

Existence and Indeterminacy of Markovian Equilibria in Dynamic Bargaining Games*

Vincent Anesi[†] John Duggan[‡]

September 10, 2017

Abstract

The paper studies stationary Markov perfect equilibria in multidimensional models of dynamic bargaining, in which the alternative chosen in one period determines the status quo for the next. We generalize a sufficient condition for existence of equilibrium due to Anesi and Seidmann (2015). We then use this existence result to show that if a weak gradient restriction holds at an alternative, then when players are sufficiently patient, there is a continuum of equilibria with absorbing sets arbitrarily close to that alternative. A sufficient condition for our gradient restriction is that the gradients of all players' utilities are linearly independent at that alternative. When the dimensionality of the set of alternatives is high, this linear independence condition holds at almost all alternatives, and equilibrium absorbing sets are dense in the set of alternatives. This implies that constructive techniques, which are common in the literature, fail to identify many plausible outcomes in dynamic bargaining games.

1 Introduction

Most formal political analyses of legislative policymaking, until recently, have used models in which legislative interaction ends once a proposal is passed (e.g., Romer and Rosenthal 1978, Baron and Ferejohn 1989, and Banks and Duggan 2000, 2006). As pointed out by Baron (1996) and later by Kalandrakis (2004), however, most legislatures have the authority to change existing laws by enacting new legislation; so that laws continue in effect only in the absence of new legislation. To explore this dynamic feature of legislative policymaking, these authors have introduced an alternative model that casts the classical spatial collective-choice problem into a dynamic bargaining framework. Each period begins with a status quo policy inherited from the previous period, and a legislator is chosen randomly to propose any feasible policy, which is then subject to an up or down vote.

*We are grateful to the editor and two referees for comments that greatly improved the paper.

[†]School of Economics, University of Nottingham. Email: vincent.anesi@nottingham.ac.uk.

[‡]Department of Political Science and Department of Economics, University of Rochester. Email: dugg@ur.rochester.edu.

If the proposal is voted up, then it is implemented in that period and becomes the next period's status quo; if it is voted down, then the ongoing status quo is implemented and remains in place until the next period. This process continues ad infinitum.

The problem immediately encountered in this framework is that existence results for stationary Markov perfect equilibria provided in the extant game-theoretic literature do not apply. The consequence has been a fast growing body of literature consisting of work that explicitly constructs stationary Markovian equilibria for bargaining games with an endogenous status quo, and then analyzes the properties of policy outcomes implied by these constructions (e.g., Baron, 1996; Kalandrakis 2004, 2010, 2014; Bowen and Zahran, 2012; Nunnari, 2014; Richter, 2014; Baron and Bowen, 2014; Zápál, 2014; and Anesi and Seidmann, 2015). These analyses are an important development in the study of legislative dynamics; but almost all either assume that the space of alternatives is unidimensional, or focus on pie-division settings where each bargainer's utility only depends on her own share of the pie. There are no known conditions that guarantee the existence of a stationary Markovian equilibrium for more general multidimensional choice spaces.¹

In this paper, we provide gradient restrictions at a given alternative that are sufficient for existence of a stationary Markov perfect equilibrium in pure strategies with an absorbing set close to that alternative, when players are sufficiently patient. In fact, our gradient restriction holds at *every* interior alternative when there is private good that can be allocated across the players, and it delivers a *continuum* of stationary Markov perfect equilibria with distinct absorbing points close to that alternative. Thus, in a large class of models of dynamic bargaining with an endogenous quo, equilibria are indeterminate—despite the fact that stationary Markovian strategies depend on the history of play only through the current status quo, sharply constraining the ability to punish and reward players for past behavior. Our results have important implications for applied bargaining models of legislative policymaking, where the norm is to construct a particular equilibrium in closed form and to analyze the properties of this equilibrium selection as parameters are varied: when players are sufficiently patient, the implicit equilibrium selection made in such analyses may be a restrictive one, with the danger that insights derived from those analyses are driven by the equilibrium selection, rather than equilibrium incentives in general. In the absence of further justifications for such a selection, the multiplicity of equilibria we highlight suggests limits on the usefulness of these constructions in predicting the policy outcomes and understanding the dynamics and comparative statics of legislative bargain-

¹An exception is Duggan and Kalandrakis (2012), who establish existence of stationary Markovian equilibria in pure strategies for general environments. They modify the basic framework by adding noise to the status quo transition and assuming preference shocks in each period. This paper concentrates on existence conditions that do not rely on such noise.

ing. Studies of dynamic bargaining with an endogenous status quo thus face an important equilibrium refinement issue.

Our analysis allows the feasible set of alternatives to be any nonempty subset of multidimensional Euclidean space, and we rely only on smoothness of utilities to apply techniques from differential topology; we do not impose any functional form restrictions or assume the existence of a private good. The bargaining protocol is the standard one, described above, and we permit the voting rule to be any non-collegial rule, i.e., no player has a veto. Our main indeterminacy result is that when players are sufficiently patient, a continuum of stationary Markov perfect equilibria in pure strategies can be constructed with absorbing sets close to any alternative at which a weak gradient restriction holds. This gradient restriction holds if the gradients of all players' utilities are linearly independent at the given alternative; and, generically, this in turn holds at almost all alternatives when the dimensionality of the set of alternatives is greater than or equal to the number of players. Linear independence of all players' gradients is sufficient for our condition, but not necessary: we provide a more general condition requiring only that linear independence of the gradients of players belonging to a given "oversized" coalition (a decisive coalition that remains decisive if any one member is removed), and that linear independence is maintained if a player outside the coalition is switched with one inside. Given a high-dimensional set of alternatives, our linear independence condition holds generically outside a set of alternatives with measure zero, with the implication that equilibria typically abound in such models.

Though above indeterminacy result makes use of equilibria in which Pareto inefficient alternatives are proposed and passed, we show in our companion working paper (Anesi and Duggan 2017) that when players are sufficiently patient, a second gradient restriction leads to a continuum of stationary Markov perfect equilibria in pure strategies with Pareto optimal absorbing points close to a given Pareto optimal alternative. Thus, the use of Pareto inefficient alternatives is not essential for the indeterminacy result, and refining away equilibria with inefficient outcomes still leaves a continuum of equilibria. We emphasize that both gradient restrictions are easily verified in economic environments and, together, cover many applications encountered in the literature, including pie-division settings and, more generally, the large class of economies with a private good component.

The analysis of equilibrium indeterminacy develops constructive techniques due to Anesi and Seidmann (2015), who establish existence, but not indeterminacy, of stationary Markov perfect equilibria for the pie-division model. The approach rests in identifying possible absorbing sets of equilibria when players are sufficiently patient. In doing so, the authors define the concept of a "simple solution" as a list of alternatives, each associated with a decisive coalition supporting it, such that for every player: the player's utility takes

two values over the list of alternatives, a reward payoff and a punishment payoff; the player is included in some but not all coalitions; and the player receives her reward payoff when included in a coalition and receives her punishment payoff when excluded. The authors show that in the pie-division setting, given any simple solution and assuming sufficiently patient players, there is a stationary Markov perfect equilibrium with an absorbing set that coincides with the simple solution. To capture more general environments, our equilibrium construction uses the concept of “semi-simple solution,” which generalizes Anesi and Seidmann’s (2015) simple solutions by allowing for the possibility of multiple punishment payoffs. This gain in flexibility allows us to push their approach well beyond pie division and to shed many of the assumptions usually made in the dynamic bargaining literature, dropping convexity and compactness of the set of alternatives, and assuming only weak gradient restrictions on players’ utilities.

As mentioned earlier, existence results for stationary Markov perfect equilibria provided in the literature on stochastic games do not apply to the dynamic bargaining framework, as they rely on continuity conditions on the transition probability that are violated in the bargaining model (cf. Duggan 2017 for a more detailed discussion). Existence and characterization results for Markov perfect equilibria have been obtained in alternative frameworks of dynamic bargaining in which the policy space is finite (Anesi 2010; Diermeier and Fong 2011, 2012; and Battaglini and Palfrey 2012) or without discounting (Anesi and Seidmann 2014) or when the set of possible status quos is countable (Duggan 2017).

A roadmap of the paper is as follows. The bargaining framework and equilibrium concept are defined in Section 2. Section 3 presents the concept of semi-simple solution and establishes that when players are sufficiently patient, a semi-simple solution can be obtained as the absorbing set of a stationary Markov perfect equilibrium in pure strategies.² Section 4 shows that given any alternative, if the gradients of members of an oversized coalition are linearly independent, and if linear independence is maintained when a player outside the coalition is switched with one inside, then there is a continuum of semi-simple solutions in an arbitrarily small neighborhood of that alternative. Section 5 combines the above observations and presents our result on indeterminacy of equilibria in dynamic bargaining games. Finally, Appendices A–B contain formal proofs omitted from the text.

²We also present the weaker concept of “mixed” semi-simple solution and show that it can be supported if mixed proposal strategies are permitted in the working paper version of this paper (Anesi and Duggan 2017).

2 Dynamic Bargaining Framework

In each of an infinite number of discrete periods, indexed $t = 1, 2, \dots$, a finite set of players $N \equiv \{1, \dots, n\}$, with $n \geq 3$, must reach a collective choice from a nonempty set of alternatives, $X \subseteq \mathfrak{R}^d$, with nonempty interior. Let x^t denote the alternative chosen in period t . Bargaining takes place as follows. Each period t begins with a status-quo alternative x^{t-1} , in place from the previous period. Player i is selected with probability $p_i \in (0, 1)$ to propose a policy in X ; all players then simultaneously vote to accept or to reject the chosen proposal. It is accepted if a coalition $C \in \mathcal{D}$ of players vote to accept, and it is rejected otherwise, where $\mathcal{D} \subseteq 2^N \setminus \{\emptyset\}$ is the nonempty collection of *decisive coalitions*, which have the authority to decide policy in a given period. If proposal y is accepted, then it is implemented in period t and becomes the status quo next period (i.e., $x^t = y$); otherwise the previous status quo, x^{t-1} , is implemented and remains the status quo in period $t+1$ (i.e., $x^t = x^{t-1}$). This process continues *ad infinitum*. The initial status quo, $x^0 \in X$, is exogenously given.

We assume the voting rule \mathcal{D} is *proper*, i.e., every pair of decisive coalitions has nonempty intersection: $C, C' \in \mathcal{D}$ implies $C \cap C' \neq \emptyset$. In addition, we assume \mathcal{D} is *monotonic*, i.e., any superset of a decisive coalition is itself decisive: $C \in \mathcal{D}$ and $C \subseteq C'$ imply $C' \in \mathcal{D}$. Finally, we assume that \mathcal{D} is *non-collegial*, in the sense that no player has a veto: we have $N \setminus \{i\} \in \mathcal{D}$ for all $i \in N$. Thus, we allow for any *quota rule* defined by $\mathcal{D} = \{C : |C| \geq q\}$, the only restrictions on the quota q being $\frac{n}{2} < q < n$. For future use, we say a decisive coalition C is *oversized* if every member of the coalition is redundant: for all $i \in C$, $C \setminus \{i\} \in \mathcal{D}$. Let \mathcal{D}^* denote the collection of oversized coalitions, and note that $N \in \mathcal{D}^*$, since \mathcal{D} is non-collegial.

The preferences of each player i over lotteries over sequences of alternatives are represented by a von Neumann-Morgenstern stage utility function $u_i : X \rightarrow \mathfrak{R}$ that is twice continuously differentiable and bounded above. Say x is *Pareto optimal* if there is no $y \in X$ such that for all $i \in N$, we have $u_i(y) \geq u_i(x)$, with strict inequality for at least one member of N . Given a sequence of alternatives $\{x^t\} \in X^\infty$, player i 's payoff is the discounted sum $(1 - \delta_i) \sum_{t=1}^{\infty} \delta_i^{t-1} u_i(x^t)$, where $\delta_i \in [0, 1)$ is her discount factor.

A noteworthy special case of our general environment is that of a *mixed economy*, in which an alternative $x = (x_1, \dots, x_n, g)$ consists of a private component $(x_1, \dots, x_n) \in \mathfrak{R}_+^n$ and possibly a public component $g \in \mathfrak{R}_+^{d-n}$, $d \geq n$. Here, the set of alternatives is $X = \{x \in \mathfrak{R}^d : f(-\sum_{i \in N} x_i, g) \leq 0\}$, where $f : \mathfrak{R}_+^{d-n+1} \rightarrow \mathfrak{R}$ is a continuous, weakly monotonic function. We then require that each u_i is strictly increasing in x_i and constant in $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n$; more formally, $\frac{\partial u_i}{\partial x_i}(x) > 0$ and $\frac{\partial u_i}{\partial x_j}(x) = 0$ for all x and all $j \neq i$. We interpret x_i as an amount of a resource allocated to i , and our restriction on

utilities reflects the assumption that there are no consumption externalities in the private good. An obvious example of a mixed economy (where the public good level is fixed at zero and therefore suppressed) is the setting of *pie-division with free disposal*, in which $X = \{(x_1, \dots, x_n) \in [0, 1]^n : \sum_{i \in N} x_i \leq 1\}$ is the n -dimensional unit simplex; but mixed economies with richer policy spaces are also common in the political economy literature (e.g., Jackson and Moselle 2002).

An alternative setting is that of *pie-division with no disposal*, in which the pie must be fully divided; it does not constitute a mixed economy, as defined above. This model, which has received considerable attention in the literature on bargaining (both with and without an endogenous status quo), is also captured by our framework. To this end, define $X = \{(x_1, \dots, x_{n-1}) \in [0, 1]^{n-1} : \sum_{i=1}^{n-1} x_i \leq 1\}$ as the $(n-1)$ -dimensional unit simplex, and assume: (i) $\frac{\partial u_i}{\partial x_i}(x) > 0$ and $\frac{\partial u_i}{\partial x_j}(x) = 0$ for all x , all $i < n$, and all $j \in N \setminus \{i, n\}$; and (ii) there exists a differentiable, real-valued function v on $[0, 1]$ with $v' > 0$ such that $u_n(x) = v(1 - \sum_{i=1}^{n-1} x_i)$ for all $x = (x_1, \dots, x_{n-1})$. The set of alternatives, formulated thusly, has full dimension, and an alternative at which each player consumes a positive amount belongs to the interior of X , so calculus-based methods can be directly applied.

We focus on subgame perfect equilibria in which players use pure stationary Markov strategies, defined as follows. For any player $i \in N$, a *stationary Markov strategy* $\sigma_i = (\pi_i, \alpha_i)$ consists of a proposal strategy $\pi_i: X \rightarrow X$, where $\pi_i(x)$ is the proposal made by player i when the current status quo is x (conditional on her being selected to propose), and a voting strategy $\alpha_i: X^2 \rightarrow \{0, 1\}$, where $\alpha_i(x, y)$ is the (degenerate) probability that i votes to accept a proposal y when the current status quo is x . A *stationary Markov perfect equilibrium* is a subgame perfect equilibrium in which all players use stationary Markov strategies. We follow the standard approach of concentrating throughout on equilibria in stage-undominated voting strategies; i.e., those in which, at any voting stage, no player uses a weakly dominated strategy. Hence, we refer to a pure stationary Markov perfect equilibrium in stage-undominated voting strategies more succinctly as a *stationary bargaining equilibrium*.

Every stationary Markov strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ (in conjunction with recognition probabilities) generates a transition function $P^\sigma: X^2 \rightarrow [0, 1]$, where $P^\sigma(x, y)$ is the probability, given σ , that the alternative implemented in the next period is y , given that the alternative implemented in the current period is x .³ We say that $x \in X$ is an *absorbing point* of σ if and only if $P^\sigma(x, x) = 1$, and we denote the set of absorbing points of σ by $A(\sigma) \equiv \{x \in X : P^\sigma(x, x) = 1\}$. We will say that σ is *no-delay* if and only if: (i) $A(\sigma) \neq \emptyset$;

³As all players use pure strategies, $P^\sigma(x, \cdot)$ is a discrete probability density with $|\text{supp}(P^\sigma(x, \cdot))| \leq n$ for all $x \in X$.

and (ii) for all $x \in X$, there is $y \in A(\sigma)$ such that $P^\sigma(x, y) = 1$. In words, a strategy profile is no-delay if an absorbing point is implemented in every period (both on and off the equilibrium path).

3 Semi-simple Solutions and Existence of Equilibria

Anesi and Seidmann (2015) show that in pie-division settings with free disposal, each simple solution (as defined in the introduction) identifies alternatives that are absorbing points of stationary bargaining equilibria for the corresponding bargaining games when players are sufficiently patient. Simple solutions always exist in the pie division model with free disposal, but there remains the possibility that a more general concept has greater applicability in general environments yet is still sufficient for the construction of stationary bargaining equilibria. In this section, we push the approach of Anesi and Seidmann (2015) further to obtain the existence of stationary bargaining equilibria corresponding to sets of alternatives exhibiting a more general structure. To this end, we propose the weaker concept of semi-simple solution and show that when players are sufficiently patient, every semi-simple solution can be supported as the absorbing set of a stationary bargaining equilibrium. This generalizes the result of Anesi and Seidmann (2015), and as we show in the following section, it allows us to obtain a continuum of stationary bargaining equilibria near any alternative satisfying a general gradient restriction. In turn, this will imply indeterminacy of stationary bargaining equilibria in a broad class of dynamic bargaining games.

Definition 1. *A set of alternatives $S \subset X$ is a **semi-simple solution** if*

(i) *there is a one-to-one mapping $\rho: S \rightarrow N$ such that for all $x \in S$,*

$$\rho(x) \in C(x) \equiv \left\{ i \in N : u_i(x) = \max_{z \in S} u_i(z) \right\} \in \mathcal{D},$$

(ii) *for all $i \in N$, u_i is not constant on S .*

In the definition of semi-simple solution, for any player i , we interpret $\max_{z \in S} u_i(z)$ as the player's reward payoff, and payoffs below this value as punishments. Thus, part (i) requires that for each $x \in S$, there is a decisive coalition that supports x to implement the reward payoff of player $\rho(x)$; and since $u_{\rho(x)}$ is not constant on S , by part (ii), it follows that there is also a decisive coalition willing to implement a punishment payoff for the player. Semi-simple solutions are a generalization of Anesi and Seidmann's (2015) simple solutions,⁴ the critical difference between the two concepts being that, in the definition of a

⁴See Section 3 of our companion working paper for a proof.

semi-simple solution, a player i does not have a single punishment payoff, but instead this can vary across alternatives that do not solve $\max_{z \in S} u_i(z)$. The first part of Definition 1 implies that a semi-simple solution contains at most n alternatives, and the second part implies that the collection $\{C(x) : x \in S\}$ of decisive coalitions has empty intersection. It is known from the social choice literature that the number of coalitions must therefore meet or exceed the Nakamura number of the voting rule; for the quota rule case, this means that any simple solution S must satisfy $|S| \geq \lceil \frac{n}{n-q} \rceil$.⁵

Next, we generalize Anesi and Seidmann's (2015) result to support any semi-simple solution as the absorbing set of a stationary bargaining equilibrium. Thus, if one can show that a semi-simple solution exists, then necessarily the dynamic bargaining game (with discount factors close enough to one) will possess a stationary bargaining equilibrium. However, existence of a semi-simple solution is not generally guaranteed. For instance, it is readily checked that no semi-simple solution can exist in settings with a unidimensional set of alternatives, $n \geq 3$, majority rule, and single-peaked utilities, as in Baron (1996), Kalandrakis (2014), and Zápal (2014).⁶

Theorem 1. *Let S be a semi-simple solution. There is a threshold $\bar{\delta} \in (0, 1)$ such that if $\min_{i \in N} \delta_i > \bar{\delta}$, then there exists a no-delay stationary bargaining equilibrium σ such that $A(\sigma) = S$.*

The proof of Theorem 1 generally follows the lines of the construction of Anesi and Seidmann (2015). To convey the idea of the proof, consider the special case in which $n = 5$, \mathcal{D} is majority rule, and assume that $\delta_i \approx 1$ for each player i . Let $\{\bar{x}^1, \bar{x}^2, \bar{x}^3, \bar{x}^4\}$ be a semi-simple solution with payoffs as depicted below,

	1	2	3	4	5
\bar{x}^1	1	1	1	0	3
\bar{x}^2	1	1	0	1	2
\bar{x}^3	1	0	1	1	1
\bar{x}^4	0	1	1	1	0

so that $C(\bar{x}^1) = \{1, 2, 3, 5\}$, $C(\bar{x}^2) = \{1, 2, 4\}$, $C(\bar{x}^3) = \{1, 3, 4\}$ and $C(\bar{x}^4) = \{2, 3, 4\}$. (To fulfill Definition 1, define the mapping ρ on $\{\bar{x}^1, \bar{x}^2, \bar{x}^3, \bar{x}^4\}$ by: $\rho(\bar{x}^k) = k$ for all $k = 1, 2, 3, 4$.) For each $i \in N$, let C_i be the decisive coalition defined by: $C_i \equiv C(\bar{x}^i)$ if $i \in N \setminus \{5\}$, and $C_5 \equiv C(\bar{x}^1)$. We say that coalition C_i "forms" in period t if alternative \bar{x}^k is

⁵The *Nakamura number* of a general, non-collegial voting rule \mathcal{D} is the size of the smallest collection of decisive coalitions having empty intersection. In the foregoing, $\lceil r \rceil$ is the smallest integer greater than or equal to real number r .

⁶Also note that a semi-simple solution would not exist if \mathcal{D} were collegial, as in Nunnari (2014).

implemented, i.e., if $x^t = \bar{x}^k$. Observe that the players in C_i , and only these players, receive their reward payoff when C_i forms. Now consider a stationary Markov strategy profile σ with the following features. In every period t , the current status quo reveals whether a coalition formed in period $t - 1$ or not; then events unfold as follows: (a) if no coalition in $\{C_j\}_{j \in N}$ formed in period $t - 1$, then every proposer i offers to form coalition C_i , and this proposal is accepted by all the members of C_i ; (b) if a coalition in $\{C_j\}_{j \in N}$, say C_i , formed in $t - 1$, then every proposer j offers to form coalition C_i again in t by proposing \bar{x}^i ; (c) if a coalition in $\{C_j\}_{j \in N}$, say C_i , formed in $t - 1$, then any (off-the-equilibrium-path) proposal that does not involve forming C_i is rejected by the members of C_i . Observe that on the path of play prescribed by σ , one of the alternatives in $\{\bar{x}^1, \bar{x}^2, \bar{x}^3, \bar{x}^4, \bar{x}^5\}$ is implemented in the first period and then never amended: in the first period, some coalition C_i forms with probability p_i and then forms in all future periods with probability one. The same occurs following any period in which no coalition formed. Hence, each player's maximum utility in each period t is given by her reward payoff, which she obtains if and only if she is a member of the coalition that forms in t .

Suppose that coalition C_i formed in period $t - 1$ and that, contrary to (b) above, the proposer selected in period t does not offer to form C_i again. By accepting such a proposal, each member j of C_i would face the risk of not being a member of the coalition that will form and, therefore, of not receiving her reward payoff in all future periods. This implies that, in period t , it is always profitable for the farsighted members of coalition C_i to oppose any proposal that does not involve forming C_i again, as prescribed by (c). As C_i is a majority coalition, it is thus impossible for any proposer j to prevent C_i from forming in period t . Therefore, proposer j cannot profitably deviate from passing, as prescribed by (b). By the same logic, if no coalition formed in period $t - 1$, then no player j can improve on accepting an offer to form coalition $C_i \ni j$, thereby obtaining her reward payoff with probability one in all future periods. It is therefore optimal for proposer i to (successfully) offer to form C_i , as prescribed by (a).

4 Multiplicity of Semi-simple Solutions

To leverage Theorem 1, it remains to find conditions under which semi-simple solutions exist. Our approach is to exploit restrictions on the gradients of players' utility functions at an interior alternative x that are sufficient for existence of semi-simple solutions near x . To provide a preliminary intuition, consider the case of the pie division with free disposal. In this setting, one can easily satisfy the conditions of Definition 1 at any interior alternative $x = (x_1, \dots, x_n)$ by transferring shares of the pie between players: for example, given

sufficiently small $\epsilon > 0$, we can construct \bar{x}^i by transferring $(n - 1)\epsilon$ from player i to the other players, increasing the consumption of every other player by ϵ . Then $\{\bar{x}^1, \dots, \bar{x}^n\}$ is a semi-simple solution in which each player i receives her unique punishment payoff from \bar{x}^i and her reward payoff from \bar{x}^j for $j \neq i$. Of course, a continuum of semi-simple solutions can thus be obtained arbitrarily close to x by varying ϵ . The crux of the proof of our first existence result is to translate this simple argument into more general environments, where we cannot freely adjust the players' payoffs by re-allocating a transferable private good. This is where our first gradient restriction, called (C1), kicks in: it allows us to apply the logic of the pie-division model with free disposal to construct a continuum of semi-simple solutions in any open neighborhood of an alternative satisfying the condition. Condition (C1) can fail in pie division models with no disposal (this occurs when \mathcal{D} is a quota rule with $q = n - 1$), but at the end of this section, we state a result that yields a continuum of semi-simple solutions near any Pareto optimal alternative satisfying a gradient restriction that is satisfied in such settings.

More precisely, condition (C1) requires that the gradients of the members of some oversized coalition C^* are linearly independent at an interior alternative x ; moreover, there is some member j of the coalition such that if we switch j with any non-member k , the gradients of members of the resulting coalition $(C^* \setminus \{j\}) \cup \{k\}$ remain linearly independent. The condition is obviously implied if the gradients $\{\nabla u_i(x) : i \in N\}$ of all players are linearly independent, for in this case, the second requirement is in fact vacuously satisfied.

(C1) There is an oversized coalition $C^* \in \mathcal{D}^*$ such that (i) the gradients $\{\nabla u_i(x) : i \in C^*\}$ are linearly independent, and (ii) there exists $j \in C^*$ such that for all $k \in N \setminus C^*$, the gradients $\{\nabla u_i(x) : i \in (C^* \setminus \{j\}) \cup \{k\}\}$ are linearly independent.

Intuitively, when condition (C1) holds at x , we can obtain all values of the utility profile $u = (u_i)_{i \in C^*}$ in some open neighborhood of $u(x) \in \mathfrak{R}^{|C^*|}$ by arbitrarily small variations of x , i.e., the Jacobian of u at x has full row rank. As mentioned above, a sufficient condition for (C1) that is satisfied in many economic applications is that the collection $\{\nabla u_i(x) : i \in N\}$ of all players' gradients is linearly independent. The latter condition is satisfied whenever the set of alternatives has a private good component; *a fortiori*, given our formulation of mixed economies, (C1) holds at any alternative x in the interior of X . In particular, (C1) holds in the model of pie-division with free disposal, as studied in Baron and Bowen (2014), Richter (2014), and Anesi and Seidmann (2015). Letting m^* denote the size of the smallest oversized coalition, i.e., $m^* = \min\{|C| : C \in \mathcal{D}^*\}$, condition (C1) holds as well in the model of pie division with no disposal as long as $m^* < n$, for then if a coalition of

size m^* includes player n , it must exclude some player $i < n$, and the members' gradients will be linearly independent. More generally, (C1) holds if for every size m^* coalition C , the gradients $\{\nabla u_i(x) : i \in C\}$ are linearly independent; and for a quota rule, this reduces to the requirement that the gradients of every coalition with $q + 1$ members are linearly independent at x . In multidimensional settings with $d \geq m^*$, Schofield's (1980) Singularity Theorem A establishes that for generic profiles of utility functions, (C1) holds at every alternative outside a union of manifolds of dimension $m^* - 1$ or less; that is, (C1) generically holds on a closed set of alternatives with measure zero.⁷

The next result establishes a continuum of semi-simple solutions in non-collegial, dynamic bargaining games under (C1). In fact, we show that given any interior alternative satisfying (C1), we can find a continuum of semi-simple solutions arbitrarily close to that alternative.

Theorem 2. *Let x be any interior point of X at which (C1) is satisfied. Every open neighborhood U of x contains a continuum of semi-simple solutions.*

For a sketch of the proof, consider a simple three-player majority-voting game, let x be any interior point of X , and for simplicity assume that condition (C1) is satisfied at x . Since the only oversized coalition is $N = \{1, 2, 3\}$, condition (C1) implies that $\{\nabla u_i(x) : i = 1, 2, 3\}$ is linearly independent. Our approach, in the context of this example, is to find a set of alternatives $\{\bar{x}^1, \bar{x}^2, \bar{x}^3\}$ such that for each alternative $k = 1, 2, 3$, the coalition supporting \bar{x}^k is $C(\bar{x}^k) = N \setminus \{k\}$, whereas player k receives her punishment payoff from \bar{x}^k . Then each $C(\bar{x}^k)$ is decisive, and we fulfill Definition 1 by specifying the mapping ρ so that $\rho(\bar{x}^k) = k + 1$ for $k = 1, 2$, and $\rho(\bar{x}^3) = 1$. To this end, define f as the function that maps vectors of alternatives $(x^1, x^2, x^3) \in X^3$ to corresponding utility vectors $(u_i(x^j))_{i,j \in N} \in \mathfrak{R}^9$. The argument is depicted in Figure 1, where we place $(u_1(x), u_2(x), u_3(x))$ at the center of the simplex in \mathfrak{R}^3 . Condition (C1) implies that the Jacobian of f has full row rank at x . By the local submersion theorem (e.g., Guillemin and Pollack, 1974), therefore, we can select alternatives $\bar{x}^1, \bar{x}^2, \bar{x}^3$ near x so as to give each player i her punishment payoff at \bar{x}^i and the remaining players their reward payoffs, e.g., for sufficiently small $\epsilon > 0$, we can set $u_i(\bar{x}^i) = u_i(x) - \epsilon$, whereas $u_j(\bar{x}^i) = u_j(x) + \epsilon$ for each $j \in N \setminus \{i\}$. Thus, $\{\bar{x}^1, \bar{x}^2, \bar{x}^3\}$ is a semi-simple solution. It is readily checked that we can use the same argument for a continuum of values of ϵ that each yield a different semi-simple solution.

⁷Schofield's (1980) result holds if we give the space of twice continuously differentiable utility profiles the Whitney topology. Condition (C1) fails at alternatives x such that the rank of $\{\nabla u_i(x) : i \in C^*\}$ has rank $m^* - r$ for $r \geq 1$. Setting $w = d$ and $z = m^*$ in Schofield's theorem, the claim follows. Smale (1974) establishes a similar result for the case of exchange economies.




Figure 1: Mapping to utility vectors

The preceding argument illustrates the proof approach in a particular example, where the only oversized coalition is the entire set of players, N itself. This is necessarily the case if the voting rule is a quota rule with $q = n - 1$, but more generally (as in the case of majority rule with $n \geq 5$) the gradient restriction imposed in (C1) will be weaker when satisfied using a smaller coalition $C^* \subsetneq N$. Returning to the example of Figure 1, maintain majority rule, but now assume $n = 5$, and assume that (C1) is satisfied with $C^* = \{1, 2, 3, 4\}$. In this case, the selection of four alternatives $\{x^1, \bar{x}^2, \bar{x}^3, \bar{x}^4\}$ proceeds as above and satisfies part (i) of Definition 1—but we must address the possibility that a player outside C^* , namely player 5, is indifferent over $\{x^1, \bar{x}^2, \bar{x}^3, \bar{x}^4\}$. In case player 5 is not indifferent, then we set $\bar{x}^1 = x^1$; and otherwise, we perturb x^1 as follows. Following the argument of Figure 1, we can, by varying x^1 in an arbitrarily small open set, vary the payoffs of players in $(C^* \setminus \{1\}) \cup \{5\}$ in an open set around $(u_2(x^1), u_3(x^1), u_4(x^1), u_5(x^1))$. Thus, we can perturb x^1 to \bar{x}^1 such that the utilities of players 2, 3, and 4 are unchanged (preserving their reward and punishment payoffs), and such that player 5's utility changes; moreover, we can make the perturbation small enough that player 1's payoff from \bar{x}^1 remains less than her reward payoff. In both cases, we obtain a set $\{\bar{x}^1, \bar{x}^2, \bar{x}^3, \bar{x}^4\}$ such that part (i) of Definition 1 is preserved, and such that player 5 is not indifferent over $\{\bar{x}^1, \bar{x}^2, \bar{x}^3, \bar{x}^4\}$, fulfilling part (ii); that is, the set is a semi-simple solution. In the general case, if multiple players outside C^* are indifferent over the alternatives in the provisional solution, then we iterate this procedure for each one.

Remark on pie division with no disposal and $m^* = n$. As mentioned above, Theo-

rem 2 does not cover the model of pie division with no disposal when $m^* = n$. Nevertheless, the result can be adapted to cover that environment by adding the assumption that x is Pareto optimal and weakening the gradient restriction of (C1) so that the gradients of every size $n - 1$ coalition are linearly independent.

(C1*) (i) x is Pareto optimal, and (ii) for all $i \in N$, the gradients $\{\nabla u_j(x) : j \in N \setminus \{i\}\}$ are linearly independent.

Although (C1) fails in the model of pie division with no disposal when $m^* = n$, condition (C1*) holds at all interior alternatives in this environment, capturing the settings of Kalandrakis (2004 and 2010) and Bowen and Zahran (2012). These and the previous examples show that conditions (C1) and (C1*) are easy to check and that they apply to a range of economic environments of interest, and the following variant of Theorem 2 shows that (C1*) also leads to a continuum of semi-simple solutions near any alternative satisfying the condition.⁸

Theorem 2*. *Let x be any interior point of X at which (C1*) is satisfied. Every open neighborhood U of x contains a continuum of semi-simple solutions.*

5 Indeterminacy of Stationary Bargaining Equilibria

Theorem 1, combined with Theorem 2, immediately yields an equilibrium existence result for the dynamic bargaining game: as discount factors become close to one, absorbing points of stationary bargaining equilibria exist near every alternative that satisfies (C1). Of greater importance are the implications of these results for the predictive power of stationary bargaining equilibria in this class of games: when players are sufficiently patient, the dynamic bargaining game admits a *continuum* of equilibria. The next result establishes indeterminacy of stationary bargaining equilibria when any alternative satisfies (C1) and players are sufficiently patient.

Theorem 3. *Let x be any interior point of X at which (C1) is satisfied. For every open neighborhood U of x , there exists $\hat{\delta} \in (0, 1)$ such that if $\min_{i \in N} \delta_i > \hat{\delta}$, then there is a continuum of semi-simple solutions in U corresponding to absorbing sets of no-delay stationary bargaining equilibria with discount factors $\delta_1, \dots, \delta_n$.*

Note that the theorem is not an immediate corollary of Theorems 1 and 2. Indeed, the threshold $\bar{\delta} = \bar{\delta}_S$ identified in Theorem 1 was only shown to apply to a given semi-simple

⁸The proof of this theorem can be found in the working paper version (Anesi and Duggan 2017).

solution S , not to the continuum of semi-simple solutions in U described in Theorem 2. Henceforth, let $\bar{\delta}_S$ be the threshold associated with S , and let \mathcal{S} be the continuum of semi-simple solutions from Theorem 2. It remains to be established that there is a subcontinuum, say \mathcal{S}^* , of that continuum such that $\hat{\delta} = \sup \{\bar{\delta}_S : S \in \mathcal{S}^*\} < 1$. To this end, let $U \subseteq X$ be an arbitrary open neighborhood of x . For each natural number k , set $\mathcal{S}^k = \{S \in \mathcal{S} : \bar{\delta}_S < 1 - \frac{1}{k}\}$, and note that $\mathcal{S} = \bigcup_{k=1}^{\infty} \mathcal{S}^k$. Thus, \mathcal{S} is the union of countably many sets, and since \mathcal{S} is a continuum, some set \mathcal{S}^k is also a continuum.⁹ Then we set $\mathcal{S}^* = \mathcal{S}^k$ and $\hat{\delta} = 1 - \frac{1}{k}$ to complete the proof.

The political economy literature on bargaining games with an endogenous status quo has devoted considerable attention to the set A^* of *dynamically stable* alternatives, i.e., the alternatives that can be supported as long run outcomes of stationary bargaining equilibria. Formally, we define A^* to consist of every alternative x for which there exists $\hat{\delta} \in (0, 1)$ such that if $\min_{i \in N} \delta_i > \hat{\delta}$, then there is a stationary bargaining equilibrium σ for discount factors $\delta_1, \dots, \delta_n$ such that $x \in A(\sigma)$. In terms of predicting bargaining outcomes, the characterization of dynamically stable alternatives is only informative if A^* is “small” relative to the set of alternatives. This is typically not the case in the dynamic bargaining game with a high-dimensional set of alternatives. As discussed above, when $d \geq m^* - 1$, for generic utility profiles, condition (C1) is satisfied on a set of alternatives with full measure, and thus the dynamically stable alternatives are dense in the set of alternatives.

Corollary 1. *If the set of alternatives at which (C1) holds is dense in $\text{int}X$, then the set A^* of dynamically stable alternatives is dense in $\text{int}X$.*

This observation is reminiscent of the cycling results in the social choice literature (e.g., McKelvey, 1979). Just as the top cycle is generically dense in the set of alternatives in sufficiently high dimensional spaces, we find that long-run bargaining outcomes for any such environment are highly indeterminate. Whereas McKelvey’s chaos theorem evokes the picture of collective choices moving arbitrarily through the set of alternatives over time, however, our results establish the possibility that collective choices via dynamic bargaining can come to rest at arbitrary locations in the set of alternatives.

We conclude that as players become patient, stationary bargaining equilibria may not only be indeterminate, but when the space of alternatives is high-dimensional, every interior alternative can be approximated by equilibrium absorbing points. Note that the indeterminacy is not created by the possibility of Pareto inefficient equilibria: our companion working paper (Anesi and Duggan 2017) provides a second gradient restriction

⁹More precisely, if we had $|\mathcal{S}^k| < |\mathcal{S}|$ for all k , then it would follow from König’s Theorem (e.g., Holz et al., 1999) that $|\mathcal{S}| \leq \sum_{k=1}^{\infty} |\mathcal{S}^k| < \prod_{k=1}^{\infty} |\mathcal{S}| = |\mathcal{S}|$, a contradiction.

under which near a given Pareto optimal alternative, we can find a continuum of semi-simple solutions involving only Pareto optimal alternatives. In the absence of *a priori* bounds on the players' discount factors, our results demonstrate difficulties for the prediction and analysis of outcomes in dynamic settings that must be addressed in future work on dynamic bargaining with an endogenous status quo.

A Proof of Theorem 1

Let $S = \{\bar{x}^1, \dots, \bar{x}^m\}$, $m \leq n$, be a semi-simple solution. We will proceed in six steps. Step 1 defines the threshold $\bar{\delta} = \bar{\delta}(\bar{x}^1, \dots, \bar{x}^m)$. Steps 2 and 3 construct σ and verify that it is a no-delay stationary Markov strategy profile. Step 4 derives players' continuation values from the definition of σ . Finally, Steps 5 and 6 use these continuation values to establish that σ is a stationary bargaining equilibrium.

Step 1: Definition of $\bar{\delta}(\bar{x}^1, \dots, \bar{x}^m)$. Let $p^{\min} \equiv \min_{i \in N} p_i$ and, for each $i \in N$, let

$$\bar{\delta}_i(\bar{x}^1, \dots, \bar{x}^m) \equiv \frac{u_i^{\sup} - \max_h u_i(\bar{x}^h)}{u_i^{\sup} - p^{\min} \min_h u_i(\bar{x}^h) - (1 - p^{\min}) \max_h u_i(\bar{x}^h)},$$

where $u_i^{\sup} > \max_h u_i(\bar{x}^h)$ is an upper-bound for $u_i(X)$ (recall that the u_i 's are bounded above). By condition (ii) in the definition of a semi-simple solution,

$$p^{\min} \min_h u_i(\bar{x}^h) + (1 - p^{\min}) \max_h u_i(\bar{x}^h) < \max_h u_i(\bar{x}^h)$$

and, therefore, $\bar{\delta}_i(\bar{x}^1, \dots, \bar{x}^m) \in (0, 1)$ for every $i \in N$. This in turn implies that

$$\bar{\delta}(\bar{x}^1, \dots, \bar{x}^m) \equiv \max_{i \in N} \bar{\delta}_i(\bar{x}^1, \dots, \bar{x}^m) \in (0, 1).$$

Moreover, as u_i is continuous for each i , $\bar{\delta}_i$ is continuous in $(\bar{x}^1, \dots, \bar{x}^m)$. Hence, $\bar{\delta}$ is a continuous function of $(\bar{x}^1, \dots, \bar{x}^m)$.

Henceforth, we assume that $\min_{i \in N} \bar{\delta}_i > \bar{\delta}(\bar{x}^1, \dots, \bar{x}^m)$.

Step 2: Definition of stationary Markov strategy profile σ . The definition of proposal strategies relies on an n -tuple of alternatives $(\hat{x}^1, \dots, \hat{x}^n) \in \{\bar{x}^1, \dots, \bar{x}^m\}^n$, defined as follows. From condition (i) in Definition 1, there exists a one-to-one mapping $\rho: \{\bar{x}^1, \dots, \bar{x}^m\} \rightarrow N$ such that for all $k = 1, \dots, m$, $\rho(\bar{x}^k) \in \{i \in N: u_i(\bar{x}^k) = \max_h u_i(\bar{x}^h)\}$. For each $i \in \rho(\{\bar{x}^1, \dots, \bar{x}^m\})$, we define \hat{x}^i as the alternative \bar{x}^k in the semi-simple solution such that $\bar{x}^k = \rho^{-1}(i)$. If $N \setminus \rho(\{\bar{x}^1, \dots, \bar{x}^m\}) \neq \emptyset$ then, for each $i \notin \rho(\{\bar{x}^1, \dots, \bar{x}^m\})$, we select \hat{x}^i among the maximizers of u_i on $\{\bar{x}^1, \dots, \bar{x}^m\}$. Observe that, by condition (ii) in

Definition 1, we have $u_i(\hat{x}^i) = \max_h u_i(\bar{x}^h) > \min_h u_i(\bar{x}^h)$, for all $i \in N$. Moreover, since ρ is one-to-one, we also have

$$\sum_{j \in N} p_j u_i(\hat{x}^j) \leq p^{\min} \min_h u_i(\bar{x}^h) + (1 - p^{\min}) \max_h u_i(\bar{x}^h),$$

for all $i \in N$.

We are now in a position to define $\sigma = (\sigma_1, \dots, \sigma_n)$. For each $i \in N$, σ_i prescribes the following behavior to player i :

(a) In the proposal stage of any period t with ongoing status quo x , (conditional on her being selected to make a proposal) she proposes

$$\phi^i(x) \equiv \begin{cases} x & \text{if } x \in \{\bar{x}^1, \dots, \bar{x}^m\}, \\ \hat{x}^i & \text{otherwise;} \end{cases}$$

(b) in the voting stage of any period t with ongoing status quo x , she accepts proposal y if and only if $W_i(y) > W_i(x)$,¹⁰ where

$$W_i(z) \equiv (1 - \delta_i)u_i(z) + \delta_i \sum_{j \in N} p_j u_i(\phi^j(z)) , \text{ for all } z \in X .$$

It is easy to see that σ is stationary Markov.

Step 3: Verification that σ is no-delay with $A(\sigma) = \{\bar{x}^1, \dots, \bar{x}^m\}$. It follows immediately from part (a) in the definition of σ that every element of the semi-simple solution is absorbing, that is, $\{\bar{x}^1, \dots, \bar{x}^m\} \subseteq A(\sigma)$. What remains to be established, therefore, is that any status quo $x \notin \{\bar{x}^1, \dots, \bar{x}^m\}$ is immediately amended to an alternative in $\{\bar{x}^1, \dots, \bar{x}^m\}$ with probability one. To see this, observe that each proposer $i \in N$ offers alternative \hat{x}^i in $\{\bar{x}^1, \dots, \bar{x}^m\}$ when the status quo is $x \notin \{\bar{x}^1, \dots, \bar{x}^m\}$. Moreover, by definition of a semi-simple solution, there is a decisive coalition $C(\hat{x}^i) \in \mathcal{D}$ such that

$$\begin{aligned} W_k(\hat{x}^i) &= u_k(\hat{x}^i) \\ &= \max_h u_k(\bar{x}^h) \\ &> (1 - \delta_k)u_k^{\sup} + \delta_k [p^{\min} \min_h u_k(\bar{x}^h) + (1 - p^{\min}) \max_h u_k(\bar{x}^h)] \\ &\geq (1 - \delta_k)u_k(x) + \delta_k \sum_{j \in N} p_j u_k(\hat{x}^j) \\ &= W_k(x), \end{aligned}$$

¹⁰These voting strategies are, in a sense, simpler than those used in Anesi and Seidmann's (2015) result with simple solutions, where indifferent voters reject proposals in some cases but not in others. This difference in voting behavior is immaterial for the derivation of the result; but assuming that indifferent voters always reject eases the exposition in our more general framework.

for all $k \in C(\hat{x}^i)$. (The first inequality follows from Step 1: $\delta_k > \bar{\delta}_k(\bar{x}^1, \dots, \bar{x}^m)$ for all $k \in N$.) From part (b) in the definition of σ , all players in the decisive coalition $C(\hat{x}^i)$ accept $\hat{x}^i \in A(\sigma)$, which is therefore implemented.

Step 4: Continuation values. We denote by $V_i(x|\sigma)$ player i 's expected discounted payoff from implementing alternative x in a given period. Suppose first that $x \in \{\bar{x}^1, \dots, \bar{x}^m\}$. It follows immediately from part (a) in the definition of σ that

$$V_i(x|\sigma) = u_i(x) = (1 - \delta_i)u_i(x) + \delta_i \sum_{j \in N} p_j u_i(\phi^j(x)) = W_i(x).$$

Suppose now that $x \notin \{\bar{x}^1, \dots, \bar{x}^m\}$. Each player i receives $(1 - \delta_i)u_i(x)$ in the current period. Then, in the next period, player j is selected with probability p_j and, as shown in Step 3, successfully proposes $\hat{x}^j = \phi^j(x)$. Hence,

$$V_i(x|\sigma) = (1 - \delta_i)u_i(x) + \delta_i \sum_{j \in N} p_j u_i(\phi^j(x)) = W_i(x), \quad (1)$$

for all $i \in N$ and all $x \in X$.

Step 5: Verification that players do not cast stage-dominated votes. Consider an arbitrary voting stage, in which a proposal y has been made to amend the current status quo x . Coupled with (1), part (b) in the definition of σ guarantees that each player i only accepts y if $V_i(y|\sigma) > V_i(x|\sigma)$, and only rejects y if $V_i(y|\sigma) \leq V_i(x|\sigma)$.

Step 6: Verification that σ is a stationary bargaining equilibrium. It follows from Step 5 (and the one-shot deviation principle) that, in any voting stage, no player can profitably deviate from σ . To complete the proof of Theorem 1, we must therefore show that there is no profitable (one-shot) deviation from σ in any proposal stage. Suppose first that the current status quo x belongs to $\{\bar{x}^1, \dots, \bar{x}^m\}$. In this case, σ prescribes proposer i to maintain x . If she deviates by proposing to change x to any other alternative $y \neq x$, then her proposal will be rejected. Indeed, if y also belongs to $\{\bar{x}^1, \dots, \bar{x}^m\}$, then part (i) in Definition 1 implies that there is a decisive coalition $C(x) \in \mathcal{D}$ such that $W_i(x) = u_i(x) \geq u_i(y) = W_i(y)$ for all $i \in C(x)$; if y does not belong to $\{\bar{x}^1, \dots, \bar{x}^m\}$, then similarly part (i) in Definition 1 implies that, for each member i of the decisive coalition

$C(x)$,

$$\begin{aligned}
W_i(x) &= u_i(x) \\
&= \max_h u_i(\bar{x}_h) \\
&> (1 - \delta_i)u_i^{\sup} + \delta_i [p^{\min} \min_h u_i(\bar{x}_h) + (1 - p^{\min}) \max_h u_i(\bar{x}_h)] \\
&\geq (1 - \delta_i)u_i(y) + \delta_i \sum_{j \in N} p_j u_i(\hat{x}^j) \\
&= (1 - \delta_i)u_i(y) + \delta_i \sum_{j \in N} p_j u_i(\phi^j(y)) \\
&= W_i(y) .
\end{aligned}$$

(The first inequality follows from Step 1: $\delta_i > \bar{\delta}_i(\bar{x}^1, \dots, \bar{x}^m)$ for all $i \in N$.) Hence, all members of the decisive coalition reject y in both cases. It is therefore impossible for any proposer to profitably deviate from σ when the current status quo is in $\{\bar{x}^1, \dots, \bar{x}^m\}$.

Suppose now that the current status quo x does not belong to $\{\bar{x}^1, \dots, \bar{x}^m\}$. If proposer i plays according to σ , then she successfully proposes \hat{x}^i (recall Step 3), thus obtaining a dynamic payoff of $V_i(\hat{x}^i|\sigma) = u_i(\hat{x}^i) = \max_h u_i(\bar{x}_h)$. Because she could simply propose x itself, instead of proposing an alternative that is rejected, it follows that if she has a profitable deviation, then she can profit from making a successful proposal y . If y belongs to $\{\bar{x}^1, \dots, \bar{x}^m\}$, then the deviation is not profitable since $V_i(\hat{x}^i|\sigma) = \max_h u_i(\bar{x}_h) \geq u_i(y) = V_i(y|\sigma)$. If y does not belong to $\{\bar{x}^1, \dots, \bar{x}^m\}$, then

$$V_i(\hat{x}^i|\sigma) = \max_h u_i(\bar{x}_h) > (1 - \delta_i)u_i(y) + \delta_i \sum_{j \in N} p_j u_i(\phi^j(y)) = V_i(y|\sigma),$$

where as above, the inequality follows from $\delta_i > \bar{\delta}_i(\bar{x}^1, \dots, \bar{x}^m)$. Hence, the deviation is again unprofitable. This completes the proof of the theorem.

B Proof of Theorem 2

Let x be an interior point of X that satisfies (C1) using coalition $C^* \in \mathcal{D}^*$, i.e., (i) the gradients $\{\nabla u_i(x) : i \in C^*\}$ are linearly independent, and (ii) there exists $j \in C^*$ such that for all $k \in N \setminus C^*$, the gradients $\{\nabla u_i(x) : i \in (C^* \setminus \{j\}) \cup \{k\}\}$ are linearly independent. For simplicity, enumerate the members of C^* as $\{1, \dots, m\}$, and assume without loss of generality that player $j = 1$ fulfills part (ii) of (C1). Now, let $U \subseteq X$ be an open

neighborhood of x . Define the mapping $f: X^m \rightarrow \mathfrak{R}^{m^2}$ by

$$f(x^1, \dots, x^m) = \begin{pmatrix} u_1(x^1) \\ \vdots \\ u_1(x^m) \\ \vdots \\ u_m(x^1) \\ \vdots \\ u_m(x^m) \end{pmatrix},$$

where for each player $i = 1, \dots, m$, there are m rows giving player i 's payoff from alternatives x^1, \dots, x^m . The derivative of f at arbitrary $(x^1, \dots, x^m) \in X^m$ is the $m^2 \times md$ matrix

$$Df(x^1, \dots, x^m) = \begin{bmatrix} Du_1(x^1) & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & Du_1(x^m) \\ \\ Du_2(x^1) & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & Du_2(x^m) \\ \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \\ Du_m(x^1) & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & Du_m(x^m) \end{bmatrix},$$

where we view $Du_i(x^j)$ as a $1 \times d$ row matrix. By assumption, the matrix $Df(x^1, \dots, x^m)$ has full row rank at (x, \dots, x) . Moreover, we have $d \geq m$, since the players' gradients are linearly independent, and therefore $dm \geq m^2$.

Let $y = (y^1, \dots, y^m) = f(x, \dots, x)$, where $y^i = (u_i(x), \dots, u_i(x))$ is then the m -fold copy of player i 's utility from x . By the local submersion theorem (e.g., Guillemin and Pollack, 1974), we can choose an arbitrarily small open set $\tilde{U} \subseteq U$ containing x such that

the image $\tilde{V} \equiv f(\tilde{U}^m)$ is an open set containing y . Therefore, there exists $\epsilon > 0$ such that

$$y_\epsilon = \begin{pmatrix} u_1(x) - \epsilon \\ u_1(x) + \epsilon \\ \vdots \\ u_1(x) + \epsilon \\ \\ u_2(x) + \epsilon \\ u_2(x) - \epsilon \\ u_2(x) + \epsilon \\ \vdots \\ u_2(x) + \epsilon \\ \\ \vdots \\ \\ u_m(x) + \epsilon \\ \vdots \\ u_m(x) + \epsilon \\ u_m(x) - \epsilon \end{pmatrix}$$

belongs to \tilde{V} . Moreover, by part (ii) of (C1), we can choose \tilde{U} sufficiently small that for all $\tilde{x} \in \tilde{U}$ and all $k = m + 1, \dots, n$, the gradients $\{\nabla u_i(\tilde{x}) : i = 2, \dots, m, k\}$ are linearly independent.

Since $y_\epsilon \in \tilde{V}$, there is a vector $(x_\epsilon^1, \dots, x_\epsilon^m) \in \tilde{U}^m$ such that $f(x_\epsilon^1, \dots, x_\epsilon^m) = y_\epsilon$. We claim that for all $i = 1, \dots, m$, u_i is not constant on $\{x_\epsilon^1, \dots, x_\epsilon^m\}$. Indeed, for each $k = 1, \dots, m$ and each $i \in C^* \setminus \{k\}$, we have

$$u_i(x_\epsilon^k) = u_i(x) + \epsilon = \max_h u_i(x_\epsilon^h),$$

which is strictly greater than $u_i(x_\epsilon^i) = u_i(x) - \epsilon = \min_h u_i(x_\epsilon^h)$, as claimed. To fulfill part (ii) of Definition 1, it remains for us to confront the possibility that players outside C^* are indifferent over the m alternatives, i.e., for some $i \in N \setminus C^*$, u_i is constant on $\{x_\epsilon^1, \dots, x_\epsilon^m\}$. To address this problem, we iteratively perturb x_ϵ^1 , and we construct these perturbations recursively to break any indifferences over $\{x_\epsilon^1, \dots, x_\epsilon^m\}$ among players $m + 1, \dots, n$, while continuing to punish player 1 and maintaining the reward payoffs of players $2, \dots, m$. To begin, set $z^0 = x_\epsilon^1$.

Step 1: If player $m + 1$ is indifferent over $\{z^0, x_\epsilon^2, \dots, x_\epsilon^m\}$, then we form the coalition

$C^1 = (C^* \setminus \{1\}) \cup \{m+1\}$, which has size m . Define the mapping $g^1: X \rightarrow \mathfrak{R}^m$ by

$$g^1(z) = \begin{pmatrix} u_2(z) \\ \vdots \\ u_m(z) \\ u_{m+1}(z) \end{pmatrix},$$

which gives the vector of payoffs to players in C^1 . Since $z^0 \in \tilde{U}$, the gradients $\{\nabla u_i(z^0) : i \in C^1\}$ are linearly independent, so that $Dg^1(z^0)$ has full row rank. By the local submersion theorem, we can choose an open set $\tilde{U}^1 \subseteq \tilde{U}$ containing z^0 such that $\tilde{V}^1 \equiv g^1(\tilde{U}^1)$ is open. Moreover, we can choose \tilde{U}^1 sufficiently small that for all $z \in \tilde{U}^1$, we have

$$u_1(x) > u_1(z).$$

Since \tilde{V}^1 is open and contains $g^1(z^0)$, there exists $\epsilon^1 > 0$ such that

$$y^1 = \begin{pmatrix} u_2(x_\epsilon^1) \\ \vdots \\ u_m(x_\epsilon^1) \\ u_{m+1}(z^0) + \epsilon^1 \end{pmatrix}$$

belongs to \tilde{V}^1 . Since $y^1 \in \tilde{V}^1$, there is an alternative $z^1 \in \tilde{U}^1$ such that $g^1(z^1) = y^1$. If player $m+1$ is not indifferent over $\{z^0, x_\epsilon^2, \dots, x_\epsilon^m\}$, then set $z^1 = z^0$.

In general, assume we are given $z^{k-1} \in \tilde{U}$ such that (a) $u_1(x) > u_1(z^{k-1})$, (b) for all $i = 2, \dots, m$, we have $u_i(z^{k-1}) = u_i(x_\epsilon^1)$, and (c) for all $i = m+1, \dots, m+k-1$, u_i is not constant on $\{z^{k-1}, x_\epsilon^2, \dots, x_\epsilon^m\}$. Then we proceed as follows.

Step k: If player $m+k$ is indifferent over $\{z^{k-1}, x_\epsilon^2, \dots, x_\epsilon^m\}$, then we form the coalition $C^k = (C^* \setminus \{1\}) \cup \{m+k\}$, which has size m . Define the mapping $g^k: X \rightarrow \mathfrak{R}^m$ by

$$g^k(z) = \begin{pmatrix} u_2(z) \\ \vdots \\ u_m(z) \\ u_{m+k}(z) \end{pmatrix},$$

which gives the vector of payoffs to players in C^k . Since $z^{k-1} \in \tilde{U}$, the gradients $\{\nabla u_i(z^{k-1}) : i \in C^k\}$ are linearly independent, so that $Dg^k(z^{k-1})$ has full row rank. By the local submersion theorem, we can choose an open set $\tilde{U}^k \subseteq \tilde{U}$ containing z^{k-1} such that $\tilde{V}^k \equiv g^k(\tilde{U}^k)$ is open. Moreover, by (a), we can choose \tilde{U}^k sufficiently small that for all $z \in \tilde{U}^k$, we have

$$u_1(x) > u_1(z).$$

Finally, by (c), we can choose \tilde{U}^k small enough that for all $z \in \tilde{U}^k$ and all $i = m + 1, \dots, m + k - 1$, u_i is not constant on $\{z, x_\epsilon^2, \dots, x_\epsilon^m\}$. Since \tilde{V}^k is open and contains $g^k(z^{k-1})$, there exists $\epsilon^k > 0$ such that

$$y^k = \begin{pmatrix} u_2(z^{k-1}) \\ \vdots \\ u_m(z^{k-1}) \\ u_{m+k}(z^{k-1}) + \epsilon^k \end{pmatrix}$$

belongs to \tilde{V}^k . Since $y^k \in \tilde{V}^k$, there is an alternative $z^k \in \tilde{U}^k$ such that $g^k(z^k) = y^k$, and by (b), we have $u_i(z^k) = u_i(x_\epsilon^1)$ for all $i = 2, \dots, m$. If player $m + k$ is not indifferent over $\{z^{k-1}, x_\epsilon^2, \dots, x_\epsilon^m\}$, then set $z^k = z^{k-1}$.

After Step $n - m$, we define the m -tuple $(\bar{x}_\epsilon^1, \dots, \bar{x}_\epsilon^m) = (z^{n-m}, x_\epsilon^2, \dots, x_\epsilon^m)$, and we claim that $\{\bar{x}_\epsilon^1, \bar{x}_\epsilon^2, \dots, \bar{x}_\epsilon^m\}$ is a semi-simple solution. Indeed, define the mapping $\rho: \{\bar{x}_\epsilon^1, \dots, \bar{x}_\epsilon^m\} \rightarrow N$ such that $\rho(\bar{x}_\epsilon^i) = i + 1$ for all $i = 1, \dots, m - 1$ and $\rho(\bar{x}_\epsilon^m) = 1$. For each $k = 1, \dots, m$, the coalition $C(\bar{x}_\epsilon^k)$ of players supporting \bar{x}_ϵ^k includes $C^* \setminus \{k\}$, and since C^* is oversized, this implies $C(\bar{x}_\epsilon^k) \in \mathcal{D}$. Of course, it follows that ρ is a one-to-one selection from the coalitions $C(\bar{x}_\epsilon^k)$. Thus, the set $\{\bar{x}_\epsilon^1, \dots, \bar{x}_\epsilon^m\}$ satisfies parts (i) and (ii) of Definition 1, i.e., it is a semi-simple solution contained in the open set U , as claimed.

Following the argument for $\epsilon > 0$, we can similarly construct vectors $y_\gamma \in \tilde{V}$ and $(x_\gamma^1, \dots, x_\gamma^m) \in U^m$ for all $\gamma \in (0, \epsilon)$. By construction, $\gamma_1 \neq \gamma_2$ implies $u_1(x_{\gamma_1}^2) = u_1(x) + \gamma_1 \neq u_1(x) + \gamma_2 = u_1(x_{\gamma_2}^2)$ and, therefore, $\{x_{\gamma_1}^1, \dots, x_{\gamma_1}^m\} \neq \{x_{\gamma_2}^1, \dots, x_{\gamma_2}^m\}$. We conclude that there is a continuum of semi-simple solutions contained in U .

References

- [1] V. Anesi (2010) "Noncooperative Foundations of Stable Sets in Voting Games," *Games and Economic Behavior* 70, 488-493.
- [2] V. Anesi and J. Duggan (2017) "Existence and Indeterminacy of Markovian Equilibria in Dynamic Bargaining Games," working paper.
- [3] V. Anesi and D.J. Seidmann (2014) "Bargaining over an Endogenous Agenda," *Theoretical Economics* 9, 445-482.
- [4] V. Anesi and D.J. Seidmann (2015) "Bargaining in Standing Committees with an Endogenous Default," *Review of Economic Studies* 82, 825-867.
- [5] J.S. Banks and J. Duggan (2000) "A Bargaining Model of Collective Choice," *American Political Science Review* 94, 73-88.

- [6] J.S. Banks and J. Duggan (2006) “A General Bargaining Model of Legislative Policy-making” *Quarterly Journal of Political Science* 1, 49-85.
- [7] D.P. Baron (1996) “A Dynamic Theory of Collective Goods Programs,” *American Political Science Review* 90, 316-330.
- [8] D.P. Baron and T.R. Bowen (2014) “Dynamic Coalitions,” working paper.
- [9] D.P. Baron and J.A. Ferejohn (1989) “Bargaining in Legislatures,” *American Political Science Review* 85, 137-164.
- [10] M. Battaglini and T. Palfrey (2012) “The Dynamics of Distributive Politics,” *Economic Theory* 49, 739-777.
- [11] T.R. Bowen and Z. Zahran (2012) “On Dynamic Compromise,” *Games and Economic Behavior* 76, 391-419.
- [12] D. Diermeier and P. Fong (2011) “Legislative Bargaining with Reconsideration,” *Quarterly Journal of Economics* 126, 947-985.
- [13] D. Diermeier and P. 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.
- [14] J. Duggan (2017) “Existence of Stationary Bargaining Equilibria,” *Games and Economic Behavior* 102, 111-126.
- [15] J. Duggan and T. Kalandrakis (2012) “Dynamic Legislative Policy Making,” *Journal of Economic Theory* 147, 1653-1688.
- [16] V. Guillemin and A. Pollack (1974) *Differential Topology*, Prentice Hall, Englewood Cliffs, New Jersey.
- [17] M. Holz, K. Steffens and E. Weitz (1999) *Introduction to Cardinal Arithmetic*, Birkhäuser Verlag, Basel.
- [18] M.O. Jackson and B. Moselle (2002) “Coalition and Party Formation in a Legislative Voting Game,” *Journal of Economic Theory* 103, 49-87.
- [19] A. Kalandrakis (2004) “A Three-Player Dynamic Majoritarian Bargaining Game,” *Journal of Economic Theory* 116, 294-322.
- [20] T. Kalandrakis (2010) “Minimum Winning Coalitions and Endogenous Status Quo,” *International Journal of Game Theory* 39, 617-643.

- [21] T. Kalandrakis (2014) “Pareto Efficiency in the Dynamic One-Dimensional Bargaining Model,” working paper.
- [22] L. Loomis and S. Sternberg (1968) *Advanced Calculus*, Addison-Wesley: Reading, MA.
- [23] R.D. McKelvey (1979) “General Conditions for Global Intransitivities in Formal Voting Models,” *Econometrica* 47, 1086-1112.
- [24] S. Nunnari (2014) “Dynamic Legislative Bargaining with Veto Power: Theory and Experiments,” working paper.
- [25] M. Richter (2014) “Fully Absorbing Dynamic Compromise,” *Journal of Economic Theory* 152, 92-104.
- [26] T. Romer and H. Rosenthal (1978) “Political Resource Allocation, Controlled Agendas, and the Status Quo,” *Public Choice* 33, 27-44.
- [27] N. Schofield (1980) “Generic Properties of Simple Bergson-Samuelson Welfare Functions,” *Journal of Mathematical Economics*, 7: 175–192.
- [28] S. Smale (1974) “Global Analysis and Economics III: Pareto Optima and Price Equilibria,” *Journal of Mathematical Economics* 1, 107-117.
- [29] J. Zápal (2014) “Simple Markovian Equilibria in Dynamic Spatial Legislative Bargaining,” working paper.