition of a consistent choice set, $C \cap R(y) \setminus P(x)$ is nonempty for all $x \in C$ and for all $y \in X$. Thus, for any $B \subseteq X$ such that $C \subseteq B$,

$$R(y) \cap B \not\subseteq P(x) \cap B \quad \text{for all } x \in C, \text{ for all } y \in X. \quad (11)$$

Choosing $B = X = QU^0$ in (11), by the definition of quasi-covering, $C \subseteq QU(QU^0) = QU^1$. Let $C \subseteq \bigcap_{\beta < \alpha} QU^{\beta}$ for given α . Choosing $B = \bigcap_{\beta < \alpha} QU^{\beta}$ in (11), by the definition

of quasi-covering, $C \subseteq QU(\bigcap_{\beta < \alpha} QU^\beta) = QU^\alpha$. By transfinite induction, $C \subseteq \bigcap_\alpha QU^\alpha = UQU$. \square

As UQU is a well defined set of the topological space X , Theorem 15 automatically implies that *a consistent choice set exists*.

Finally, we tie the existence results concerning consistent choice sets to the existence issue of policy programs that satisfy the one-deviation property. By Theorem 5, we have shown that a terminating policy program with the one-deviation property does exist and that the set of outcomes that are implementable via any such program is contained in the ultimate uncovered set.

COROLLARY 16. *There is a terminating policy program that meets the one-deviation property that implements outcomes in UQU . Moreover, UQU contains all outcomes that can be implemented via any terminating policy program that meets the one-deviation property.*

Thus, it is without loss of generality to focus on the ultimate quasi-uncovered set UQU if one is interested in the welfare consequences of dynamic political decision making.

External stability, deep covering, and covering sets

We now connect the quasi-uncovered set and its iterations to the deep uncovered set (Duggan 2013) and to the covering set (Dutta 1988).¹⁷ Given a subset B of X , we say that x *deep covers* y in B if

$$R(x) \cap B \subseteq P(y) \cap B \quad \text{and} \quad x, y \in B.$$

The set of *deep uncovered* outcomes in B is

$$DU(B) := \{x \in X : x \text{ is not deep covered by } y \in B \text{ in } B\}.$$

The difference between quasi-covering under B and deep covering in B is that the latter concerns only outcomes in B , but the former accounts for covering relations over outcomes in X . This has profound consequences on the associated uncovered sets. While quasi-covering under B leaves more room for an outcome in B to be covered (i.e., by an outcome not in B), it may also allow an outcome not in B to be uncovered under B . Hence, in general, a clear relationship does not exist between the quasi-uncovered set $DU(B)$ and the deep uncovered set $DU(B)$.

However, we next observe that iteration of the quasi-uncovered set leads to more selective sets of outcomes than does iteration of the deep uncovered set. Define the iterations of the deep uncovered set as follows. Set $DU^0 = X$ and let $DU^\alpha = DU(\bigcap_{\beta < \alpha} DU^\beta)$ for all ordinals α . By construction, $DU(DU^\alpha) = DU^{\alpha+1}$ for all ordinals α .

¹⁷For different notions of covering, see Duggan (2013) and Penn (2006).

Duggan (2013) shows that the deep uncovered set of a compact set is closed and nonempty. Hence, by transfinite induction, DU^α is closed and nonempty, and the *ultimate deep uncovered* set $UDU = \bigcap_\alpha DU^\alpha$ is nonempty.¹⁸ The proof of the next proposition is relegated to Section A.1.

PROPOSITION 17. *We have $QU^\alpha \subseteq DU^\alpha$ for any ordinal α .*

By Proposition 17, since $UQU = \bigcap_\alpha QU^\alpha$ and $UDU = \bigcap_\alpha DU^\alpha$, it immediately follows that also $UQU \subseteq UDU$. Thus, by Theorem 15, we conclude that the ultimate deep uncovered set UDU contains as a subset all consistent choice sets and, hence, all outcomes that can be implemented via a terminating policy program that meets the one-deviation property. However, when the inclusion $UQU \subseteq UDU$ is *strict*, Theorem 15 also implies that UDU is *not* a consistent choice set. In such a case, UDU contains elements that cannot be implemented via a terminating policy program that meets the one-deviation property, by Theorem 5.

When X is a finite set, the ultimate uncovered set is *externally stable* in the sense that for any outcome x not in UQU , there is another outcome y in UQU that covers x in $UQU \cup \{x\}$. This is deemed as a desirable property of the set since it guarantees that the play always converges back to the set.

A problematic feature of general infinite domains is that external stability of the ultimate uncovered set does not carry over to there, as discussed by Duggan (2013). It is, therefore, noteworthy that our existence result concerning policy processes that meet the one-deviation property (Theorem 15) does not have to rely on the external stability of UQU . What suffices for the one-deviation property is that any x in UQU is not quasi-covered in UQU by any element in X . This questions the importance of external stability as a necessary reflection of strategically sophisticated voting behavior.

However, external stability does have an interpretation in terms of the one-deviation principle. We now argue that *if* a set happens to be externally and internally stable in the sense of the *covering set* of Dutta (1988), then it is a consistent choice set.

A covering set D is defined by the property that $x \in X$ is not deep covered in $D \cup \{x\}$ if and only if $x \in D$. We now argue that any covering set is a consistent choice set (see Section A.1 for a proof).

PROPOSITION 18. *Let $D \subseteq X$ be a covering set. Then D is a consistent choice set.*

By construction, the characterizing feature of a covering set is that it is stable, both externally and internally: no element x in D is deep covered in $D \cup \{x\}$ and every element x not in D is deep covered in $D \cup \{x\}$. Proposition 18 points out that these properties can be used to construct a policy process that meets the one-deviation property. However, the converse is not true even when X is a finite set. There are consistent choice sets that are not covering sets. That is, a policy process that meets the one-deviation property need not reflect external stability even in the finite case (see the working paper version of this paper).

¹⁸The ultimate uncovered set is analyzed in the finite case, e.g., by Miller (1980), Dutta (1988), and Laslier (?L98). In that case, all the notions of covering coincide.

5. NONCOOPERATIVE INTERPRETATION

We now argue that the one-deviation property of a policy program (2) characterizes subgame perfect equilibria of many natural legislative bargaining procedures. We construct a game form in which all the *strong* subgame perfect equilibria, i.e., those that are robust against coalitional deviations, are described by (2).

Consider the situation where there is an odd number n of voters choosing a policy in X . Denote the (complete, transitive) preferences of voter i with asymmetric preferences P_i . The associated majority relation M is defined by

$$x M y \quad \text{if } |\{i: x P_i y\}| > \frac{n}{2}.$$

Denote by \bar{M} the *completion* of M , i.e. $x \bar{M} y$ if not $y M x$ for all $x, y \in X$.

Consider the following dynamic plurality voting procedure Γ^P . Voting takes place in discrete periods $t = 0, 1, \dots$. At period $t = 0$, there is an initial status quo policy x_0 . At period t , the voters cast their votes simultaneously from $\{\text{STOP}\} \cup X$. If STOP receives a larger number of votes than any $y \in X$, then the status quo x_t is implemented. If the outcome $y \in X$ receives a larger number of votes than $z \in X \setminus \{y\}$ or STOP , then y becomes the new status quo x_{t+1} at stage $t + 1$. In the case of a tie, a winner is chosen by some deterministic rule. The game continues until a policy is implemented. Never implementing a policy is strictly worse for all the agents than implementing any of the policies in X in finite time. There is no discounting.

Denote by H the set of possible histories of status quos. For simplicity, let us assume that the dynamic voting strategy ϕ is conditioned on H . Summarize the voting act of the agents after history h when x is the status quo by $\phi(h, x) \in \{\text{STOP}\} \cup X$ such that

$$\phi(h, x) = \begin{cases} \text{STOP} & \text{if STOP is the winner of the plurality voting at } (h, x) \\ y & \text{if } y \in X \text{ is the winner of the plurality voting at } (h, x). \end{cases}$$

Denote by $\phi^\tau(h)$ the τ th iteration of ϕ starting from h . Let $\bar{\phi}(h) = (\phi^0(h), \phi^1(h), \dots)$. Then, as in Section 2, $\mu[\bar{\phi}(h)]$ is the outcome that is implemented when the agents adhere to strategy ϕ after history h , i.e., for all $(h, x) \in H$,

$$\mu[\bar{\phi}(h, x)] = \begin{cases} x & \text{if STOP is the winner of the plurality voting at } (h, x) \\ \mu[\bar{\phi}(h, x, y)] & \text{if } y \in X \text{ is the winner of the plurality voting at } (h, x). \end{cases}$$

In particular, $\mu[\bar{\phi}(h, \phi(h))] = \mu[\bar{\phi}(h)]$ for all h .

As before, our focus is on *terminating* equilibria that implement an outcome in finite time after any history. Our solution concept is the *strong subgame perfect equilibria* (SSPE), requiring that a strategy profile be immune to profitable *coalitional* deviations, where each member of the coalition strictly benefits from the deviation. No restrictions are set on the form of the coalition. In the case of a terminating strategy profile and since any infinite terminal history is worse for all players than any finite terminal history, it is sufficient to check that there is no profitable one-time deviation for any coalition (at a single instant of time). Thus, after any history, there should be no coalition whose members would strictly benefit from a deviation, given the continuation play.

It is now easy to see that any SSPE is characterized by (1). In an SSPE, the strategy must be such that a majority coalition cannot profitably make a one-time deviation to the play path. Hence

$$\text{not } \mu[\bar{\phi}(h, x, a)] \bar{M} \mu[\bar{\phi}(h, x)] \quad \text{for all } a \in \{\text{STOP}\} \cup X. \quad (12)$$

Equivalently,

$$\mu[\bar{\phi}(h, x)] \bar{M} \mu[\bar{\phi}(h, x, a)] \quad \text{for all } a \in \{\text{STOP}\} \cup X. \quad (13)$$

For sufficiency, suppose that ϕ satisfies (13) and that ϕ arises from a voting strategy in which any voter chooses $\phi(h)$ at history h . Then a deviation from the equilibrium path requires a deviation by a majority coalition. But since ϕ satisfies (12), there is no majority coalition whose members strictly prefer a one-time deviation. Thus ϕ is a SSPE. The proof of the following proposition is relegated to Section A.1.

PROPOSITION 19. *Strategy ϕ arises from a terminating SSPE voting strategy of the dynamic plurality procedure Γ^P if and only if there is a terminating policy program σ that satisfies the one-deviation property with respect to the binary relation \bar{M} such that $\phi = \sigma$.*

The “if” part of the proof suggests the following hypothesis. Any sequential moves legislative bargaining game that has the property that, after any nonterminal history, any majority coalition can force, by mutually coordinated action, the implementation of any given outcome, has a subgame perfect equilibrium that corroborates a terminating policy program that meets the one-deviation restriction. An example of such a game is the nondiscounted version of the classic Baron and Ferejohn (1989) model (either version). In related work, Anesi and Seidmann (forthcoming) and Anesi (2006, 2010, 2012) demonstrate how Markov perfect equilibria (that are dependent only on the current status quo) are characterized by vNM stable sets in various amendment games.

However, as our one-deviation restriction is conditioned only on the histories of status quos, noncooperative agenda amendment games may also entertain equilibria that it does not characterize. These equilibria are sensitive to the recognition protocol that specifies how the amendment proposal is made in each stage.

6. RELATION TO SOME OTHER MODELS

In this section, we connect our results to the solution concepts of Bernheim and Slavov (2009), Konishi and Ray (2003), and Vartiainen (2011). The solution of Bernheim and Slavov (2009) can only be applied to a “real time” blocking model, where infinitely lived agents gain per period utility from a state that may be affected by the society. However, the solutions of Konishi and Ray (2003) and Vartiainen (2011) also fit well to the current scenario, where a single social decision is made once and for all after a coalitional contemplation.

6.1 Bernheim and Slavov (2009)

Bernheim and Slavov (2009) expand the idea of Condorcet dominance to a setting where political decisions are made repeatedly. The set X of outcomes is now interpreted as the set of social *states* that may change in dates $t = 0, 1, \dots$. Policy making is now an ongoing process where the individuals gain benefits from the policy choices in each period t . There is a set $\{1, \dots, n\}$ of agents. Each agent $i = 1, \dots, n$ is endowed with a per period utility function $u_i: X \rightarrow \mathbb{R}$.

Letting H denote the set of all possible finite paths of social alternatives—the set of histories—a dynamic policy program is now a function $\sigma: H \rightarrow X$, capturing the transitions from one history to another. Let H be the set of all histories of states (x_0, \dots, x_t) , where x_0 is exogenously given. The transition from history h to the next history $(h, \sigma(h))$ is induced by a winning majority coalition whose members stand to benefit from it. As before, $\sigma^t(h)$ denotes the t th iteration of σ , starting from the history h .

We now focus on policy programs that are *absorbing* in the sense that after all histories, the policy path converges in finite time to an absorbing state in which it stays permanently. That is, for any history h , there is an integer T_h such that $\sigma^t(h) = x$ for all $t > T_h$. The absorbing state of the policy program σ that starts from history h is then well defined for all h and is denoted by $\alpha[\sigma(h)]$.

We assume that there is no discounting: the intertemporal payoffs are evaluated by the limit-of-the-means criterion. For an absorbing policy program, it holds then that the intertemporal payoff of agent i from policy σ at history $h \in H$ is given by $u_i(\alpha[\bar{\sigma}(h)])$.

The majority relation $M \subset X \times X$ is defined by

$$x M y \quad \text{if } |\{i: u_i(x) \geq u_i(y)\}| > \frac{n}{2}. \quad (14)$$

An absorbing policy program σ is a *dynamic Condorcet winner* (DCW) if

$$\alpha[\bar{\sigma}(h, x)] M \alpha[\bar{\sigma}(h, x, y)] \quad \text{for all } (h, x) \in H, \text{ for all } y \in X. \quad (15)$$

The key difference between DCW and our one-time deviation criterion is that there the relationship is never-ending, whereas in our model, the game ends when the one-shot decision is made. In other words, in our context, the majority can permanently commit to the current state, but in Bernheim and Slavov (2009) it cannot. Nothing prevents the program from moving from the current state to a “path of states” state that is worse for a majority coalition than the status quo outcome. As a consequence, for a DCW, it does *not* need to hold that

$$\alpha[\bar{\sigma}(h, x)] M x, \quad (16)$$

which makes our necessary condition (2) vacuous. The existence result of Bernheim and Slavov (2009) hinges on this additional leeway. Moreover, because their solution does not demand (16), it cannot be applied to the one-time decision making situation, which is the theme of this paper.

We now argue that the one-deviation restriction is a strictly stronger condition than DCW. To this end, replace a policy program σ^C as constructed in (3)–(5) with an ever-continuing program that is otherwise equivalent but has the property that $\sigma^C(q_x, y) =$

STOP is replaced by an infinite repetition of y (for a precise construction, see Section A.2). Then σ^C has C as the set of absorbing states.

By Lemma 4, σ^C meets (15) and, hence, is a DCW. By Theorem 5, we can state the following proposition.

PROPOSITION 20. *For any consistent choice set C , defined with respect to the majority relation M , there is a DCW with C as the set of absorbing states.*

However, the converse is not true. To see this, recall the sufficient condition (adjusted to our finite, no-discounting case) of Bernheim and Slavov (2009, Theorem 3(ii)): If there is a Condorcet cycle w^1, w^2, \dots, w^K such that $w^k M w^{k+1}$ for all $k = 1, \dots, K - 1$ and $w^K P w^1$, then each state w^k can be supported as an absorbing state of an absorbing DCW.¹⁹

The following example demonstrates that some absorbing states of a DCW cannot be implemented via a terminating policy program that meets the one-deviation restriction. Denote $u_S(x) = (u_i(x))_{i \in S}$ for any $S \subseteq \{1, \dots, n\}$.

EXAMPLE 21. There are three agents $\{1, 2, 3\}$ and four states $X = \{x, y, z, w\}$. The payoffs are

$$u_{\{1,2,3\}}(x) = (1, 3, 4)$$

$$u_{\{1,2,3\}}(y) = (3, 4, 1)$$

$$u_{\{1,2,3\}}(z) = (4, 1, 2)$$

$$u_{\{1,2,3\}}(w) = (2, 2, 3).$$

These payoffs induce a majority relation over states such that $y M x$, $z M y$, $w M z$, $x M w$, $x M z$, and $y M w$, the case depicted in Example 7. Thus the unique consistent choice set is $\{x, y, z\}$. However, since any state in X is a member of a Condorcet cycle, any of them can be supported as an absorbing state of an absorbing DCW. \diamond

6.2 Konishi and Ray (2003) and Vartiainen (2011)

We now relate the one-deviation property (1) to the equilibrium condition of Konishi and Ray (2003) and Vartiainen (2011). The key finding is that the one-deviation property is stronger, as it is not conditioned on the identities of the active coalitions and requires robustness against *all* feasible coalitional deviations. Consequently, the one-deviation property is not only independent of the details of the coalition formation technology—a desirable property from the viewpoint of institution-free predictions—but it also has more cutting power than the equilibrium condition of Vartiainen (2011), when applied to the majority choice context (we prove this by an example). Moreover, the current setup proves the existence of the solution far beyond the finite case studied in Vartiainen (2011).

¹⁹Bernheim and Slavov (2009) actually say something stronger: that such w^k can be supported as an absorbing state of a DCW that always converges to an absorbing state in at most two periods.

There is a set $\{1, \dots, n\}$ of agents, making a one-shot decision from the set X of social alternatives. Agent i has a utility function $u_i: X \rightarrow \mathbb{R}$. Each coalition $S \subseteq \{1, \dots, n\}$ has an opportunity to move the play from the status quo outcome x to any outcome in the set $F_S(x) \subseteq X \cup \{\text{STOP}\}$ such that $\text{STOP} \in F_S(x)$ for all S . Coalitional actions may be contingent on past coalitional actions and, hence, a history consists of a sequence such as $x_0, S_1, x_1, \dots, S_k, x_k$, where the status quo $x_j \in X$ is induced by the coalition $S_j \subseteq \{1, \dots, n\}$ for all $j = 1, \dots, k$. If, under status quo x , all coalitions agree on STOP , then x is implemented.

In the current context, the process of coalition formation specifies the next status quo and the coalition that induces it. The process of coalition formation is then characterized by a pair of functions (σ, I) , defined for all past histories $\hat{h} = (x_0, S_0, x_1, S_2, \dots, x_{k-1}, S_k)$ and current status quos x such that $I(\hat{h}, x) \subseteq \{1, \dots, n\}$ and $\sigma(\hat{h}, x) \in F_{I(\hat{h}, x)}(x)$. Letting $(I^t, \sigma^t)(\hat{h}, x) = (I, \sigma)(\hat{h}, x), (I^0, \sigma^0)(\hat{h}, x), \dots, (I^{t-1}, \sigma^{t-1})(\hat{h}, x)$ for all $t = 1, \dots$, we can now describe agent i 's payoff from a terminating process of coalition formation simply as $u_i(\mu[\bar{\sigma}(\hat{h}, x)])$, where $\bar{\sigma}(\hat{h}, x) = (\sigma^0(\hat{h}, x), \sigma^1(\hat{h}, x), \dots)$ is a sequence of status quos that originates from the history (\hat{h}, x) and where $\mu[\bar{\sigma}(\hat{h}, x)]$ is the final element of that sequence.

If coalition S chooses a coalitional action $a \in F_S(x)$ after history (\hat{h}, x) , then

$$\mu[\bar{\sigma}(\hat{h}, x, S, a)] = \begin{cases} \mu[\bar{\sigma}(\hat{h}, x, S, y)] & \text{if } a = y \in X \\ x & \text{if } a = \text{STOP}. \end{cases}$$

For any coalition S , use the notation $u_S(x) > u_S(y)$ if $u_i(x) > u_i(y)$ for all $i \in S$ and use $u_S(x) \geq u_S(y)$ if $u_i(x) \geq u_i(y)$ for all $i \in S$. Given a terminating process of coalition formation (σ, I) and a history (\hat{h}, x) , we say that a coalition S has a *weakly preferred* move $a \in F_S(x)$ from x if $u_S(\mu[\bar{\sigma}(\hat{h}, x, S, a)]) \geq u_S(x)$. Furthermore, a move a is *efficient* for the coalition S if there is no $b \in F_S(x)$ such that $u_S(\mu[\bar{\sigma}(\hat{h}, x, S, b)]) > u_S(\mu[\bar{\sigma}(\hat{h}, x, S, a)])$.

Given these notions, we may now specify the solution concept of Konishi and Ray (2003), as defined by Vartiainen (2011).²⁰ A terminating process of coalition formation (σ, I) is a *dynamic equilibrium process of coalition formation* (DEPCF) if it satisfies, for all *histories of coalitional moves*, i.e., $(\hat{h}, x) \in \bigcup_{k=0}^{\infty} (X \times \{\text{all coalitions}\})^k \times X$, the following conditions:

- (i) If $\sigma(\hat{h}, x) \in X$, then $\sigma(\hat{h}, x)$ is an efficient and weakly preferred move from x for the coalition $I(\hat{h}, x)$.
- (ii) If $\sigma(\hat{h}, x) = \text{STOP}$, then STOP is an efficient move from x for *all* coalitions S .

Vartiainen (2011) characterizes DEPCF and proves its existence when X is a finite set.²¹

²⁰To interpret the results of this subsection in the infinitely repeated blocking model of Konishi and Ray (2003), exchange the notion of absorbing state with the act of stopping, as in Section 6.1.

²¹In the formulation of Vartiainen (2011), the notation is simplified by encoding the identity of the active coalition directly into the definition of an outcome. A history dependent process of coalition formation can then be described by the function σ alone. Despite this superficial simplification, histories of the active coalitions are used by the solution at each step and they are crucial for the existence result.

From the viewpoint of institution-free predictions, DEPCF is not an unproblematic concept, however. First, as also discussed by Konishi and Ray (2003, Appendix A), part (i) of the definition relies on an implicit commitment assumption. Once a majority coalition $I(h, x)$ activates, no other coalition can do anything to affect the next status quo. What if there is another majority coalition that prefers another outcome, given the continuation play paths from the possible histories after the next step?²² The ideal way to circumvent this problem would be to require efficiency for *all* the relevant coalitions.

The second problem, which is related to the first, concerns the level of information embedded into the notion of history. That the identities of the active coalitions may affect the process is a strong assumption and requires extensive knowledge on individual voters' preferences. It also rules out some commonly used voting technologies, e.g., secret ballots. A more satisfactory concept would be insensitive to unreasonable details by being dependent only on the history of commonly observed variables, such as the sequence of status quos. Under such circumstances, it would not matter which one of the coalitions induces the next status quo.

Both of the above concerns are avoided if the solution concept is required to be independent of the identity of the blocking coalition. We now argue that the one-deviation property (1) can be interpreted as such a solution. Formally, define a strengthened version of DEPCF for all *histories of status quos*, i.e., $(h, x) \in \bigcup_{k=0}^{\infty} X^k$:

- (i*) If $\sigma(h, x) \in X$, then $\sigma(h, x)$ is an efficient move for *all* coalitions S and is a weakly preferred move from x for *some* coalition S .
- (ii*) If $\sigma(h, x) = \text{STOP}$, then STOP is an efficient move from x for *all* coalitions S .

Note that this version of DEPCF uses no information on the identities of the active coalitions anywhere. Hence, the function σ alone specifies the coalition formation process.

Let us call a terminating coalition formation process σ that meets (i*) and (ii*) as a *DEPCF**. Since DEPCF* relies on weaker informational conditions and requires stronger efficiency properties, any DEPCF* is a DEPCF. We now argue that DEPCF* is equivalent to the one-deviation property in the current context.

Let us define a majority relation M as in (14). The next proposition, whose proof is found in Section A.1, states the equivalence between the one-deviation restriction and DEPCF* when applied to endogenous majority choice. The key observation is that the second part of (i*), that $\sigma(h, x)$ should be a weakly preferred move from x for some S whenever $\sigma(h, x) \in X$, becomes redundant when DEPCF* is applied to the majority choice context.

PROPOSITION 22. *Let $F_S(x) = X$ for all majority coalitions S and let $F_S(x) = \{\text{STOP}\}$ for all other coalitions. Then σ is a DEPCF* if and only if it satisfies the one-deviation property with respect to the majority relation M .*

By Theorem 15, a *DEPCF** exists. The proof of Vartiainen (2011) concerning the existence of a DEPCF in a finite setup critically hinges on the assumption that the histories

²²For a criticism of the assumption and an alternative approach, see Xue (1998).

can also be conditioned on the *identities* of the active coalitions. Hence, the existence of a DEPCF* is a novel result, even when restricted to a finite set X .²³

We now demonstrate that DEPCF* has strictly more cutting power than DEPCF. Since any DEPCF* is a DEPCF, it suffices, by Theorem 5 and Proposition 22, to find a situation where the implemented outcomes of a DEPCF do not form a consistent choice set. A detailed description of the DEPCF used in the next example can be found in Section A.3.

EXAMPLE 23. Reconsider the situation depicted in Example 21 and interpret the states as feasible social alternatives. There is a DEPCF that implements *any* of the alternatives $\{x, y, z, w\}$. \diamond

Recall from Example 7 that the unique consistent choice set in this case is $\{x, y, z\}$. Since, by Proposition 22, a DEPCF* permits only implementation of this set of outcomes, w is ruled out as a robust prediction when moving from DEPCF to DEPCF*.

The observation that *any* of the alternatives $\{x, y, z, w\}$ can be implemented in a DEPCF relies on knowledge of the identities of active coalitions.²⁴ The DEPCF that is used punishes any coalition that deviates from implementing an element in $\{x, y, z, w\}$ and it punishes any coalition that fails to obey the punishment scheme.

In particular, in the constructed DEPCF, it is reliable for the coalitions to implement w , since a deviation to x by the coalition $\{2, 3\}$ can be punished by allowing the coalition $\{1, 2\}$ to move the game further to y , which is then implemented. This construction, which renders the original deviation unprofitable for the deviating coalition, crucially depends on the assumption that the punishment can be made contingent on the identity of the coalition: if it cannot, then $\{1, 2\}$ will make the original deviation from w to x , so as to induce y . To prevent such a deviation, coalition $\{1, 2\}$ must be punished in a different way, by implementing x . Thus, without knowing in detail which coalition made the deviation and what the preferences are of the members of the active coalition, one cannot tailor the punishment correctly and one cannot keep it optimal for all the coalitions not to deviate. In the absence of detailed knowledge of the coalition formation technology or of the preferences of the agents, one should, therefore, be cautious in assuming that w can be implemented.

Comparing Proposition 22 to (15) and (16), it then follows that DEPCF* also refines DCW in the model of Bernheim and Slavov (2009). That DEPCF* has strictly more cutting power follows from Examples 21 and 23.

7. CONCLUDING REMARKS

We study farsighted political decision making when the voting acts may be conditioned on the history. We abstract from the details of the voting procedure and assume that individual preferences are aggregated by a social preference (e.g., majority) ordering.

²³The result relies on the assumed properties of the coalitional choice sets $(F_S)_S$. With general choice sets, the existence of a DEPCF* cannot be guaranteed.

²⁴Or that any element of $\{x, y, z, w\}$ can be supported as an absorbing state of the real-time blocking model of Konishi and Ray (2003).

Choices are made on the basis of binary comparisons: the current status quo may be challenged with another outcome and the status quo is implemented if it is not defeated by any challenger. The key aspect of the model is farsightedness and that blocking behavior may be conditioned on the past blockings. The solution we apply is the standard one-deviation principle.

We establish a reduced form characterization of policy programs that meet the one-deviation property. Our main result is the general existence of the solution. The conditions under which the existence is shown encompass virtually all scenarios of interest, including the case of Euclidean preferences. Also, as the characterization is directly in terms of the underlying majority relation, the solution is readily comparable to other concepts in the social choice literature.

Our results contribute to the voting literature in three dimensions. First, we show that the one-deviation property, which is almost a synonym for sequential rationality in noncooperative models, is also a natural way to model collective decision making in the canonical social choice scenario.

Second, we show that the prediction problem—the key challenge of coalitional analysis—can be solved in a reliable way in the political choice scenario. Indeed, our solution is optimal for all majority coalitions at all times, even under transitory moves, and, hence, it is independent of many of the details of the institutional structure that constrain coalitional behavior. Compared to Bernheim and Slavov (2009), Konishi and Ray (2003), and Vartiainen (2011), the current solution is more robust, has more cutting power, and requires less information on the part of the modeler. We also prove the existence of the solution in a more general domain.²⁵

Third, we develop a new covering method that has good existence and stability properties. The iterated version of our uncovered set is used to prove the existence of our solution.

APPENDIX

A.1 *Omitted proofs*

PROOF OF PROPOSITION 8. If: Let the set $V \subseteq X$ satisfy internal and external stability. By external stability, there is a simple policy program $\sigma: X \rightarrow X \cup \{\text{STOP}\}$ such that $\sigma(x) = \text{STOP}$ for all $x \in V$ and such that $\sigma(x) \in R(x) \cap V$ for all $x \notin V$. Then $V = \{x \in X : \sigma(x) = \text{STOP}\}$. We show that σ satisfies the one-deviation property. Let x be the status quo. If $x \notin V$, then obediently following σ implements $\sigma(x)$. Stopping at x implements x , which, since $\sigma(x) \in R(x)$, is not a profitable deviation. Any other one-time deviation implements some $y \in V$. By internal stability, $y \notin P(\sigma(x))$, which implies that the deviation is not profitable. If $x \in V$, then obediently following σ implements x . A one-time deviation implements some $y \in V$. By internal stability, $y \notin P(x)$ and, hence, the deviation is not profitable.

Only if: Let the function $\sigma: X \rightarrow X \cup \{\text{STOP}\}$ characterize a simple policy program that meets the one-deviation property. We show that the set $V = \{x \in X : \sigma(x) = \text{STOP}\}$

²⁵However, the focus of Konishi and Ray (2003) and Vartiainen (2011) is in the general coalitional model à la Chwe (1994).

satisfies internal and external stability. Let $x \in V$. Suppose that $y \in P(x)$ for some $y \in V$. But then, since $\sigma(y) = \text{STOP}$, a one-time deviation to y would be profitable when x is the status quo. Let $x \notin V$, i.e., $\sigma(x) \neq \text{STOP}$. Identify the outcome $y \in V$ that is implemented. Since stopping the game when x is the status quo is not profitable, $y \in R(x)$, as desired. \square

PROOF OF PROPOSITION 17. The claim holds for $\alpha = 0$. By transfinite induction, it suffices to show that if the claim holds for the ordinals $\beta < \alpha$, then it holds also for the ordinal α . Assume that $QU^\beta \subseteq DU^\beta$ for all $\beta < \alpha$. Let $x \notin DU^\alpha$. Then there is $y \in \bigcap_{\beta < \alpha} DU^\beta$ such that y deep covers x in $\bigcap_{\beta < \alpha} DU^\beta$ and, hence, y quasi-covers x under $\bigcap_{\beta < \alpha} DU^\beta$. By assumption, $\bigcap_{\beta < \alpha} QU^\beta \subseteq \bigcap_{\beta < \alpha} DU^\beta$. But then y quasi-covers x under $\bigcap_{\beta < \alpha} QU^\beta$ and, hence, $x \notin QU(\bigcap_{\beta < \alpha} QU^\beta) = QU^\alpha$. \square

PROOF OF PROPOSITION 18. Take $x \in D$ and let $y \in X$. We find an element z in D such that $z \in R(y) \setminus P(x)$. If $y \notin P(x)$, then $x = z$ qualifies as such an element. Thus let $y \in P(x)$. If $y \in D$, then since D is a covering set, it follows that $R(y) \cap D \not\subseteq P(x) \cap D$. Thus, there is $z \in D$ such that $z \in R(y) \setminus P(x)$.

Thus let $y \in P(x) \setminus D$. Since D is a covering set, there is $z \in D$ such that $R(z) \cap D \subseteq P(y) \cap D$. Since $z \in P(y)$, we are done if $z \notin P(x)$. Suppose that $z \in P(x)$. Since $x, z \in D$ and D is a covering set, necessarily $R(z) \cap D \not\subseteq P(x) \cap D$. Thus there is $w \in D$ such that $w \in R(z) \setminus P(x)$. Since $R(z) \cap D \subseteq P(y) \cap D$ and $w \in R(z) \cap D$, we have that $w \in P(y)$. Thus, $w \in R(y) \setminus P(x)$, as desired. \square

PROOF OF PROPOSITION 22. First we claim that under the assumptions made in the proposition, (i*) can be stated in the following reduced form: If $\sigma(h, x) \in X$, then $\sigma(h, x)$ is an efficient move for all coalitions S . That is, the condition that $\sigma(h, x)$ must be a weakly preferred move from x for some coalition S becomes redundant. To see this, suppose that $\sigma(h, x)$ is not a weakly preferred move from x for any majority coalition S . Then $u_S(\mu[\bar{\sigma}(h, x)]) \not\geq u_S(x)$ for all majority coalitions S . Then $|\{i: u_i(\mu[\bar{\sigma}(h, x)]) \geq u_i(x)\}| < n/2 \leq |\{i: u_i(x) > u_i(\mu[\bar{\sigma}(h, x)])\}|$. Hence, there is a majority coalition S' such that $u_{S'}(x) > u_{S'}(\mu[\bar{\sigma}(h, x)])$. But then $\sigma(h, x)$ is not efficient for S' .

Thus we can conclude that DEPCF* is equivalent to the requirement that $\sigma(h, x) \in X \cup \{\text{STOP}\}$ is an efficient move for all majority coalitions S for all (h, x) . We claim that this is equivalent to σ satisfying the one-deviation property. That $\sigma(h, x)$ is efficient for a majority coalition S is equivalent to $u_S(\mu[\bar{\sigma}(h, x, a)]) \not\geq u_S(\mu[\bar{\sigma}(h, x)])$ for all $a \in X \cup \{\text{STOP}\}$. That this holds for all majority coalitions S is equivalent to

$$|\{i: u_i(\mu[\bar{\sigma}(h, x, a)]) > u_i(\mu[\bar{\sigma}(h, x)])\}| < n/2 \leq |\{i: u_i(\mu[\bar{\sigma}(h, x)]) \geq u_i(\mu[\bar{\sigma}(h, x)])\}|$$

for all $a \in X \cup \{\text{STOP}\}$. This, in turn, is equivalent to $\mu[\bar{\sigma}(h, x)] M \mu[\bar{\sigma}(h, x, a)]$ for all $a \in X \cup \{\text{STOP}\}$, as desired. \square

TABLE 1. ???.

(a, b)	S	c	\bar{S}
(x, y)	$\{1, 2\}$	z	$\{1, 3\}$
(y, z)	$\{1, 3\}$	x	$\{2, 3\}$
(z, w)	$\{2, 3\}$	y	$\{1, 2\}$
(w, x)	$\{2, 3\}$	y	$\{1, 2\}$
(z, x)	$\{2, 3\}$	y	$\{1, 2\}$
(w, y)	$\{1, 2\}$	z	$\{1, 3\}$

A.2 Construction of policy program σ^C for the model of Bernheim and Slavov (2009)

Construct a partitioning $\{H_x\}_{x \in C}$ of H as in (4). Let the policy program σ^C satisfy, for all $h \in H_x$,

$$\sigma^C(h, y) = \begin{cases} y & \text{if } y \in C \setminus P(x) \\ z(x, y) & \text{if } y \notin C \setminus P(x). \end{cases}$$

The only difference between this program and the one defined in (3)–(5) concerns the choice $\sigma^C(h, y)$ when $y \in C \setminus P(x)$. Since

$$\sigma^C(h, y, \dots, y) = y$$

and

$$(h, y, \dots, y) \in H_y,$$

it follows that the policy program σ^C is absorbing in a very strong sense: starting from any history (h, y) , the program starts repeating the status quo in at most one period (Bernheim and Slavov 2009 call this property *stationarity*). By an argument analogous to that of Lemma 4, this policy program is DCW and has C as the set of absorbing states.

A.3 Construction of a DEPCF in Example 21

We now construct a DEPCF that implements any of the alternatives $\{x, y, z, w\}$ in Example 21. This construction requires that if $u_S(b) > u_S(a)$ for a majority coalition S , then there is c and a majority coalition S' such that $u_{S'}(c) > u_{S'}(b)$ and $u_S(c) \not> u_S(a)$. Table 1 shows how to choose c and \bar{S} for each combination of (a, b) and S .

We now specify the process of coalition formation (σ, I) and a partitioning of the histories $\bigcup_{k=0}^{\infty} (X \times \{\text{majority coalitions}\})^k$ into phases $\{H_{a,S}\}_{a \in \{x,y,z,w\}, S \text{ majority coalition}}$. The process of coalition formation (σ, I) is constructed as follows: If h belongs to the phase $H_{a,S}$, and the status quo is b , two alternatives exist:

- If $u_S(b) \not> u_S(a)$, then $\sigma(h, b) = \text{STOP}$.
- If $u_S(b) > u_S(a)$, then $\sigma(h, b) = c$ and $I(h, b) = \bar{S}$ such that $u_S(c) \not> u_S(a)$ and $u_{\bar{S}}(c) > u_{\bar{S}}(b)$ (see Table 1).

Phases $\{H_{a,S}\}_{a \in \{x,y,z,w\}, S \text{ majority coalition}}$ are constructed as follows: If h belongs to the phase $H_{a,S}$, and the status quo is b , there are three alternatives:

- If $u_S(b) \not\geq u_S(a)$ and a majority coalition S' induces a status quo $d \neq \text{STOP}$, then $(h, b, S', d) \in H_{b,S'}$.
- If $u_S(b) > u_S(a)$ and a majority coalition \bar{S} induces a status quo $d \neq c$, then $(h, b, \bar{S}, d) \in H_{c,\bar{S}}$ (see Table 1).
- In all other cases, the new history remains in the phase $H_{a,S}$.

To see how the constructed DEPCF works, note that at any phase $H_{a,S}$ and at any status quo b , it takes at most two steps for the process to implement an outcome ξ such that $u_S(\xi) \not\geq u_S(a)$. In particular, if $b = a$, then a is implemented instantaneously. Consider the consequences of a deviation. If (a) applies, then a deviation by S' to $d \neq \text{STOP}$ results in a phase change to $H_{b,S'}$ and implementation of an outcome ξ such that $u_{S'}(\xi) \not\geq u_{S'}(b)$. Without a deviation by S' , b is implemented. Hence, the deviation is not profitable. If (b) applies, then a deviation by \bar{S} to $d \neq c$ results in a phase change to $H_{c,\bar{S}}$ and implementation of an outcome ξ such that $u_{\bar{S}}(\xi) \not\geq u_{\bar{S}}(c)$. Without a deviation by \bar{S} , c becomes implemented. Hence, the deviation is not profitable.

The set of implemented outcomes of this DEPCF is $\{x, y, z, w\}$. As argued in Example 21, the unique set of the consistent choice set is $\{x, y, z\}$.

REFERENCES

- Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin (2012), "Dynamics and stability of constitutions, coalitions, and clubs." *American Economic Review*, 102, 1446–1476. [188, 190]
- Anesi, Vincent (2006), "Committees with farsighted voters: A new interpretation of stable sets." *Social Choice and Welfare*, 27, 595–610. [188, 190, 205]
- Anesi, Vincent (2010), "Noncooperative foundations of stable sets in voting games." *Games and Economic Behavior*, 70, 488–493. [188, 190, 198, 205]
- Anesi, Vincent (2012), "A new old solution for weak tournaments." *Social Choice and Welfare*, 39, 919–930. [188, 190, 198, 205]
- Anesi, Vincent and Daniel Seidmann (forthcoming), "Bargaining over an endogenous agenda." *Theoretical Economics*. [188, 189, 190, 198, 205]
- Banks, Jeffrey S. (1985), "Sophisticated voting outcomes and agenda control." *Social Choice and Welfare*, 1, 295–306. [191]
- Baron, David P. and John A. Ferejohn (1989), "Bargaining in legislatures." *American Political Science Review*, 83, 1181–1206. [205]
- Bell, Colin E. (1981), "A random voting graph almost surely has a Hamiltonian cycle when the number of alternatives is large." *Econometrica*, 49, 1597–1603. [187]
- Bernheim, B. Douglas and Sita N. Slavov (2009), "A solution concept for majority rule in dynamic settings." *Review of Economic Studies*, 76, 33–62. [187, 188, 189, 191, 192, 193, 205, 206, 207, 210, 211, 213]

- Chwe, Michael S.-Y. (1994), "Farsighted coalitional stability." *Journal of Economic Theory*, 63, 299–325. [188, 189, 211]
- Diermeier, Daniel and Pohan Fong (2012), "Characterization of the von Neumann–Morgenstern stable set in a non-cooperative model of dynamic policy-making with a persistent agenda setter." *Games and Economic Behavior*, 76, 349–353. [188, 190]
- Duggan, John (2006), "Endogenous voting agendas." *Social Choice and Welfare*, 27, 495–530. [188, 191]
- Duggan, John (2013), "Uncovered sets." *Social Choice and Welfare*, 41, 489–535. [190, 194, 198, 199, 202, 203]
- Duggan, John and Jeffrey S. Banks (2008), "A dynamic model of democratic elections in multidimensional policy spaces." *Quarterly Journal of Political Science*, 3, 269–299. [188]
- Dutta, Bhaskar (1988), "Covering sets and a new Condorcet correspondence." *Journal of Economic Theory*, 44, 63–80. [189, 190, 202, 203]
- Dutta, Bhaskar, Matthew O. Jackson, and Michel Le Breton (2001), "Strategic candidacy and voting procedures." *Econometrica*, 69, 1013–1037. [188, 191]
- Dutta, Bhaskar, Matthew O. Jackson, and Michel Le Breton (2002), "Voting by successive elimination and strategic candidacy." *Journal of Economic Theory*, 103, 190–218. [188, 191]
- Dutta, Bhaskar, Matthew O. Jackson, and Michel Le Breton (2004), "Equilibrium agenda formation." *Social Choice and Welfare*, 23, 21–57. [188, 191]
- Dutta, Bhaskar, Matthew O. Jackson, and Michel Le Breton (2005), "The Banks set and the uncovered set in budget allocation problems." In *Social Choice and Strategic Decisions* (David Austen-Smith and John Duggan, eds.), 163–199, Springer, Berlin. [190]
- Fishburn, Peter C. (1977), "Condorcet social choice functions." *SIAM Journal of Applied Mathematics*, 33, 469–489. [189]
- Konishi, Hideo and Debraj Ray (2003), "Coalition formation as a dynamic process." *Journal of Economic Theory*, 110, 1–41. [187, 188, 189, 191, 193, 197, 205, 207, 208, 209, 210, 211]
- McKelvey, Richard D. (1979), "General conditions for global intransitivities in formal voting models." *Econometrica*, 47, 1085–1112. [187, 188]
- Miller, Nicholas R. (1980), "A new solution set to for tournaments and majority voting: Further graph-theoretical approaches to the theory of voting." *American Journal of Political Science*, 24, 68–96. [189, 191, 203]
- Penn, Elizabeth M. (2006), "Alternate definitions of the uncovered set and their implications." *Social Choice and Welfare*, 27, 83–87. [188, 190, 202]
- Penn, Elizabeth M. (2008), "A distributive n -amendment game with endogenous agenda formation." *Public Choice*, 136, 201–213. [191, 198]

Penn, Elizabeth M. (2009), "A model of farsighted voting." *American Journal of Political Science*, 53, 36–54. [188]

Ray, Debraj (2007), *A Game Theoretic Perspective on Coalition Formation*. Oxford University Press, Oxford. [188]

Ray, Debraj and Rajiv Vohra (1997), "Equilibrium binding agreements." *Journal of Economic Theory*, 73, 30–78. [189]

Roberts, Kevin (2007), "Condorcet cycles? A model of intertemporal voting." *Social Choice and Welfare*, 29, 383–404. [188]

Rubinstein, Ariel (1979), "A note about the 'nowhere denseness' of societies having an equilibrium under majority rule." *Econometrica*, 47, 511–514. [187]

Schofield, Norman (1983), "Generic instability of majority rule." *Review of Economic Studies*, 50, 695–705. [187]

Shepsle, Kenneth A. and Barry R. Weingast (1984), "Uncovered sets and sophisticated voting outcomes with implications for agenda institutions." *American Journal of Political Science*, 28, 49–74. [191]

Vartiainen, Hannu (2011), "Dynamic coalitional equilibrium." *Journal of Economic Theory*, 146, 672–698. [187, 188, 189, 191, 193, 197, 205, 207, 208, 209, 211]

Xue, Licun (1998), "Coalitional stability under perfect foresight." *Economic Theory*, 11, 603–627. [188, 209]

Submitted 2012-4-24. Final version accepted 2012-12-10. Available online 2012-12-10.