

Existence of Stationary Bargaining Equilibria*

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

February 15, 2016

Abstract

This paper addresses the question of existence of stationary Markov perfect equilibria in a class of dynamic games that includes many known bargaining models and models of coalition formation. General sufficient conditions for existence of equilibria are currently lacking in a number of interesting environments, e.g., models with non-convexities, consumption lower bounds, or an evolving state variable. The main result establishes existence of equilibrium under compactness and continuity conditions, without the structure of convexity or strict comprehensiveness used in the extant literature. The proof requires a precise selection of voting equilibria following different proposals using a generalization of Fatou's lemma.

1 Introduction

Many applications of game-theoretic analysis are predicated on the existence of equilibrium. In simple models, equilibria may be solved for directly, but in more complex models, a general existence result can indicate structure conducive to equilibria and can guide the search for a solution. In less tractable models, the explicit construction of equilibria may not be possible, and in such cases, a general existence result can serve to underpin characterization results by ensuring they are non-vacuous. This paper establishes existence of stationary Markov perfect equilibria in a class of dynamic games that grows out of the bargaining literature in economics and political science, originating with the seminal work of Rubinstein (1982). In economics, a number of

*This paper supersedes the working paper version, which is titled "Coalitional Bargaining Equilibria." I thank Paulo Borelli for helpful comments.

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 following 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 prior question of existence has been left open:

- legislative bargaining in a multi-dimensional policy space in which stage utilities are non-concave or the set of alternatives is non-convex,
- coalition formation in NTU environments in which strict comprehensiveness fails,
- collective dynamic programming, where policies are chosen each period and affect the value of a discrete state variable that evolves endogenously over time.

Three main applications corresponding to the above environments motivate the general existence result. The first is a legislature that must choose a tax level that determines the set of equilibrium allocations of an economy, with bargaining coming to an end once agreement is reached. Because the correspondence of Walrasian allocations generally has closed but not convex graph, extant results do not deliver existence of equilibrium in the bargaining game. More generally, it may be that stage payoffs are induced by subsequent play of an unmodeled game and violate the usual convexity conditions. The second is coalition formation in public good economies where individuals have property rights over endowments and consumption is non-negative, violating strict comprehensiveness. The third is a group of agents who repeatedly negotiate over a level of public investment (or a level of extraction of a common pool resource), where current choices determine the future capital stock (or amount of resource available) and possibly future preferences over consumption. This paper considers a general model of dynamic bargaining that unifies all of the above examples, and the main result yields stationary Markov perfect equilibria in each of them.

In the general model, one agent proposes an outcome or remains silent, other agents simultaneously respond, and the game then may or may not continue in this fashion. Stage games are parameterized by a countable set of states. Outcomes may belong to an abstract finite set, or they 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. Formally, it is assumed 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 general: 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 strategies determining future states and outcomes over an infinite horizon along the path of play. The equilibrium concept is stationary Markov perfect equilibrium in stage-undominated voting strategies.

A key to the general existence result for such stationary bargaining equilibria is the possibility of mixing on the part of the agents, including not just mixing over proposals, but, importantly, mixed voting as well. Ray and Vohra (1999) and Banks and Duggan (2000, 2006) 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, even for simple finite examples, the analysis of voting strategies must be more nuanced. The necessity of mixed voting strategies also arises in Herings and Houba’s (2013) analysis of a three-player, three-alternative bargaining game with no discounting.¹

¹See the discussion of their Examples 3.2–3.4.

Before discussing mixed voting in more detail, the next example illustrates the need for mixed proposal strategies in a simple, finite model.

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

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

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 could 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, it is shown 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); and each agent accepts her favorite outcome if proposed, accepts her second-favorite if she herself proposed it, rejects her second-favorite if it is proposed by another agent, and rejects her least favorite outcome. \square

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 in some cases, delay *must* occur in equilibrium, highlighting the possible role of non-

convexities in the phenomenon of inefficient delay.²

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 it fails with positive probability, so that agent 1 does not propose x in a no-delay equilibrium. Since agent 1 will not propose y , her worst outcome, in a no-delay equilibrium, it must be that agent 1 proposes z with probability one, and this passes. Note that agent 3 can obtain a payoff of at least .7 as proposer by proposing y , the favorite outcome of agent 2. 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.³ 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 *needed* for existence in the general bargaining framework, we 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 .2.⁴ We first show 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 $\delta v_2 \geq .7$. We examine two cases. In case 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. Since agent 1 will not propose y and $\delta v_2 \geq .7$, it must be that $\delta v_2 \leq (.9)[(.8)\delta v_2 + (.2)(.7)]$, but then $\delta v_2 < .7$, a contradiction. In the remaining case, agent 1 remains silent or proposes z (or mixes between the two), and therefore $\delta v_2 < .7$, a contradiction that proves the claim.

²See also Herings and Houba (2013) for analysis of a three-player, three-alternative bargaining game with no discounting in which delay may occur in equilibrium.

³Technically, 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).

⁴The example is robust, as agent 1 could well propose with small probability.

Next, we show 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. Therefore, x must pass with probability strictly between zero and one if proposed, so agent 2 must use a non-degenerate mixed voting strategy. \square

Thus, mixed voting strategies are needed to obtain equilibrium existence in any class of environments containing these simple finite examples. The approach of this paper is to find a fixed point in the space of potential equilibrium proposal strategies augmented by continuation values and proposer payoffs; we then exploit the special structure of the dynamic bargaining model to back out mixed strategy voting equilibria following different proposals to fulfill the proposers' optimality conditions while generating the targeted continuation values and proposer payoffs. Equilibrium voting strategies cannot be derived directly in the fixed point argument because voting strategies are conditioned on proposals, raising the problem of finding a suitable topology on the space of voting strategies. This difficulty is circumvented by locating a fixed point in a more manageable space, submerging voting strategies via a generalization of Fatou's lemma in the fixed point argument. It is worth noting that existence of stationary Markov 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,⁵ but their results do not yield stationarity. Existence results for stationary Markov perfect equilibria provided in the literature on stochastic games are partial,⁶ and they rely on a critical continuity condition on the transition probability 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 from proposal subgames to voting subgames is discontinuous with respect to the strong topology on the range of the transition probability used in the stochastic games literature.

⁵See also Fudenberg and Levine (1983), Hellwig and Leininger (1987), and Hellwig, Leininger, Reny, and Robson (1990).

⁶For example, Nowak and Raghavan (1992) use correlation, Nowak (1985) allows ϵ -best responses, Duggan (2012) assumes a noisy component of the state, and Barelli and Duggan (2014) use semi-Markov strategies.

The framework 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) along with Eraslan and Merlo’s (2002) majority-rule version, the model of one-dimensional policies with sidepayments 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. In the latter models, bargaining ends once acceptance is reached. Anesi (2006,2010) considers a model of bargaining over a finite set of alternatives with an endogenous status quo, i.e., the alternative chosen in one period becomes the status quo in the next, and this process is repeated ad infinitum. His model is covered by this paper, but because it is finite, equilibrium existence is not an issue. Duggan and Kalandrakis (2012) prove existence of stationary bargaining equilibrium in a general bargaining model, but they add noise to the status quo transition and assume idiosyncratic preference shocks in each period, taking the model outside the framework of this paper. We can, however, incorporate a similar structure by dropping preference shocks and restricting status quo alternatives to a countable set.

On the coalition formation side, this paper 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. 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—a feature that takes it outside the framework of this paper.⁷ Those authors prove existence of equilibrium under the rejector-becomes-proposer protocol for NTU games satisfying strict comprehensiveness, an assumption that allows them to restrict attention to pure voting strategies. Strict comprehensiveness is a weak assumption in private good economies but is easily violated in public good environments with consumption lower bounds,⁸ leading to the kinds of challenges illustrated in Example 2 and creating a role for mixed voting. Modifying the protocol of Ray and Vohra (1999) so that the selection of proposer is exogenous (but depends on the set of players who

⁷See also Chatterjee, Dutta, Ray, and Sengupta (1993), Bloch (1996), and Krishna and Serrano (1996). Serrano (2005) and Ray (2007) provide surveys of this literature.

⁸Suppose 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 geometric terms, the Pareto frontier of the set of utility imputations has a “flat” portion.

have not left the game), the existence result of this paper yields a stationary bargaining equilibrium in their model.⁹

In the remainder of the paper, Section 2 presents the general dynamic bargaining model. Section 3 contains the formal statement and proof of the existence theorem, and Section 4 provides an informal discussion of the proof. Proofs of auxiliary results are located in the appendix.

2 Dynamic Bargaining Framework

The model is given by the list $(S, N, M, X, \{X_C\}_{C \subseteq N}, \{u_i\}_{i \in N}, p, \{\delta_i\}_{i \in N})$. Here, S is a nonempty, countable set of states with the discrete topology; N is a countable set of agents, also with the discrete topology; $M: S \rightrightarrows N$ specifies for each state s a nonempty, finite set $M(s)$ of $m(s)$ active agents; X is a nonempty, compact metric space of outcomes; $X_C: S \rightrightarrows X$ specifies for each coalition $C \subseteq N$ and state s a closed set $X_C(s)$ consisting of outcomes feasible for coalition C in state s ;¹⁰ $u_i: X \times S \times \{0, 1\} \rightarrow \mathfrak{R}$ is a bounded and continuous stage payoff function for agent each i , where $a \in \{0, 1\}$ is a voting outcome described below, and $u_i(x, s, a)$ is the payoff to agent i from outcome x in state s with voting outcome a ; $p: S \times X \times S \times \{0, 1\} \rightarrow [0, 1]$ is a continuous state transition probability function, where $p(s'|x, s, a)$ is the probability that given outcome x in state s with voting outcome a , the state transitions to s' ; and $\delta_i: S \times S \rightarrow [0, 1]$ specifies a discount factor $\delta_i(s, s')$ for each agent i that gives the weight of tomorrow's payoffs relative to today's when we transition from state s to state s' .

Given an initial state s_1 , the dynamic bargaining game among the agents is governed by the following recursively defined protocol. 1) Each period t begins with a state s_t . 2) A proposer $i(s_t) \in M(s_t)$ determined by the state proposes any outcome x in X . 3) The active agents $M(s_t)$ simultaneously decide whether to accept or reject this proposal. 4) If there is a coalition $C \subseteq M(s_t)$ such that $x \in X_C(s_t)$ and all members of C accept x , then the proposal passes, i.e., $a = 1$, and x is the outcome in period t , and

⁹The technique of selecting voting equilibria cannot be used in the general rejector-becomes-proposer model, as convexity of voting equilibrium outcomes is not ensured. I conjecture that existence is regained if public randomization occurs before agents respond to proposals.

¹⁰We allow $C = \emptyset$, in which case the bargaining protocol of the model implies that any outcome in $X_\emptyset(s)$ can be implemented by the proposer without acceptance by any agent; however, this case does not play a special role, as such outcomes could be included in the set of outcomes feasible for the proposer.

payoffs $u_i(x, s_t, 1)$ accrue to the active agents; otherwise, the proposal fails, i.e., $a = 0$, and the active agents receive default payoffs $u_i(x, s_t, 0)$; and all inactive agents receive a stage payoff of zero. Finally, 5) the game transitions to period $t+1$, a new state s_{t+1} is drawn with probability $p(s_{t+1}|x, s_t, a)$, and the protocol is repeated with the following period's payoffs discounted by $\delta_i(s_t, s_{t+1})$. Thus, given sequences $\mathbf{s} = (s_t)_{t=1}^\infty$, $\mathbf{x} = (x_t)_{t=1}^\infty$, and $\mathbf{a} = (a_t)_{t=1}^\infty$ of states, proposals, and voting outcomes, the discounted payoff in period $t \geq 2$ for agent i is

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

and the discounted sum of payoffs is

$$U_i(\mathbf{s}, \mathbf{x}, \mathbf{a}) = u_i(s_1, x_1, a_1) + \sum_{t=2}^{\infty} U_{i,t}(\mathbf{s}, \mathbf{x}, \mathbf{a}).$$

Given the above structure, it is conceivable that every outcome $x \in X$ would be accepted if proposed, although the proposer could prefer that no proposal were accepted. To address this possibility, we assume that the proposer has the option to refrain from proposing, remaining silent and maintaining the status quo instead. To formalize this assumption, we assume there is a default outcome $q \in X$ that is an isolated point, i.e., $\{q\}$ is open in X , and that q is always feasible for all coalitions, i.e., for all s and all C , we have $q \in X_C(s)$. Consistent with this interpretation, we assume that stage payoffs and the state transition are generated by the default when a proposal is rejected: for all i , all s , and all x , we have $u_i(x, s, 0) = u_i(q, s, 1)$ and $p(\cdot|x, s, 0) = p(\cdot|q, s, 1)$.¹¹ This restriction allows payoffs in the case of rejection to depend arbitrarily on the current state, but not on the proposal itself, conforming with bargaining protocols in the literature. This assumption allows us to write, henceforth,

$$\begin{array}{l} u_i(x, s) = u_i(x, s, 1) \\ u_i(q, s) = u_i(x, s, 0) \end{array} \quad \text{and} \quad \begin{array}{l} p(s'|x, s) = p(s'|x, s, 1) \\ p(s'|q, s) = p(s'|x, s, 0), \end{array}$$

¹¹The assumption that q is isolated is without loss of generality. Given a model in which the default $q \in X$ 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 $u'_i(q', s, a) = u_i(q, s, a)$ and $p'(\cdot|q', j, s, a) = p(\cdot|q, j, s, a)$. Thus, the re-defined default q' enters in these functions exactly as the original default and fulfills our assumptions without affecting the strategic structure of the game.

suppressing the outcome of voting from the arguments of stage payoffs and the transition probability. Furthermore, we assume that once an active agent becomes inactive, she remains inactive thereafter. Formally, for all agents i , all natural numbers T , all sequences $(x_0, \dots, x_{T-1}) \in X^T$ of outcomes, and all sequences $(s_0, \dots, s_T) \in S^{T+1}$ of states, if $i \in M(s_0) \setminus M(s_1)$ and $\prod_{t=1}^T p(s_t | x_{t-1}, s_{t-1}) > 0$, then $i \notin M(s_t)$. This is largely without loss of generality, as an agent can remain active yet receive a zero payoff for arbitrarily long (or infinite) sequences of states.¹²

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, we impose the joint restriction on δ_i and p that there exists $\hat{T} < \infty$ such that

$$\sup \left[\prod_{t=1}^{\hat{T}-1} p(s_{t+1} | x_t, s_t) \delta_i(s_t, s_{t+1}) - \prod_{t=1}^{\hat{T}-1} p(s_{t+1} | x_t, s_t) \right] < 0,$$

where the supremum is over sequences of outcomes $(x_1, \dots, x_{\hat{T}-1}) \in X^{\hat{T}-1}$ and states $(s_1, \dots, s_{\hat{T}}) \in S^{\hat{T}}$, such that $\prod_{t=1}^{\hat{T}-1} p(s_{t+1} | x_t, s_t) > 0$. To parse this condition, note that if each discount factor $\delta_i(s, s')$ were equal to one, then the lefthand side would equal zero. Clearly, this is ruled out. The inequality says that over any span of \hat{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 $\bar{\delta} = \sup\{\delta_i(s, s') \mid i \in N, s, s' \in S, \delta_i(s, s') < 1\}$, 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, and obviously the discounting assumption is satisfied if each agent i discounts payoffs over time according to a fixed discount factor $\delta_i \in [0, 1)$. The scope of the above-defined framework is discussed in relation to the literature at the end of the section.

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

¹²Technically, this assumption allows us to restrict attention in the existence proof to a finite number of continuation values in any state, a restriction that is important in the application of Fatou's lemma, Lemma 3.2.

agent i is a mapping $\pi_i: S_i \rightarrow \mathcal{P}(X)$, where $\pi_i(s)(Y)$ denotes the probability that i proposes an outcome belonging to the Borel measurable set Y in state s . A *stationary voting strategy* for i 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* for i 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, we 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 condition: for all i and all s ,

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

where $\alpha(x, s)$ is shorthand for the probability that the proposal of x in state s is accepted, i.e., it is defined by

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

Note that by the assumption that the proposer can impose the default, $\alpha(x, s) = 1$ if $x = q$. Given σ , define agent i '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 *stationary 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

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

and 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 an agent is indifferent, i.e., $U_i(x, s; \sigma) = U_i(q, s; \sigma)$. This definition implicitly imposes the requirement that voting strategies are stage undominated 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). The restriction imposed is somewhat stronger, as it applies even if the proposal is not feasible for any coalition and even if an agent is a “dummy voter,” but that extra restrictiveness is inconsequential. The refinement on voting strategies could be dropped if, instead, voting were specified as sequential, a common approach that does not affect the set of equilibrium outcomes.

The above framework is sufficiently general to encompass the standard Rubinstein (1982) model of two-player, alternating-offer bargaining and subsequent work on legislative bargaining. The alternating offer protocol is obtained by assuming three states: an active state in which player 1 proposes, an active state in which player 2 proposes, and a terminal state; then the transition probability switches between the two active states until an agreement is reached, after which the game moves to the terminal state. The n -player distributive model of Baron and Ferejohn (1989) and spatial models of Banks and Duggan (2000,2006) feature a randomly recognized proposer, and these are obtained by assuming n active states and a terminal state, and by specifying the transition probability so that draws of proposers are independent across time until agreement is reached, after which the terminal state is reached. The models of Merlo and Wilson (1985) and Eraslan and Merlo (2002) are similarly obtained, but with each active state having two components, the identity of the proposer and the size of the pie. The papers in this literature impose convexity conditions on feasible sets and stage payoffs to obtain existence of equilibrium in pure voting strategies, but the framework of this paper does not use that structure. Thus, the model captures non-convex problems, such as a legislature that chooses a tax rate and

a Walrasian equilibrium of an economy; more formally, if $E(t)$ denotes the set of Walrasian equilibria given tax rate t , then we can set $X = \text{graph}(E)$ equal to the graph of the Walrasian equilibrium correspondence to capture legislative bargaining in an economic environment.

Models of coalition formation in NTU games are such that a proposer identifies a coalition C and a vector of payoffs in the feasible set $V(C) \subseteq \mathfrak{R}^C$, where we assume each $V(C)$ is compact. Here, we specify $X_C(s)$ to consist of pairs (C, y) such that $y \in V(C)$. To obtain the protocol of Okada (1996), a state must specify the set of active agents and the current proposer, and the transition is such that: if a proposal is rejected, then the set of active agents remains the same, and a new proposer is drawn from a fixed distribution over the active agents; if a proposal $x = (C, y)$ is accepted, then the coalition C leaves the game with payoffs y , a new proposer is drawn from the smaller set of active agents, and the process is repeated until all agents have left the game. Although Okada (1996) assumes a TU structure, the framework easily extends to the NTU setting to provide a version of Ray and Vohra (1999) with exogenous proposer selection, allowing public good economies in which consumption lower bounds are present and strict comprehensiveness is violated.

The special cases described above all possess a terminal state that is, typically in applications, reached in equilibrium with probability one. We can also obtain collective dynamic programming problems such that in each period: a state is given, a proposal is made and voted on, the outcome determines the distribution of the next period's state, and so on, ad infinitum. As an example of this structure, we obtain Anesi's (2006,2010) model of bargaining over a finite set of alternatives with an endogenous status quo: here, a state corresponds to the current status quo alternative at the beginning of a period, and the transition probability is specified so that if a proposal is accepted, then it becomes the new state at the beginning of the next period; and if a proposal is rejected, then the state remains unchanged. More generally, we can allow the set of alternatives to be a continuum, as long as the set of possible status quo alternatives is a countable subset, and the transition from outcomes to the status quo is continuous; that is, there is a countable set $Q \subseteq X$ such that for all $x \in X$, $\sum_{y \in Q} p(y|x) = 1$, and for all $y \in Q$, the probability $p(y|x)$ that y is the status quo following outcome x is continuous in x . This formulation allows the status quo to approximate the outcome from the previous period, providing a version of Duggan and Kalandrakis (2012) that drops preference shocks while restricting the status

quo to a countable set. These special cases use the state to represent a status quo alternative that remains in place when agreement is not reached, but it can be used more generally to model the state of a discrete system, such as a capital stock level, production capacity, regulatory statutes, or legal codes.

3 Existence of Stationary Bargaining Equilibria

For the statement of the following theorem, we 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 we assume u_i and p are jointly continuous in their arguments. We write $v_{i,s}$ to denote agent i 's continuation value calculated at the beginning of a period in state s , and $v = (v_{i,s})_{i \in N, s \in S} \in \mathfrak{R}^{N \times S}$ for 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 stationary 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 stationary bargaining equilibria, along with upper hemicontinuity of the correspondence of equilibrium continuation values.

Theorem 3.1 *The correspondence $E: \Gamma \rightrightarrows \mathfrak{R}^{N \times S}$ has non-empty, closed values and is upper hemicontinuous.*

The remainder of this section consists of the proof of the theorem. The argument relies on two lemmas. The first, a special case of Theorem 2.4 of Tan and Wu (2002), extends an interval-valued correspondence with closed graph from a compact subset of a metric space to the entire metric space.

Lemma 3.1 *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 $y \in Y$, $\Phi(y) = \phi(y)$.*

The second lemma extends the version of Fatou's lemma on upper hemicontinuity 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 $\Phi: X \times Y \rightrightarrows \mathfrak{R}^k$ have nonempty values and closed graph. Let $\int \Phi(x, y) \mu(dx)$ consist of all integrals of measurable selections from $\Phi(\cdot, y)$. The latter authors prove that the correspondence of integrals, $\int \Phi(x, \cdot) \mu(dx): Y \rightrightarrows \mathfrak{R}^k$, has closed graph. For our arguments, however, we need to allow the probability measure μ to vary, and as a consequence we add the assumption of convex values. Another technical extension, which is needed for the application of this paper, is that we consider integrals of selections ϕ from $\Phi(\cdot, y)$ composed with a continuous function, as in $f(x, \phi(x), y)$, that is linear in its second argument. The proof of the lemma is provided in the appendix.

Lemma 3.2 *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 $f: X \times [0, 1]^k \times Y \rightarrow \mathfrak{R}^n$ 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.

We now formally prove Theorem 1. We 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

$$\underline{u} = \frac{\hat{T}}{1 - \delta} \cdot \inf_{i, x, s} u_i(x, s) \quad \text{and} \quad \bar{u} = \frac{\hat{T}}{1 - \delta} \cdot \sup_{i, x, s} u_i(x, s),$$

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

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

with elements $\theta = (\pi, w, v)$. As usual, $\mathcal{P}(X(s))$ is imbedded in the vector space M of signed measures with the weak* topology (as the topological

dual of the space of bounded, continuous, real-valued functions on X), which is Hausdorff and locally convex. As is well-known, $\mathcal{P}(X(s))$ is compact in the weak* topology. Of course, we imbed $[\underline{u}, \bar{u}]$ in the real line with the Euclidean topology. Then the product topology on $(M^S) \times (\mathfrak{R}^S) \times (\mathfrak{R}^{N \times S})$ makes it a locally convex, Hausdorff topological space, and Θ is a non-empty, compact, convex subset of this space. Finally, let $\Theta^+ = \Theta \times \Gamma$ be Θ augmented by the parameters of the model. Denote a generic element of Θ^+ by $\theta^+ = (\pi, w, v, \gamma)$.

We 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 stationary bargaining equilibrium in the model parameterized by γ ; and conversely, each stationary bargaining equilibrium corresponds to a fixed point; and the correspondence of fixed points has closed graph. We will define F as a product correspondence $F = P \times W \times V$, and 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 function simulates an agent's dynamic payoff, under the assumption that v represents future payoffs, and is continuous in its arguments.

Definition of P : For each state s and each coalition C , define the correspondences

$$\begin{aligned} A_C(s, \theta^+) &= \{x \in X_C(s) \mid \text{for all } i \in C, U_i(x, s, \theta^+) \geq U_i(q, s, \theta^+)\} \\ A_C^\circ(s, \theta^+) &= \{x \in X_C(s) \mid \text{for all } i \in C, U_i(x, s, \theta^+) > U_i(q, s, \theta^+)\}, \end{aligned}$$

and as well define the correspondences

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

and note that $q \in A(s, \theta^+)$. Intuitively, given parameters θ^+ , the set $A^\circ(s, \theta^+)$ consists of the outcomes that would necessarily be accepted if proposed in state s , and $A(s, \theta^+)$ consists of the outcomes that could possibly be accepted if proposed in state s . Since $X(s)$ has the relative topology induced by X , continuity of U_i implies that for all $s \in S$, the correspondence $A(s, \cdot)$ has closed graph in $\Theta^+ \times X(s)$, $A^\circ(s, \cdot)$ has open graph in $\Theta^+ \times X(s)$, and for

all θ^+ , $\text{clos}A^\circ(s, \theta^+) \subseteq A(s, \theta^+)$. Furthermore, the correspondence $A(s, \cdot)$ is actually upper hemi-continuous. Since $A^\circ(s, \cdot)$ has open graph, the correspondence is lower hemi-continuous as a function of θ^+ , and Aliprantis and Border's (2006) Lemma 17.29 implies that the extended-real-valued mapping $\sup\{U_{i(s)}(x, s, \theta^+) \mid x \in A^\circ(s, \theta^+)\}$ is lower semi-continuous. Then, as the pointwise maximum of two lower semi-continuous functions, it follows that

$$f(s, \theta^+) = \max \left\{ \sup\{U_{i(s)}(x, s, \theta^+) \mid x \in A^\circ(s, \theta^+)\}, U_{i(s)}(q, s, \theta^+) \right\}.$$

is lower semi-continuous as a function of θ^+ . Since the proposer can, in equilibrium, obtain any outcome strictly acceptable to some coalition for which it is feasible, and can always obtain the default by choosing the default, we interpret $f(s, \theta^+)$ as the "security value" of the proposer in state s given parameters θ^+ . Now define

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

as the set of possibly acceptable proposals that meet or exceed the proposer's security value. This set is non-empty. Indeed, in case $\sup\{U_{i(s)}(x, s, \theta^+) \mid x \in A^\circ(s, \theta^+)\} \geq U_{i(s)}(q, s, \theta^+)$, then $A^\circ(s, \theta^+) \neq \emptyset$, and 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 $U_{i(s)}(\bar{x}, s, \theta^+) \geq \sup\{U_{i(s)}(x, s, \theta^+) \mid x \in A^\circ(s, \theta^+)\}$, and therefore $\bar{x} \in \hat{P}(s, \theta^+)$. In case $\sup\{U_{i(s)}(x, s, \theta^+) \mid x \in A^\circ(s, \theta^+)\} < U_{i(s)}(q, s, \theta^+)$, then $q \in \hat{P}(s, \theta^+)$. Furthermore, by continuity of U_i and lower semi-continuity of f , $\hat{P}(s, \cdot)$ has closed graph in $\Theta^+ \times X(s)$. Then define $P: \Theta^+ \rightrightarrows (\prod_{s \in S} \mathcal{P}(X(s)))$ by

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

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

Definition of W : Let $\text{supp}(\pi_s)$ denote the support of π_s , i.e., the smallest closed set Y such that $\pi_s(Y) = 1$, 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). By Aliprantis and Border's (2006) Lemma 17.29, the mapping

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

is upper semi-continuous as a function of θ^+ . Define the (possibly empty) set

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

For each state s , since $f(s, \cdot)$ is lower semi-continuous and $g(s, \cdot)$ is upper semi-continuous in θ^+ , the correspondence $\hat{W}(s, \cdot)$ has closed, in fact, compact graph in $\Theta^+ \times [\underline{u}, \bar{u}]$. Since the projection mapping from $\text{graph}(f)$ to Θ^+ is continuous, the set

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

is compact. To see that $\hat{\Theta}(s) \neq \emptyset$, choose any $\theta^+ = (\pi, w, v, \gamma)$ such that π_s puts probability one on a payoff maximizing outcome in $X(s)$ for proposer $i(s)$ in model γ . By Lemma 3.1, we can extend $\hat{W}(s, \cdot)$ from $\hat{\Theta}(s)$ to a correspondence (still denoted $\hat{W}(s, \cdot)$) on Θ^+ that has non-empty, convex values and has closed graph. Then define the correspondence $W: \Theta^+ \rightrightarrows [\underline{u}, \bar{u}]^S$ by

$$W(\theta^+) = \prod_{s \in S} \hat{W}(s, \theta^+),$$

which has non-empty, convex values and has closed graph. \square

Definition of V : Given state $s \in S_i$, each agent j 's payoff depends on the probability that agent i 's proposals pass. The unambiguous cases are that 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. The critical case is when agent i proposes $x \in A(s, \theta^+) \setminus A^\circ(s, \theta^+)$, for then the equilibrium conditions on voting strategies impose no restrictions on the probability that x passes. However, if $x \in \text{supp}(\pi_s)$, then the proposal should, to be consistent, generate the payoff w_s for agent i , and this provides a restriction on voting strategies. Indeed, 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 $U_i(x, s, \theta^+) > U_i(q, s, \theta^+)$, we must have

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

More generally, when $U_i(x, s, \theta^+) \neq U_i(q, s, \theta^+)$, 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, when $U_i(x, s, \theta^+) = U_i(q, s, \theta^+)$,

this is undefined. Next, define the correspondence $\hat{A}: X \times S \times \Theta^+ \rightrightarrows [0, 1]$ by

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

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^+ , the correspondence $\hat{A}(\cdot, s, \theta^+)$ 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 any 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 $\hat{A}(\cdot, s, \theta^+)$. Note that the selection does not necessarily satisfy the conditions for equilibrium in voting subgames: it is possible, e.g., that $\hat{a}(x, s, \theta^+) < 1$ yet $x \in A^\circ(s, \theta^+)$. This discrepancy is repaired after the fixed point argument.

Define $\hat{V}(s, \theta^+)$ 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 $\hat{a}(\cdot, s, \theta^+): X \rightarrow [0, 1]$ from $\hat{A}(\cdot, s, \theta^+)$, we specify that the vector $v' = (v'_{i,s})_{i \in N}$ of continuation values belongs to $\hat{V}(s, \theta^+)$ if and only if for all $j \in M(s)$, we have

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

and for all $j \in N \setminus M(s)$, $v'_{j,s} = 0$. Note that $\hat{V}(s, \theta^+)$ 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 $\{x \in X \mid U_{i(s)}(x, s, \theta^+) = U_{i(s)}(q, s, \theta^+)\}$, which is measurable. Furthermore, since $\hat{A}(\cdot, s, \theta^+)$ is convex-valued, convexity of $\hat{V}(s, \theta^+)$ follows. That $\hat{V}(s, \cdot)$ has closed graph in $\Theta^+ \times [\underline{u}, \bar{u}]^N$ follows from Lemma 3.2. Indeed, to apply that result, let X be the set of outcomes, let $Y = ([\underline{u}, \bar{u}]^S) \times ([\underline{u}, \bar{u}]^{N \times S}) \times \Gamma$, let $k = 1$, and let $\Phi = \hat{A}(\cdot, s, \cdot)$. Note that the countable product of metric spaces is metrizable in the product topology (see Theorem 3.36 of Aliprantis and Border (2006)), so Y is metric. Enumerate the active agents in $M(s)$ as $1, \dots, n$, and 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 $x \in X$, all $y = (w, v, \gamma) \in Y$, and all $a \in [0, 1]$.¹³ Let the correspondence F consist of integrals of f with respect to $\mu = \pi_s$, so that

¹³Note that the definition of $f_i(x, a, y)$ makes use of the simulated dynamic payoff $U_i(x, s, \theta^+)$, but the latter depends only on $y = (w, v, \gamma)$ and not on proposal strategies π .

$\hat{V}(s, \theta^+) = F(y, \mu)$. Then closed graph of $\hat{V}(s, \cdot)$ follows immediately from the lemma. Finally, define $V: \Theta^+ \rightrightarrows [\underline{u}, \bar{u}]^{N \times S}$ by

$$V(\theta^+) = \prod_{s \in S} \hat{V}(s, \theta^+),$$

which, following the above argument, has non-empty, convex values and has closed graph. \square

Fixed points of F : 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, standard continuity arguments imply that the correspondence from parameters γ to the set of fixed points of $F(\cdot, \gamma)$ has closed graph. \square

Claim: For all $(w, v, \gamma) \in [\underline{u}, \bar{u}]^S \times [\underline{u}, \bar{u}]^{N \times S} \times \Gamma$, 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 stationary bargaining equilibrium $\sigma^* = (\pi^*, \alpha^*)$ with continuation values $v = (v_i(s; \sigma^*))_{i \in N, s \in S}$ and proposer payoffs

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

for each $s \in S$.

To prove the claim, let (w, v, γ) be given. We first prove the “only if” direction. To this end, consider any π such that (π, w, v) is a fixed point of $F(\cdot, \gamma)$. For all $i \in N$ and all $s \in S_i$, we have $\pi_s \in \mathcal{P}(\hat{P}(s, \theta^+))$, so that $\text{supp}(\pi_s) \subseteq \hat{P}(s, \theta^+)$, and therefore $f(s, \theta^+) \leq g(s, \theta^+)$. It follows that $w_s \in \hat{W}(s, \theta^+) = [f(s, \theta^+), g(s, \theta^+)]$, and furthermore, for all $x \in \text{supp}(\pi_s)$, we have $U_i(x, s, \theta^+) \geq w_s \geq f(s, \theta^+) \geq U_i(q, s, \theta^+)$. Then each $v_{j,s}$ is determined by a selection $\hat{\alpha}(\cdot, s, \theta^+)$ 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 :

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

Next, we specify voting strategies α^* to satisfy the conditions of equilibrium by considering three cases of proposals, and we then define proposal strategies π^* to complete the specification of the equilibrium.

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 she strictly prefers x to rejection, i.e.,

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

which means that x will pass with probability one if proposed, i.e., $\alpha^*(x, s) = 1$. Note that it is possible the selection $\hat{\alpha}(\cdot, s, \theta^+)$ actually specifies that x fail with positive probability, i.e., $\hat{\alpha}(x, s, \theta^+) < 1$. 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 we consider $x \in A^\circ(s, \theta^+)$, we have $f(s, \theta^+) \geq U_i(x, s, \theta^+)$. But if $x \in \text{supp}(\pi_s)$, then we have

$$f(s, \theta^+) \geq U_i(x, s, \theta^+) \geq w_s \geq f(s, \theta^+),$$

which implies $U_i(x, s, \theta^+) = w_s$. Thus, $U_i(x, s, \theta^+) > U_i(q, s, \theta^+)$ would imply $\hat{\alpha}(x, s, \theta^+) = 1$ by construction of \hat{A} . We conclude that $\hat{\alpha}(x, s, \theta^+) < 1$ is possible only if $U_i(x, s, \theta^+) \leq U_i(q, s, \theta^+)$, and since $x \in \text{supp}(\pi_s)$, 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 we define equilibrium proposal strategies, below, we correct the inconsistency highlighted here by specifying that with probability $1 - \hat{\alpha}(x, s, \theta^+)$, the agent proposes 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 each agent j vote in the obvious way if she has a strict preference: if $U_j(x, s, \theta^+) > U_j(q, s, \theta^+)$, then $\alpha_j^*(x, s) = 1$; and if $U_j(x, s, \theta^+) < U_j(q, s, \theta^+)$, then $\alpha_j^*(x, s) = 0$. There is some coalition C with $x \in A_C(s, \theta^+)$, and for each such coalition C , 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 $\hat{\alpha}(x, s, \theta^+)$. 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, i.e., $\alpha_j^*(x, s) = 0$. 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 $\hat{\alpha}(x, s, \theta^+)$. We thus 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 $A(s, \theta^+)$. We specify voting strategies as in Case 1, so x fails with probability one, i.e., $\alpha^*(x, s) = 0$. It is possible that $\hat{\alpha}(x, s, \theta^+) > 0$, 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 immaterial.

To specify proposal strategies, consider any agent i and state $s \in S_i$. We specify 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 , so $U_i(x, s, \theta^+) = U_i(q, s, \theta^+)$, we require that the proposer place probability $1 - \hat{\alpha}(x, s, \theta^+)$ on q , and otherwise, the agent mixes according to π_s . Formally, define $\pi_i^*(s)$ so that for all measurable $Y \subseteq X \setminus \{q\}$,

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

and

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

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

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 $w_s \geq f(s, \theta^+) \geq U_i(x, s, \theta^+)$, so the expected payoff from proposing x does not exceed w_s . In Case 2, the acceptance probability $\hat{\alpha}(x, s, \theta^+)$ is chosen so that if the inequality $U_i(x, s, \theta^+) > U_i(q, s, \theta^+)$ holds, 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 stationary bargaining equilibrium.

For the “if” direction, consider a stationary 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 claim. Note by optimality of proposal strategies, we have $w_s \geq U_i(q, s; \sigma^*)$ for all $i \in N$ and all $s \in S_i$. Define π by modifying

π^* so that for all i and all $s \in S_i$, the agent proposes $q \in X(s)$ whenever the original proposal strategy dictates a proposal of $x \in X \setminus X(s)$, i.e., $\pi_s(\{q\}) = \pi_i^*(s)(\{q\} \cup (X \setminus X(s)))$ and for all measurable $Y \subseteq X(s) \setminus \{q\}$, $\pi_s(Y) = \pi_i^*(s)(Y)$. To establish that $(\pi, w, v) \in F(\pi, w, v, \gamma)$, we alter α^* in three ways. First, adjust each α_j^* so that the agent accepts q if it is proposed. Second, following any 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., $\alpha^*(x, s)U_i(x, s; \sigma^*) + (1 - \alpha^*(x, s))U_i(q, s; \sigma) = w_s$. Third, following a proposal of x by i such that $U_i(x, s; \sigma^*) < U_i(q, s; \sigma^*)$, adjust voting strategies so that x fails with probability one. Note that 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, and the probability of a proposal x with $U_i(x, s; \sigma^*) < U_i(q, s; \sigma^*)$ is zero, so these modifications do not affect the agents' continuation values. Let $\theta^+ = (\pi, w, v, \gamma)$, and let $\hat{\alpha}(x, s, \theta^+)$ denote the probability that proposal x is accepted in state s , given the above alterations. Then we have $\hat{\alpha}(x, s, \theta^+) \in \hat{A}(x, s, \theta^+)$ for all x and all s , i.e., the acceptance probability $\hat{\alpha}(\cdot, s, \theta^+)$ 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 $F(\cdot, \gamma)$. \square

Finally, the above claim establishes a correspondence between fixed points of $F(\cdot, \gamma)$ and the stationary bargaining equilibria in model γ , immediately delivering existence of equilibria and non-empty values of the correspondence E . Closed graph follows as well, using the facts that $E(\gamma)$ is just the projection of the fixed points of $F(\cdot, \gamma)$ onto $[\underline{u}, \bar{u}]^{N \times S}$ and that F has compact range. Since E has compact range, this implies upper hemicontinuity, completing the proof of the theorem.

4 Discussion of Proof

The proof consists of a fixed point argument for a particular correspondence, where an element of the domain reflects a particular the state of the game, and the correspondence updates the state of the game to reflect the agents' optimality conditions; of course, the construction must be such that a fixed point determines a stationary bargaining equilibrium. An essential component of the state of the game is the agents' proposal strategies, π , but mixed

voting implies the possibility of delay, so that the information in π is not sufficient to evaluate the objective function of the proposer in (1). Continuation values are required for this, and so we add the vector $v = (v_{i,s})_{i \in N, s \in S}$ to the description of the state of the game, where $v_{i,s}$ represents the expected discounted payoff to agent i at the beginning of a period in state s . However, more is needed, for the proposer’s objective function also depends on the acceptance probability $\alpha(x, s)$, and the latter is indeterminate: it is equal to one when x is feasible for a coalition that strictly prefers x to the default; it is equal to zero when x is not feasible for a coalition that weakly prefers x to the default; but it may range between zero and one in the remaining case that indifferent agents are pivotal.

One possible response is to explicitly include the agents’ acceptance strategies in the description of the state of the game, expanding the domain to include the vector $\alpha = (\alpha_i)_{i \in N}$ 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 restrict the variation of $\alpha_i(x)$ as we range over x or would fix a measure on X and require us to ignore sets of measure zero proposals. But the acceptance strategies α_i are endogenous, and there is no way to a priori restrict the variation of an agent’s acceptance strategy as a function of the proposal; and we are not given an exogenous probability measure on proposals that allows us to ignore “small” subsets of proposals. Thus, these maneuvers are 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. To develop this, first note that we do not require the individual agents’ acceptance strategies to evaluate the proposer’s objective function, but only the overall probability of acceptance, $\alpha(x, s)$. The key insight of the argument is that although these acceptance probabilities are conditioned on the proposal x and live in a complex, infinite-dimensional space, they can be reduced to a dimensionality equal to the cardinality of the state space. Specifically, suppose we know the expected payoff to the proposer in each state, denoted w_s . Given vectors $w = (w_s)_{s \in S}$ and v of proposer payoffs of continuation values,

the expected payoff to the proposer $i = i(s)$ if x is accepted in state s is

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

and the expected payoff to the proposer if x is rejected in state s is

$$U_i(q, s; v) = u_i(q, s) + \sum_{s' \in S} p(s'|q, s) \delta_i(s, s') v_{i, s'}.$$

In the most relevant case, where $U_i(x, s; v) \geq w_s \geq U_i(q, s; v)$, with at least one inequality strict, the expected payoff to the proposer must satisfy

$$w_s = \alpha(x, s) U_i(x, s; v) + (1 - \alpha(x, s)) U_i(q, s; v),$$

and thus we can back out the acceptance probability as

$$\alpha(x, s) = \frac{w_s - U_i(q, s; v)}{U_i(x, s; v) - U_i(q, s; v)}. \quad (2)$$

Other cases can be dealt with by bounding this probability to be between zero and one, the only difficulty arising when $U_i(x, s; v) = U_i(q, s; v)$.

In the latter case, the proposer is indifferent between the proposal being accepted or rejected, and the proposer's expected payoff does not imply any restriction on the acceptance probability. These arguments suggest a correspondence $\hat{A}(\cdot, s; w, v): X \rightarrow [0, 1]$ of acceptance probabilities such that $\hat{A}(x, s; w, v)$ is given by (2) whenever $U_{i(s)}(x, s; w, v) \neq U_{i(s)}(q, s; v)$, while imposing bounds to make the probability well-defined, and whenever indifference holds with $x \neq q$, we set $\hat{A}(x, s; w, v) = [0, 1]$, reflecting the fact that acceptance probabilities are unrestricted.¹⁴ Every measurable selection $\hat{\alpha}(\cdot, s): X \rightarrow [0, 1]$ from this correspondence yields an expected payoff of w_s to the proposer in state s , a fact that will be used to back out acceptance probabilities after a fixed point (π^*, v^*, w^*) is determined. Importantly, the correspondence has closed graph and convex values. Note that we are not assured that the selection of acceptance probabilities is consistent with the equilibrium conditions in voting subgames, a matter that is addressed after the fixed point argument.

Thus, we describe the state of the game by a triple (π, w, v) , and we define a product correspondence $F = P \times W \times V$ from the space of such

¹⁴For technical reasons, the construction in the proof assumes that when $x = q$, the proposal is accepted with probability one; this is without loss of generality.

triples into itself. To update proposal strategies, we define the set $A(v)$ of the outcomes that are feasible for a coalition and weakly preferred to rejection by members of the coalition; and we define the set $A^\circ(v)$ of the outcomes that are feasible for a coalition and strictly preferred to rejection by members of the coalition. Intuitively, outcomes in $A^\circ(v)$ must be accepted with probability one if proposed, while a proposal is accepted with positive probability only if it belongs to $A(v)$. These sets imply upper and lower bounds on the proposer's expected payoff,

$$\max_{y \in A(v)} U_{i(s)}(y, s; v) \quad \text{and} \quad \max \left\{ \sup_{y \in A^\circ(v)} U_{i(s)}(y, s; v), U_{i(s)}(q, s; v) \right\},$$

respectively, where the lower bound reflects the proposer's option of choosing the default. We refer to the lower bound as the proposer's "security value," and we define the P component of the correspondence so that $P(\pi, w, v)$ consists of all mixed proposal strategy profiles π' such that in each state s , the mixture π'_s places probability one on elements of $A(v)$ that give the proposer an expected payoff at least as great as her security value. This component identifies, in a sense, the set of candidate equilibrium mixed proposal strategies: we can without loss of generality restrict proposals to belong to $A(v)$, and in equilibrium, the proposer will never propose an alternative that offers less than her security value. Straightforward arguments show that the correspondence P has closed graph and nonempty, convex values.

In equilibrium, the proposer in state s must be indifferent between all proposals in the support of her proposal strategy π_s , so all proposals in the support of π_s must yield the same payoff to the proposer. Given a triple (π, w, v) , the W component of the correspondence then identifies the set of possible proposer payoffs in each state. Specifically, in state s , the proposer must obtain at least her security value, giving the lowest possible proposer payoff, but identifying the highest possible proposer payoff is somewhat more complex. It is possible that in equilibrium, the proposer mixes over outcomes that offer distinct dynamic payoffs, e.g., $U_i(x, s; v) > U_i(y, s; v) \geq U_i(q, s; v)$, yet is indifferent due to variation in the probability of rejection; that is, it is possible that x and y are both optimal because x is accepted with lower probability. Let $f(s; v)$ denote the security value of the proposer, and let

$$g(s; \pi, v) = \min\{U_{i(s)}(x, s; v) \mid x \in \text{supp}(\pi_s)\}$$

denote the minimum dynamic payoff achieved over the support of π_s . Then the set of proposer payoffs that can be achieved by using acceptance probabilities to "level off" dynamic utilities in the support of π_s is just the

interval $[f(s, v), g(s; \pi, v)]$. It may be that this interval is empty for some triples (π, w, v) , but can define the W component by

$$W(\pi, w, v) = [f(s, v), g(s; \pi, v)]$$

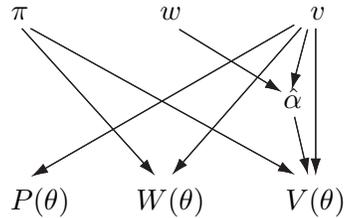
when the interval is nonempty, extending it to the full domain in an upper hemicontinuous way by Lemma 3.1.

The updating of continuation values requires great care, because these will depend not only on proposal strategies, but also on acceptance probabilities, which are not explicitly given in the domain of the correspondence. Instead, we infer them from w and v using the correspondence $\hat{A}(\cdot, s; w, v)$ described above. Specifically, to update the continuation values in state s , we enumerate the active agents as $1, \dots, n$, and we collect the vectors $(v'_{1,s}, \dots, v'_{n,s})$ such that there exists a measurable selection $\hat{\alpha}(\cdot, s): X \rightarrow [0, 1]$ from $\hat{A}(\cdot, s; w, v)$ satisfying

$$v'_{j,s} = \int_X [\hat{\alpha}(x, s)U_j(x, s; v) + (1 - \hat{\alpha}(x, s)U_j(q, s; v)]\pi_s(dx)$$

for all $j = 1, \dots, n$. The set $\hat{V}(s; \pi, v)$ of such vectors is nonempty and convex, and we use a generalization of Fatou's lemma, Lemma 3.2, to establish closed graph of the correspondence $(\pi, w, v) \mapsto \hat{V}(s; \pi, v)$. Finally, we define $V(\pi, w, v) = \prod_{s \in S} \hat{V}(s; \pi, v)$ as the product of these sets, and V also has closed graph and nonempty, convex values.

This completes the specification of the correspondence F , pictured below, where $\theta = (\pi, w, v)$. Each component is specified so that it has nonempty, convex values and closed graph, and F inherits these properties. By Glicksberg's fixed point theorem, F admits a fixed point, say (π^*, w^*, v^*) , and there-with a selection $\hat{\alpha}(\cdot, s)$ from $\hat{A}(\cdot, s; w^*, v^*)$ that generates the continuation values v^* , but the construction does not ensure that the selection



$\hat{\alpha}(\cdot, s)$ will be consistent with the equilibrium conditions in voting subgames, so minor adjustments are required to fulfill the conditions for equilibrium. In particular, it is possible that a proposed outcome x passes with probability less than one although it is feasible for a coalition of agents who strictly prefer x to rejection, and so it should 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, the problem can only arise 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

A Example 1

In this section of the appendix, it is shown that the unique stationary bargaining equilibrium discounted continuation value in Example 1 is $.7$. Consider any stationary bargaining equilibrium. We first show 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, we show 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.

B Proof of Lemma 3.2

To prove Lemma 3.2, 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 , it follows that $\text{supp}(\nu^m) \subseteq \text{graph}(\Phi(\cdot, y^m))$. 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

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

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 $(x, a) \in (X \times [0, 1]) \setminus T$, $P_{\alpha|\mathcal{T}}(\cdot, x, a)$ 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

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

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

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

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 (2012) “Dynamics and Stability of Constitutions, Coalitions, and Clubs,” *American Economic Review*, 102: 1446–76.
- [2] C. Aliprantis and K. Border (2006) *Infinite Dimensional Analysis*, 3rd ed., Springer: New York.
- [3] V. Anesi (2006) “Committees with Farsighted Voters: A New Interpretation of Stable Sets,” *Social Choice and Welfare*, 27: 595–610.
- [4] V. Anesi (2010) “Noncooperative Foundations of Stable Sets in Voting Games,” *Games and Economic Behavior*, 70: 488–493.
- [5] R. Aumann (1976) “An Elementary Proof that Integration Preserves Uppersemicontinuity,” *Journal of Mathematical Economics*, 3: 15–18.
- [6] J. Banks and J. Duggan (2000) “A Bargaining Model of Collective Choice,” *American Political Science Review*, 94: 73–88.
- [7] J. Banks and J. Duggan (2006) “A General Bargaining Model of Legislative Policy-making,” *Quarterly Journal of Political Science*, 1: 49–85.
- [8] D. Baron (1996) “A dynamic theory of collective goods programs,” *American Political Science Review*, 90: 316–330.
- [9] P. Billingsley (1968) *Convergence of Probability Measures*, Wiley: New York.
- [10] P. Borelli and J. Duggan (2014) “A Note on Semi-Markov Perfect Equilibria in Discounted Stochastic Games,” *Journal of Economic Theory*, 151: 596–604.

- [11] T. Börgers (1989) “Perfect Equilibrium Histories of Finite- and Infinite-horizon Games,” *Journal of Economic Theory*, 47: 218–227.
- [12] T. Börgers (1991) “Upper Hemicontinuity of the Correspondence of Subgame-perfect Equilibrium Outcomes,” *Journal of Mathematical Economics*, 20: 89–106.
- [13] J. Duggan (2012) “Noisy Stochastic Games,” *Econometrica*, 80: 2017–2045.
- [14] J. Duggan and T. Kalandrakis (2012) “Dynamic Legislative Policy Making,” *Journal of Economic Theory*, 147: 1653–1688.
- [15] R. Dudley (2002) *Real Analysis and Probability*, Cambridge University Press: New York.
- [16] H. Eraslan and A. Merlo (2002) “Majority Rule in a Stochastic Model of Bargaining,” *Journal of Economic Theory*, 103: 31–78.
- [17] D. Fudenberg and D. Levine (1983) “Subgame-perfect Equilibria of Finite and Infinite-horizon Games,” *Journal of Economic Theory*, 31: 251–268.
- [18] C. Harris (1985a) “Existence and Characterization of the Perfect Equilibrium in Games of Perfect Information,” *Econometrica*, 53: 613–628.
- [19] C. Harris (1985b) “A Characterization of the Perfect Equilibria of Infinite-horizon Games,” *Journal of Economic Theory*, 37: 99–125.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] J.-J. Herings and H. Houba (2013) “The Condorcet Paradox Revisited,” Unpublished manuscript.

- [24] J.-J. Herings and A. Predtetchinski (2009) “Bargaining with Non-convexities,” Unpublished manuscript.
- [25] T. Kalandrakis (2004a) “Equilibria in Sequential Bargaining Games as Solutions to Systems of Equations,” *Economics Letters*, 84: 407–411.
- [26] T. Kalandrakis (2004b) “A Three-player Dynamic Majoritarian Bargaining Game,” *Journal of Economic Theory*, 116: 294–322.
- [27] T. Kalandrakis (2009) “Minimum Winning Coalitions and Endogenous Status Quo,” *International Journal of Game Theory*, 39: 617–643.
- [28] H. Konishi and D. Ray (2003) “Coalition Formation as a Dynamic Process,” *Journal of Economic Theory*, 1–41.
- [29] V. Krishna and R. Serrano (1996) “Multilateral Bargaining,” *Review of Economic Studies*, 63: 61–80.
- [30] A. Merlo and C. Wilson (1995) “A Stochastic Model of Sequential Bargaining with Complete Information,” *Econometrica*, 63: 371–399.
- [31] 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.
- [32] 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.
- [33] A. Okada (1996) “A Non-cooperative Coalitional Bargaining Game with Random Proposers,” *Games and Economic Behavior*, 16: 97–108.
- [34] D. Ray (2007) *A Game-Theoretic Perspective on Coalition Formation*, Oxford University Press: New York.
- [35] D. Ray and R. Vohra (1999) “A Theory of Endogenous Coalition Structures,” *Games and Economic Behavior*, 26: 286–336.
- [36] P. Rogers (1969) “Non-zero sum stochastic games,” Operations Research Center Report no. 69-8, University of California, Berkeley.
- [37] A. Rubinstein (1982) “Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 50: 97–109.

- [38] R. Serrano (2005) “Fifty Years of the Nash Program: 1953–2003,” *Investigaciones Económicas*, 29: 219–258.
- [39] M. Sobel (1971) “Noncooperative Stochastic Games,” *Annals of Mathematical Statistics*, 42: 1930–1935.
- [40] K. Tan and Z. Wu (2002) “An Extension Theorem and Duals of Gale-Mas-Collel’s and Shafer-Sonnenschein’s Theorems,” D. O’Regan and R. Agarwal (eds) *Set Valued Mappings with Applications in Nonlinear Analysis*, Series in Mathematical Analysis and Applications, Taylor & Francis: New York.
- [41] N. Yannelis (1990) “On the Upper and Lower Semicontinuity of the Aumann Integral,” *Journal of Mathematical Economics*, 19: 373–389.