

Coalitional Bargaining Equilibria

John Duggan

Working Paper No. 62

May 2011

W. ALLEN
WALLIS
Institute of
POLITICAL
ECONOMY

UNIVERSITY OF
ROCHESTER

Coalitional Bargaining Equilibria

John Duggan
Department of Political Science
and Department of Economics
University of Rochester

April 25, 2011

Abstract

This paper takes up the foundational issue of existence of stationary subgame perfect equilibria in a general class of coalitional bargaining games that includes many known bargaining models and models of coalition formation. General sufficient conditions for existence of equilibria are currently lacking in many interesting environments: bargaining models with non-concave stage utility functions, models with a Pareto optimal status quo alternative and heterogeneous discount factors, and models of coalition formation in public good economies with consumption lower bounds. This paper establishes existence of stationary equilibrium under compactness and continuity conditions, without the structure of convexity or comprehensiveness used in the extant literature. The proof requires a precise selection of voting equilibria following different proposals. The result is applied to obtain equilibria in models of bargaining over taxes, coalition formation in NTU environments, and collective dynamic programming problems.

1 Introduction

This paper establishes existence of stationary subgame perfect equilibria in a general model that captures bargaining in economic environments, coalition formation in NTU environments, and collective dynamic programming problems. The framework grows out of the bargaining literature in economics and political science, originating with the seminal work of Rubinstein (1982). In economics, a number of papers following Chatterjee, Dutta, Ray, and Sengupta (1993) have investigated models of coalition formation built on bargaining protocols extending Rubinstein's, while in political science, a literature on legislative bargaining building on Baron and Ferejohn (1989) has emerged. The focus of both strands has been on the characterization of stationary equilibria, but there are quite natural environments for which the question of existence has been left open:

- legislative bargaining in a one-dimensional policy space (as in Banks and Duggan (2000)) in which stage utilities are single-peaked but not concave,
- legislative bargaining in which the status quo is a Pareto optimal policy (as in Banks and Duggan (2006)) and the discount factors of the agents are heterogeneous,
- coalition formation in public good economies where individuals have property rights over endowments (as in Ray and Vohra (1999)) but consumption is non-negative.

Three main applications motivate the general existence result. The first is a legislature that must choose taxes that determine the set of equilibrium allocations of an economy. Because the equilibrium correspondence has closed but not convex graph, extant results do not apply, but in the framework of this paper taxes can be generated by Markovian equilibrium behavior of legislators in a dynamic bargaining game. The example extends easily to bargaining over parameters of general games. The second is coalition formation in an NTU model in which proposers are determined stochastically, as in Binmore (1987), rather than according to the rejector-becomes-proposer protocol that is common in the coalition formation literature. The third is a model of collective dynamic programming with a countable set of states, in which a state-dependent decision-maker proposes a collective action, which is then voted on by the agents. As in the standard single-agent model, current states and actions determine future states, but now the identity of the decision-maker can change over time, and collective actions require the approval of at least one decisive coalition.

I analyze a general model of coalitional bargaining such that in each period, one agent proposes an outcome or remains silent, other agents simultaneously respond, and the game then may or may not continue. Stage games are parameterized by a countable set of states. Outcomes may be divisions of a pie, choices of policy in a multidimensional space, allocations of private and public goods in an economy, or choices of achievable vectors of utilities for coalitions. I assume only that the set of outcomes is a compact metric space and that the agents have continuous stage utilities—no convexity conditions are imposed. The process governing the timing of proposals is also quite general: agents may make alternating or sequential proposals, or the selection of a proposer may be random, with recognition probabilities possibly changing over time with the state. And the rule by which proposals are accepted or rejected is quite arbitrary: acceptance may require the assent of a majority of agents, or, for that matter, each coalition may have authority to implement proposals from some set of feasible outcomes that depends on the coalition. The latter assumption captures supermajority voting rules, economic models in which agents have property rights over their endowments, and more generally TU and NTU environments common in the literature on coalition formation. These feasible sets may vary over time as a function of the state, capturing shocks to production technology or endowments. In contrast to much of the literature on legislative bargaining and coalition formation, where the game ends once an outcome is determined, meaningful interaction may be ongoing, with equilibrium outcomes determining future states and outcomes over an infinite horizon. The equilibrium concept is stationary Markov perfect equilibrium in stage-undominated voting strategies.

A key to existence of coalitional bargaining equilibria is the possibility of mixing on the part of the agents, an approach taken by Ray and Vohra (1999) and Banks and Duggan (2000, 2006). Those papers deliver existence results in which proposers may randomize over the coalition proposed to, and they are quite general in some respects, but they impose sufficiently strong conditions to ensure that equilibria can be found in which the agents use pure voting strategies. In Ray and Vohra (1999), the restriction to pure voting strategies is achieved by the possibility of side payments in TU games or strict comprehensiveness in NTU games, whereas Banks and Duggan (2000) rely on concavity of the agents' stage utilities and the assumption of a bad "status quo," and Banks and Duggan (2006) use concavity and a constraint qualification on the indifference contours of the agents. To obtain a more general existence result for these environments, including simple finite examples, the analysis of voting strategies must be more nuanced. The next example uses a finite model to illustrate the need for mixed proposal strategies.

Example 1 Assume there are three agents and three outcomes with utilities as in the table below. Assume that each period in which the game is played, an agent is randomly drawn (with probability $1/3$) to propose an outcome or simply remain silent. A proposal passes if two agents accept, in which case the game ends with the proposed outcome and corresponding payoffs. If a proposal fails, or the proposer remains silent, then we move to the next period and the process is repeated. Payoffs are discounted by a common factor $\delta = .9$ each period in which delay occurs. A stationary equilibrium in pure strategies consists of a proposal for each agent and a response strategy to each possible proposal. It is not an equilibrium for each agent to propose her favorite outcome and for that proposal to pass with probability one: in that case, agent 1 would propose x , but agent 2's discounted continuation value would be $\delta v_2 = (.9/3)[.7 + 2 + 0] = .81 > .7$, which exceeds the payoff from x . Thus, voting for x is dominated for agents 2 and 3, so it would not pass. And it is not an equilibrium for each agent to propose her second-favorite outcome and for that proposal to pass with probability one: in that case, agent 1 would propose z and receive a payoff of $.7$, but the payoff from remaining silent is $\delta v_1 = .81 > .7$, so the agent would not make a proposal. In the appendix, I show that in every stationary equilibrium, each agent's discounted continuation value must be $.7$, and it can be checked that there is no pure strategy profile that generates these continuation values. Thus, the conditions for equilibrium lead to mixing, as in the following: each agent proposes with probability $.39$ her second-favorite outcome, which passes with probability one, and with probability $.61$ she proposes her favorite outcome, which fails with probability one. There is actually a continuum of stationary equilibria in which agents use pure voting strategies, for instead of proposing her favorite outcome, an agent could equivalently remain silent, or she could mix over the two options. \square

	1	2	3
x	2	.7	0
y	0	2	.7
z	.7	0	2

Concavity assumptions in the standard legislative bargaining framework imply that proposals are always accepted in equilibrium — delay cannot occur. The next example, a continuation of the first, shows that delay *must* occur in equilibrium, highlighting the possible role of non-convexities in the phenomenon of inefficient delay.

Example 1 (cont.) Suppose there is an equilibrium with no delay, so no agent remains silent and every proposal made in equilibrium passes with probability one. We have already observed that it is not the case that for every agent, if the agent proposes her favorite outcome, then it passes with probability one. So assume that if agent 1 proposes x , then agent 2 rejects x with positive probability. Then it must be that agent 1 proposes z with probability one, and this passes. But then $\delta v_3 \geq (.3)[2 + 0 + .7] > .7$, contradicting the result, shown in the appendix, that $\delta v_3 = .7$ in equilibrium. Therefore, all equilibria produce delay with positive probability. \square

For finite environments, existence of stationary equilibria in mixed proposal and voting strategies follows from a standard result on existence of mixed strategy equilibria in finite stochastic games.¹ Although the agents use pure voting strategies in the equilibrium of Example 1, the existence result for finite stochastic games does not rule out the possibility of mixed voting strategies. To see that mixed voting strategies are indeed needed for existence in the coalitional bargaining framework, I

¹Technically, we model voting as simultaneous and eliminate stage-dominated voting strategies. But we can equivalently model voting as sequential and apply the existence result of Rogers (1969) and Sobel (1971).

reconsider the running example.

Example 2 Modify Example 1 by assuming agent 1 is recognized to propose with probability .8 and agent 2 is recognized with probability .2.² I claim it cannot be the case that when agent 1 proposes x , it fails with probability one. Indeed, suppose it did. Then agent 2 rejects x , so $\delta v_2 \geq .7$. If agent 1 proposes x with positive probability in equilibrium, then the agent's payoff from proposing is δv_1 , and since the agent could propose z and receive a payoff of .7, we have $\delta v_1 \geq .7$. Therefore, agent 2 must propose x with positive probability, but then agent 2's expected payoff from proposing is .7, so $\delta v_2 \leq (.9)[(.8)\delta v_2 + (.2)(.7)]$, which implies $\delta v_2 < .7$, a contradiction. Likewise, agent 1 does not remain silent, and we conclude that the agent proposes z with probability one, and therefore $\delta v_2 < .7$, a contradiction that proves the claim. Next, I claim that it cannot be the case that when agent 1 proposes x , it passes with probability one. Suppose it did. Then agent 1 will propose x , and since agent 2 accepts it, $\delta v_2 \leq .7$. Then $\delta v_3 \leq (.9)[(.8)(0) + (.2)(2)] < .7$. It follows that agent 3 will accept y when proposed, so agent 2 will propose y . But then $\delta v_2 \geq (.9)[(.8)(.7) + (.2)(2)] > .7$, a contradiction. \square

Mixed voting strategies must therefore be used to obtain equilibrium existence in any class of environments containing these simple finite examples. The approach of this paper exploits the special structure of the coalitional bargaining model to select mixed strategy voting equilibria following different proposals precisely so as to maintain the proposers' optimality conditions. A key to this approach is that because agents who are indifferent between accepting and rejecting a proposal can mix arbitrarily, the set of voting equilibria is convex. Because voting strategies are conditioned on proposals, which take values in a compact metric space, this problem is potentially infinite-dimensional and leads to the question of a suitable topology on the space of voting strategies. I circumvent this difficulty by applying a version of Fatou's lemma to reduce the problem to a finite-dimensional one,³ submerging voting strategies in the construction of the fixed point correspondence. It is worth noting that existence of stationary subgame perfect equilibrium in the general model does not follow from the extant game-theoretic literature. Work by Harris (1985a,b), Börgers (1989,1991), and Harris, Reny, and Robson (1995) establishes existence of subgame perfect equilibria in perfect information games and of correlated subgame perfect equilibria in games of "almost perfect" information,⁴ and in fact these results apply to the coalitional bargaining model, but their results do not yield stationarity. The literature on stochastic games focuses on stationarity, and the coalitional bargaining model can indeed be formulated as a stochastic game, but fully stationary subgame perfect equilibria are not established for general games,⁵ and that literature relies on a critical continuity condition on the transition probability of the game that is violated in the bargaining model: the issue is that the outcome chosen by the proposer is directly voted on by the agents, unmediated by any random noise; this deterministic transition probability from proposal subgames to voting subgames is discontinuous with respect to the strong topology used in the literature.

²The example is robust, in the sense that agent 1 could well propose with small probability.

³Or, if the set of states is countably infinite, to a problem of countable dimensionality.

⁴See also Fudenberg and Levine (1983), Hellwig and Leininger (1987), and Hellwig, Leininger, Reny, and Robson (1990).

⁵For example, correlation is used (Nowak and Raghavan (1992)), or stationarity is relaxed (Chakrabarti (1999)), or ϵ -best responses are allowed (Nowak (1985)).

The model of this paper covers the literature on legislative bargaining, including the distributive model of Baron and Ferejohn (1989), the unanimity bargaining model with stochastic pie of Merlo and Wilson (1995) and Eraslan and Merlo’s (2002) majority-rule version, the one-dimensional model of Jackson and Moselle (2002), the general spatial models of Banks and Duggan (2000,2006), and Kalandrakis’s (2004a) version with proposer selection following an arbitrary Markov chain. On the coalition formation side, it generalizes Okada’s (1996) TU model with random proposers to the NTU setting, and it extends the NTU model of Herings and Predtetchinski (2009) to arbitrary voting rules. Much of the literature on coalitions formation, such as Chatterjee, Dutta, Ray, and Sengupta (1993), Bloch (1996), Krishna and Serrano (1996), and Ray and Vohra (1999) assume that the identity of next period’s proposer is endogenous — it is the first agent, if any, to reject the current proposal — that takes it outside the framework of this paper.⁶ Ray and Vohra (1999) prove existence of equilibrium with this protocol for NTU games satisfying strict comprehensiveness, a weak assumption in private good economies but one that is easily violated in public good environments with consumption lower bounds.⁷ In the absence of strict comprehensiveness, however, existence is problematic, and the technique of selecting voting equilibria cannot be used because we may lose the key property of convexity of voting equilibrium outcomes: an example in the concluding section illustrates how the set of equilibrium outcomes in a voting subgame can be non-convex in that model, undermining the fixed point argument used in this paper and pointing to a potentially deep problem of equilibrium existence. Nevertheless, I conjecture that existence is regained if public randomization is allowed before agents respond to proposals.

In the remainder of the paper, Section 2 presents the general coalitional bargaining model, and Section 3 considers a number of special cases to illustrate its flexibility. I describe how to obtain finite environments with a moving status quo, the legislative bargaining model of Baron and Ferejohn (1989) and its successors, and in particular an application to bargaining over taxes in an economy or parameters of a more general game, a generalization of Okada’s (1996) TU coalition formation model to the NTU case, along with a version of the model with coalitional externalities and incomplete contracts, and finally a model of collective dynamic programming in which agents vote on collective choices that influence the evolution of the state. Section 4 contains the formal statement and proof of the existence theorem, Section 5 gives an overview of the proof approach, and Section 6 concludes. Statement and proof of auxiliary results are contained in the appendix.

2 Coalitional Bargaining Framework

The model is given by the list $(N, S, X, \{X_C\}_{C \subseteq N}, \{u_i\}_{i \in N}, q, p, \{\delta_i\}_{i \in N})$, where N is a finite set of n agents, S is a countable set of states with the discrete topology, X is a compact metric space of outcomes, $X_C: S \rightrightarrows X$ specifies for each coalition C and state s a closed set $X_C(s)$ consisting of outcomes feasible for coalition C in state s , $u_i: X \times S \rightarrow \Re$ is a bounded and continuous stage payoff function for agent i , q is a fixed default outcome belonging to X , $p: S \times X \times S \rightarrow [0, 1]$

⁶See Serrano (2005), Bandyopadhyay and Chatterjee (2006), and Ray (2007) for recent surveys of this literature.

⁷Suppose there are two agents, one public good, and one private good. Consider the Pareto optimal allocation in which one agent is given all of the private good, and she chooses her optimal level of public good production. If that level is positive, then the utility to the second agent could easily exceed the utility from his endowment, but we cannot reduce his utility in a way that benefits the first agent. In other words, the Pareto frontier of the set of utility imputations has a “flat” portion.

is a continuous state transition probability function, and $\delta_i: S \times S \rightarrow [0, 1]$ specifies the discount factor $\delta_i(s, s')$ for agent i that gives the weight of tomorrow's payoffs relative to today's when we transition from state s to state s' . The coalitional bargaining game among the agents is governed by the following protocol: 1) each period t begins with a state s ; 2) an agent $i(s)$ determined by the state proposes any outcome x in X ; 3) the agents simultaneously decide whether to accept or reject this proposal; 4) if there is a coalition $C \subseteq N$ such that $x \in X_C(s)$ and all members of C accept x , then x is the outcome in period t , and payoffs $u_i(x, s)$ accrue to the agents; otherwise, the agents receive payoffs $u_i(q, s)$ from the default outcome $q \in X$; finally, 5) the game transitions to period $t + 1$, a new state s' is drawn with probability $p(s'|x, s)$, and the protocol is repeated with the following period's payoffs discounted by $\delta_i(s, s')$. Thus, given sequences $\mathbf{s} = (s_t)_{t=1}^\infty$ and $\mathbf{x} = (x_t)_{t=1}^\infty$ of states and outcomes, the discounted payoff in period $t \geq 2$ for agent i is

$$U_{i,t}(\mathbf{s}, \mathbf{x}) = \left(\prod_{k=2}^t \delta_i(s_{k-1}, s_k) \right) u_i(s_k, x_k),$$

and the discounted sum of payoffs is

$$U_i(\mathbf{s}, \mathbf{x}) = u_i(s_1, x_1) + \sum_{t=2}^{\infty} U_{i,t}(\mathbf{s}, \mathbf{x}).$$

Since each u_i is bounded below, I normalize stage payoffs so that $\min_{i,s,x} u_i(s, x) \geq 0$ without loss of generality.

Assume that the sets of feasible outcomes are closed, and therefore compact, and that they are monotonic. That is, for states s and all coalitions C and C' with $C \subseteq C'$, we have $X_C(s) \subseteq X_{C'}(s)$. Assume also that the status quo q is an isolated point, i.e., $\{q\}$ is open in X , and that q is always feasible for all coalitions, i.e., $q \in X_C(s)$ for all s and all C , so that a proposer always has the option to “pass” and obtain the default for the current period. The assumption that the default is isolated is without loss of generality. Given a model in which $q \in X$ is not isolated, we can modify the model by appending an artificial element q' to X to obtain a new set $X' = X \cup \{q'\}$ of outcomes, extending the metric on X so that q' is isolated in X' , and extending utilities and the transition probability so that $u'_i(q', s) = u_i(q, s)$ and $p'(s'|q', j, s) = p(s'|q, j, s)$. Thus, the re-defined default q' enters in these functions exactly as the previous default and fulfills our assumptions without affecting the strategic structure of the game. The former default is still an element of X' and replicates q' and is not isolated, but the existence of such an outcome is compatible with the other assumptions.

Discount factors are specified quite generally, allowing them to depend on the states from which, and to which, the game transitions. This is helpful in capturing environments where the time between some decisions is inconsequential, so no discounting occurs. So that dynamic payoffs are well-defined, I impose the joint restriction on δ_i and p that there exists $T < \infty$ such that

$$\sup \prod_{t=1}^{T-1} p(s_{t+1}|x_t, s_t) \delta_i(s_t, s_{t+1}) < \prod_{t=1}^{T-1} p(s_{t+1}|x_t, s_t),$$

where the supremum is over sequences $(x_1, \dots, x_{T-1}) \in X^{T-1}$ and $(s_1, \dots, s_T) \in S^T$ such that $\prod_{t=1}^{T-1} p(s_{t+1}|x_t, s_t) > 0$. To parse this condition, note that if each discount factor $\delta_i(s, s')$ were

equal to one, then the lefthand side would equal the righthand side. Clearly, this is ruled out. The inequality says that over any span of T periods with positive probability, regardless of the states and outcomes determined over that span, payoffs are discounted at some point along that sequence. Furthermore, define $\bar{\delta} = \max\{\delta_i(s, s') \mid i \in N, s, s' \in S, \delta_i(s, s') < 1\}$, and assume $\bar{\delta} < 1$. Thus, there may be transitions that are essentially instantaneous, but there do not exist transitions that are arbitrarily close to instantaneous; this is automatically satisfied if S is finite. Obviously, the discounting assumption is satisfied if each agent i discounts payoffs over time according to a fixed discount factor $\delta_i \in [0, 1)$.

To define proposal strategies for an agent i , let S_i be the subset of $s \in S$ such that $i = i(s)$, and let $\mathcal{P}(X)$ be the space of Borel probability measures on X endowed with the weak* topology. A *stationary proposal strategy* for agent i is a measurable mapping $\pi_i: S_i \rightarrow \mathcal{P}(X)$, where $\pi_i(s)(Y)$ denotes the probability that i proposes an outcome in the Borel measurable set Y in state s . A *stationary voting strategy* is a Borel measurable mapping $\alpha_i: X \times S \rightarrow [0, 1]$, where $\alpha_i(x, s)$ is the probability that i votes to accept proposal x in state s . A *stationary strategy* is then a pair $\sigma_i = (\pi_i, \alpha_i)$. Let $\sigma = (\sigma_1, \dots, \sigma_n)$ denote a stationary strategy profile. Note one departure from stationarity as defined in Banks and Duggan (2000,2006), Ray and Vohra (1999), and others: because the state contains the identity of the proposer, I allow voting strategies $\alpha_j(x, s)$ to implicitly depend on the proposer. This is innocuous when, as in previous work, voting equilibria are essentially unique and in pure strategies. In the current framework, however, the selection of mixed strategy voting equilibria when some agents are indifferent can depend on the proposer.

Given a profile σ of stationary strategies, we can calculate the expected discounted sum of payoffs for agent i when state s obtains at the beginning of a period. These continuation values, $v_i(s; \sigma)$, satisfy the following recursive relationship: for all i and all s ,

$$v_i(s; \sigma) = \int_X \left[\alpha(x, s) \left[u_i(x, s) + \sum_{s' \in S} p(s'|x, s) \delta_i(s, s') v_i(s'; \sigma) \right] + (1 - \alpha(x, s)) \left[u_i(q, s) + \sum_{s' \in S} p(s'|q, s) \delta_i(s, s') v_i(s'; \sigma) \right] \right] \pi_{i(s)}(s)(dx),$$

where $\alpha(x, s)$ is shorthand for the probability that the proposal of x in state s is accepted. It is defined formally as

$$\alpha(x, s) = \sum_{C \subseteq N: x \in X_C(s)} \left(\prod_{j \in C} \alpha_j(x, s) \right) \left(\prod_{j \notin C} (1 - \alpha_j(x, s)) \right).$$

Note that by the assumption that the proposer can impose the default, $\alpha(x, s) = 1$ if $x = q$ and the proposer votes to accept. Given σ , define the agent's *dynamic payoff* from outcome x in state s as

$$U_i(x, s; \sigma) = u_i(x, s) + \sum_{s' \in S} p(s'|x, s) \delta_i(s, s') v_i(s'; \sigma),$$

incorporating not only the current stage payoff but expected future payoffs as well. Note that dynamic payoffs are continuous in (x, s) .

A *coalitional bargaining equilibrium* is a stationary strategy profile σ such that agents propose optimally given the voting strategies of others and such that agents use stage-undominated voting

strategies following proposals. Formally, the requirement on proposal strategies is that for all i and all $s \in S_i$, $\pi_i(s)$ put probability one on solutions to

$$\max_{x \in X(s)} \alpha(x, s)U_i(x, s; \sigma) + (1 - \alpha(x, s))U_i(q, s; \sigma),$$

and, using the assumption that the feasible sets $X_C(s)$ are monotonic, the requirement on voting strategies is that for all i and all s ,

$$\alpha_i(x, s) = \begin{cases} 1 & \text{if } U_i(x, s; \sigma) > U_i(q, s; \sigma) \\ 0 & \text{if } U_i(x, s; \sigma) < U_i(q, s; \sigma), \end{cases}$$

with no restriction on votes when agents are indifferent, i.e., $U_i(x, s; \sigma) = U_i(q, s; \sigma)$. The implicit requirement of stage undominated voting strategies is standard and used to preclude Nash equilibria of the voting game in which, for example, all agents vote reject and none are pivotal (so rejection is trivially a best response). This refinement could be dropped if, instead, I specified that voting were sequential, a common approach that does not affect the set of equilibrium outcomes.

3 Special Cases

The model set forth in the previous section has been described in a parsimonious way. Rather than explicitly building in complex structure, we capture special cases of interest by exploiting hidden generality of the model. This section shows how perhaps unexpected structure can be obtained by suitable specification of the coalitional bargaining model.

3.1 Finite Environments with Moving Status Quo

It is well-known that stationary subgame perfect equilibria exist in finite stochastic games, and that result applies (with minor adjustments) to the current framework. Thus, the contribution of this paper is equilibrium existence in dynamic coalitional games with infinite sets of feasible outcomes. Nevertheless, it is noteworthy that the approach to existence taken here applies equally well to finite versions of the model. For example, let P denote a finite set of “positions,” let $X = P \cup \{q\}$ be the set of positions augmented by an abstract default outcome q , and let the set of states be $S = N \times P$, where a state $s = (i, x)$ specifies the proposer in the current period and the current position. Given $s = (i, x)$, the set $X_C(s)$ represents the set of positions to which coalition C has authority to move from x , and given outcome z , stage utilities are defined by

$$u_i(z, (j, x)) = \begin{cases} r_i(z) & \text{if } z \in X \setminus \{q\}, \\ r_i(x) & \text{if } z = q, \end{cases}$$

where $r_i: X \rightarrow \mathfrak{R}$ is arbitrary. Given outcome z in state (i, x) , transition probabilities are such that if $z \in X \setminus \{q\}$, then the current position moves to x , and the next proposer is selected according to a time-invariant probability distribution on the set of agents. That is, state (j, z) is selected with probability p_j , where $\sum p_k = 1$. And if the proposer chooses the default or makes an unsuccessful proposal, then the current position stays at x , and a new proposer is again randomly drawn, so

(j, x) is selected with probability p_j . Thus, the state (i, x) keeps track of the status quo position, x , and if the members of an authorized coalition accept a new position, then it becomes part of the status quo in the next period; otherwise, x is maintained. Note that the technical role of the default q is to modulate stage utilities and transition probabilities — it does not produce utility, but it controls when utility is derived from the current status quo.

Fixing discount factors $\delta_i \in [0, 1)$ for each agent and instantiating to majority rule with the specification

$$X_C(s) = \begin{cases} X & \text{if } |C| > \frac{n}{2}, \\ \{q\} & \text{else,} \end{cases}$$


we obtain a finite version of the endogenous status quo bargaining models of Baron (1996) and Kalandrakis (2004b,2009). Or, assuming $X_C(s)$ is independent of the proposer i , the primitives of this model are as in Konishi and Ray (2003). Although their analysis is cooperative, as they analyze the properties of a class of Markov chains consistent with coalitional incentives of farsighted agents, a coalitional bargaining equilibrium essentially generates a process of coalition formation (PCF) in the sense of those authors.⁸ Finally, given a state (i, x) and a failed proposal, we can specify that the next proposer is an agent who has not previously proposed and that the transition is not discounted until either a proposal passes or there are no remaining unused agents, giving us an agenda-setting protocol as in Acemoglu, Egorov, and Sonin (2009).

3.2 General Legislative Bargaining

The legislative bargaining model of Baron and Ferejohn (1989) is subsumed by specifying the set of states as $S = N \cup \{\tau\}$, where $i \in S$ indicates the proposer in a given period and $\tau \in S$ is a terminal state. Let $X = \Delta \cup \{q\}$, where Δ is the unit simplex in \mathfrak{R}^n , with $x \in \Delta$ representing an allocation of a “dollar” among the n agents. Again specify majority voting. Stage utilities from outcome $x \in X$ in state s are specified as

$$u_i(x, s) = \begin{cases} \frac{x_i}{1-\delta} & \text{if } x \in \Delta \text{ and } s \in N \\ 0 & \text{else,} \end{cases}$$

where stage utilities are discounted by a common factor $\delta \in [0, 1)$ that is independent of the state. Transition probabilities are specified so that the terminal state τ is absorbing. In state $i \in S$, if $x \in X \setminus \{q\}$ is proposed and receives the support of a majority of agents, then payoffs are accrued as specified above, and the game transitions to the terminal state, and the agents receive zero payoffs thereafter. If no allocation is



⁸There is a slight wedge between their equilibrium concept and the one I use, as the non-cooperative analysis allows for the possibility of delay, even if there is some coalition that strictly prefers to move to a different outcome that is feasible for the coalition.

passed in the current period, then the agents receive the default payoff of zero, and the state transitions to $j \in N$ with probability p_j , where $\sum p_j = 1$. The transition probability $p(s'|z, i)$ is depicted above.

The coalitional bargaining model also captures the more general versions of the Baron-Ferejohn model as analyzed in Banks and Duggan (2000,2006). In the earlier of those papers, the latter authors assume a set A of alternatives that is a compact, convex subset of finite-dimensional Euclidean space, generalizing the unit simplex, and they assume stage utilities $u_i(a)$ are continuous, concave, and positive on A , whereas the default payoff is zero. The voting rule is given by an arbitrary nonempty collection \mathcal{D} of decisive coalitions that is monotonic in the sense that $C \in \mathcal{D}$ and $C \subseteq C'$ implies $C' \in \mathcal{D}$. Discount factors may differ across agents. We capture this model by appending an abstract default q to the set A of alternatives to obtain the set $X = A \cup \{q\}$ of outcomes, and we extend stage utilities so that $u_i(x, j) = u_i(x)/(1 - \delta_i)$ and $u_i(x, \tau) = 0$ for $x \in A$ and $u_i(q, j) = u_i(q, \tau) = 0$. In the later paper, the authors allow for default payoffs to be given by an arbitrary alternative, say $a_q \in A$, and they assume a common discount factor and a constraint qualification on the indifference contours of the agents that generalizes the customary condition of strict quasi-concavity. Then we specify the default payoff as $u_i(q, j) = u_i(a_q)$. Those papers establish existence of stationary bargaining equilibria in which agents use pure voting strategies, and those results can be obtained as special cases of the main theorem of this paper.⁹ The convexity conditions employed by Banks and Duggan (2000,2006) are necessitated by the focus on pure voting strategies. The more general setting of the current paper captures models in which those assumptions are violated and for which no results on existence of stationary equilibria are currently known. For example, in some applications it is natural to assume single-peakedness (or quasi-concavity more generally), whereas cardinal restrictions such as concavity are less convincing. Moreover, in applications where payoffs are given by an alternative that is not worse than every other conceivable alternative for every agent (which likely covers a great portion of situations of interest), there may be no reason to assume agents discount future payoffs at the same rate, and so we would like to allow for heterogeneous discount factors.

Now consider legislative bargaining among a group of legislators who must choose a tax system $t \in T$, a compact metric space, which then determines a set of equilibrium allocations $E(t) \subseteq \mathbb{R}^{km}$ in an economy with m consumers (which may or may not include the legislators) and k commodities. Assume that the range of E is compact and that the correspondence has nonempty values and closed graph. Then define the set of outcomes as $X = \{q\} \cup \text{graph}(E)$, and let each legislator i have stage utility

$$u_i(x, s) = \begin{cases} \frac{u_i(t, a, s)}{1 - \delta_i} & \text{if } x = (t, a) \in \text{graph}(E) \text{ and } s \in N \\ u_i(t_q, a_q, s) & \text{if } x = q \text{ and } s \in N \\ 0 & \text{if } s = \tau, \end{cases}$$

where t_q is the status quo tax system and $a_q \in E(t_q)$ the status quo equilibrium allocation, and $u_i(t, a, s)$ is jointly continuous in its arguments. The voting rule is given by an arbitrary collection \mathcal{D} of decisive coalitions. Note that a proposal consists of not only a tax system t , but also an equilibrium allocation $a \in E(t)$ consistent with the tax system, allowing the proposer to resolve indeterminacy of economic equilibria arbitrarily. We could for the most part suppress the

⁹With the observation that in the framework of Banks and Duggan (2000,2006), all proposals weakly acceptable to all members of a decisive coalition must pass in equilibrium.

announcement of the equilibrium allocation, instead relying on an appropriate selection of equilibrium allocations following tax systems proposed by a given agent. A conceptual subtlety arises, however, if a proposer mixes over two outcomes (t, a) and (t, a') with the same tax system, for then consumers must observe the legislator's announcement of the allocation (or some other sunspot) to select equilibria with the correct probabilities.

Because the correspondence has compact graph, the assumptions of the model are fulfilled here, and therefore a coalitional bargaining equilibrium exists. This example can be extended to any situation in which a group of agents bargains over parameters of a game, as long as the parameter space is compact metric and the equilibrium correspondence has nonempty values and compact graph. In fact, the main theorem of this paper establishes that the correspondence of coalitional bargaining equilibrium continuation value vectors has closed graph, so we can obtain existence of equilibrium in an issue-by-issue version of the model. Suppose for example that the set of outcomes is the square $X = [0, 1]^d$ and that the dimensions must be decided on in order. Specifically, a committee $C_h \subseteq N$ with jurisdiction on dimension h bargains over the coordinate x_h given previous decisions x_1, \dots, x_{h-1} on other coordinates. Then the bargaining game among members of committee C_d given x_1, \dots, x_{d-1} is a one-dimensional version of the Baron-Ferejohn model parameterized by (x_1, \dots, x_{d-1}) . The correspondence $E^d: [0, 1]^{d-1} \rightrightarrows \mathfrak{R}^n$ of coalitional bargaining equilibrium continuation values has nonempty values and compact graph, and therefore the bargaining game in committee C_{d-1} admits equilibria for each (x_1, \dots, x_{d-2}) , and the set of equilibrium payoff vectors of this game has nonempty values and compact graph, with each vector $(v_1, \dots, v_n) \in E(x_1, \dots, x_{d-2})$ of continuation values predicated on a selection of equilibria for the bargaining game in C_d , and so on. Finally, the bargaining game for the first committee C_1 admits an equilibrium, which embeds equilibrium selections in all later subgames.

3.3 Coalition Formation in NTU Environments

Okada (1996) generalizes the legislative bargaining model of Baron and Ferejohn (1989) by considering an abstract coalitional bargaining environment with transferable utility. Let v be a TU coalition function such that $v(\{i\}) = 0$, $v(C \cup C') \geq v(C) + v(C')$ for disjoint coalitions, and $v(N) > 0$. Define

$$V(C) = \left\{ x \in \mathfrak{R}_+^{|C|} \mid \sum_{i \in C} x_i \leq v(C) \right\}$$

as the payoff vectors feasible for C . In any period, some pairwise disjoint collection \mathcal{C} of coalitions has formed, and the agents $A = N \setminus \bigcup \mathcal{C}$ remain active. An agent i is drawn from the set of active agents with probability p_i^A and makes a proposal (C, x) such that $x \in V(C)$. The members of C vote to accept or reject the proposal, and if all members accept, then each $i \in C$ receives stage payoff $x_i/(1 - \delta)$, and each $i \notin C$ receives a stage payoff of zero. The game moves to the next period, with the collection $\mathcal{C}' = \mathcal{C} \cup \{C\}$ of formed coalitions and continues until no active players remain, at which point the game ends. Stage utilities are discounted by the common factor $\delta \in [0, 1)$ after each period. Okada (1996) analyzes pure strategy stationary equilibria of this game and characterizes when such equilibria exist, but pure strategy equilibria sometimes fail to exist in this setting.


This model is obtained as a special case of the coalitional bargaining model by defining S to consist of all pairs (i, \mathcal{C}) , where i represents the proposer for the current period and \mathcal{C} is a pairwise disjoint collection of coalitions representing the coalitions that have formed, along with a terminal state τ . Define $X = \{q\} \cup \bigcup \{V(C) \mid C \subseteq N\}$, where q is an abstract default outcome, and for each C and (i, \mathcal{C}) , define

$$X_C(i, \mathcal{C}) = \begin{cases} V(C) \cup \{q\} & \text{if } C \cap \bigcup \mathcal{C} = \emptyset, \\ \{q\} & \text{else.} \end{cases}$$

Define stage utilities from outcome x in state s as

$$u_i(x, s) = \begin{cases} \frac{x_i}{1-\delta} & \text{if } i \in C \text{ and } x \in V(C) \text{ and } s \neq \tau, \\ 0 & \text{else.} \end{cases}$$

Transition probabilities are specified so that the terminal state τ is absorbing. In state $s = (i, \mathcal{C})$, if $x \in X_C(s)$ is proposed and is accepted by all members of C , and if there exists $j \in A' = N \setminus (C \cup \bigcup \mathcal{C})$, then the game transitions to state $(j, \mathcal{C} \cup \{C\})$ with probability $p_j^{A'}$; in case there exists no such agent j , then the game transitions to the terminal state. If no allocation is passed in the current period, then agents receive the default payoff of zero, and the state transitions to (j, \mathcal{C}) , where $j \in A = N \setminus \bigcup \mathcal{C}$ is selected as proposer with probability p_j^A . The transition probability $p(s'|z, (i, \mathcal{C}))$ is depicted below.



An implication of the main theorem of the current paper is that when agents are permitted to mix, a coalitional bargaining equilibrium indeed exists. In fact, the result applies in considerably more general environments: it is sufficient that feasible outcomes are given by an NTU coalition function V such that $V(C) \subseteq \mathbb{R}^{|C|}$ is nonempty and compact for each coalition C and such that $V(\{i\}) = \{0\}$ for each agent i . This captures, for example, coalitional bargaining in private good exchange economies, but is also encompasses environments in which strict comprehensiveness is violated, including public good economies with consumption lower bounds. Ray and Vohra (1999) prove existence in general TU settings, and they extend the result to NTU environments satisfying strict comprehensiveness in the discussion following the proof of their Theorem 2.1. Their proof approach is predicated on the rejector-becomes-proposer protocol, however, so a contribution of the existence result here is to allow for stochastically determined proposers, or rotating protocols, or more general models in which the next proposer is determined as a function of the current state. Moreover, Ray and Vohra's proof strategy requires strict comprehensiveness, which is unavailable in some economic environments of interest. Herings and Predtetchinski (2009) prove existence of pure strategy stationary equilibria in NTU environments satisfying (not necessarily strict) comprehensiveness, although they assume voting is by unanimity rule. Thus, the result of the current paper

shows that existence (in mixed strategies) is maintained for arbitrary voting rules, even without comprehensiveness.

Ray and Vohra (1999) actually consider a model that is built on a partitional coalition function, which allows limited externalities between coalitions, and in this respect is more general than the coalitional bargaining model. In the TU context, given a coalition C and a partition \mathcal{C} of $N \setminus C$, the value of C is written $v(C; \mathcal{C})$, so the value of the coalition can depend on which other coalitions form (though not the allocations selected by the other coalitions), and when a proposer offers an allocation to the coalition C , it must be contingent on the collection of other coalitions that subsequently form. In effect, an agent proposes a menu of allocations, one for each possible partition \mathcal{C} . The approach of the current paper does not directly extend to partition function environments such as this, because it requires that the payoffs of the agents at the end of the game depend on proposals (menus) accepted in earlier periods. This is only possible in the coalitional bargaining framework if coalitional agreements in early periods are encoded in the state variable, but because the set of states is countable, they cannot generally be used to carry forward past agreements, which may vary over a continuous space, to the end of the game. The problem is that in this version of the partition function model, coalitions are able to write complex contracts that are contingent on future coalitions that form.

We can, however, capture a model of externalities in which coalitions cannot write binding contracts that commit them in advance to contingent allocations. To do so, assume an initial phase of coalition formation, in which a randomly drawn agent proposes to a set of agents that they form a coalition—without specifying an allocation of payoffs. Then, the latter agents respond, and so on, and this process continues until all agents have selected into groups (possibly singletons). The offer to form a coalition can also include a voting rule, or “constitution,” used to choose an allocation in the second phase, as long as the set of possible constitutions is finite. In the second phase, the formed coalitions sequentially allocate payoffs from their feasible sets among their members; because agents have already selected into coalitions, these allocations will depend in equilibrium on the coalitions that have formed. Because the set of possible coalition structures is finite, we can encode the outcomes of the first phase in the state variable. Say at the end of the formation phase, the state is $s = \{C_1, \dots, C_m\}$. Then the members of coalition C_1 choose an allocation $x_1 \in V(C; s)$, and the state transitions to $s' = \{C_1, \dots, C_m\} \times \{C_1\}$, where the second component records the coalitions that have agreed upon allocations, then coalition C_2 chooses an allocation, and so on, until all coalitions have chosen allocations and the game ends. In this way, we can capture environments characterized by coalitional externalities and incomplete contracts.

3.4 Collective Dynamic Programming

The coalitional bargaining framework extends well beyond the previous examples. In the legislative bargaining and coalition formation environments, there is a terminal state that is reached when, respectively, a majority first accepts a proposal or all coalitions have formed. We can also extend the standard dynamic programming framework with a finite state space S and compact metric space A of actions, where $A(s)$ is the set of actions feasible in state s , to allow for collective choice. Now, instead of a single decision maker, we assume a finite set of agents and interpret actions $a \in A$ as collective actions, such as allocations of goods for current consumption and a level of

public investment. Assume each agent i has continuous stage utility $u_i(a, s)$ defined on action-state pairs, and let $p(s'|a, s)$ denote the transition probability function. It may be that each state s determines a unique decision maker $i(s)$ who chooses any feasible action $a \in A(s)$, but we can allow for majority voting or more complex voting rules. Letting $\mathcal{D}(s)$ denote the monotonic collection of decisive coalitions in state s , we assume that i proposes an action $a \in A(s)$ and the agents then vote to either accept the proposal a or reject it in favor of an exogenously specified default $d(s) \in A(s)$. In the public investment example, this could simply be that the agents consume their endowments and invest nothing. If the members of a decisive coalition accept a , then the agents accrue stage utility $u_i(a, s)$, discounted by δ_i^{t-1} in period t , and the game moves to the next period with state s' drawn from $p(\cdot|a, s)$; otherwise, payoffs are $u_i(d(s), s)$ and the transition probability is $p(\cdot|d(s), s)$.

We obtain the collective dynamic programming model as a special case of the coalitional bargaining framework in a straightforward way, appending an abstract default outcome q to A and specifying the feasible outcomes $X_C(s)$ for coalition C as the set $A(s) \cup \{q\}$ if $C \in \mathcal{D}(s)$ and as $\{q\}$ otherwise. We extend stage utilities and the transition probability so that q plays the same role as the default $d(s)$ in each state, i.e., $u_i(q, s) = u_i(d(s), s)$ and $p(s'|q, s) = p(s'|d(s), s)$. This model deviates from the standard framework in two ways. In the standard framework, the decision maker chooses an action in state s knowing that she chooses optimally in every subsequent period, giving rise to the well-known Bellman equation. But in the collective dynamic programming model, decision maker i must anticipate that future choices may be determined by other agents, a distortion that can affect the current decision by agent i . This wedge is present even in the model where the decision making agent $i(s)$ may unilaterally choose any collective action $a \in A(s)$. When i 's proposal is subject to approval by a decisive coalition, there is a second wedge between the collective framework and the standard one due to "political constraints." In effect, agent i 's problem is to choose a optimally subject to political constraints, which are endogenous, as well as the usual feasibility constraints.

4 Existence of Coalitional Bargaining Equilibria

For the statement of the following theorem, parameterize stage payoff functions and the transition probability on states by the elements γ of a metric space Γ , as in $u_i(x, s, \gamma)$ and $p(s'|x, s, \gamma)$, and assume u_i and p are jointly continuous in their arguments. In this section, $v_{i,s}$ will denote agent i 's continuation value calculated at the beginning of a period in state s , and $v = (v_{i,s})_{i \in N, s \in S} \in \mathfrak{R}^{N \times S}$ a vector of continuation values. It is understood that $\mathfrak{R}^{N \times S}$, and other explicitly defined product spaces, are endowed with the product topology. Define the correspondence $E: \Gamma \rightrightarrows \mathfrak{R}^{N \times S}$ such that $E(\gamma)$ consists of vectors v such that in the model parameterized by γ , there exists a coalitional bargaining equilibrium $\sigma = (\pi, \alpha)$ with continuation values $v = (v_i(s; \sigma))_{i \in N, s \in S}$. The next result establishes existence of coalitional bargaining equilibria, along with upper hemicontinuity of the correspondence of equilibrium continuation values.

Theorem 4.1 *The correspondence $E: \Gamma \rightrightarrows \mathfrak{R}^{N \times S}$ has non-empty, closed values and is upper hemicontinuous.*

The rest of this section consists of the existence proof; see the next section for an in-depth discussion of the approach. We will use w_s to denote the expected discounted payoff of the proposer in state s and $w = (w_s)_{s \in S} \in \mathfrak{R}^S$ for a profile of proposer payoffs. Using boundedness of u_i , define

$$\bar{u} = \frac{T}{1 - \delta} \cdot \sup_{i,x,s} u_i(x, s),$$

so that we can assume $v_{i,s}, w_s \in [0, \bar{u}]$ for all i and s . I use the notation $\pi = (\pi_s)_{s \in S} \in \mathcal{P}(X)^S$ for a profile of mixed proposal strategies. Let $X(s) = \bigcup_{C \subseteq N} X_C(s)$ be the feasible outcomes in state s , and define the nonempty, convex, compact product space

$$\Theta = \left(\prod_{s \in S} \mathcal{P}(X(s)) \right) \times ([0, \bar{u}]^S) \times ([0, \bar{u}]^{N \times S}),$$

with elements $\theta = (\pi, w, v)$. Finally, let $\Theta^+ = \Theta \times \Gamma$ be this space augmented by the parameters of the model. Denote a generic element of Θ^+ by $\theta^+ = (\pi, w, v, \gamma)$.

I will define a correspondence $F: \Theta^+ \rightrightarrows \Theta$ such that for all $\gamma \in \Gamma$, $F(\cdot, \gamma)$ has a fixed point $\theta^* = (\pi^*, w^*, v^*) \in F(\theta^*, \gamma)$; each fixed point θ^* corresponds to a coalitional bargaining equilibrium in the model parameterized by γ ; and the correspondence of fixed points has closed graph. Write F as a product correspondence $F = P \times W \times V$. For the construction of the component correspondences, it will be useful to define an analogue of dynamic payoffs as

$$U_i(x, s, \theta^+) = u_i(x, s, \gamma) + \sum_{s' \in S} p(s'|x, s, \gamma) \delta_i(s, s') v_{i,s'}$$

for all i and s . This simulates an agent's dynamic payoff under the assumption that v represents future payoffs and is continuous in its arguments.

I first define P . For each state s and agent i , define the correspondences $A_i(s, \cdot), A_i^\circ(s, \cdot): \Theta^+ \rightrightarrows X(s)$ by

$$\begin{aligned} A_i(s, \theta^+) &= \{x \in X(s) \mid U_i(x, s, \theta^+) \geq U_i(q, s, \theta^+)\} \\ A_i^\circ(s, \theta^+) &= \{x \in X(s) \mid U_i(x, s, \theta^+) > U_i(q, s, \theta^+)\}, \end{aligned}$$

and for each coalition C , define $A_C(s, \cdot), A_C^\circ(s, \cdot): \Theta^+ \rightrightarrows X(s)$ by

$$\begin{aligned} A_C(s, \theta^+) &= \{x \in X_C(s) \mid \text{for all } i \in C, x \in A_i(s, \theta^+)\} = \bigcap_{i \in C} A_i(s, \theta^+) \\ A_C^\circ(s, \theta^+) &= \{x \in X_C(s) \mid \text{for all } i \in C, x \in A_i^\circ(s, \theta^+)\} = \bigcap_{i \in C} A_i^\circ(s, \theta^+). \end{aligned}$$

Then define $A(s, \cdot), A^\circ(s, \cdot): \Theta^+ \rightrightarrows X(s)$ by

$$A(s, \theta^+) = \bigcup_{C \subseteq N} A_C(s, \theta^+) \quad \text{and} \quad A^\circ(s, \theta^+) = \bigcup_{C \subseteq N} A_C^\circ(s, \theta^+),$$

and note that $q \in A(s, \theta^+)$. Intuitively, if continuation values are given by θ^+ , then $A^\circ(s, \theta^+)$ consists of the outcomes that would necessarily pass if proposed in equilibrium, and $A(s, \theta^+)$ consists

of the outcomes that could possibly pass in s if proposed in equilibrium. Endowing $X(s)$ with the relative topology induced by X , continuity of U_i implies that for all $s \in S$, $A(s, \cdot)$ has closed graph, $A^\circ(s, \cdot)$ has open graph, and for all θ^+ , $\text{clos}A^\circ(s, \theta^+) \subseteq A(s, \theta^+)$. Furthermore, because $X_C(s)$ is closed (and therefore compact), $A(s, \cdot)$ is actually upper hemi-continuous.

To define P , for each s , let $\hat{u}(s, \cdot): \Theta^+ \rightarrow \overline{\mathfrak{R}}$ be the mapping defined by

$$\hat{u}(s, \theta^+) = \sup\{U_{i(s)}(x, s, \theta^+) \mid x \in A^\circ(s, \theta^+)\},$$

where $\sup \emptyset = -\infty$. Intuitively, this is the payoff that the proposer in state s can guarantee in equilibrium by proposing outcomes strictly acceptable to some coalition. Note that if $\hat{u}(s, \theta^+) > -\infty$, then by compactness of the closure of $A^\circ(s, \theta^+)$ and continuity of U_i , there exists $\bar{x} \in A(s, \theta^+)$ such that $U_{i(s)}(\bar{x}, s, \theta^+) = \hat{u}(s, \theta^+)$. Furthermore, since $A^\circ(s, \cdot)$ has open graph, it is lower hemi-continuous, and Aliprantis and Border's (2006) Lemma 17.29 then implies that $\hat{u}(s, \cdot)$ is lower semi-continuous. Since the proposer can always impose the default, define the "security value"

$$f(s, \theta^+) = \max\{\hat{u}(s, \theta^+), U_{i(s)}(q, s, \theta^+)\}.$$

As each U_i is continuous, and as the pointwise maximum of two lower semi-continuous functions is lower semi-continuous, we see that $f(s, \cdot)$ is lower semi-continuous; and since S is countable with the discrete topology, it follows that f is jointly lower semi-continuous in (s, θ^+) . Define

$$\hat{P}(s, \theta^+) = \{x \in A(s, \theta^+) \mid U_{i(s)}(x, s, \theta^+) \geq f(s, \theta^+)\}.$$

This set is non-empty. Indeed, if $\hat{u}(s, \theta^+) \geq U_{i(s)}(q, s, \theta^+)$, then $\bar{x} \in \hat{P}(s, \theta^+)$, where \bar{x} is described above. Otherwise, set $q \in \hat{P}(s, \theta^+)$. Furthermore, by continuity of U_i and lower semi-continuity of f , \hat{P} has closed graph. Then define $\hat{P}: S \times \Theta^+ \rightrightarrows \mathcal{P}(X)$ by

$$P(s, \theta^+) = \mathcal{P}(\hat{P}(s, \theta^+)).$$

By Aliprantis and Border's (2006) Theorem 17.13, this correspondence has non-empty, convex values and has closed graph.

To define W , let $\text{supp}(\pi_i(s))$ denote the support of $\pi_i(s)$, i.e., the smallest closed set Y such that $\pi_i(s)(Y) = 1$, and note that the correspondence $\text{supp}: \mathcal{P}(X) \rightrightarrows X$ is lower hemi-continuous. (See Aliprantis and Border's (2006) Theorem 17.14.) Now define $g: S \times \Theta^+ \rightarrow \mathfrak{R}$ by

$$g(s, \theta^+) = \min\{U_{i(s)}(s, x, \theta^+) \mid x \in \text{supp}(\pi_s)\},$$

which is well-defined by compactness of X and continuity of $U_{i(s)}$. By Aliprantis and Border's (2006) Lemma 17.14, g is upper semi-continuous. Define the (possibly empty-valued) correspondence $\hat{W}: S \times \Theta^+ \rightrightarrows [0, \bar{u}]$ by

$$\hat{W}(s, \theta^+) = [f(s, \theta^+), g(s, \theta^+)].$$

Clearly, \hat{W} is convex-valued. By Lemma A.1, this correspondence has closed, in fact, compact graph. Since projections of compact sets are compact, the set

$$\hat{\Theta} = \{(s, \theta^+) \in S \times \Theta^+ \mid f(s, \theta^+) \leq g(s, \theta^+)\}$$

is compact. To see that $\hat{\Theta} \neq \emptyset$, choose any $\theta^+ = (\pi, w, v, \gamma)$ such that π_s puts probability one on a payoff maximizing outcome in $X(s)$ for the proposer $i(s)$ in model γ . Thus, by Lemma A.3, we can extend \hat{W} from $\hat{\Theta}$ to a correspondence $W: S \times \Theta^+ \rightrightarrows [0, \bar{w}]$ that has non-empty, convex values and has closed graph.

Finally, I define V . Given state $s \in S_i$, each agent j 's expected payoff depends on the probability that agent i 's proposals pass. If i proposes $x \in A^\circ(s, \theta^+)$, then the proposal must pass in equilibrium, and if $x \notin A(s, \theta^+)$, then the proposal must fail. If agent i proposes $x \in A(s, \theta^+) \setminus A^\circ(s, \theta^+)$ such that $U_i(x, s, \theta^+) \geq w_s \geq U_i(q, s, \theta^+)$ in equilibrium, then the payoff from proposing x must equal w_s , i.e., the probability, say \hat{a} , that x is accepted must satisfy

$$w_s = \hat{a}U_i(x, s, \theta^+) + (1 - \hat{a})U_i(q, s, \theta^+),$$

or, assuming $U_i(x, s, \theta^+) > U_i(q, s, \theta^+)$,

$$\hat{a} = \frac{w_s - U_i(q, s, \theta^+)}{U_i(x, s, \theta^+) - U_i(q, s, \theta^+)}.$$

More generally, when $U_i(x, s, \theta^+) > U_i(q, s, \theta^+)$ but w_s is unrestricted, define

$$\hat{a}(x, s, \theta^+) = \max \left\{ 0, \min \left\{ 1, \frac{w_s - U_i(q, s, \theta^+)}{U_i(x, s, \theta^+) - U_i(q, s, \theta^+)} \right\} \right\},$$

which is continuous in (x, s, θ^+) . Of course, this is not defined when $U_i(x, s, \theta^+) = U_i(q, s, \theta^+)$. Next, define the correspondence $\hat{A}: X \times S \times \Theta^+ \rightrightarrows [0, 1]$ by

$$\hat{A}(x, s, \theta^+) = \begin{cases} \{1\} & \text{if } x = q \\ \{\hat{a}(x, s, \theta^+)\} & \text{if } U_{i(s)}(x, s, \theta^+) > U_{i(s)}(q, s, \theta^+) \\ [0, 1] & \text{else,} \end{cases}$$

and note that \hat{A} has non-empty, convex values. Moreover, \hat{A} has closed graph because $\hat{a}(x, s, \theta^+)$ is continuous, U_i is continuous, and q is isolated. Given s and θ^+ , this correspondence gives the acceptance probabilities, as a function of the outcome proposed in s , that are consistent with the proposer's payoff w_s in θ^+ . Then agent j 's continuation value in s is determined by the precise way that acceptance probabilities depend on proposals, i.e., by a selection from the $\hat{A}(\cdot, s, \theta^+)$ correspondence. Note that the selection is not necessarily satisfy the conditions for equilibrium in voting subgames: it is possible, for example, that $\hat{a}(x, s, \theta^+) < 1$ yet $x \in A^\circ(s, \theta^+)$. This discrepancy is repaired after the fixed point argument.

Define $V(s, \theta^+)$ to be the set of possible continuation value vectors in state s induced by measurable selections from $\hat{A}(\cdot, s, \theta^+)$ as follows: given each measurable section $\hat{a}(\cdot, s, \theta^+)$, we specify that the vector $v' = (v'_{i,s})_{i \in N, s \in S}$ of continuation values defined by

$$v'_{j,s} = \int_X [\hat{a}(x, s, \theta^+)U_j(x, s, \theta^+) + (1 - \hat{a}(x, s, \theta^+))U_j(q, s, \theta^+)]\pi_s(dx), \quad j \in N, s \in S,$$

belongs to $V(s, \theta^+)$. Note first that $V(s, \theta^+)$ is non-empty. Indeed, we obtain a measurable selection from $\hat{A}(\cdot, s, \theta^+)$ by taking any function constant at 1/2 on the set $\{x \in X \mid U_{i(s)}(x, s, \theta^+) =$

$U_{i(s)}(q, s, \theta^+)$, a measurable set. Furthermore, since $\hat{A}(\cdot, s, \theta^+)$ is convex-valued, so is $V(s, \theta^+)$. That V has closed graph follows from a version of Fatou's lemma presented in Lemma A.4. Indeed, to apply that result, let X be the set of outcomes, let $Y = S \times \Theta^+$, let $k = 1$, and let $\Phi = \hat{A}$. Let $f = (f_1, \dots, f_n)$ be defined by

$$f_i(x, a, y) = aU_i(x, s, \theta^+) + (1 - a)U_i(q, s, \theta^+)$$

for all $x \in X$, $y = (s, \theta^+) \in Y$, and $a \in [0, 1]$, and let the correspondence F consist of integrals of f with respect to $\mu = \pi_s$. The desired result then follows immediately from the lemma.

These components together define F , a correspondence with non-empty, convex values and closed graph. By Glicksberg's theorem, for each $\gamma \in \Gamma$, $F(\cdot, \gamma)$ has a fixed point θ^* . Furthermore, the correspondence from parameters γ to the set of fixed points of $F(\cdot, \gamma)$ has closed graph. The next lemma establishes a correspondence between fixed points of $F(\cdot, \gamma)$ and the coalitional bargaining equilibria in model γ , immediately delivering existence of equilibria and non-empty values of the correspondence E . Closed graph follows as well, since $E(\gamma)$ is just the projection of the fixed points of $F(\cdot, \gamma)$ onto $[0, \bar{u}]^{N \times S}$. Then since E has compact range, it is upper hemicontinuous, as required.

Lemma 4.1 *For all $(w, v, \gamma) \in [0, \bar{u}]^S \times [0, \bar{u}]^{N \times S} \times \Gamma$, there exists $\pi \in \prod_{s \in S} \mathcal{P}(X(s))$ such that $(\pi, w, v) \in F(\pi, w, v, \gamma)$ if and only if there is a coalitional bargaining equilibrium $\sigma^* = (\pi^*, \alpha^*)$ with continuation values $v = (v_i(s; \sigma^*))_{i \in N, s \in S}$ and proposer payoffs*

$$w_s = \int_X [\alpha^*(x, s)U_{i(s)}(x, s; \sigma^*) + (1 - \alpha^*(x, s))U_{i(s)}(q, s; \sigma^*)] \pi_{i(s)}^*(s)(dx), \quad s \in S.$$

Let (w, v, γ) be given. I first prove the “only if” direction. Consider any π such that (π, w, v) is a fixed point of $F(\cdot, \gamma)$. For all $i \in N$ and all $s \in S_i$, it follows by construction that $\text{supp}(\pi_s) \subseteq \hat{P}(s, \theta^+)$, and therefore that $f(s, \theta^+) \leq g(s, \theta^+)$. It then follows that $W(s, \theta^+) = \hat{W}(s, \theta^+)$. Thus, for all $x \in \text{supp}(\pi_s)$, we have $U_i(x, s, \theta^+) \geq w_s \geq \max\{\hat{u}(s, \theta^+), U_i(q, s, \theta^+)\}$. Then each $v_{j,s}$ is determined by a selection $\hat{a}(\cdot, s, \theta^+)$ such that acceptance probabilities entail that every proposal x in the support of π_s yields expected payoff w_s to the proposer in state s :

$$w_s = \hat{a}(x, s, \theta^+)U_i(x, s, \theta^+) + (1 - \hat{a}(x, s, \theta^+))U_i(q, s, \theta^+).$$

Next, we specify voting strategies α^* to satisfy the conditions of equilibrium by considering three different types of proposal: x is indifferent between proposing x or imposing the default. When we define equilibrium proposal strategies, below, we correct the inconsistency highlighted here by specifying that with probability $1 - \hat{a}(x, s, \theta^+)$, the agent propose q instead of x .

Case 1: In state $s \in S_i$, agent i proposes x in $A^\circ(s, \theta^+)$. We specify that each agent j accepts x if and only if x belongs to $A_j^\circ(s, \theta^+)$, i.e.,

$$\alpha_j^*(x, s) = \begin{cases} 1 & \text{if } x \in A_j^\circ(s, \theta^+) \\ 0 & \text{else,} \end{cases}$$

which means that x will pass with probability one if proposed, i.e., $\alpha^*(x, s) = 1$. Note that it is possible the selection $\hat{a}(\cdot, s, \theta^+)$ actually specifies that x fail with positive probability, i.e.,

$\hat{a}(x, s, \theta^+) < 1$. This can create an inconsistency in the calculation of the agents' continuation values if π_s puts positive probability on such outcomes, but this can occur only under special conditions. Since $x \in A^\circ(s, \theta^+)$, we have $\hat{u}(s, \theta^+) \geq U_i(x, s, \theta^+)$. But if $x \in \text{supp}(\pi_s)$, then we have

$$\hat{u}(s, \theta^+) \geq U_i(x, s, \theta^+) \geq w_s \geq \max\{\hat{u}(s, \theta^+), U_i(q, s, \theta^+)\} \geq \hat{u}(s, \theta^+),$$

which implies $U_i(x, s, \theta^+) = w_s$. Thus, $U_i(x, s, \theta^+) > U_i(q, s, \theta^+)$ would imply $\hat{a}(x, s, \theta^+) = 1$ by construction of \hat{A} . We conclude that $\hat{a}(x, s, \theta^+) < 1$ is possible only if $U_i(x, s, \theta^+) \leq U_i(q, s, \theta^+)$, and we already have the opposite inequality. Thus, the problem described above can only arise if $U_i(x, s, \theta^+) = U_i(q, s, \theta^+)$, i.e., if agent i is indifferent between proposing x or imposing the default. When I define equilibrium proposal strategies, below, I correct the inconsistency highlighted here by specifying that with probability $1 - \hat{a}(x, s, \theta^+)$, the agent propose q instead of x .

Case 2: In state $s \in S_i$, agent i proposes x in $A(s, \theta^+)$ but not in $A^\circ(s, \theta^+)$. We specify that for each agent j , if $x \in A_j^\circ(s, \theta^+)$, then j accepts x with probability one; if $x \notin A_j(s, \theta^+)$, then j rejects x . Then for every coalition C with $x \in A_C(s, \theta^+)$, there is some $j \in C$ with $U_j(x, s, \theta^+) = U_j(q, s, \theta^+)$. By choosing the vote probabilities of indifferent voters correctly, we can ensure that the probability x is passes is indeed $\hat{a}(x, s, \theta^+)$. To elaborate, consider any coalition C such that $x \in A_C(s, \theta^+)$, and let $C_0 = \{j \in C \mid U_j(x, s, \theta^+) = U_j(q, s, \theta^+)\}$ denote the members of C indifferent between accepting x and rejection, and let $C_1 = \{j \notin C \mid U_j(x, s, \theta^+) = U_j(q, s, \theta^+)\}$ denote the indifferent agents who do not belong to C . To complete the specification of voting strategies in this case, we specify that each $j \in C_1$ reject x with probability one. Now, if all members of C_0 vote to reject x , then it will fail; and if all members of C_0 accept x , then it will pass. Thus, by the intermediate value theorem, there exists $c \in (0, 1)$ such that if all members of C_0 accept x with probability c , then it passes with probability $\hat{a}(x, s, \theta^+)$. We therefore specify that $\alpha_j^*(x, s) = c$ for all $j \in C_0$, obtaining the desired acceptance probability.

Case 3: In state s , agent i proposes x outside $A(s, \theta^+)$. Then each agent j accepts x if and only if $x \in A_j(s, \theta^+)$, which means x fails with probability one, i.e., $\alpha^*(x, s) = 0$. It is possible that $\hat{a}(x, s, \theta^*) > 0$, but since $\text{supp}(\pi_s) \subseteq \hat{P}(s, \theta^+)$, we have $\pi_s(A(s, \theta^+)) = 1$, so outcomes outside $A(s, \theta^+)$ are never proposed in equilibrium. Thus, the discrepancy does not affect the agents' continuation values and is not problematic.

To specify proposal strategies, consider any agent i and state $s \in S_i$. We stipulate that the agent mixes according to π_s , modified to correct the discrepancy in Case 1 above. When the agent is indifferent between imposing the default and proposing an outcome $x \in A^\circ(s, \theta^+)$ in the support of π_s , we require that the proposer place probability $1 - \hat{a}(x, s, \theta^+)$ on q , and otherwise, the agent mixes according to π_s . Formally, we define $\pi_i^*(s)$ so that for all measurable $Y \subseteq X \setminus \{q\}$,

$$\pi_i^*(s)(Y) = \pi_s(Y \setminus A^\circ(s, \theta^+)) + \int_{Y \cap A^\circ(s, \theta^+)} \hat{a}(x, s, \theta^+) \pi_s(dx)$$

and

$$\pi_i^*(x)(\{q\}) = \pi_s(\{q\}) + \int_{Y \cap A^\circ(s, \theta^+)} (1 - \hat{a}(x, s, \theta^+)) \pi_s(dx).$$

This maintains the continuation values generated from the fixed point, so $v = (v_i(s; \sigma^*))_{i \in N, s \in S}$, and we have $U_i(x, s; \sigma^*) = U_i(x, s; \theta^+)$ for all i, x , and s . Moreover, the proposers' expected payoffs are w_s , as in the statement of the lemma.

To see optimality of π_i^* , we must show that no proposal yields an expected payoff greater than w_s . In Case 1, above, a proposal x passes with probability one, and since $x \in A^\circ(s, \theta^+)$, we have $w_s \geq \hat{u}(s, \theta^+) \geq U_i(x, s, \theta^+)$, so the expected payoff from proposing x does not exceed w_s . In Case 2, the acceptance probability $\hat{a}(x, s, \theta^+)$ is chosen so that if $U_i(x, s, \theta^+) > U_i(q, s, \theta^+)$, then the expected payoff from proposing x is exactly w_s . Indeed, recall that $w_s \geq U_i(q, s, \theta^+)$. If $U_i(x, s, \theta^+) > w_s \geq U_i(q, s, \theta^+)$, then the expected payoff from proposing x is w_s by construction; and if $U_i(x, s, \theta^+) = U_i(q, s, \theta^+)$, then the acceptance probability is unrestricted, but then we have $w_s \geq U_i(q, s, \theta^+) = U_i(x, s, \theta^+)$, so proposing x is not a profitable deviation. In Case 3, proposals are rejected with probability one, and since $w_s \geq U_i(q, s, \theta^+)$, no profitable deviation is possible. Therefore, (π^*, α^*) comprises a coalitional bargaining equilibrium.

For the “if” direction, consider a coalitional bargaining equilibrium $\sigma^* = (\pi^*, \alpha^*)$ with $v = (v_i(s; \sigma^*))_{i \in N, s \in S}$ and proposer payoffs $w = (w_s)_{s \in S}$ as in the statement of the lemma. Note by optimality of proposal strategies, we have $w_s \geq U_i(q, s; \sigma^*)$ for all $i \in N$ and all $s \in S_i$. We modify σ^* in two ways. First, adjust each α_j^* so that the agent accepts q if it is proposed. Second, following a proposal of x by agent i such that $U_i(x, s; \sigma^*) > U_i(q, s; \sigma^*)$, let the agents mix so that: (i) x passes with probability one if $w_s \geq U_i(x, s)$, (ii) x fails with probability one if $w_s = U_i(q, s)$, and (iii) if $U_i(x, s; \sigma^*) > w_s > U_i(q, s; \sigma^*)$, then the proposer’s expected payoff is exactly w_s , i.e., $\alpha^*(x, s)U_i(x, s; \sigma^*) + (1 - \alpha^*(x, s))U_i(q, s; \sigma) = w_s$. The conditions (i)–(iii) necessarily hold for all $x \in \text{supp}(\pi_i^*(s))$, except perhaps on a set of $\pi_i^*(s)$ -measure zero, so our modifications do not affect the agents’ continuation values. Letting $\pi = (\pi_i^*(s))_{i \in N, s \in S_i}$ and $\theta^+ = (\pi, w, v, \gamma)$, it follows that $\alpha^*(x, s) \in \hat{A}(x, s, \theta^+)$ for all x and all s , i.e., the acceptance probability $\alpha^*(\cdot, s)$ is a selection from $\hat{A}(\cdot, s, \theta^+)$, which implies $v \in V(\theta^+)$. Furthermore, we have $U_i(x, s; \theta^+) = U_i(x, s; \sigma^*)$ for all i , all x , and all s , and this implies $\pi \in P(\theta^+)$, and finally $w \in W(\theta^+)$. Therefore, (π, w, v) is a fixed point of $F(\cdot, \gamma)$, completing the proof.

5 Overview of Proof

To convey the approach to existence, I take the argument of Banks and Duggan (2000) as a starting point. There, a proposer is randomly drawn and a vote held; if the proposal passes, then the game ends, and otherwise the game continues to the next period and the process is repeated. Assume for now that stage payoffs are strictly positive (so delay is Pareto inefficient) and concave. The proof of existence in Banks and Duggan (2000) takes the form of a fixed point argument in the domain of profiles of proposal strategies, $\pi = (\pi_1, \dots, \pi_n)$. Given a profile π , assuming these proposal strategies are used in the future and there is no delay (without loss of generality in this model), each agent’s continuation value is calculated directly: letting p_j denote the probability that j is recognized to propose, it is $v_i(\pi) = \sum_{j \in N} p_j \int_x u_i(x) \pi_j(dx)$. Given these continuation values, and assuming that each agent accepts when indifferent (also without loss of generality), we can calculate for each agent the set $A_i(\pi)$ of alternatives the agent would accept conditional on being proposed, and then we specify the set $A(\pi)$ of alternatives that would pass if proposed. The equilibrium condition on proposals is that each agent propose a utility-maximizing element of $A(\pi)$ if

recognized, engendering a correspondence from the space of mixed proposal strategy profiles into itself, below,

$$\pi = (\pi_i)_{i \in N} \longrightarrow (v_i(\pi))_{i \in N} \longrightarrow A(\pi) \longrightarrow (\mathcal{P}(\arg \max_{y \in A(\pi)} u_i(y)))_{i \in N}$$

and the existence proof consists in verifying the conditions of Glicksberg’s theorem. Compactness (in the weak* topology) and convexity of the domain is not problematic. That the correspondence has convex values follows directly from the construction, and a simple continuity argument establishes that it has closed values.

The crux of the argument is twofold: showing that the correspondence has nonempty values and that it is upper hemi-continuous. Nonemptiness follows from concavity of stage payoffs. Indeed, letting $E[\pi]$ denote the mean of the probability measure $\sum_{j \in N} p_j \pi_j$, concavity yields $u_i(E[\pi]) \geq v_i(\pi)$ for each agent, and therefore $E[\pi] \in A(\pi)$. In fact, because stage payoffs are positive, we have $u_i(E[\pi]) > \delta_i v_i(\pi)$, so each agent would strictly accept this expected outcome if it is proposed. The proof of upper hemi-continuity follows from an application of the theorem of the maximum. For this, it is important that the correspondence $A(\pi)$ is continuous. Upper hemi-continuity of $A(\pi)$ is straightforward to verify, but the argument for lower hemi-continuity uses the fact that the expected value $E[\pi]$ belongs to the set $A^\circ(\pi)$ of proposals that strictly pass; by continuity, the correspondence $A^\circ(\pi)$ has open graph, and concavity implies that $A(\pi)$ lies in the closure of $A^\circ(\pi)$. Therefore, $A(\pi)$ is indeed lower hemi-continuous. The approach of Banks and Duggan (2006) is similar, but instead of assuming a bad status quo, that paper allows for an arbitrary status quo and imposes a common discount factor to obtain the result.

When concavity is dropped, risk aversion is no longer sufficient for nonemptiness of $A(\pi)$. To address this, we can append a default outcome available to the proposer in case no other proposal is weakly preferable to delay for a decisive coalition, but the argument for lower hemi-continuity of $A(\pi)$ no longer goes through, as illustrated in Figure 1. Here, consider a sequence $\{\pi^m\}$ of proposal strategy profiles converging to a limit π , with an agent’s discounted continuation value decreasing along the sequence and her acceptance correspondence expanding discontinuously at the limit, violating lower hemicontinuity. A single agent’s acceptance set will always be nonempty (it will contain all utility maximizing outcomes), but this is not necessarily true of coalitional acceptance sets $A_C(\pi)$, which may be empty along the sequence with coalitionally acceptable outcomes appearing in the limit—another way of violating lower hemicontinuity. The problem in these examples is that in the limit, an agent is indifferent between accepting x and rejecting it, and we have naively specified that the agent vote to accept when indifferent.

The logic of Ray and Vohra (1999) still relies on lower hemicontinuity of coalitional acceptance sets in the NTU framework, but the authors obtain this property by assuming strict comprehensiveness of the coalitional function V . In their model, once a coalition has formed, the remaining agents play a “smaller” coalition formation game, continuing until all agents have selected into (possibly singleton) coalitions. The existence of equilibrium in the game with n agents therefore requires existence of equilibria in all smaller games, which is assumed in an induction argument. The argument for the n -player game uses a fixed point argument in which the domain consists of the expected payoff, w_i , to each agent when proposing and for each agent a probability distri-




Figure 1: Lower hemi-continuity problem

bution, γ_i , over coalitions. Given the rejector-becomes-proposer protocol, this allows the authors to calculate the value of proposing to each coalition, and then to calculate the expected payoff, w_i^j , to i when j proposes, assuming: j proposes according to π_j , and equilibrium payoffs when i is excluded from the forming coalition are given by induction. The latter is needed, for when agent i proposes, the agent can effectively determine the next proposer j and can obtain the payoff w_i^j . Finally, this determines the agent's optimal payoff as proposer and the set of optimal proposals (either proposing to a coalition or delegating to another agent), as depicted below.

$$\begin{array}{ccc}
 (\gamma, w) & \longrightarrow & (\max_{y \in A_C(w)} u_i(y))_{i,C} & \longrightarrow & (w_i^j)_{i,j \in N} \\
 & & & \searrow & \\
 & & & & (\mathcal{P}(\arg \max_{C,j} \{w_i^C, w_i^j\}), \max_{C,j} \{w_i^C, w_i^j\})_{i \in N}
 \end{array}$$

This correspondence satisfies the conditions of Kakutani's theorem and therefore admits a fixed point, which corresponds to a stationary equilibrium of the model. Without strict comprehensiveness, however, acceptance sets may violate lower hemicontinuity in a way that creates a discontinuity in the proposer's optimal payoff to a coalition, as depicted in Figure 2, where agent 1 cannot form a coalition with agent 2 along the sequence but in the limit can obtain a positive payoff with the agent.

Clearly, the assumptions of concavity and positive utility play an important role in the argument of Banks and Duggan (2000), and strict comprehensiveness is critical for the approach of Ray and Vohra (1999), but none of these conditions are available in the general framework of the current paper.¹⁰ One approach to the problems outlined above is to specify that agents reject when indifferent, but this creates difficulties for upper hemi-continuity of $A(\pi)$, as can be seen by modifying the example in Figure 1 so that the agent's discounted continuation increases (rather than decreases) to $\delta_i v_i(\pi)$. To salvage the fixed point argument, we must specify the votes of indifferent voters in a more nuanced way. Another possibility is to explicitly include the agents'

¹⁰Although stage payoffs are normalized to be non-negative, I make no assumption about payoffs from the default outcome, which are zero in much of the legislative bargaining literature.




Figure 2: Lower hemi-continuity problem

acceptance strategies in the definition of the correspondence, expanding the domain to include the vector $\alpha = (\alpha_1, \dots, \alpha_n)$ of acceptance strategies. But acceptance decisions are conditioned on the proposal made, so each $\alpha_i \in \{0, 1\}^X$ lives in a function space, and it is not clear how this space should be topologized when the set X of outcomes is infinite: the product topology does not give sequential compactness and is not useful in the current context, and other common topologies would fix a measure on X and require us to ignore sets of measure zero, a maneuver unavailable here.

The approach in this paper is to circumvent acceptance strategies by means of a Fatou's lemma argument, establishing existence of a fixed point and backing voting strategies out in a way consistent with equilibrium. A key insight is that to calculate the agents' continuation values, it is only necessary to know their proposal strategies and the probability that any proposed outcome passes — we do not need the individual acceptance probabilities of the agents. Moreover, although these acceptance probabilities are conditioned on the outcome proposed and live in a complex, infinite-dimensional space, they can be reduced to a dimensionality equal to the cardinality of the state space. Specifically, if we know the continuation values and the expected payoff to the proposer in each state, then in a given state we can calculate the proposer's payoffs if the proposal is rejected, say a , and if the proposal is accepted, say c . Moreover, if the proposer's expected payoff from proposing, say b , satisfies $a < b < c$, then the probability that the proposal is accepted must be $\frac{b-a}{c-a}$ to reconcile proposer payoffs with these other quantities. This observation, detailed below, suggests that given continuation values and proposer payoffs, we can sufficiently pin down voting strategies to update the agents' continuation values. In fact, the argument is complicated by the contingency that the proposer is indifferent between acceptance and rejection, in which case $a = c$ in the above story, and an acceptance probability is not uniquely pinned down.¹¹ The fixed point argument therefore involves integration over all selections of acceptance probabilities, and this is where Fatou's lemma plays a critical role in proving upper hemicontinuity of the correspondence. Thus, I expand the domain of the correspondence to include continuation values and proposer payoffs as well as the agents' mixed proposal strategies, then deduce the existence of a fixed point, and then back out voting strategies consistent with equilibrium.

To elaborate, return to the setting of Banks and Duggan (2000) but without assuming concavity or a bad status quo. Suppose that an agent is recognized as proposer, that continuation values are

¹¹This is irrelevant for updating the proposer's continuation values, but some other agents may not be indifferent, so the specification of acceptance probability is significant.

$v = (v_1, \dots, v_n)$, that her expected payoff from proposing is w_i , and that she uses a mixed proposal strategy π_i . Consider one play of the coalitional bargaining game induced by v , in which the agents act as though future payoffs were given by these continuation values, and construct the set $A(v)$ of outcomes that give the members of one or more decisive coalitions a dynamic payoff at least equal to the dynamic payoff of rejection; these are the proposals that could conceivably pass in an equilibrium of the induced game. And construct the set $A^\circ(v)$ of proposals that the members of a decisive coalition strictly prefer to rejection; these outcomes would surely pass in equilibrium if proposed. Of course, a proposer can always choose to impose the default. This gives us an upper bound and lower bound,

$$\max_{y \in A(v)} \frac{u_i(y)}{1 - \delta_i} \quad \text{and} \quad \max \left\{ \sup_{y \in A^\circ(v)} \frac{u_i(y)}{1 - \delta_i}, u_i(q) + \frac{\delta_i v_i}{1 - \delta_i} \right\},$$

on the expected payoff of the proposer in any equilibrium of the induced game. We refer to the lower bound as the proposer's "security value." Assume:

- (i) π_i puts probability one on outcomes in $A(v)$ that meet or exceed the agent's security value,
- (ii) w_i is greater than or equal to the agent's security value,
- (iii) the payoff of the worst outcome in the support of π_i is at least w_i .

Then the problem is to specify acceptance probabilities $\alpha(x)$ for every outcome that are consistent with π_i being optimal and with the incentives in voting subgames.

In order that π_i is an optimal proposal, it must be that agent i is indifferent over all of the alternatives in the support of π_i : they must all yield an expected payoff of w_i . So it suffices to specify acceptance probabilities such that all outcomes in the support of π_i yield an expected payoff of w_i , and no outcomes yield a higher payoff. Given an outcome x in the support of π_i , suppose for the sake of argument that the proposer strictly prefers x to rejection, i.e.,

$$\frac{u_i(x)}{1 - \delta_i} > u_i(q) + \frac{\delta_i v_i}{1 - \delta_i}.$$

If w_i is above agent i 's dynamic payoff from x , then we specify that x passes with probability one, and if w_i is below the payoff from q , we specify that x fails with probability one. Otherwise, w_i is between agent i 's dynamic payoff from x and q , and then we must have

$$w_i = \alpha(x) \left(\frac{u_i(x)}{1 - \delta_i} \right) + (1 - \alpha(x)) \left(u_i(q) + \frac{\delta_i v_i}{1 - \delta_i} \right),$$

so the probability that x passes when proposed by agent i must be


$$\alpha(x) = \frac{(1 - \delta_i)w_i - u_i(q) - \delta_i v_i}{u_i(x) - u_i(q) - \delta_i v_i}.$$

This pins down the probability that x passes when proposed by agent i , unless the proposer is indifferent between x passing and remaining silent. In that case, the acceptance probability is

indeterminate. This construction yields a correspondence $\hat{A}^i(\cdot; v, w): X \rightrightarrows [0, 1]$ of acceptance probabilities such that for every selection α^i from $\hat{A}^i(\cdot; v, w)$, the mixed proposal strategy π_i is optimal and yields an expected payoff of w_i to agent i from proposing. Importantly, the correspondence $\hat{A} = \hat{A}^1 \times \cdots \times \hat{A}^n$ has convex values and closed graph. Note that we are not assured that the selection of acceptance probabilities is consistent with the equilibrium conditions in voting subgames, a point to which we return later.

The fixed point correspondence $F = P \times W \times V$ maps from triples (π, w, v) to product sets of such triples. We then update the agents' continuation values as follows: for every agent i and every selection α^i from \hat{A}^i , we calculate new continuation values $\hat{v} = (\hat{v}_1, \dots, \hat{v}_n)$ by integrating over proposals with respect to the mixed proposal strategies $\pi = (\pi_1, \dots, \pi_n)$, and we collect these in the set $V(\pi, w, v)$. This correspondence inherits convex values from \hat{A} , and we use a version of Fatou's lemma to prove that it has closed graph. To be more explicit, we consider sequences $\{(\pi^m, w^m, v^m)\}$ and $\{\hat{v}^m\}$ such that $\hat{v}^m \rightarrow \hat{v}$, $(\pi^m, w^m, v^m) \rightarrow (\pi, w, v)$, and $\hat{v}^m \in V(\pi^m, w^m, v^m)$ for all m . Each \hat{v}^m corresponds to some selection α^m from the correspondence $\hat{A}(\cdot; \pi^m, w^m, v^m)$. By assumption, the sequence of updated continuation values converges to \hat{v} , but we do not know whether the sequence $\{\alpha^m\}$ converges—we have not topologized the space of acceptance strategies, and there is no need to do so. By Lemma A.4, the limiting continuation values \hat{v} will indeed correspond to some selection from $A(\cdot; \pi, w, v)$, even if the sequence $\{\alpha^m\}$ of selections does not converge to α in an interesting topology. For the application of the lemma, however, it is critical that \hat{A} have convex values: here, convexity begets closed graph.

The definitions of the components P and W are less involved, with the updated proposal strategies taking v as input and updated proposer values taking π and v ; this is pictured below, where $\theta = (\pi, w, v)$. These components are specified so that they have nonempty, convex values and closed graph and so that for every fixed point (π, w, v) of the correspondence, conditions (i)–(iii) are satisfied. Continuation values are generated by a selection from $\hat{A}(\cdot; v, w)$ as described above, but minor adjustments are required, as the construction does not ensure that this selection will be consistent with the equilibrium conditions in voting subgames. In particular, it is possible that a proposed outcome x passes with probability less than one although a decisive coalition of agents strictly prefer x to rejection, and so it must pass with probability one in equilibrium. In the proof, we show this can happen only under restrictive conditions, and we modify the proposal and voting strategies derived from the fixed point argument to correct this problem. Specifically, this is only possible if the proposer is indifferent between x and imposing the default, and so we specify that the agent propose q whenever x would have been proposed and failed, preserving the agents' continuation values while satisfying the conditions for equilibrium



6 Concluding Discussion

The main result of this paper establishes existence of coalitional bargaining equilibria, along with upper hemi-continuity of the equilibrium correspondence, in a general model of bargaining and

coalition formation. The theorem improves known results by dropping all convexity and comprehensiveness: it imposes only a metric space structure on the set of outcomes and assumes only compactness and continuity conditions. As applications, we obtain existence of equilibrium in models of bargaining over parameters of a game (including issue-by-issue bargaining among committees), coalition formation in general NTU environments with stochastically determined proposers, and collective dynamic programming with a metric space of actions and countable state space.

The existence result relies on the assumption that in case a proposal x is rejected in state s , the transition to next period's state is governed by $p(\cdot|q, s)$, which is independent of the identity of the agents who reject the proposal. In particular, the order of voting and the identity of the first agent to reject are irrelevant. An alternative protocol of interest in the literature on coalition formation, but not covered by the current framework, is that in which a proposer makes an offer to a particular coalition, the offer is considered sequentially by members of that coalition, and the first member to reject (if any) makes a counter-offer. This "rejector-becomes-proposer" protocol presents especially difficult problems for existence in the general environments considered here. When strict comprehensiveness is assumed, as in Ray and Vohra (1999), a proposer can make transfers to break indifference among coalition members, so we can restrict attention to pure voting strategies, and the set equilibrium outcomes in a voting subgame is trivially convex. In general, however, indifferent voters cannot always be induced to accept a proposal, and mixed voting equilibria are unavoidable. In the proof of Theorem 1, it is important that the agents' continuation values depend only on the probability that proposals pass or fail (not on who rejects), and that in any voting equilibrium, this set of overall acceptance probabilities is convex: if a proposal can pass with some intermediate probability between zero and one, then it must be that some voters are indifferent between accepting and rejecting, and we can specify that they mix arbitrarily to convexify the set of acceptance probabilities.

But when the first coalition member to reject becomes the next proposer, the agents' continuation values depend on who rejects first. Given mixed voting strategies of the agents, we must describe an outcome of the voting game as a probability distribution over the event that the proposal passes and, for each member i , the event that i is the first to reject. Here, indifference cannot be used to convexify the set of voting outcomes. The problem is illustrated in the first panel of Figure 3, where I consider the possibility that outcome x is proposed to the coalition $\{1, 2\}$, and the agents respond in order. If agent 1 is the first to reject, then the agent makes a counter offer and all agents receive their continuation values, an outcome I represent simply by q_1 . I represent the outcome that agent 2 is the first to reject by q_2 . Imagine that agent 2 is indifferent between accepting and rejecting, so the agent may mix arbitrarily in equilibrium; assume agent 1 strictly prefers to accept if agent 2 will accept, and the agent strictly prefers to reject if agent 2 will reject. Then there are two pure strategy equilibria, one, (accept, accept), in which the proposal passes, and one, (reject, reject), in which agent 1 is the first to reject.

I claim that the equally weighted convex combination of these distributions cannot be obtained in any voting equilibrium. If the probability that agent 1 is the first to reject is one half, then it must be that the agent's vote is determined by a coin toss, as in the second panel of Figure 3; and if the proposal passes with probability one half, then it must be that agent 2 always accepts. But then agent 1 has a strict preference to accept, and it is not a best response to mix. Because of this non-convexity, the application of Fatou's lemma to prove closed graph of the correspondence $V(s, \theta^+)$




Figure 3: Nonconvex voting equilibrium outcomes

in the proof of Theorem 1 does not go through. The current proof approach can be salvaged by considering correlated voting equilibria, allowing the members of the coalition to observe a public randomization device before responding, to obtain convexity of the set of voting outcomes. This approach, while an extension of the notion of stationary equilibrium used in the bargaining literature, is a well-known remedy to non-existence of stationary equilibria in stochastic games and subgame equilibria in general extensive form games; I conjecture that it will yield existence of stationary subgame perfect equilibrium (with correlated voting equilibria) in the coalitional bargaining model.

A Technical Appendix

Example 1 *The unique coalitional bargaining equilibrium discounted continuation value is .7.* Consider any coalitional bargaining equilibrium. I first claim that for each agent i , $\delta v_i \geq .7$. By symmetry, we can focus on agent 1, so suppose $\delta v_1 < .7$. Then agent 1 must accept z when proposed, and therefore agent 3 proposes z with probability one. It cannot be that agent 2 accepts x with probability one when proposed by agent 1, for then agent 1 would propose x , and then $\delta v_1 \geq (.3)[2 + 0 + .7] > .7$, a contradiction. Therefore $\delta v_2 \geq .7$, and it follows that outcome y must be realized with positive probability when agent 2 proposes, so agent 3 must accept y with positive probability, which implies $\delta v_3 \leq .7$. On the other hand, it cannot be that agent 3 accepts y with probability one, for then agent 2 would propose it, and then $\delta v_3 \geq (.3)[0 + .7 + 2] > .7$, a contradiction. Therefore, $\delta v_3 = .7$. Letting p denote the probability that agent 2 proposes y or remains silent, we have $.7 = \delta v_3 \geq (.3)[0 + p(.7) + 2]$, which implies $p < .5$. Therefore, the probability that agent 2 proposes x is $1 - p > .5$. Since agent 1 accepts x , we then have $\delta v_1 \geq (.3)[.7 + (.5)(2) + .7] > .7$, a contradiction that establishes the claim. Next, I claim that for each agent i , $\delta v_i \leq .7$. Focusing on agent 1, suppose $\delta v_1 > .7$. Then agent 1 rejects z when it is proposed, and since the agent's payoff from remaining silent exceeds the payoff from z , agent 1 does not propose z . Then $\delta v_3 < .7$, and agent 3 must accept y if it is proposed, so agent 2 proposes y , and it passes. Furthermore, since agent 3's payoff from proposing y is $.7 > \delta v_3$, the agent will not propose z and will not remain silent. But then $\delta v_1 \leq (.3)[2 + 0 + 0] < .7$, a contradiction. \square

Two very elementary lemmas follow.

Lemma A.1 *Let X be a topological space, let $f: X \rightarrow \mathfrak{R}$ be lower semi-continuous, and let $g: X \rightarrow$*

\mathfrak{R} be upper semi-continuous. Define $\phi: X \rightrightarrows \mathfrak{R}$ by $\phi(x) = [f(x), g(x)]$ for all $x \in X$. Then ϕ has closed graph.

Proof: Take $\{(x^\alpha, y^\alpha)\}$ such that $y^\alpha \in \phi(x^\alpha)$ for all α and $(x^\alpha, y^\alpha) \rightarrow (x, y)$. Then $f(x^\alpha) \leq y^\alpha \leq g(x^\alpha)$ for all α . By assumption, $f(x) \leq \liminf f(x^\alpha) \leq y \leq \limsup g(x^\alpha) \leq g(x)$, so $y \in \phi(x)$. ■

Lemma A.2 *Let X be a topological space, let $\phi: X \rightrightarrows [a, b]$ be a non-empty, convex-valued correspondence with closed graph and compact range. Then there exist $f: X \rightarrow \mathfrak{R}$ and $g: X \rightarrow \mathfrak{R}$ such that f is lower semi-continuous, g is upper semi-continuous, and for all $x \in X$, $\phi(x) = [f(x), g(x)]$.*

Proof: By Aliprantis and Border's (2006) Lemma 17.30, the function $g(x) := \max \phi(x)$ is upper semi-continuous. Similarly, $h(x) := \max\{-c \mid c \in \phi(x)\}$ is upper semi-continuous, which implies that $f(x) := -h(x) = \min \phi(x)$ is lower semi-continuous, as required. ■

The next lemma extends an interval-valued correspondence with closed graph from a compact subset of a metric space to the entire metric space.

Lemma A.3 *Let X be a metric space, let Y be a non-empty, compact subset of X . Let $\phi: Y \rightrightarrows [a, b]$ be a non-empty, convex-valued correspondence with closed graph and compact range. Then there exists $\Phi: X \rightrightarrows [a, b]$ such that Φ has non-empty, convex values, has closed graph, and extends ϕ , i.e., for all $y \in Y$, $\Phi(y) = \phi(y)$.*

Proof: The proof is constructive. Let f and g be as in the statement of Lemma A.2. For all $x \in X$, let $M(x)$ be the set of solutions to

$$\min_{y \in Y} d(x, y),$$

where d , our metric, is continuous. By the theorem of the maximum, M is upper hemi-continuous with non-empty, compact values. Now define

$$g^*(x) = \max_{y \in M(x)} g(y)$$

for all $x \in X$. By Aliprantis and Border's (2006) Lemma 17.30, g^* is upper semi-continuous. Similarly,

$$h(x) := \max_{y \in M(x)} -f(y)$$

is upper semi-continuous, so

$$f^*(x) := -h(x) = \min_{y \in M(x)} f(y)$$

is lower semi-continuous. Define $\Phi(x) = [f^*(x), g^*(x)]$ for all $x \in X$. This correspondence clearly has non-empty, convex values, and it has closed graph by Lemma A.1. ■

The final lemma extends the version of Fatou's lemma on upper hemi-continuity of integrals of correspondences due to Aumann (1976) and Yannelis (1990). For a simplified statement of their result, let X and Y be metric spaces, with X compact, let (X, Σ, μ) be a measure space with Σ the completion of the Borel σ -algebra, and let $F: X \times Y \rightrightarrows \mathfrak{R}^k$ have nonempty values and closed graph. Let $\int F(x, y)\mu(dx)$ consist of all integrals of measurable selections from $F(\cdot, y)$. Then the correspondence of integrals, $\int F(x, \cdot)\mu(dx): Y \rightrightarrows \mathfrak{R}^k$, has closed graph. For our arguments, we need to allow the probability measure μ to vary, and as a consequence we add the assumption of convex values. A technical extension, which is useful in the application of this paper, is that I consider integrals of selections from $F(\cdot, y)$ composed with a continuous function $f(x, \phi(x), y)$ that is linear in its second argument.

Lemma A.4 *Let X and Y be metric spaces, and assume X is compact. Let $\Phi: X \times Y \rightrightarrows [0, 1]^k$ be a correspondence with non-empty, convex values and closed graph. Let $f: X \times [0, 1] \times Y \rightarrow \mathfrak{R}^n$ be continuous, and assume that for all $x \in X$ and all $y \in Y$, $f(x, a, y)$ is affine linear in $a \in [0, 1]^k$. Then the correspondence $F: Y \times \mathcal{P}(X) \rightrightarrows \mathfrak{R}^n$ defined by*

$$F(y, \mu) = \left\{ \int_X f(x, \phi(x), y)\mu(dx) \mid \phi \text{ is a Borel mble selection from } \Phi(\cdot, y) \right\}$$

for all $(y, \mu) \in Y \times \mathcal{P}(X)$ has closed graph.

Proof: Consider a sequence $\{(\mu^m, y^m, c^m)\}$ in $\mathcal{P}(X) \times Y \times [0, 1]^k$ such that $c^m \in F(y^m, \mu^m)$ for all m and such that $(\mu^m, y^m, c^m) \rightarrow (\mu, y, c)$. Thus, for each m , there exists a measurable selection ϕ^m from $\Phi(\cdot, y^m)$ such that

$$c^m = \int_X f(x, \phi^m(x), y^m)\mu^m(dx).$$

Let \mathcal{X} and \mathcal{A} denote the Borel sigma-algebras on X and $[0, 1]^k$, respectively, and let \mathcal{S} denote the Borel sigma-algebra on $X \times [0, 1]^k$. Note that $\mathcal{S} = \mathcal{X} \otimes \mathcal{A}$. (See Aliprantis and Border's (2006) Theorem 4.44.) Define the probability measure ν^m on $(X \times [0, 1]^k, \mathcal{S})$ as follows: given Borel measurable $S \in \mathcal{S}$, let

$$\nu^m(S) = \mu^m(\{x \in X \mid (x, \phi^m(x)) \in S\}).$$

Since $\Phi(\cdot, y^m)$ has closed graph for each m , we have $\text{supp}\nu^m \subseteq \text{graph}\Phi(\cdot, y^m)$. By a change of variables,

$$c^m = \int_{X \times [0, 1]^k} f(x, a, y^m)\nu^m(d(x, a)).$$

Furthermore, since $X \times [0, 1]^k$ is compact, $\{\nu^m\}$ must have a weak* convergent subsequence (still indexed by m for simplicity) with limit, say, ν . Since Φ has closed graph and the support of a probability measure varies lower hemicontinuously in the weak* topology, we have

$$\text{supp}(\nu) \subseteq \limsup \text{supp}(\nu^m) \subseteq \limsup \text{graph}(\Phi(\cdot, y^m)) \subseteq \text{graph}(\Phi(\cdot, y)).$$

Using continuity of f , Billingsley's (1968) Theorem 5.5 implies that

$$c = \lim c^m = \lim \int_{X \times [0,1]^k} f(x, a, y^m) \nu^m(d(x, a)) = \int_{X \times [0,1]^k} f(x, a, y) \nu(d(x, a)).$$

Note that the marginal of ν^m on X is just μ^m , and by Billingsley's (1968) Theorem 3.1, $\nu^m \rightarrow \nu$ weak* implies that the marginals of ν^m also converge weak* to the marginal of ν . Thus, the marginal of ν on X is in fact μ .

Fixing y , define the random variable ξ on the probability space $(X \times [0, 1]^k, \mathcal{S}, \nu)$ by $\xi(x, a) = x$, and define the random variable α on $(X \times [0, 1]^k, \mathcal{S}, \nu)$ by $\alpha(x, a) = a$. Let $\mathcal{T} = \{\{Z\} \times [0, 1]^k \mid Z \in \mathcal{X}\}$ be the sigma-algebra of events conditioning on information about x . By Dudley's (2002) Theorem 10.2.2, there exists a conditional distribution for α given \mathcal{T} , $P_{\alpha|\mathcal{T}} : \mathcal{A} \times X \times [0, 1]^k \rightarrow [0, 1]$ such that (i) there exists $T \in \mathcal{T}$ such that $\nu(T) = 0$ and for all $(x, a) \in (X \times [0, 1]^k) \setminus T$, $P_{\alpha|\mathcal{T}}(\cdot, x, a)$ is a probability measure on \mathcal{A} , and (ii) for all $A \in \mathcal{A}$, $P_{\alpha|\mathcal{T}}(A, \cdot)$ is a version of the probability of A , conditional on \mathcal{T} , and is \mathcal{T} -measurable, i.e., constant in a . Then, by Dudley's (2002) Theorem 10.2.1, conditional distributions $\{P_x \mid x \in X\}$ exist for ν , i.e., for all $A \in \mathcal{A}$, all $Z \in \mathcal{X}$, all $(x, a) \in X \times [0, 1]^k$, and all $T \in \mathcal{T}$,

- (a) P_x is a probability measure on $([0, 1]^k, \mathcal{A})$,
- (b) $\nu(Z \times A) = \int_Z P_x(A) \mu(dx)$,
- (c) $x \mapsto P_x(A)$ is \mathcal{X} -measurable.

Furthermore, we have $P_x(A) = P_{\alpha|\mathcal{T}}(A, x, a)$. Finally, for every integrable $g : X \times [0, 1]^k \rightarrow \mathfrak{R}$, we have

$$\int_{X \times [0,1]^k} g(x, a) \nu(d(x, a)) = \int_X \int_{[0,1]^k} g(x, a) P_x(da) \mu(dx).$$

If $g : X \times [0, 1]^k \rightarrow \mathfrak{R}^n$ and each component g_i is integrable, then the latter observation extends straightforwardly.

Since $X \times [0, 1]^k$ is compact and f is continuous, each component f_i is integrable. As a consequence of the preceding observations, we have

$$\int_{X \times [0,1]^k} f(x, a, y) \nu(d(x, a)) = \int_X \int_{[0,1]^k} f(x, a, y) P_x(da) \mu(dx) = \int_X f(x, E[a|x], y) \mu(dx),$$

where $E[a|x] = \int_{[0,1]^k} a P_x(da)$ is Borel measurable, and where the second equality relies on linearity of $f(x, a, y)$ in a . Also, since $\text{supp} \nu \subseteq \text{graph} \Phi(\cdot, y)$, we have

$$1 = \int_{X \times [0,1]^k} I_{\text{graph} \Phi(\cdot, y)}(x, a) \nu(d(x, a)) = \int_X \int_{[0,1]^k} I_{\text{graph} \Phi(\cdot, y)}(x, a) P_x(da) \mu(dx),$$

and it follows that for μ -almost every x , $P_x(\Phi(x, y)) = 1$. Since $\Phi(x, y)$ is convex, we have $E[a|x] \in \Phi(x, y)$. We can then construct ϕ by splicing $E[a|\cdot]$ with an arbitrary measurable selection on the measure zero set such that $E[a|x] \notin \Phi(x, y)$ without affecting the value of the integral

$\int f(x, E[a|x], y) d\mu$. Thus, $E[a|\cdot]: X \rightarrow [0, 1]^k$ yields a measurable selection ϕ from Φ satisfying

$$c = \int_X f(x, \phi(x), y) \mu(dx)$$

and therefore $c \in F(y, \mu)$, as required. ■

References

- [1] D. Acemoglu, G. Egorov, and K. Sonin (2009) “Dynamics and Stability of Constitutions, Coalitions, and Clubs,” Unpublished manuscript.
- [2] C. Aliprantis and K. Border (2006) *Infinite Dimensional Analysis*, 3rd ed., Springer: New York.
- [3] R. Aumann (1976) “An Elementary Proof that Integration Preserves Uppersemicontinuity,” *Journal of Mathematical Economics*, 3: 15–18.
- [4] S. Bandyopadhyay and K. Chatterjee (2006) “Coalition Theory and Its Applications: A Survey,” *The Economic Journal*, 116: F136–F155.
- [5] J. Banks and J. Duggan (2000) “A Bargaining Model of Collective Choice,” *American Political Science Review*, 94: 73–88.
- [6] J. Banks and J. Duggan (2006) “A General Bargaining Model of Legislative Policy-making,” *Quarterly Journal of Political Science*, 1: 49–85.
- [7] D. Baron (1996) “A dynamic theory of collective goods programs,” *American Political Science Review*, 90: 316–330.
- [8] P. Billingsley (1968) *Convergence of Probability Measures*, Wiley: New York.
- [9] T. Börgers (1989) “Perfect Equilibrium Histories of Finite- and Infinite-horizon Games,” *Journal of Economic Theory*, 47: 218–227.
- [10] T. Börgers (1991) “Upper Hemicontinuity of the Correspondence of Subgame-perfect Equilibrium Outcomes,” *Journal of Mathematical Economics*, 20: 89–106.
- [11] S. Chakrabarti (1999) “Markov Equilibria in Discounted Stochastic Games,” *Journal of Economic Theory*, 85: 294–327.
- [12] R. Dudley (2002) *Real Analysis and Probability*, Cambridge University Press: New York.
- [13] H. Eraslan and A. Merlo (2002) “Majority Rule in a Stochastic Model of Bargaining,” *Journal of Economic Theory*, 103: 31–78.
- [14] D. Fudenberg and D. Levine (1983) “Subgame-perfect Equilibria of Finite and Infinite-horizon Games,” *Journal of Economic Theory*, 31: 251–268.

- [15] C. Harris (1985a) “Existence and Characterization of the Perfect Equilibrium in Games of Perfect Information,” *Econometrica*, 53: 613–628.
- [16] C. Harris (1985b) “A Characterization of the Perfect Equilibria of Infinite-horizon Games,” *Journal of Economic Theory*, 37: 99–125.
- [17] C. Harris, P. Reny, and A. Robson (1995) “The Existence of Subgame-Perfect Equilibrium in Continuous Games with Almost Perfect Information: A Case for Public Randomization,” *Econometrica*, 63: 507–544.
- [18] M. Hellwig, W. Leininger (1987) “On the Existence of Subgame-perfect Equilibrium in Infinite-action Games of Perfect Information,” *Journal of Economic Theory*, 43: 55–75.
- [19] M. Hellwig, W. Leininger, P. Reny, A. Robson (1990) “Subgame-perfect Equilibrium in Continuous Games of Perfect Information: An Elementary Approach to Existence and Approximation by Discrete Games,” *Journal of Economic Theory*, 52: 406–422.
- [20] J.-J. Herings and A. Predtetchinski (2009) “Bargaining with Non-convexities,” Unpublished manuscript.
- [21] T. Kalandrakis (2004a) “Equilibria in Sequential Bargaining Games as Solutions to Systems of Equations,” *Economics Letters*, 84: 407–411.
- [22] T. Kalandrakis (2004b) “A Three-player Dynamic Majoritarian Bargaining Game,” *Journal of Economic Theory*, 116: 294–322.
- [23] T. Kalandrakis (2009) “Minimum Winning Coalitions and Endogenous Status Quo,” *International Journal of Game Theory*, 39: 617–643.
- [24] H. Konishi and D. Ray (2003) “Coalition Formation as a Dynamic Process,” *Journal of Economic Theory*, 1–41.
- [25] V. Krishna and R. Serrano (1996) “Multilateral Bargaining,” *Review of Economic Studies*, 63: 61–80.
- [26] A. Merlo and C. Wilson (1995) “A Stochastic Model of Sequential Bargaining with Complete Information,” *Econometrica*, 63: 371–399.
- [27] A. Nowak (1985) “Existence of Equilibrium Stationary Strategies in Discounted Non-Cooperative Stochastic Games with Uncountable State Space,” *Journal of Optimization Theory and Applications*, 45: 591–602.
- [28] A. Nowak and T. Raghavan (1992) “Existence of Stationary Correlated Equilibria with Symmetric Information for Discounted Stochastic Games,” *Mathematics of Operations Research*, 17: 519–526.
- [29] A. Okada (1996) “A Non-cooperative Coalitional Bargaining Game with Random Proposers,” *Games and Economic Behavior*, 16: 97–108.
- [30] D. Ray (2007) *A Game-Theoretic Perspective on Coalition Formation*, Oxford University Press: New York.

- [31] D. Ray and R. Vohra (1999) “A Theory of Endogenous Coalition Structures,” *Games and Economic Behavior*, 26: 286–336.
- [32] P. Rogers (1969) “Non-zero sum stochastic games,” Operations Research Center Report no. 69-8, University of California, Berkeley.
- [33] R. Serrano (2005) “Fifty Years of the Nash Program: 1953–2003,” *Investigaciones Económicas*, 29: 219–258.
- [34] N. Yannelis (1990) “On the Upper and Lower Semicontinuity of the Aumann Integral,” *Journal of Mathematical Economics*, 19: 373–389.