A Newton Collocation Method for Solving Dynamic Bargaining Games

John Duggan and Tasos Kalandrakis

Working Paper No. 60
August 2009

W. ALLEN
WALLIS
Institute of
POLITICAL
ECONOMY

UNIVERSITY OF
ROCHESTER

A Newton Collocation Method for Solving Dynamic Bargaining Games

John Duggan* Tasos Kalandrakis†

August 18, 2009

Abstract

We develop and implement a collocation method to solve for an equilibrium in the dynamic legislative bargaining game of Duggan and Kalandrakis (2008). We formulate the collocation equations in a quasi-discrete version of the model, and we show that the collocation equations are locally Lipchitz continuous and directionally differentiable. In numerical experiments, we successfully implement a globally convergent variant of Broyden's method on a preconditioned version of the collocation equations, and the method economizes on computation cost by more than 50% compared to the value iteration method. We rely on a continuity property of the equilibrium set to obtain increasingly precise approximations of solutions to the continuum model. We showcase these techniques with an illustration of the dynamic core convergence theorem of Duggan and Kalandrakis (2008) in a nine-player, two-dimensional model with negative quadratic preferences.

*Department of Political Science and Department of Economics, University of Rochester

†Department of Political Science, University of Rochester

1 Introduction

To examine strategic incentives in ongoing collective choice problems, we consider a class of dynamic bargaining games in which a sequence of proposals and votes generates policy outcomes over time. The status quo policy evolves endogenously, with today's policy determining tomorrow's status quo and forward-looking players anticipating the future consequences of their decisions. Specifically, we take up the legislative bargaining framework of Duggan and Kalandrakis (2008). In the general model of that paper, players possess some (perhaps small) degree of uncertainty about the future state of the game: there is noise in the transition from today's outcome to tomorrow's status quo, and players' preferences are subject to transitory preference shocks each period. The analytic derivation of equilibrium solutions is prohibitive in this framework, owing to the complexity of strategies (which are conditioned on the realized status quo and preference shocks), but the ability to compute equilibria is nevertheless essential for the development of our understanding of this type of dynamic social interaction. Computed equilibria can provide possibility theorems, and computational results, when used systematically, can also provide the germs of general theorems. Finally, in empirical work, the structural estimation of model parameters or the calibration of a model to observed data, as the case may be, both rely on the ability to compute equilibria.

We propose a method to compute equilibria in the dynamic bargaining framework. We formulate stationary equilibria as solutions to a system of functional equations, the unknowns of which are essentially the future expected utilities (or "dynamic utilities") of the players. We solve these functional equations using a collocation method, a method for solving functional equations that belongs in the general family of projection methods (Judd (1998), Chapter 11). In particular, we posit a finite-dimensional representation of the unknown equilibrium dynamic utility functions as linear combinations of Chebyshev polynomials, and we seek coefficients for these representations that solve the equilibrium functional equations exactly at a finite number of points (that coincide with the roots of the Chebyshev polynomials). This version of the collocation method is theoretically justified by the Chebyshev Interpolation Theorem (Judd (1998); Rivlin (1990)) and by the theoretical results of Duggan and Kalandrakis (2008), which guarantee smoothness properties and a priori bounds for the equilibrium expected utilities and their derivatives: we can approximate the unknown functions to an a priori specified, but arbitrary, level of precision by adding higher degree Chebyshev polynomials to represent these functions.

Having transformed the original equilibrium functional equations to a system of a finite number of equations in an equal number of unknown coefficients, we face the practical problem of solving this system of equations. Our computational analysis then focuses on the "quasi-discrete" model, in which the uncertainty in the model is continuous (and so the state space of the bargaining game is infinite), but given the status quo, only a finite number of alternatives may be proposed. This allows us to exactly solve the optimization problem of a proposer by an exhaustive grid search. Since the collocation equations are intermediated by the players' best response behavior, and because they involve the integration over future uncertainty in the model, those equations are nonlinear and potentially ill-behaved.

We establish, however, that they are locally Lipschitz continuous and directionally differentiable, i.e., they belong in a class of nonlinear equations for which various generalizations of Newton's method and the convergence properties thereof have been studied in recent years (Ip and Kyparisis (1992); Martinez and Qi (1995); Pang (1990, 1991); Qi (1993); Xu and Chang (1997)). Moreover, we show that we can obtain an equilibrium of the model with a continuum of feasible alternatives by taking the limit of equilibria of a sequence of quasi-discrete approximations.

We implement a version of Broyden's method (Broyden (1965)) to solve for an equilibrium. Our version of Broyden's method "preconditions" the collocation equations to obtain faster rates of convergence, and we show that this algorithm outperforms two alternatives, a pseudo-Newton method and simple value iteration, in a series of experiments. Finally, we apply the method to illustrate the core convergence result of Duggan and Kalandrakis (2008) by specifying configurations of ideal points approaching the canonical setting in which one alternative belongs to the majority core; for each configuration, we compute the invariant distribution (representing the long run distribution over alternatives) generated by a stationary equilibrium; and we show that these invariant distributions pile mass near the limiting core point. Interestingly, convergence appears to be faster when the players are more patient.

Many situations of interest in political economy possess the structure of a dynamic bargaining game: some player proposes an alternative, that proposal is considered by other players and possibly agreed to, and the game possibly continues into future periods. One branch of this literature considers environments in which bargaining ends once agreement is reached, with play continuing into the future only if a proposal is rejected.[1] We focus, instead, on a class of models in which bargaining continues ad infinitum, whether there is agreement in a period or not. Baron (1996) analyzes the one-dimensional version of the model with single-peaked utilities, Kalandrakis (2004, 2007) studies the canonical divide-the-dollar environment, Cho (2005) considers policy making in a stage game that emulates aspects of parliamentary government, and Battaglini and Coate (2007) characterize stationary equilibria in a model of public good provision and taxation with identical legislators and a stock of public goods that evolves over time. With general stage payoffs and feasible set of alternatives, Duggan and Kalandrakis (2008) assume that stage payoffs and the transition to next period's status quo are subject to (arbitrarily small) shocks, adding uncertainty about the future state of the game. In that paper, we establish existence and a number of desirable technical properties of stationary equilibria; we also examine the ergodic properties of equilibria, and we provide a core convergence result for long run equilibrium policy outcomes as the noise in the model goes to zero and the model becomes close to admitting a core alternative.[2]

Little work has been done on computation of equilibrium in this class of bargaining games. Baron and Herron (2003) give a numerical calculation of equilibrium in a three-

[1]See Rubinstein (1982), Binmore (1987), Baron and Ferejohn (1989), Banks and Duggan (2000, 2006), and others.

[2]At a further distance from our paper is work on finite-state dynamic voting games, such as Acemoglu, Egorov, and Sonin (2008) and Diermeier and Fong (2008).

player, finite-horizon version of the model, and Penn (2009) provides numerical illustrations of her model. Closest to the current paper is the work of Duggan, Kalandrakis, and Manjunath (2008), who consider a special case of the model of Duggan and Kalandrakis (2008) to examine the effect of the presidential veto in a US-like political system. But the approach of the current paper differs from the former in several respects. First, Duggan, Kalandrakis, and Manjunath (2008) use function approximation instead of function interpolation, so that in their case equilibrium is not obtained as a solution to a system of collocation equations. Second, they work with a continuous proposal space for the legislators and use continuous optimization methods to solve for legislators' optimal proposals, whereas we implement our techniques in a model where the space of possible proposals at each status quo is finite. Third, those authors use a version of value iteration to obtain an equilibrium. In this paper, we consider value iteration as one possible solution method, but we implement and provide theoretical justification for the use of Newton and Newton-like methods. These differences, especially the last two, amount to significant gains on computation time.

In what follows, we first present the model, define our equilibrium concept, and provide background results on the model in Section 2. In Section 3, we formally describe the collocation method and define the collocation equations. In Section 4, we provide theoretical results for the quasi-discrete model, establishing smoothness properties of the collocation equations and our approximation result for the continuum model. In Section 5, we describe in detail our implementation of Broyden's method for solving the collocation equations. In Section 6, we provide the results of our numerical experiments and our illustration of core convergence. Section 7 concludes, and the appendix contains the proof of our smoothness result.

2 Dynamic Bargaining Framework

In this section, we first present the bargaining model, and we then define our equilibrium concept, a refinement of stationary Markov perfect equilibrium, and review the known foundational results for the model.

2.1 Bargaining Model

We consider a finite set N of players, $i = 1, \ldots, n$, who determine policy over an infinite horizon, with periods indexed $t = 1, 2, \ldots$. Interaction proceeds as follows in each period. A status quo policy $q \in \mathbb{R}^d$ and a vector $\theta = (\theta_1, \ldots, \theta_n) \in \mathbb{R}^{nd}$ of preference parameters are realized and publicly observed. A player $i \in N$ is drawn at random, with probabilities $p_1, \ldots, p_n$, and proposes any policy $y \in X \cup \{q\}$, where $X \subseteq \mathbb{R}^d$ represents the set of feasible policies. The players vote simultaneously to accept y or reject it in favor of the status quo q. The proposal passes if a coalition $C \in \mathscr{D}$ of players vote to accept, and it fails otherwise, where $\mathscr{D}$ is a nonempty collection $\mathscr{D} \subseteq 2^N \setminus \{\emptyset\}$ of *decisive coalitions* satisfying only the minimal monotonicity requirement that if one coalition is decisive, and we add players to that coalition, then the larger coalition is also decisive. Formally, we assume that if $C \in \mathscr{D}$

and $C \subseteq C' \subseteq N$, then $C' \in \mathscr{D}$. The policy outcome for period t, denoted x_t, is y if the proposal passes and is q otherwise. Each player j receives stage utility $u_j(x_t)+\theta_j \cdot x_t$, where $\theta_j \in \mathbb{R}^d$ is a utility shock for the player. Finally, the status quo q' for period $t+1$ is drawn from the density $g(\cdot|x_t)$, a new vector $\theta' = (\theta'_1, \ldots, \theta'_n)$ of preference shocks is drawn from the density $f(\cdot)$, and the above procedure is repeated in period $t+1$. Payoffs in the dynamic game are given by the expected discounted sum of stage utilities, as is standard, and we denote the discount factor of player i by $\delta_i \in [0, 1)$.

We impose a number of regularity conditions on the policy space. We assume that the set of feasible policies, X, is cut out by a finite number of functions $h_\ell \colon \mathbb{R}^d \to \mathbb{R}$, indexed by $\ell \in K$. We partition K into inequality constraints, K^{in}, and equality constraints, K^{eq}, and we assume that

$$X = \{x \in \mathbb{R}^d : h_\ell(x) \geq 0, \ell \in K^{in}, h_\ell(x) = 0, \ell \in K^{eq}\}.$$

We further assume that X is compact, and that h_ℓ is r-times continuously differentiable for all $\ell \in K$, where $r \geq \max\{2, d\}$.[3] For technical reasons, we impose the weak condition that for all $x \in X$, $\{Dh_\ell(x) : \ell \in \overline{K}(x)\}$ is linearly independent, where $\overline{K}(x)$ is the subset of $\ell \in K$, including equality constraints, such that $h_\ell(x) = 0$. These assumptions allow us to capture quite general manifolds. An important special case is the *quasi-discrete model,* in which the policy space $X \subseteq \mathbb{R}^d$ is finite.[4] Even if the space of interest is a continuum, this special class of model plays an important role in our computational analysis, where we make use of limits of equilibria of quasi-discrete models.

We assume $u_i \colon \mathbb{R}^d \to \mathbb{R}$ is r-times continuously differentiable. The presence of preference shocks in the model captures uncertainty about the players' future policy preferences. For example, in the important special case of negative quadratic stage utility, where $u_i(x) = -||x_j - x_t||^2$ and x_j is player i's unperturbed ideal point, the preference shock θ_i is equivalent to a perturbation of the player's ideal point x_i. We assume that the vector $\theta = (\theta_1, \ldots, \theta_n)$ is distributed according to a density f with support contained in the set $\Theta = [\underline{\theta}, \overline{\theta}]^{nd} \subseteq \mathbb{R}^{nd}$, and we further assume a compact set $\tilde{X} = [\underline{x}, \overline{x}]^d \subseteq \mathbb{R}^d$ with $X \subseteq \tilde{X}$ and a bound b_f such that for all $i \in N$, all $\theta \in \Theta$, and all $x \in \tilde{X}$, we have $|u_i(x) + \theta_i \cdot x| f(\theta) \leq b_f$. The noise on the status quo captures the idea that players are uncertain about the way policy decisions today will be implemented in the future. We assume that the density $g \colon \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$, with values $g(q|x)$, is jointly measurable in (q, x), and that for all x, the support of the density $g(\cdot|x)$ lies in $\tilde{X}$. We do not assume that the support of $g(\cdot|x)$ lies in X, though of course we allow it. Furthermore, we assume a bound b_g such that for all q, we have: $g(q|x)$ is r-times continuously differentiable in x; if $r < \infty$, then all derivatives of order $1, \ldots, r$ are bounded in norm by b_f, and the r-th derivative of $g(q|x)$ with respect to x is Lipschitz continuous with modulus b_f; and if $r = \infty$, then derivatives of all orders $1, 2 \ldots$ are bounded in norm by b_f.

A strategy in the game consists of two components, one giving the proposals of a player when recognized to propose and the other giving the votes of the player after a proposal

[3] Of course, we allow $r = \infty$.

[4] We obtain a finite X using suitably "oscillating" equality constraints. We can, for example, isolate a grid on $[0, 1]^d$ by using trigonometric functions, as in $\{x \in \mathbb{R}^d : \sin(2\pi x_i \alpha) = 0, i = 1, \ldots, d\}$, for appropriate α.

is made. While these choices can in principle depend arbitrarily on histories, we seek subgame perfect equilibria in which players use stationary Markov strategies, which we denote $\sigma_i = (\pi_i, \alpha_i)$. Our main focus will be on pure strategies.[5] Thus, player i's proposal strategy is a measurable mapping $\pi_i\colon \mathbb{R}^d \times \Theta \to \mathbb{R}^d$, where $\pi_i(q,\theta)$ is the policy proposed by i given status quo q and utility shocks θ; and player i's voting strategy is a measurable mapping $\alpha_i\colon \mathbb{R}^d \times \mathbb{R}^d \times \Theta \to \{0,1\}$, where $\alpha_i(y,q,\theta) = 1$ if i accepts proposal y given status quo q and utility shocks θ and $\alpha_i(y,q,\theta) = 0$ if i rejects. We let $\sigma = (\sigma_1,\ldots,\sigma_n)$ denote a stationary strategy profile. We may equivalently represent voting strategies by the set of feasible proposals a player would vote to accept. We define this *acceptance set* for i as $A_i(q,\theta;\sigma) = \{y \in X \cup \{q\} : \alpha_i(y,q,\theta_i) = 1\}$. Letting C denote a coalition of players, we then define

$$A_C(q,\theta;\sigma) \;=\; \bigcap_{i\in C} A_i(q,\theta;\sigma) \quad \text{and} \quad A(q,\theta;\sigma) \;=\; \bigcup_{C\in\mathscr{D}} A_C(q,\theta;\sigma)$$

as the *coalitional acceptance set* for C and the *collective acceptance set,* respectively. The latter consists of all policies that would receive the votes of all members of at least one decisive coalition and would, therefore, pass if proposed.

Given strategy profile σ, we define $v_i(x;\sigma)$ as player i's discounted continuation value at the beginning of period $t+1$ from policy outcome x in period t. We then define i's "policy-specific" dynamic payoff as

$$U_i(x;\sigma) = u_i(x) + \delta_i v_i(x;\sigma). \tag{1}$$

Then the discounted payoff to player i from implementing policy x in the current period given preference shock θ_i is $U_i(x;\sigma) + \theta_i \cdot x$. We focus on voting strategies that are "deferential," i.e., players vote to accept when indifferent between a proposed policy and the status quo, which allows us to then consider only no-delay equilibria, meaning no player ever proposes a policy that is rejected. (In lieu of that, the player can just as well propose the status quo.) Our measurability assumptions on strategies imply that continuation values are also measurable, and therefore they satisfy

$$v_i(x;\sigma) = \int_q \int_\theta \sum_j p_j[U_i(\pi_j(q,\theta);\sigma) + \theta_i \cdot \pi_j(q,\theta)] f(\theta) g(q|x) d\theta dq \tag{2}$$

for all policies x.[6] Note that we can restrict the domain of U_i to the compact set $\tilde{X} \subset \mathbb{R}^d$, as players are restricted to propose in $X \cup \{q\} \subseteq \tilde{X}$ in any period with status quo q, and the distribution of the status quo has support restricted to $\tilde{X}$.

[5] Duggan and Kalandrakis's (2008) Theorem 2 establishes that this is without loss of generality, as any equilibrium in stationary mixed strategies is essentially equivalent to some equilibrium in pure strategies. See Subsection 2.2 for further explanation.

[6] Note that continuation values v_i are "ex ante," in the sense that they are calculated by integrating over q and θ. The dynamic utilities U_i differ from those of Duggan and Kalandrakis (2008) by subtracting out the current period's preference shock. Also note that we do not normalize dynamic payoffs by $(1-\delta_i)$ as we do in our earlier paper.

2.2 Stationary Bargaining Equilibrium

With this formalism established, we can now define a class of stationary Markov perfect equilibria of special interest. Intuitively, we require that players always propose optimally and that they always vote in their best interest. It is well-known that the latter requirement is unrestrictive in simultaneous voting games, however, as arbitrary outcomes can be supported by Nash equilibria in which no voter is pivotal. To address this difficulty, we follow the standard approach of refining the set of Nash equilibria in voting subgames by requiring that players delete votes that are dominated in the stage game. Thus, we say a strategy profile σ is a *stationary bargaining equilibrium* if the following conditions hold:

- for every status quo q, every shock θ, and every player i, $\pi_i(q, \theta)$ solves

$$\begin{aligned} \max_y \, & U_i(y; \sigma) + \theta_i \cdot y \\ \text{s.t.}\, & y \in A(q, \theta; \sigma), \end{aligned} \tag{3}$$

- for every status quo q, every shock θ, every proposal y, and every player i,

$$\alpha_i(y, q, \theta_i) = \begin{cases} 1 & \text{if } U_i(y; \sigma) + \theta_i \cdot y \geq U_i(q; \sigma) + \theta_i \cdot q \\ 0 & \text{else.} \end{cases} \tag{4}$$

Thus, as required by subgame perfection, proposers choose optimally after all histories and the votes of players are, furthermore, consistent with the usual dominance criterion and deferential voting.

Duggan and Kalandrakis (2008) provide foundational results on stationary bargaining equilibria. Chief among those results is the following theorem, which establishes existence and a number of desirable regularity properties of equilibria.[7]

Theorem 1 *There exists a stationary bargaining equilibrium, σ, possessing the following properties.*

1. *Continuation values are differentiable: for every player i, $v_i(x; \sigma)$ is r-times continuously differentiable as a function of x.*

2. *Proposals are almost always strictly best: for every status quo q, almost all shocks θ, every player i, and every $y \in A(q, \theta; \sigma)$ distinct from the proposal $\pi_i(q, \theta)$, we have $U_i(\pi_i(q, \theta); \sigma) + \theta_i \cdot \pi_i(q, \theta) > U_i(y; \sigma) + \theta_i \cdot y$.*

3. *Proposal strategies are almost always continuously differentiable: for every status quo q, almost all shocks θ, and every player i such that $\pi_i(q, \theta) \neq q$, $\pi_i(q, \theta)$ is continuously differentiable in an open set around (q, θ).*

[7] Duggan and Kalandrakis (2008) use the term "pure stationary legislative equilibrium" for the concept we consider here. They state only continuity of equilibrium proposal strategies, rather than the condition of continuous differentiability stated in part 3. See the working paper version, Duggan and Kalandrakis (2007), for the statement and proof of the stronger result we give here.

4. *Binding voters, if any, are almost never redundant: for every status quo q, almost all shocks θ, and every player i, if $\pi_i(q,\theta) \neq q$ and there exists j such that $U_j(\pi_i(q,\theta);\sigma) + \theta_j \cdot \pi_i(q,\theta) = U_j(q;\sigma) + \theta_j \cdot q$, then*

$$\{\ell \in N : U_\ell(\pi_i(q,\theta);\sigma) + \theta_\ell \cdot \pi_i(q,\theta) \geq U_\ell(q;\sigma) + \theta_\ell \cdot q\} \setminus \{j\} \notin \mathscr{D}.$$

Part 1 of Theorem 1 implies that dynamic utilities, $U_i(x;\sigma) = u_i(x) + \delta_i v_i(x;\sigma)$, are also r-times differentiable, allowing us in principle to employ first order conditions to characterize optimal proposals. Optimal proposals are essentially strict, by part 2, and equilibrium proposal strategies are continuously differentiable almost everywhere, by part 3, permitting, in principle, the application of calculus techniques in computing comparative statics. Finally, part 4 of the theorem informs us that, generically, if a voter is indifferent between a proposal and the status quo, then the player is pivotal, in the sense that the remaining players willing to vote for the proposal are not decisive, i.e., the coalition of players who accept the proposal is minimally decisive. While Theorem 1 holds in the general model, its implications for the quasi-discrete model, in which the set X of alternatives is finite, take a simple form. There, dynamic utilities are almost always injective: for every status quo q, almost all shocks θ, every player i, and all $x, y \in X \cup \{q\}$, we have $U_i(x;\sigma) + \theta_i \cdot x \neq U_i(y;\sigma) + \theta_i \cdot y$. Thus, the conditions in parts 2 and 4 are trivially satisfied. Finally, part 3 can be strengthened: equilibrium proposal strategies are such that for every q, almost all θ, and all i with $\pi_i(q,\theta) \neq q$, π_i is in fact constant in an open set around (q,θ).

Theorem 2 of Duggan and Kalandrakis (2008) considers the possibility that proposers mix over optimal proposals and that voters mix when indifferent between a proposal and the status quo. The result establishes that when we broaden our definition of equilibrium in this way, we do not introduce new equilibrium behavior in any meaningful sense: every mixed strategy equilibrium is equivalent (up to a measure-zero set of status quos and preference shocks) to a stationary bargaining equilibrium as defined above. Moreover, *every* stationary bargaining equilibrium satisfies the properties in parts 1–4 of Theorem 1. This allows us to focus on pure strategies without loss of generality and increases the scope for computation of equilibrium. In the earlier paper, we also show that the correspondence of stationary bargaining equilibria has closed graph, a desirable continuity property that facilitates our analysis of quasi-discrete approximations of the continuum model in Subsection 4.2; and we provide general conditions under which equilibrium strategies induce an invariant distribution over the policy space that represents the long run policy outcomes of the system.

Finally, Duggan and Kalandrakis (2008) prove a core convergence theorem for models "close" to the canonical social choice model, in which players have quadratic stage utilities, the voting rule is strong (if a coalition is not decisive, then its complement is), there is a unique core policy, and the player located at the core has positive probability of proposing. Theorem 6 of that paper establishes that in this case, the equilibrium invariant distributions collapse to the point mass on the limiting core policy. We take up this result in Section 6.3, where we demonstrate the numerical methods developed in this paper and depict convergence to the core in a two-dimensional, majority-rule version of the bargaining model.

3 The Collocation Method

In this section, we develop a collocation method to solve for an equilibrium of the bargaining game. In particular, in Subsection 3.1, we formulate the problem of finding an equilibrium as the solution of a system of functional equations. In Subsection 3.2, we describe how we approximate these functional equations with a finite-dimensional system of equations, the collocation equations. Before we attempt to solve the collocation equations, we must be able to numerically evaluate them, and we discuss issues related to the numerical evaluation of the collocation equations in Subsection 3.3.

3.1 Equilibrium Functional Equations

We start with the observation that an equilibrium is fully characterized by the corresponding policy-specific dynamic payoffs $U_1, \ldots, U_n$ of the n players. Indeed, given $U = (U_1, \ldots, U_n)$, define

$$A(q, \theta; U) \;=\; \bigcup_{C \in \mathscr{D}} \big[\bigcap_{i \in C} A_i(q, \theta; U_i) \big],$$

where

$$A_i(q, \theta; U_i) \quad = \quad \{y \in X \cup \{q\} : U_i(y) + \theta_i \cdot y \geq U_i(q) + \theta_i \cdot q\} \tag{5}$$

gives the acceptance set of player i after eliminating stage-dominated voting strategies given dynamic payoffs U_i. Furthermore, define the policy $\pi_i(q, \theta; U)$ to be the player's optimal proposal given dynamic payoff U_i and voting behavior in (5), i.e., it solves

$$\begin{aligned} &\max_y U_i(y) + \theta_i \cdot y \\ &\text{s.t.}\, y \in A(q, \theta; U). \end{aligned} \tag{6}$$

By an application of Theorem I.3.1 in Mas-Colell (1985), the above optimization problem has a unique solution for almost all shocks θ_i, pinning down the optimal proposals $\pi_i(q, \theta; U)$ almost everywhere. Then we can express equilibrium dynamic payoffs as solutions to the functional equation

$$U_i(x) \quad = \quad u_i(x) + \delta_i \int_q \int_\theta \sum_{h \in N} p_h \big[U_i(\pi_h(q, \theta; U)) + \theta_i \cdot \pi_h(q, \theta; U) \big] f(\theta) g(q|x) d\theta dq, \tag{7}$$

and we can focus our search for an equilibrium on computing equilibrium dynamic payoff functions $U = (U_1, \ldots, U_n)$ solving the functional equations (7).

Note that since both the stage utility functions u_i and, by property 1 of Theorem 1, the continuation values v_i are r-times continuously differentiable with derivatives that are uniformly bounded in $\tilde{X}$, we conclude that any solution U to the functional equations (7) must belong in the space of r-times continuously differentiable functions $C^r(\tilde{X}, \mathbb{R}^n)$. Furthermore, Lemma 5 of Duggan and Kalandrakis (2008) establishes that the solutions to (7) must lie in an a priori specified compact space with derivatives satisfying a uniform bound across $\tilde{X}$.

3.2 Collocation Equations

The numerical solution of the functional equations (7) on a computer must necessarily involve a representation of the infinite-dimensional objects U in finite dimensions. A standard approach is to proceed by choosing a finite-dimensional subspace of the function space of the candidate solutions U, and then restricting the search for approximate solutions to this subspace. Our choice of basis is a finite set of m Chebyshev polynomials,

$$\{T_1, \ldots, T_m\}.$$

Specifically, we specify a number m_ℓ of the univariate Chebyshev polynomials of degree 0 through $m_{\ell-1}$ for each dimension $\ell = 1, \ldots, d$, and then we obtain the basis $\{T_1, \ldots, T_m\}$ of $m = \prod_{\ell=1}^{d} m_\ell$ polynomials using tensor products of these univariate polynomials. Once the basis $\{T_1, \ldots, T_m\}$ is fixed, we seek solutions for the expected payoff functions U_i that take the form

$$U(x; c_i) = \sum_{j=1}^{m} c_{i,j} T_j(x), \tag{8}$$

where $c_i = (c_{i,1,}, \ldots, c_{i,m}) \in \mathbb{R}^m$ is a vector of *collocation coefficients* corresponding to player i. We write $c = (c_1, \ldots, c_n) \in \mathbb{R}^{nm}$ for a vector specifying the coefficients of all players.

The choice of the Chebyshev polynomial basis is appealing for a number of reasons, including the fact, noted above, that any solutions U to (7) are continuously differentiable with bounded derivatives. Nevertheless, it is unlikely that the actual solutions reside in the subspace spanned by this basis for any finite m. Thus, instead of satisfying equations (7) for all $x \in \tilde{X}$, the collocation method ensures that these equations are satisfied finitely many times, specifically at a finite number of judiciously chosen points in the domain of U_i. Thus, we choose m *collocation nodes*

$$\{\tilde{x}_1, \ldots, \tilde{x}_m\} \subset \tilde{X},$$

and we seek to find collocation coefficients $c = (c_1, \ldots, c_n) \in \mathbb{R}^{nm}$ for the n players so that for every player $i = 1, \ldots, n$, the following collocation equations are satisfied at each of the m collocation nodes $\tilde{x}_k$, $k = 1, \ldots, m$:

$$U(\tilde{x}_k; c_i) = u_i(\tilde{x}_k) + \delta_i \int_q \int_\theta \sum_{h \in N} p_h \big[U(\pi_h(q, \theta; c); c_i) + \theta_i \cdot \pi_h(q, \theta; c) \big] f(\theta) g(q|\tilde{x}_k) d\theta dq. \tag{9}$$

The function $\pi_h(q, \theta; c)$ in (9) is identical to the solution $\pi_h(q, \theta; U)$ of the optimization problem in (6) when $U = (U(\cdot; c_1), \ldots, U(\cdot; c_n))$. Given our choice of the Chebyshev polynomial basis, there is an elegant theory that dictates the location of the collocation nodes $\tilde{x}_k$ at the roots of the Chebyshev polynomials.[8] Since any solution U belongs in $C^r(\tilde{X}, \mathbb{R}^n)$, and since we have an a priori established bound on the derivatives of U, the Chebyshev

[8] We obtain these nodes by tensor products of the roots of the univariate bases. Judd (1998) refers to this version of the collocation method in which the collocation nodes $\tilde{x}_k$ coincide with the roots of the polynomial basis as *orthogonal collocation*.

Interpolation Theorem guarantees that we can approximate the function U up to arbitrary precision, using the combination of the Chebyshev polynomial basis and the roots of the corresponding Chebyshev polynomials as collocation nodes (Judd (1998)), by increasing the degree of approximation and the corresponding number of collocation nodes.

We have thus reduced the problem of solving the functional equations (7) for equilibrium expected payoffs U to solving the nm collocation equations in (9) for the nm collocation coefficients $c \in \mathbb{R}^{nm}$. More compactly, we seek a solution to the equations $F(c) = 0$, where $F : \mathbb{R}^{nm} \to \mathbb{R}^{nm}$ is a function defined from the collocation equations (9) as

$$F_{i,k}(c) = U(\tilde{x}_k; c_i) - \big[u_i(\tilde{x}_k) + \delta_i \int_q \int_\theta \sum_{h \in N} p_h \big[U(\pi_h(q, \theta; c); c_i) + \theta_i \cdot \pi_h(q, \theta; c)\big] f(\theta) g(q|\tilde{x}_k) d\theta dq\big], \tag{10}$$

and where the index i, k corresponds to the k-th collocation node, $\tilde{x}_k$, and player i. In the next subsection, we turn to the practical question of evaluating this collocation function.

3.3 Collocation Function Evaluation

The evaluation of the collocation function F requires us to tackle two computational issues. First, we must be able to evaluate the integrals with respect to the status quo q and the preference shocks θ. The integral with respect to the status quo q is d-dimensional. Thus, when the dimension of the policy space d is of small or moderate size, we can extend unidimensional quadrature techniques to perform this integration. In particular, in the applications we consider, we assume that each coordinate of the status quo is drawn independently, so that the density $g(q|x)$ is a product of densities. Thus, we can use Gaussian quadrature along each dimension with weight function given by the density of the coordinate of the status quo that corresponds to this dimension. The required d-dimensional nodes and weights are easily obtained from the unidimensional ones using tensor products (Judd (1998), Miranda and Fackler (2002)). In practice, for each collocation node $\tilde{x}_k$, we specify a total of α quadrature nodes $q_{k,j} \in Q_k$ and corresponding weights $\omega_{k,j}$, $j = 1, \ldots, \alpha$, using $g(q|\tilde{x}_k)$ as the weight function, and then we compute

$$\int_q \Phi(q, c) g(q|\tilde{x}_k) dq \quad \approx \quad \sum_{q_{k,j} \in Q_k} \Phi(q_{k,j}, c)\omega_{k,j},$$

where $\Phi(q, c) = \int_\theta \sum_{h \in N} p_h \big[U(\pi_h(q, \theta; c); c_i) + \theta_i \cdot \pi_h(q, \theta; c)\big] f(\theta) d\theta$.

Integration with respect to the preference shocks θ is more challenging, as the associated integral is nd-dimensional. The Guassian quadrature approach we described above would be impractical in this case, as the required number of quadrature nodes becomes prohibitive. Instead, we switch to a quasi-Monte Carlo integration method using a Sobol sequence of β quasi-random numbers (Press et al. (1992)), $\theta^\ell \in \Theta$, $\ell = 1, \ldots, \beta$, so that for every quadrature node $q_{k,j}$, we compute

$$\int_\theta \phi(\theta, q_{k,j}, c) d\theta \quad \approx \quad \beta^{-1} \sum_{\ell=1}^{\beta} \phi(\theta^\ell, q_{k,j}, c),$$

where $\phi(\theta, q, c) = \sum_{h \in N} p_h \big[U(\pi_h(q, \theta; c); c_i) + \theta_i \cdot \pi_h(q, \theta; c)\big] f(\theta)$.

In addition to the above integrations, the second major numerical issue involved in evaluating the collocation function F is the computation of the optimal proposals $\pi_h(q, \theta; c)$. This entails solving the optimization problem of player h for each collocation node $\tilde{x}_k$, each quadrature node $q_{k,j}$, and each θ^ℓ in the Sobol sequence, i.e., we must solve a total of $n \times m \times \alpha \times \beta$ optimization problems in order to evaluate the collocation function F once. Note that for arbitrary values of the collocation coefficients c, the players' expected utility functions need not be concave, nor is there any guarantee that these functions would be concave in equilibrium. Even if these functions were concave, the feasible set of proposals available to player h is not convex, as is obvious from (5), since it is the union of sets of proposals acceptable to each coalition. Thus, these optimization problems are quite challenging and entail the possibility of significant numerical error in the evaluation of F if, for example, a local maximum is computed instead of a global maximum. When the space of policy proposals X is a continuum, Duggan, Kalandrakis, and Manjunath (2008) employ a *Nelder-Mead* maximization algorithm to compute $\pi_h(q, \theta; c)$ after initially approximating the solution via a grid search to safeguard against the possibility of locating a local, as opposed to global, maximizer. In this paper, we take an alternative approach by assuming that the set of feasible proposals X is finite, i.e., we work in the quasi-discrete model, where we can compute optimal proposals without error by straightforward exhaustive search. Furthermore, the computation times reported in Section 6.2 suggest we can perform the players' maximization relatively efficiently, even when the feasible set X comprises a large number of points.

In the next section, we show that our use of the quasi-discrete model is justified on at least two more grounds. First, the collocation function in the quasi-discrete model is sufficiently smooth to allow us to use Newton-like methods in order to compute equilibria. Second, we show that we can recover equilibria of the continuous model by computing equilibria for a sequence of quasi-discrete models.

4 The Quasi-Discrete Model

In this section, we specialize the collocation method to the quasi-discrete model. In particular, we assume that the set of feasible proposals is a finite set, $|X| < \infty$, and we index a typical element as x_p. As discussed in the previous subsection, the main advantage of this formulation is that we are able to accurately evaluate the optimization problem of proposers, thus providing a realistic setting for the numerical evaluation of the collocation equations. It is key for the tractability of this problem, however, that the collocation equations be reasonably smooth. This precondition is no mere formality, for the definition of the collocation equations in (10) involves integration over optimal policy choices, $\pi_h(q, \theta; c)$, which are not differentiable in c. Nevertheless, we establish a sufficient level of smoothness for the applicability of Newton-like or inexact Newton methods: in particular, the collocation functions are locally Lipschitz and directionally differentiable. We then show that the solution of the quasi-discrete model allows us to compute equilibria of the continuous model as the limit of equilibria of a sequence of quasi-discrete approximations.

4.1 Smoothness of the Collocation Equations

Before proceeding to the result of this subsection, we provide some elementary definitions and facts used in the proof. We say that a function $g : \mathbb{R}^n \to \mathbb{R}^n$ is *locally Lipschitz continuous* if for every $x \in \mathbb{R}^n$, there exist $\epsilon > 0$ and a constant $M \in \mathbb{R}$ such that for all $y, z \in B_\epsilon(x)$,

$$||g(y) - g(z)|| \quad \leq \quad M||y - z||.$$

A locally Lipschitz function g is differentiable almost everywhere (by Rademacher's theorem). The *directional derivative* of g in direction s is

$$g'(x|s) \quad = \quad \lim_{\alpha \to 0^+} \frac{g(x + \alpha s) - g(x)}{\alpha},$$

provided this limit is well-defined. The smoothness properties of the collocation function rely on the properties of functions that can be represented as a continuous splicing of a finite number of continuously differentiable functions. A function $g : \mathbb{R}^n \to \mathbb{R}^n$ is *piecewise smooth* if it is continuous and there exists a finite collection of continuously differentiable functions $\{g_1, \ldots, g_m\}$ such that each g_j is defined on an open domain and for all $x \in \mathbb{R}^n$, there is a $j = 1, \ldots, m$ such that $g(x) = g_j(x)$.[9] It is known that piecewise smooth functions are locally Lipschitz and have directional derivatives (Kuntz and Scholtes (1994)).

The following theorem establishes that the collocation function F is sufficiently smooth in c in order to allow us to pursue Newton-like methods for the solution of nonlinear systems of equations. If further properties hold in a neighborhood of a solution to the equations, then convergence will also be fast.[10]

Theorem 2 *Assume X is finite and f is C^1. The collocation function $F : \mathbb{R}^{nm} \to \mathbb{R}^{nm}$ is locally Lipschitz continuous and directionally differentiable.*

The proof, located in the appendix, establishes that the collocation equations are integrals of piecewise smooth functions and that the properties of these piecewise smooth functions carry over to F. Specifically, given any status quo q, we break the integral

$$\int_\theta \sum_{h \in N} p_h[U(\pi_h(q, \theta; c); c_i) + \theta_i \cdot \pi_h(q, \theta; c)]f(\theta)d\theta \tag{11}$$

into a finite number of integrals over polyhedral subsets of preference shocks. Each subset $\Theta(q, y, h, A_{-h}; c)$ corresponds to the set of preference shocks $\theta = (\theta_1, \ldots, \theta_n)$ such that player h's optimal proposal is y and the acceptance sets of the other players are given by

[9] Here, we use the generalized definition of piecewise smoothness suggested by Kuntz and Scholtes (1995).

[10] If c^* solves $F(c) = 0$ and F has continuous directional derivatives in an open set around c^* with non-singular Jacobian, J^*, at c^*, it then follows from Theorem 4.1 of Ip and Kyparisis (1992) that Broyden's method (see Subsection 5.1) converges super-linearly to a solution in a neighborhood of (c^*, J^*).

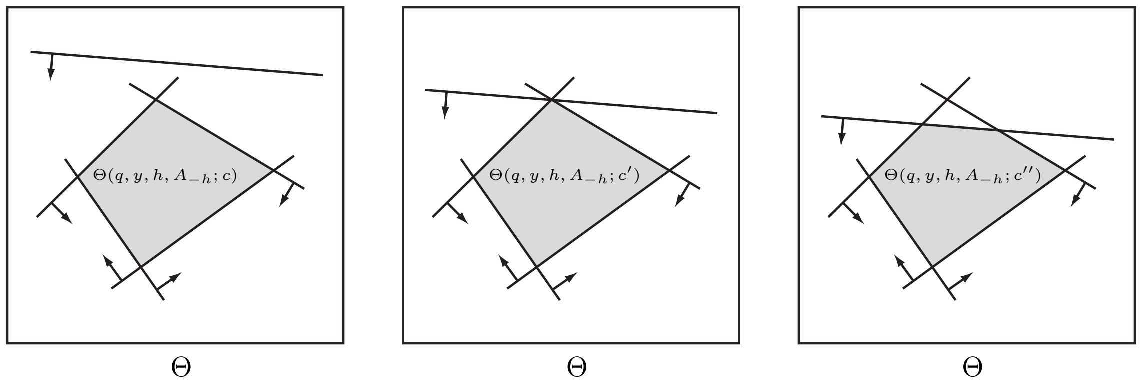

Figure 1: Non-differentiability in collocation coefficients.

A_{-h} when the status quo is q. As such, each "cell" is defined by a finite number of linear inequalities in which the collocation coefficients c enter the righthand side. The integral

$$\int_{\Theta(q,y,h,A_{-h};c)} [U(y; c_i) + \theta_i \cdot y] f(\theta) d\theta \tag{12}$$

over a particular cell is clearly continuous, but non-differentiabilities may arise because as we vary c, different sets of inequalities may become binding. In Figure 1, we depict a change from c to c' to c'', supposing for simplicity that this variation only affects the upper most linear constraint. This change has a smooth effect on (12) until c', where the upper most constraint becomes binding, and the effect is smooth thereafter.

While our decomposition of (11) suppresses the dependence of optimal proposals on the collocation coefficients (through the term $\pi_h(q, \theta; c)$), we cannot avoid non-differentiabilities that may be inherent in the structure of equilibrium. Nevertheless, we show that the integral (12) over any cell is piecewise smooth, so that (11), rewritten as the sum

$$\sum_h p_h \sum_{A_{-h}} \sum_{y \in X \cup \{q\}} \int_{\Theta(q,y,h,A_{-h};c)} [U(y; c_i) + \theta_i \cdot y] f(\theta) d\theta$$

over all possible proposers h, acceptance sets of other players A_{-h}, and possible proposals y, is also piecewise smooth. Piecewise smoothness of (11) implies that for a given status quo q, it is locally Lipschitz in c, but in order to integrate this term over q, we need some bound on its behavior across status quos. The final hurdle in the proof is to construct, for every c, a local Lipschitz constant for (11) that is uniform over q, allowing us to conclude that the integral of (11) over status quos q, and therefore the collocation equations F, is locally Lipschitz and directionally differentiable.

4.2 Quasi-Discrete Approximation

In this subsection, we show that while the quasi-discrete model affords computational tractability, it also provides a means for computation of equilibrium in the continuum model. In particular, we can recover an equilibrium of the continuum model via a sequence of computed equilibria for quasi-discrete models. Given a model with a continuum X of feasible policies, we consider an algorithm for computing stationary bargaining equilibria by means of an increasing sequence of finite grids on the policy space. To be concrete, let $\{X^\ell\}$ be a sequence of finite approximations converging to X in the Hausdorff metric. For each ℓ, define a corresponding "quasi-discrete" model that is identical to the original model except for the fact that feasible proposals (but not the status quo) are now constrained to lie in X^ℓ. The quasi-discrete model is a special case of our bargaining model, for as mentioned above, we can obtain the finite set X^ℓ of feasible policies by appropriately specifying equality constraints. Therefore, Theorem 1 yields at least one stationary bargaining equilibrium, with policy-specific dynamic payoff $U^\ell = (U_1^\ell, \ldots, U_n^\ell)$, in each quasi-discrete model.

We now establish that the sequence $\{U^\ell\}$ necessarily admits a convergent subsequence, and that the limit of any such subsequence corresponds to a stationary bargaining equilibrium of the continuum model. The notion of convergence we use is the topology of C^r-uniform convergence on compacta, a fairly strong topology (which entails a correspondingly strong convergence result).[11]

Theorem 3 *Given a model with set X of alternatives, let $\{X^\ell\}$ be a sequence of quasi-discrete models $X^\ell \subseteq \tilde{X}$ converging to X in the Hausdorff metric. Then in each quasi-discrete model X^ℓ, there exists a stationary bargaining equilibrium with policy-specific dynamic payoffs $U^\ell = (U_1^\ell, \ldots, U_n^\ell)$; and for every sequence $\{U^\ell\}$ of equilibrium dynamic payoffs of quasi-discrete models, there is an accumulation point U of $\{U^\ell\}$ and a stationary bargaining equilibrium of the model with set X alternatives with dynamic payoffs U.*

Proof For each quasi-discrete model X^ℓ, existence of a stationary bargaining equilibrium σ^ℓ follows directly from Theorem 2, and we let U^ℓ be the policy-specific dynamic payoff generated by σ^ℓ as in (1) and (2). Now consider any such sequence $\{U^\ell\}$ of equilibrium dynamic payoffs corresponding to the sequence $\{X^\ell\}$ of quasi-discrete models. For each m, let v^ℓ be as in (1) and (2). Letting b_h denote the Lebesgue measure of $\tilde{X}$, Duggan and Kalandrakis (2008) define a subset $\mathcal{V} \subseteq C^r(\mathbb{R}^d, \mathbb{R}^n)$ that consists of $v\colon \mathbb{R}^d \to \mathbb{R}^n$ satisfying the following: (i) if $r < \infty$, then the derivatives of v of order $0, 1, 2, \ldots$ are bounded in norm by $\sqrt{n} b_f b_g b_h$, and the rth derivative of v is Lipschitz continuous with modulus $\sqrt{n} b_f b_g b_h$; and (ii) if $r = \infty$, then the derivatives of v of all orders are bounded in norm by $\sqrt{n} b_f b_g b_h$. Their Lemma 5 establishes that $\mathcal{V}$ is nonempty and compact in the topology of C^r-uniform convergence on compacta, and it shows that the equilibrium continuation values v^ℓ belong to $\mathcal{V}$ for all m. Thus, there is a convergent subsequence, still indexed by m for simplicity, such that $v^\ell \to v$ with limit $v \in \mathcal{V}$. Theorem 3 of Duggan and Kalandrakis (2008) establishes closed graph of the stationary bargaining equilibrium correspondence, and it follows that there is a stationary bargaining equilibrium, say σ, with continuation value $v(\cdot; \sigma) = v$.

[11] To describe this topology, let $\hat{r}$ be a natural number and $Y \subseteq \mathbb{R}^d$, and define the norm $||f||_{\hat{r},Y}$ on $C^{\hat{r}}(\mathbb{R}^d, \mathbb{R}^n)$ as $\sup\{||\partial f(x)|| : x \in Y\}$, where ∂f is the $\hat{r}$th derivative of f. Then a sequence $\{f^m\}$ of functions converges to f in $C^r(\mathbb{R}^d, \mathbb{R}^n)$ if and only if for every $\hat{r} = 0, 1, \ldots, r$ and every compact set $Y \subseteq \mathbb{R}^d$, we have $||f^m - f||_{\hat{r},Y} \to 0$. We say $f^m \to f$ in $C^\infty(\mathbb{R}^d, \mathbb{R}^n)$ if and only if it converges in $C^r(\mathbb{R}^d, \mathbb{R}^n)$ for all $r = 0, 1, 2, \ldots$

Defining U by $U_i(x;\sigma) = u_i(x) + \delta_i v_i(x)$, we have

$$U_i^\ell \;=\; u_i + \delta_i v_i^\ell \;\rightarrow\; u_i + \delta_i v_i \;=\; U_i,$$

as required. ∎

We have framed Theorem 3 in relatively simple terms for computational purposes, fixing stage utilities, discount factors, and other parameters as the policy space becomes finer. Because the upper hemicontinuity result of Duggan and Kalandrakis (2008) is general in these respects, however, it is straightforward to extend the proof of the above result to allow these other parameters to vary.

5 Solving the Collocation Equations

The collocation equations $F(c) = 0$ are non-linear in the unknown collocation coefficients c, and hence a natural strategy to solve these equations is to apply a variant of Newton's method. Due to the smoothness properties of the collocation function established in Theorem 2, several generalizations of Newton's method are in principle applicable in our problem. In Subsection 5.1, we discuss one such method, Broyden's method, that offers several advantages compared to other alternatives. In Subsection 5.2, we discuss implementation issues and present a transformation of the collocation equations that significantly improves the performance of the resulting algorithm. In Subsection 5.3, we briefly discuss two additional algorithms for the solution of the collocation equations.

5.1 Broyden's Method

Recall that in Newton's method we start with an initial candidate c^0 for the solution and then at the $(\tau+1)$-th iteration, we obtain a new candidate solution according to the formula

$$c^{\tau+1} \;=\; c^\tau - [F'(c^\tau)]^{-1} F(c^\tau),$$

where $F'(c^\tau)$ is the Jacobian of F evaluated at the current iterate c^τ. Recently, a number of authors have studied modifications of Newton's method when F is not everywhere differentiable but satisfies weaker smoothness properties such as local Lipschitz continuity and directional differentiability, as is the case in our problem. One class of these methods directly generalizes Newton's method by using a generalized Jacobian in lieu of $F'(c^\tau)$ when F is not differentiable at c^τ (e.g., Pang (1990), Qi and Sun (1993)). A second alternative falls within the broad class of inexact or quasi-Newton methods (e.g., Ip and Kyparisis (1992), Martinez and Qi (1995), Qi (1997)). Primarily motivated by the observation that the evaluation of the Jacobian is typically very costly (even when it exists), these methods operate by providing an initial estimate of the Jacobian and then ensuring an inexact (but easy to compute) update to the Jacobian with each iteration. These methods behave quite well as long as the sequence of the updated Jacobians stays within a certain distance from the Jacobian at the solution. Any method that circumvents the need to compute the Jacobian is particularly appealing in our problem as the Jacobian of the collocation equations is very costly to compute analytically or numerically. We choose to work with a particularly robust version in the class of inexact Newton methods, namely Broyden's method (Broyden (1965)).

Broyden's method requires an initial guess of the solution c^0, just as Newton's method does. In addition, at the beginning of the algorithm we must also supply an initial guess, B^0, of the Jacobian at the solution. At the τ-th iteration, we obtain a new update of the unknown collocation coefficients as in Newton's method,

$$c^{\tau+1} \quad = \quad c^{\tau} - [B^{\tau}]^{-1} F(c^{\tau}), \tag{13}$$

and a new update of the Jacobian according to the formula

$$\begin{aligned} B^{\tau+1} \quad &= \quad B^{\tau} + \frac{(F(c^{\tau+1}) - F(c^{\tau}) - B^{\tau}(c^{\tau+1} - c^{\tau}))(c^{\tau+1} - c^{\tau})^T}{(c^{\tau+1} - c^{\tau})^T (c^{\tau+1} - c^{\tau})} \\ &= \quad B^{\tau} + \frac{F(c^{\tau+1})(c^{\tau+1} - c^{\tau})^T}{(c^{\tau+1} - c^{\tau})^T (c^{\tau+1} - c^{\tau})}, \end{aligned} \tag{14}$$

where $(c^{\tau+1} - c^{\tau})^T$ denotes the transpose of $c^{\tau+1} - c^{\tau}$. The conditions for fast convergence of Broyden's method, described following Theorem 2, are analogous to those needed for convergence of Newton's method, but in addition to requiring that the initial guess c^0 must be sufficiently close to the solution, Broyden's method also requires that the initial approximation to the Jacobian, B^0, is sufficiently close to the Jacobian at the solution. We address the non-trivial problem of supplying an accurate initial approximation B^0 for our implementation of Broyden's method in the next subsection.

In practice, both Newton and quasi-Newton algorithms are modified to prevent oscillatory behavior or divergence of the iteration sequence. A focal point of intervention in these algorithms is a modification of the Newton step in cases when that step does not lead to a decrease in the residual, $||F(c^{\tau+1})|| \leq ||F(c^{\tau})||$. If the default (quasi-) Newton step does not produce a sufficient decrease in the residual norm, then the step is adjusted by performing a line search along the originally suggested direction until a decrease is achieved. A popular globalization strategy of this form is the Armijo rule (Armijo (1966)), which we use in our implementation. In the particular version we use, the required line search along the direction suggested by (13) is performed using optimization techniques on a parabolic approximation of the residual function. Details of the Armijo rule and the particular implementation can be found in Kelley (2003).

5.2 Implementation and Preconditioning

As we already discussed, Broyden's method requires the analyst to supply an initial estimate B^0 of the Jacobian, and the performance or even eventual convergence of the algorithm hinges on the quality of this initial estimate. Assuming the initial iterate c^0 is close to the solution, a good initial value for B^0 can be obtained by computing the actual Jacobian at c^0.[12] In principle, this can be done numerically, but aside from the possibility of numerical

[12] Recall that the collocation function F of the quasi-discrete model is differentiable almost everywhere by Theorem 2.

error, the numerical evaluation of the Jacobian is impractical in our case as it requires multiple evaluations of F, which are very expensive to perform. A superior alternative is the analytic evaluation of the Jacobian, but this is also prohibitively costly in our problem. In particular, in order to compute the partial derivative $\frac{\partial F_{i,k}(c)}{\partial c_{j,\ell}}$, we must account for the indirect effect of a change in the collocation coefficient $c_{j,\ell}$ on the proposal strategies $\pi_h(\theta, q; c)$.

Nevertheless, a strategy for a second-best approximation to the Jacobian is available analytically. This approximation evaluates the derivative by ignoring any effect of changes in the collocation coefficients that is channeled through changes in proposal strategies. Evaluating only the direct effects, it is immediate from (8) and (10) that the pseudo-derivative resulting from this approach, denoted $\psi_{(i,k),(j,\ell)}(c) \approx \frac{\partial F_{i,k}(c)}{\partial c_{j,\ell}}$, takes the form

$$\psi_{(i,k),(j,\ell)}(c) = \begin{cases} T_\ell(\tilde{x}_k) - \delta_i \int_q \int_\theta \sum_{h \in N} p_h T_\ell(\pi_h(q, \theta; c)) f(\theta) d\theta g(q|\tilde{x}_k) dq & \text{if} \quad i = j \\ 0 & \text{if} \quad i \neq j. \end{cases}$$

Let $\Psi(c)$ denote the $mn \times mn$ matrix with entries corresponding to the pseudo-derivatives computed above and evaluated at c, so $\Psi(c)$ is a block-diagonal matrix with the size of each block equal to $m \times m$. As a consequence, the evaluation of the inverse of $\Psi(c)$ is relatively inexpensive.[13] Thus, we can choose c^0, set $B^0 = \Psi(c^0)$, and apply Broyden's method to the collocation function F. As discussed by Kelley (2003), this is equivalent to setting $B^0 = I_{nm}$, where I_{nm} is the $nm \times nm$ identity matrix, and then solving a *left-preconditioned* version of the collocation equations using the function

$$\Psi F(c) = [\Psi(c^0)]^{-1} F(c)$$

instead of the function F.

While the above left-preconditioned version of Broyden's method offers one feasible route for the implementation of the method, the availability of a relatively cheap approximation to the Jacobian $\Psi(c)$ of F for every c suggests an even more appealing alternative. Observe that if $\Psi(c)$ is indeed a good approximation of the Jacobian of F and is invertible, then the Jacobian of the function $[\Psi(c)]^{-1}F(c)$ is close to the identity matrix I_{nm}. The low-cost of computing $[\Psi(c)]^{-1}$ suggests that we can consider applying Broyden's algorithm to a modified function $\hat{F}$ given by

$$\hat{F}(c) = [\Psi(c)]^{-1} F(c) = c - [\Psi(c)]^{-1} S(c), \tag{15}$$

where $S(c)$ is a $nm \times 1$ column vector whose entry corresponding to i's coefficient on T_k is

$$S_{i,k}(c) = u_i(\tilde{x}_k) + \delta_i \int_q \int_\theta \sum_{h \in N} p_h(\theta_i \cdot \pi_h(q, \theta; c)) f(\theta) d\theta g(q|\tilde{x}_k) dq. \tag{16}$$

Assuming $\Psi(c^*)$ is non-singular,[14] the collocation coefficients c^* solve the collocation equations (i.e., $F(c^*) = 0$) if and only if they solve $\hat{F}(c^*) = 0$. As a result, we have transformed

[13] In fact, in the case in which players share the same discount factor, the blocks forming the block diagonal matrix $\Psi(c)$ are identical for each of the n players, and computation of the inverse of $\Psi(c)$ reduces to the computation of the inverse of one $m \times m$ matrix.

[14] It is straightforward to show that for all c, $\Psi(c)$ can be singular for at most a finite number (possibly zero) of discount factors $(\delta_1, \ldots, \delta_n)$.

the system of equations $F(c) = 0$ to a system $\hat{F}(c) = 0$ that has a lower condition number (see Judd (1998), pages 67–70 and Section 5.7) and is much closer to being linear. We thus proceed to apply Broyden's method to the function $\hat{F}$ instead of the collocation function F, using initial iterate c^0 and initial approximate Jacobian $B^0 = I_{nm}$. It turns out that this implementation of Broyden's method yields far superior performance in the numerical experiments we consider.

5.3 Two Alternative Solution Methods

In this subsection, we describe two additional methods that can be used as alternatives to Broyden's method in order to solve the collocation equations. Outside the class of Newton-like methods for solving non-linear equations, a standard approach for the solution of the collocation equations is a form of *value iteration.* This method starts with an initial guess c^0, and in the τ-th iteration it generates a new set of collocation coefficients $c^{\tau+1}$ by performing Chebyshev function interpolation on the expected utility functions evaluated at the collocation nodes, where the values of these functions are computed using the collocation coefficients, c^τ. Specifically, at the τ-th iteration, the collocation coefficients c^τ imply a vector $\hat{U}^\tau = (\hat{U}_1^\tau, \ldots, \hat{U}_n^\tau)$ of nm values for the n players' expected utilities at each of the m collocation nodes according to

$$\hat{U}_i^\tau(\tilde{x}_k) = u_i(\tilde{x}_k) + \delta_i \int_q \int_\theta \sum_{h \in N} p_h \left[U(\pi_h(q, \theta; c^\tau); c_i^\tau) + \theta_i \cdot \pi_h(q, \theta; c^\tau) \right] f(\theta) d\theta g(q|\tilde{x}_k) dq,$$

which updates $\hat{U}^\tau$ by computing "best response" proposals, $\pi_h(q, \theta; c^\tau)$. We then use the values $\hat{U}_i^\tau(\tilde{x}_k)$ to obtain $c^{\tau+1}$ by performing interpolation, i.e., by solving the linear system of equations

$$\hat{U}_i^\tau(\tilde{x}_k) = \sum_{\ell=1}^{m} c_{i,\ell}^{\tau+1} T_\ell(\tilde{x}_k), \tag{17}$$

$i = 1, \ldots, n$, $k = 1, \ldots, m$, for the unknown $c^{\tau+1}$. A form of value iteration is used by Duggan, Kalandrakis, and Manjunath (2008) to solve a variant of the model that we study. They also use a Chebyshev polynomial representation of the unknown functions to be solved for, but in their solution method the corresponding equations (17) are not solved exactly, as they use more nodes at which to evaluate the unknown functions than coefficients, i.e., they perform function approximation instead of interpolation.

A second alternative is motivated by the form of the transformed collocation function $\hat{F}$ in (15). In particular, it is apparent from the definition of the function $\hat{F}$ that collocation coefficients solving the equations $\hat{F}(c) = 0$ constitute a fixed point of the function

$$\hat{\hat{F}}(c) = -[\Psi(c)]^{-1} S(c),$$

where $S(c)$ is defined in (16). Thus, when viewed as a fixed point $c = \hat{\hat{F}}(c)$, a solution to the original collocation equations $F(c) = 0$ can then be obtained by *fixed point iteration* on the auxiliary function $\hat{\hat{F}}$, so that at the τ-th iteration we obtain

$$c^{\tau+1} \;\; = \;\; \hat{\hat{F}}(c^\tau). \tag{18}$$

An alternative way to motivate this iterative method is to construe it as a *pseudo-Newton* method. Again guided by (15), we note that the updating step described in (18) is equivalent to a Newton step for the collocation function F, since

$$c^{\tau+1} \;\; = \;\; \hat{\hat{F}}(c^\tau) = c^\tau - [\Psi(c^\tau)]^{-1}F(c^\tau),$$

where instead of the Jacobian $F'(c^\tau)$ used in the conventional Newton method, we have substituted the approximation $\Psi(c^\tau)$. Viewed as a pseudo-Newton method, the iteration (18) then admits all the globalization strategies employed for correcting Newton iterations. In particular, if at the τ-th iteration the step suggested by the pseudo Jacobian $\Psi(c^\tau)$ does not lead to a sufficient decrease in the norm of the residual $||F(c^{\tau+1})||$ compared to the residual $||F(c^\tau)||$, then we can adjust the length of the step along the same direction. In the numerical experiments we report, we implement both Broyden's method and this pseudo-Newton method using the same Armijo rule and a parabolic line search for the optimal step size at each iteration.

6 Numerical Experiments & Core Convergence

In this section, we conduct a number of numerical experiments designed to evaluate the performance of the collocation method and the techniques for solving the collocation equations for dynamic bargaining games we developed in Sections 3–5. We begin in Subsection 6.1 with a description of the numerical specification of the model parameters and other numerical specifications required to implement the algorithms. In Subsection 6.2, we discuss and compare the performance of the three methods for solving the collocation equations, and we provide an application of Theorem 3 for the purposes of approximating equilibria in a model with a continuous space of policies. As an application of these techniques, we conclude in Subsection 6.3 with an illustration of the core convergence result of Duggan and Kalandrakis (2008).

6.1 Specifications

Throughout this section, we specify models with a two-dimensional policy space ($d = 2$),[15] and we assume that the set of feasible policies is a finite grid contained in the square

$$\tilde{X} \;\; = \;\; [-1, 1]^2.$$

Given $x = (x_1, x_2)$, we set the support of the density of the status quo, $g(q|x)$, to

$$[\tfrac{9}{10}x_1 - \tfrac{1}{10}, \tfrac{9}{10}x_1 + \tfrac{1}{10}] \times [\tfrac{9}{10}x_2 - \tfrac{1}{10}, \tfrac{9}{10}x_2 + \tfrac{1}{10}] \subset \tilde{X}.$$

[15]We have written software that allows us (at increasing cost) to solve for equilibria in models with higher-dimensional policy spaces.

Each coordinate $q_i, i = 1, 2$, of the status quo is independently Beta distributed in $[\frac{9}{10}x_i - \frac{1}{10}, \frac{9}{10}x_i + \frac{1}{10}]$ with parameters $a = b = 5$.[16] This specification ensures that $g(q|x)$ is twice continuously differentiable with respect to x with Lipschitz bounded second derivative, as assumed in Subsection 2.1. We fix the distribution of the preference shocks, $f(\theta)$, to be uniform in $(-\frac{1}{20}, \frac{1}{20})^{2n}$. We solve models with $n = 9$ players and simple majority rule, so that the set of winning coalitions is given by

$$\mathcal{D} \quad = \quad \{C \in N : |C| \geq 5\}.$$

We specify uniform recognition probabilities across players, so that $p_i = \frac{1}{9}$, and set common discount factors $\delta_i = \delta$. With the exception of one set of computations in the core convergence subsection, we use a value of $\delta = 0.7$ for the discount factor. The stage utilities of all players are negative quadratic with the ideal point of player i denoted $\hat{x}_i$, so that

$$u_i(x) \quad = \quad -(x_1 - \hat{x}_{i,1})^2 - (x_2 - \hat{x}_{i,2})^2.$$

We specify a number $m = 31 \times 31 = 961$ of collocation nodes $\tilde{x}_k$, which are located at the roots of the Chebyshev polynomials as described in Section 3. Based on experimentation with alternative numbers of collocation nodes, this choice seems to strike a good compromise between achieving sufficient precision in the approximation of the expected utility functions U and computational cost. We borrow routines from the MATLAB *compecon* toolbox of Miranda and Fackler (2002) to generate the collocation nodes as well as the values of the Chebyshev polynomials at these nodes. For each collocation node $\tilde{x}_k$, we specify $\alpha = 25$ Gaussian quadrature nodes in order to perform the integration with respect to the status quo q. These nodes and the corresponding weights are specified using the MATLAB routine *qnwbeta* from the *compecon* toolbox of Miranda and Fackler (2002). We use a Sobol sequence of $\beta = 128$ quasi-random numbers as implemented by Burkardt (2007) in the MATLAB environment. We vary the size of the policy grid X, which is in all cases uniform in $[-1, 1]^2$ with sizes ranging from $7 \times 7 = 49$ to $51 \times 51 = 2{,}601$ points, although most of the computations are performed with a grid of size $35 \times 35 = 1{,}225$. Given the above, each function evaluation requires us to solve $n \times m \times \alpha \times \beta = 27{,}676{,}800$ optimization problems.

We have written MATLAB routines to perform these optimizations and the integrations required in order to evaluate the collocation functions F and $\hat{F}$ or $\hat{\hat{F}}$. These routines take advantage of MATLAB's *Parallel Computing Toolbox,* so that the players' optimization problems are executed in parallel in four processors.[17] We used Kelley (2003)'s MATLAB routine *brsola* to implement the Broyden-Armijo method with a parabolic line search to adjust the Newton step. We adapted the same code in order to implement the same line search procedure for the pseudo-Newton method described in Section 5.3.

[16] Thus, $g(q|x)$ takes the form

$$g(q|x) \quad = \quad \begin{cases} \prod_{i=1}^{2} \frac{(q_i - \frac{9}{10}x_i + \frac{1}{10})^4 (\frac{1}{10} + \frac{9}{10}x_i - q_i)^4}{B(5,5)(\frac{2}{10})^9} & \text{if} \quad q_i \in [\frac{9}{10}x_i - \frac{1}{10}, \frac{9}{10}x_i + \frac{1}{10}], i = 1, 2, \\ 0 & \text{otherwise.} \end{cases}$$

[17] All computations reported in this paper were executed on a 3.2 GHz dual Quad-Core Intel Xeon Mac-Pro machine with 8GB of memory operating under Microsoft's 64-bit Windows Vista Professional operating system.

	Broyden	Pseudo-Newton	Value Iteration
Number that converged to 10^{-5}	10	8	10
Median function evaluations needed	8.5	9.5	19
Average function evaluations needed	10.4	8.875	18.8
Min function evaluations needed	6	6	17
Max function evaluations needed	22	12	21
Number that converged to 5×10^{-5}	10	10	10
Median function evaluations needed	5	5	14
Average function evaluations needed	5.7	5	13.5
Min function evaluations needed	5	4	12
Max function evaluations needed	9	6	14

Table 1: Performance of three solution methods.

6.2 Numerical Experiments

To compare the performance of the three solution methods discussed in Section 5, we randomly drew ten sets of ideal points $\hat{x}_i$ for the nine players from the uniform distribution in $[-1, 1]^2$. We then applied each of the three procedures (Broyden's method, the pseudo-Newton method, and value iteration) to compute an equilibrium for each of the models specified by the ten sets of ideal points. In these computations, the size of the policy grid is given by $35 \times 35 = 1,225$ points. All iterations were initiated with a value $c^0 = 0$ for the collocation coefficients. We monitored convergence using the averaged L_2 norm of the residual of the transformed collocation function, i.e., $\frac{||\hat{F}(c)||_2}{\sqrt{nm}}$. The function $\hat{F}$ is evaluated in our implementation of Broyden's method, and for the purposes of comparison we evaluate the residual of that function for both the value iteration and the pseudo-Newton methods.[18] We required a residual of 10^{-5} for convergence of the iterations. In Table 1, we report on the performance of the three methods.

[18]These extra computations constitute a negligible fraction of the overall computation. Note that the cost of a function evaluation in Broyden's method is identical to the cost of evaluating the expected payoff values $\hat{U}$ in value iteration, so it is appropriate to gauge the cost of these methods in terms of function evaluations. The bulk of this computational cost arises from solving the optimization problems of the proposers.

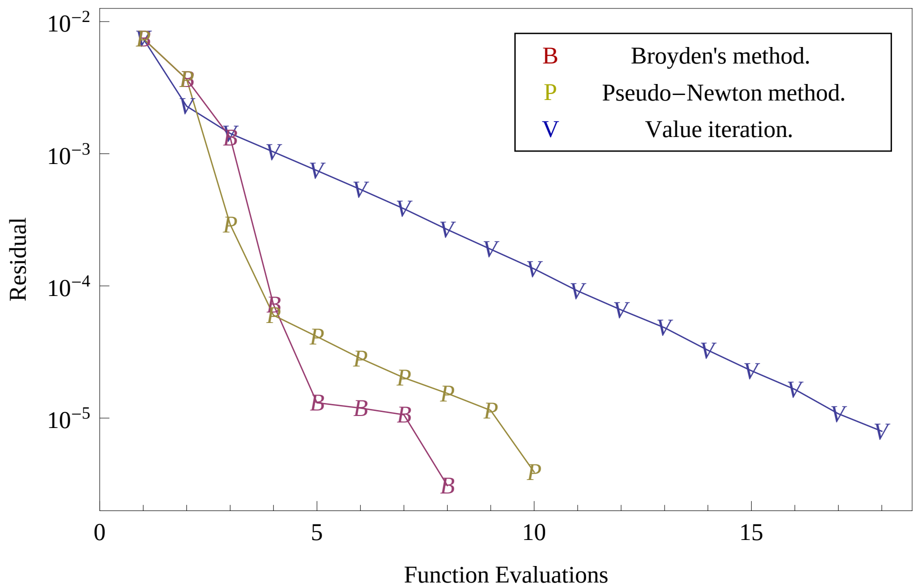

Figure 2: Convergence paths for three solution methods.

As is evident from the first row of Table 1, both Broyden's method and value iteration converged to the target tolerance level in all cases, while the pseudo-Newton method failed to converge in two cases.[19] This occurred when the line search procedure failed to find a step size that produced a sufficient reduction in the norm of the residual. According to Table 1, our implementation of Broyden's method outperforms value iteration in terms of the required function evaluations, as it required only 10.4 function evaluations on average compared to 18.8 for value iteration, thus economizing on function evaluations by at least 50%. Furthermore, in the typical case, Broyden's method converged in 8 to 9 function evaluations, while value iteration consistently required roughly double that number in the range from 17 to 21 function evaluations. Table 1 also reveals that our pseudo-Newton method is quite competitive in the initial iterations. In particular, the pseudo-Newton method outperforms the other two methods when the required tolerance level is relaxed to 5×10^{-5}, in which case it converged in all ten specifications.

In Figure 2, we display the convergence path of the three algorithms in one of the ten specifications in which the performance of the three methods (measured in terms of the number of function evaluations required for convergence) is roughly at the medians reported in Table 1. This figure is representative of the performance of the three methods. As expected, value iteration exhibits a consistent linear decrease in the residual norm. On the other hand, the other two methods exhibit super-linear convergence initially, but then slow down near the tolerance level. Although Newton-like methods perform better near the solution in the absence of numerical error in function evaluation, the slower convergence exhibited in the later iterations of Broyden's method in Figure 2 is not surprising given

[19] We discarded those two cases in calculating average function evaluations needed by pseudo-Newton.

the numerical precision that we have built into the problem. Interestingly, it is exactly near that tolerance level that the pseudo-Newton method loses its initial advantage over Broyden's method. The advantage of Broyden's method at these later iterations is due to the fact that it accumulates information across iterations that results in a more accurate representation of the Jacobian of the function $\hat{F}$ in later iterations. Our pseudo-Newton method, on the other hand, relies on the pseudo-Jacobian $\Psi(c^\tau)$ at each iteration, and that contains no information on the indirect effect of changes in the collocation coefficients c^τ on the likelihood of different values for the optimal proposals $\pi_h(q, \theta; c^\tau)$.

We performed a number of checks to test the quality of the solution produced at these tolerance levels. Given solution c^* obtained with either of the above methods, we evaluated the expected payoff of players at a large number of 50×50 test points on a uniform grid in $\tilde{X}$ that do not coincide with the 31×31 collocation nodes. We then computed the residual difference between these function values and those obtained from the collocation coefficients; namely, for each test point $z \in \tilde{X}$ and each player, we computed

$$R_i(z) = U(z; c_i^*) - \big[u_i(z) + \delta_i \int_q \int_\theta \sum_{h \in N} p_h \big[U(\pi_h(q, \theta; c^*); c_i^*) + \theta_i \cdot \pi_h(q, \theta; c^*)\big] f(\theta) g(q|z) d\theta dq\big].$$

The average L_2 norm of these residuals is roughly 7×10^{-5}, and it is less than 5×10^{-4} in the L_∞ norm. These numbers are consistent with the tolerance level of 10^{-5} used to gauge convergence of the algorithms.

It should be noted that we also attempted alternative implementations of Broyden's method, using the unconditioned function F or the left-preconditioned version ΨF discussed in Subsection 5.2, instead of the dynamically preconditioned version $\hat{F}$, but these versions of Broyden's method did not perform as well compared to either of the alternatives that we present. The conclusion we reach from this discussion and from Table 1 is that Broyden's method constitutes a viable solution method that can lead to a significant economization in computing time over value iteration when solving the quasi-discrete dynamic bargaining games that we study. Translated in terms of computer time, value iteration, with a median number of 19 function evaluations, required roughly in excess of 54 minutes to solve a typical instance of our nine player model, while the corresponding cost for Broyden's method drops to just 23 minutes. If only a rough approximation of the solution is required, then our pseudo-Newton method is a competitive alternative.

We conclude this section with an illustration of Theorem 3, which allows us to obtain an increasingly precise approximation of the equilibrium of the continuum model from a sequence of equilibrium computations of quasi-discrete models. In particular, we chose the typical specification of ideal points among the ten used to produce Table 1, on which we reported in Figure 2, and we used Broyden's method to compute an equilibrium for a sequence of six quasi-discrete models with grid sizes ranging from 7×7 to 51×51, roughly doubling the number of points in the policy grid X with each successive model. In order to ensure independence of the solutions for each model, we started Broyden's method from $c^0 = 0$ for each of the six models in the sequence.[20] For each of the six models, we attained

[20] This is not the approach we would use if our goal were to economize on computation time. In that case,

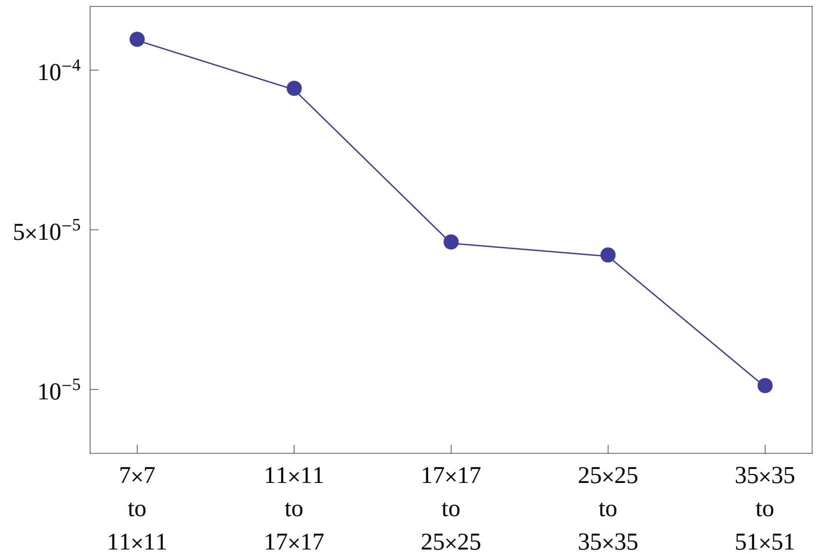

Figure 3: Change in solution with finer policy grids.

convergence to within the tolerance level of 10^{-5}.[21] In Figure 3, we depict the change in the L_2 norm of the solution obtained for each successive pair of models. As is evident from the decreasing line in Figure 3, these solutions grow successively closer to each other with each increase on the size of the policy grid. In fact, the equilibrium of the model with 35×35 grid differs from that obtained for the 51×51 grid by roughly 10^{-5}, suggesting that (at least for the specifications we use) we reasonably approximate equilibria of the continuum model by looking at these grid sizes.

6.3 Core Convergence

Every equilibrium of the dynamic bargaining game we consider admits an invariant distribution (at least one) over implemented policies, an implication of Theorem 4 of Duggan and Kalandrakis (2008). In Theorem 6 of that paper, we also show that when the noise on the status quo and the preference shocks are small, and when the stage utilities are close to quadratic and are close to admitting a core policy, and the player located at the core has positive probability of recognition, then the invariant distributions generated by equilibria of the model must be close (in the sense of weak convergence) to the unit mass on the limiting core policy. In this subsection, we implement the Broyden collocation method in

we would use the solution output (and possibly the approximated Jacobian) from application of Broyden's method in each lower grid size model in order to provide a 'hot' starting point for the next model with finer policy grid, thus significantly economizing on the number of function evaluations needed for the models with finer policy grids.

[21] Convergence to this tolerance level is generally harder in the models with coarser policy spaces, and Broyden's method converged to within 2×10^{-5} in the model with an 11×11 policy grid. We then ran a small number of value iterations to get the solution of that model to within the required tolerance level.

	$\hat{x}_1$	$\hat{x}_2$	$\hat{x}_3$	$\hat{x}_4$	$\hat{x}_5$	$\hat{x}_6$	$\hat{x}_7$	$\hat{x}_8$
$\hat{x}_{i,1}$	–0.8	0.3	–0.2	0.9	0.1	–0.15	0.3	–0.9
$\hat{x}_{i,2}$	0	0	0.2	–0.9	0.6	–0.9	0.2	–0.6

Table 2: Location of ideal points of players 1–8 in core convergence experiment.

the two-dimensional, nine-player model specified in Subsection 6.1 to provide a numerical illustration of this convergence result.

As a first step, we address the requirements of Duggan and Kalandrakis's (2008) dynamic core theorem. The specification we have chosen meets the requirements regarding the curvature of the players' stage utilities, as we assume they are quadratic. Duggan and Kalandrakis's (2008) theorem assumes a sequence of models that exhibit increasingly smaller levels of noise on preferences and the status quo. In our specification, we have assumed a relatively small level of noise, which we keep fixed in the computations that follow. These computations demonstrate that the equilibrium forces that yield the theorem take effect well before the noise on preferences and the status quo become negligible, i.e., even if fixed at the levels we have already specified. In addition to these requirements, the theorem assumes a configuration of stage preferences that becomes closer to admitting a core. For that purpose, we fix the ideal points of players 1 through 8 at the values reported in Table 2. It is straightforward to verify that, given negative quadratic stage utilities and absent preference shocks, these ideal points satisfy Plott's (1967) pairwise-symmetry conditions for the existence of a core point at the origin of the space $(0, 0)$. With the ideal points of the remaining players thus fixed, we can move the ideal point of player 9 from an arbitrary position toward the origin and monitor the effect of this move on the invariant distribution over policies induced in equilibrium.

In our numerical experiments, we varied the ideal point of player 9 from the location $\hat{x}_9 = (-0.4, -0.4)$, at which a core point does not exist, to an intermediate location $\hat{x}_9 = (-0.2, -0.2)$, and finally to the point $\hat{x}_9 = (0, 0)$, at which player 9 is located at the core (absent preference shocks). We considered two possible values for the common discount factor, a low value of $\delta = 0.3$ and the value $\delta = 0.7$ used for the computations we have already reported. We computed equilibria for each of these six configurations of ideal points and discount factors using a space of policies X given by a 51×51 uniform grid in $[-1, 1]^2$. To compute an equilibrium for that grid size, we first computed equilibria using coarser grids at a much lower computation cost, and then we gradually increased the size of the grid, using the solutions from smaller grids in order to initiate Broyden's algorithm for finer grids. At the 51×51 grid size, we required a more stringent convergence tolerance of 5×10^{-6} for termination of the algorithm, and convergence typically required only one or two iterations at this grid level, as the initial values were already quite close to the solution.[22]

Upon obtaining an equilibrium in this fashion, we simulated a long sequence of equi-

[22]Thus, consistent with Theorem 3 and the conclusion drawn from Figure 3, we can view the computed equilibria at this grid size as a good approximation of equilibrium in the continuous model.

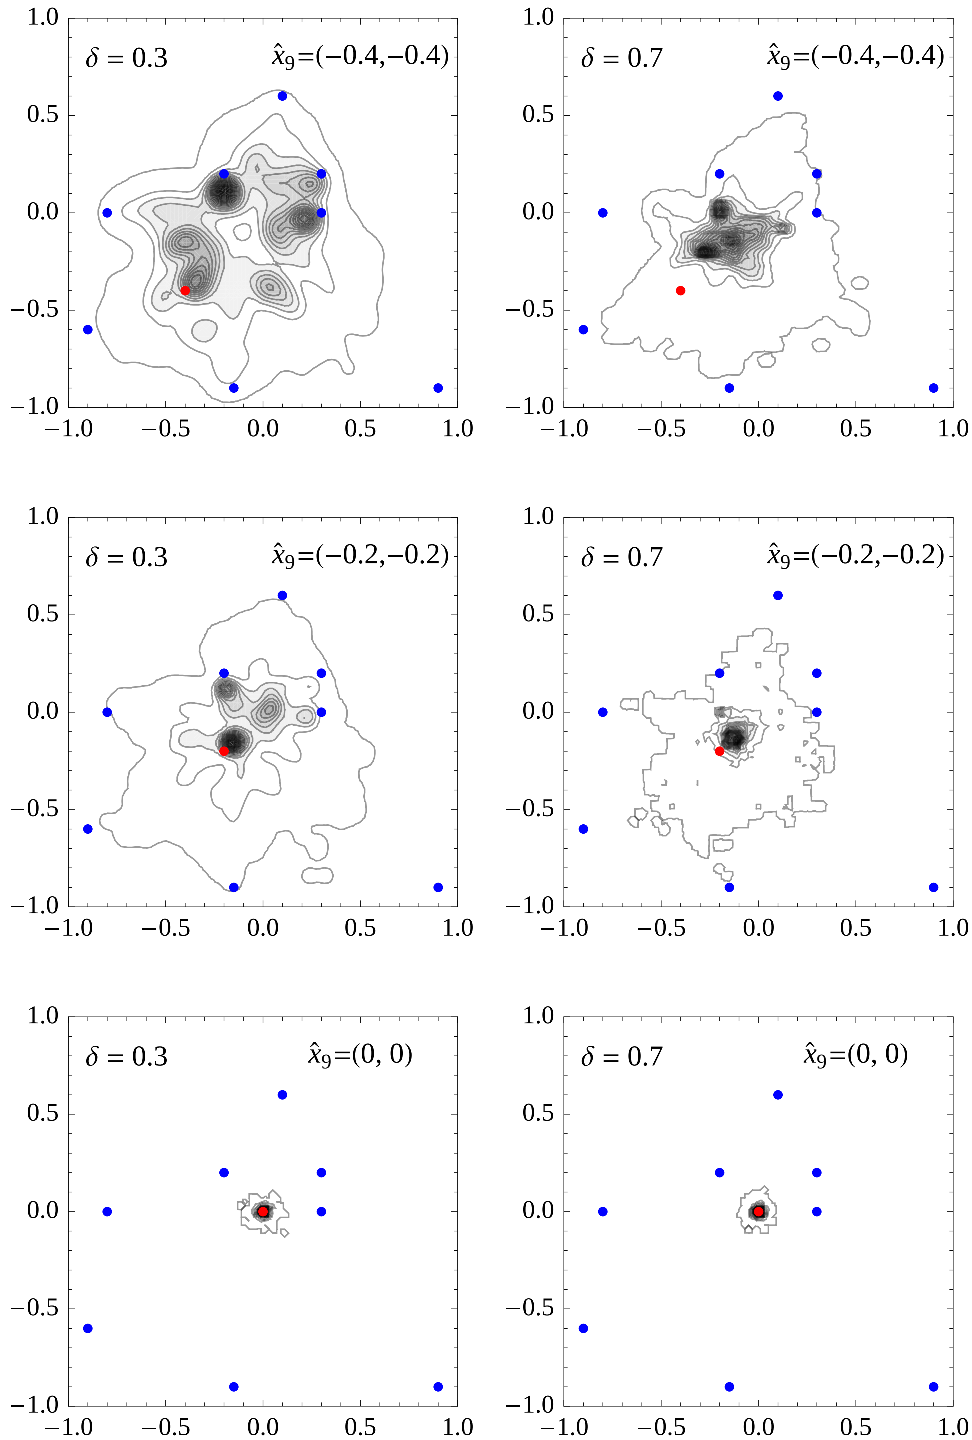

Figure 4: Core convergence.

Models with $\delta = .3$ appear on the left, those with $\delta = .7$ on the right. Long-term policies when $\hat{x}_9 \neq (0,0)$ tend to be more moderate when players are more patient. Long-term policies pile mass at the core point when $\hat{x}_9 = (0,0)$.

librium play over 20,000 periods, retaining the policy implemented in each period, and then used the last 15,000 periods as a sample from an invariant distribution induced by the equilibrium. We used this sample in order to depict this invariant distribution over policies in Figure 4 using a kernel density estimator. The first column of Figure 4 displays the invariant distribution over policies for the three locations for the ideal point of player 9 and a low discount factor $\delta = 0.3$, and the second column corresponds to a discount factor $\delta = 0.7$. The last row of these graphs corresponds to the case where player 9's location is at the core point $(0, 0)$. Moving from the first row to the last row of these graphs, we find that the equilibrium invariant distribution piles more mass near the core point for both discount factors. In the last row, the invariant distribution is obviously concentrated near the origin with virtually no policies occurring away from the limiting core point in accordance with Theorem 6 of Duggan and Kalandrakis (2008).

Observe that the invariant distribution when $\delta = 0.3$ and $\hat{x}_9 \neq (0, 0)$ is markedly more dispersed than the corresponding distribution when $\delta = 0.7$. Evidently, the strategic incentives of the players induce them to be more moderate and seek policies closer to the center of the policy space when the discount factor is larger. Such moderate policies insulate players from the risk of a large change in policy when a player with whom they disagree significantly is the proposer and can leverage the policy outcome in her favor due to the fact that a status quo is located too far away from the remaining players' preferences. The moderating effect of players' patience was also identified in the computations of Baron and Herron (2003) in a two-dimensional setting with three players symmetrically located in an equilateral triangle. Our analysis suggests that the observations of those authors are robust to the horizon of the game, and they raise the question of whether a theoretical explanation for this regularity is possible — a question we leave open.

In order to illustrate the players' voting and proposal strategies, we depict the collective acceptance set and optimal proposals for various status quo in Figure 5. The three figures on the right column of Figure 5 correspond to the equilibrium for the specification of ideal points that is the furthest from satisfying Plott's conditions so that the ideal points of players 1 to 8 are as specified in Table 2, that of player 9 is located at $(-0.4, -0.4)$, and the discount factor is set to $\delta = 0.7$. For the purposes of comparison we also depict the corresponding acceptance sets and proposal strategies when players are impatient ($\delta = 0$) in the left column of Figure 5. A number of observations emerge from this comparison. First, the collective acceptance sets when players are patient tend to include more alternatives in the center of the policy space and fewer at the extremes. Second, many players strategically compromise their proposals by offering a moderate proposal instead of choosing the feasible policy that is closest to their ideal point, even in cases when it is feasible. Third, the collective acceptance set grow small as the status quo moves to a more central location at $q = (-0.1, -0.1)$.

In Figure 6, we display the equilibrium preferences of one of the players in each of the three models with discount factor $\delta = 0.7$ to elucidate the nature of the equilibrium and the strategic incentives of the players. The first row of Figure 6 depicts the continuation value of player 7 and is in some sense a representation of the effect of future equilibrium play on that player's incentives. Note that for all three models (player 9 located at $(-0.4, -0.4)$,

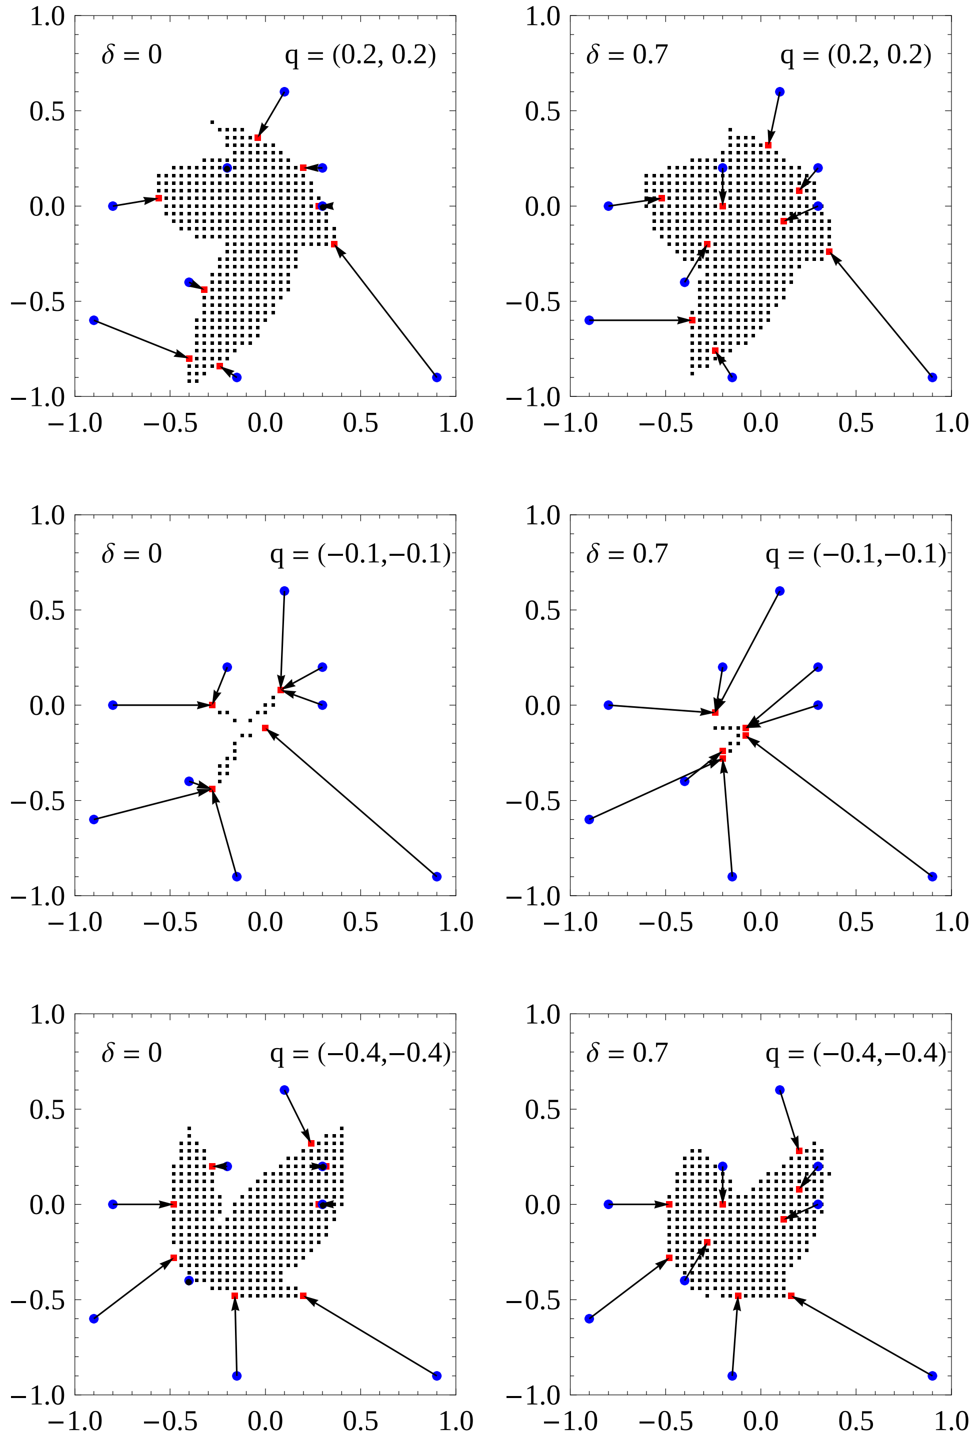

Figure 5: Social acceptance sets and proposal strategies.

Collectively acceptable alternatives are displayed in black, stage ideal points in blue, optimal proposals in red. Arrows originate from ideal points and point to the proposals of the corresponding player. Individual preference shocks θ_i are set to zero for all i.

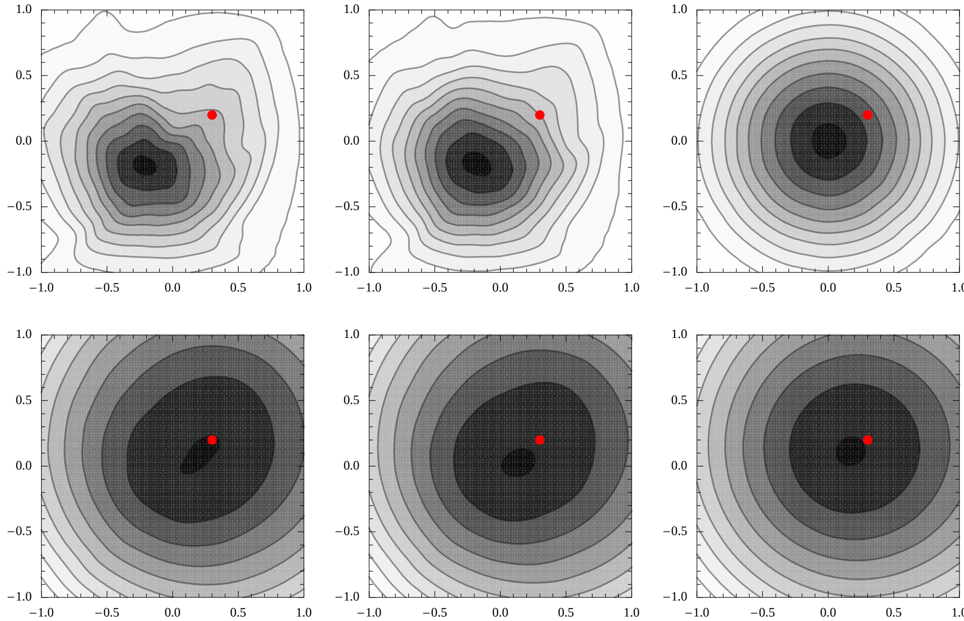

Figure 6: Continuation value and dynamic utility.

Continuation value $v_7(x)$ (first row) and dynamic utility $U_7(x)$ (second row) of player 7, for the models with player 9's stage utility ideal point located at $(-0.4, -0.4)$, $(-0.2, -0.2)$, and $(0, 0)$, respectively. The red point indicates player 7's stage ideal point. Darker areas indicate higher utility.

$(-0.2, -0.2)$, and $(0, 0)$, respectively), the policy outcomes that engender the best future distribution of policies for player 7 are those that are close to the center of the policy space, near the ideal policy of player 9. This area of the policy space concentrates most of the mass of the equilibrium invariant distribution, as is evident by the second column of Figure 4, and player 7's continuation value is maximized when the current policy is near that area of the policy space. Nevertheless, player 7 has a stage ideal point fixed at $\hat{x}_7 = (0.3, 0.2)$, which is removed from the policies that generate higher future utility for that player. The net effect of these two incentives, i.e., those concerning player 7's future expected utility versus the utility derived from present policy, is that the policy generating the highest expected discounted payoff for player 7 is a compromise between her stage ideal point and the policies that are near the center of the policy space. This is evident from the second row of Figure 6. This row displays player 7's dynamic utility $U_7(x)$, which in all cases appears to have a maximizer closer to the origin compared to the stage ideal point of the player. The dynamic utility also appears to be non-concave in the equilibrium of the model with $\hat{x}_9 = (-0.4, -0.4)$, although this non-concavity is not very pronounced, despite the fact that player 7's stage utility is strictly concave. Of course, the dynamic utility U_7 is derived endogenously as the convex combination of the stage utility and player 7's continuation value function, so the non-concavity is not surprising given the fact that player 7's continuation value is shaped by the complicated nature of future equilibrium play.

7 Conclusion

We have proposed one approach to computation of equilibrium in a general class of dynamic bargaining games, and we have provided theoretical support for this approach. Despite the fact that proposals are endogenous in our model, we have proven that we can approximate equilibrium dynamic utilities as the solution to a sufficiently smooth system of collocation equations. We have shown that a preconditioned version of Broyden's method constitutes a viable solution method for the quasi-discrete model, and that it can lead to a significant economization in computing time over value iteration. Furthermore, taking the limit of equilibria of a sequence of increasingly finer quasi-discrete models, we can compute equilibria of the continuum model. We employ these techniques in an illustration of core convergence in the model. Our computational results suggest that greater patience on the part of the players may induce greater policy moderation, opening the question of a theoretical explanation for this regularity. These contributions, taken together, should complement the future analysis of such theoretical questions as the effect of proposal power or the distribution of voting rights on equilibrium policy outcomes, and they may facilitate the empirical estimation of unobservable parameters, such as proposal probabilities or discount factors, in dynamic bargaining games.

A Proof of Theorem 2

Theorem 2 *Assume X is finite and f is C^1. The collocation function $F : \mathbb{R}^{nm} \to \mathbb{R}^{nm}$ is locally Lipschitz continuous and directionally differentiable with directional derivatives.*

Proof To any $c \in \mathbb{R}^{nm}$ we can associate policy-specific dynamic payoffs $U(y; c_i)$ as in (9), and we can consider the optimality conditions (5) and (6). We define $\Theta(q, y, h, a_{-h}; c)$ as the closure of the subset of preference shocks $\theta \in \mathbb{R}^{nd}$ for which the acceptance sets of the players other than h satisfying (5) are summarized by the vector $a_{-h} \in A = \{0,1\}^{(n-1)|X|}$, and the proposal $y \in X \cup \{q\}$ solves (6) for proposer h. Here, $a_{-h} = (a_1, \ldots, a_{h-1}, a_{h+1}, \ldots, a_n)$, and $a_{j,p} = 1$ indicates that j will vote to accept x_p if proposed, and $a_{j,p} = 0$ indicates rejection. To be more precise, the shocks $\theta \in \Theta(q, y, h, a_{-h}; c)$ are characterized by two sets of inequalities. The first set of inequalities relates to the voting incentives of players other than h. For all $x_p \in X$ and all $j \in N \setminus \{h\}$, θ must satisfy

$$U(q; c_j) + \theta_j \cdot q \quad \leq \quad U(x_p; c_j) + \theta_j \cdot x_p \tag{19}$$

if $a_{j,p} = 1$, and it must satisfy

$$U(q; c_j) + \theta_j \cdot q \quad \geq \quad U(x_p; c_j) + \theta_j \cdot x_p \tag{20}$$

if $a_{j,p} = 0$. The second set of inequalities concerns the optimality of proposing y for player h. Let $C(x_p, h, a_{-h}) = \{i \in N \setminus \{h\} : a_{i,p} = 1\}$ be the coalition of players other than h who accept proposal x_p when acceptance sets of players other than h are given by a_{-h}. Given a status quo q and potential proposal $y \in X \cup \{q\}$, let $Y(y, h, a_{-h}) = (\{x_p \in X : C(x_p, h, a_{-h}) \cup \{h\} \in \mathcal{D}\} \cup \{q\}) \setminus \{y\}$ be the subset of policies other than y that can pass by the vote of h and those of other players, given acceptance sets a_{-h}. For proposing $y \in X \cup \{q\}$ to be optimal for player h given a_{-h}, θ must be such that for all $z \in Y(y, h, a_{-h})$,

$$U(z; c_h) + \theta_h \cdot z \quad \leq \quad U(y; c_h) + \theta_h \cdot y. \tag{21}$$

Thus, the closure of the subset of preference shocks such that h proposes y when the status quo is q and the acceptance sets of other players are summarized by a_{-h} is defined as follows:

$$\Theta(q, y, h, a_{-h}; c) \quad = \quad \{\theta \in \Theta : (19), (20), \text{ and } (21) \text{ hold}\}.$$

It follows that $\Theta(y, q, h, a_{-h}; c)$ is a convex polytope defined by a finite number of inequalities that are linear in θ. Note that when $x_p \neq q$, the inequalities (19)–(21) are satisfied with equality only for a lower dimensional set of θ's, and furthermore, the intersection

$$\Theta(q, y, h, a_{-h}; c) \cap \Theta(q, y', h, a_{-h}; c)$$

is a measure-zero subset of $\mathbb{R}^{nd}$ for all distinct $y, y' \in \tilde{X}$. Hence, the probability that player h proposes y when the status quo is q is given by

$$\sum_{a_{-h} \in A} \int_{\Theta(q,y,h,a_{-h};c)} f(\theta) d\theta.$$

Based on the above, we rewrite the collocation equations as follows:

$$\begin{aligned} F_{i,k}(c) \quad = \quad & U(\hat{x}_k; c_i) - \Big[u_i(\hat{x}_k) + \delta_i \sum_{h \in N} p_h \sum_{a_{-h} \in A} \\ & \int_q \sum_{y \in X \cup \{q\}} \int_{\Theta(q,y,h,a_{-h};c)} \big[U(y; c_i) + \theta_i \cdot y\big] f(\theta) d\theta g(q|\hat{x}_k) dq\Big]. \end{aligned}$$

Note that in the above integral, the status quo q ranges over $\mathbb{R}^d$, and the probability that the realized status quo lies in the finite set X is zero; thus, we can neglect the measure-zero event that $q \in \tilde{X} \setminus X$.

For all $i \in N$, all $q \in \tilde{X}$, all $y \in X \cup \{q\}$, all $h \in N$, and all $a_{-h} \in A$, we define a function $G_i(\cdot; q, y, h, a_{-h}) : \mathbb{R}^{nm} \to \mathbb{R}$ by

$$G_i(c; q, y, h, a_{-h}) \quad = \quad \int_{\Theta(q,y,h,a_{-h};c)} \big[U(y; c_i) + \theta_i \cdot y\big] f(\theta) d\theta,$$

so that $F_{i,k}$ can be expressed equivalently as

$$F_{i,k}(c) = U(\hat{x}_k; c_i) - \Big[u_i(\hat{x}_k) + \delta_i \sum_{h \in N} p_h \sum_{a_{-h} \in A} \int_q \sum_{y \in X \cup \{q\}} G_i(c; q, y, h, a_{-h}) g(q|\hat{x}_k) dq\Big]. \tag{22}$$

The properties of F of interest will follow from the properties of G.

We prove the theorem in three steps.

Step 1. *For all $i \in N$, all $q \in \tilde{X} \setminus X$, all $y \in X \cup \{q\}$, all $h \in N$, and all $a_{-h} \in A$, there exist functions $\hat{G}_j : C_j \to \mathbb{R}$, $j = 0, \ldots, K$, such that $\{C_j\}_{j=0}^K$ is an open covering of $\mathbb{R}^{nm}$; for all $j = 0, \ldots, K$, $\hat{G}_j$ is continuously differentiable; and $G_i(\cdot; q, y, h, a_{-h})$ is piecewise smooth with representation $\{\hat{G}_0, \hat{G}_1, \ldots, \hat{G}_K\}$.*

Fix $i \in N$, $q \in \tilde{X} \setminus X$, $y \in X \cup \{q\}$, $h \in N$, and $a_{-h} \in A$. To conserve notation, we will suppress the parameters i, y, h, a_{-h} in our construction of the representation; given q, these parameters range over a finite set, and so this convenience is harmless. Because q ranges over the infinite set $\tilde{X}$, however, we will make the dependence of K, C_j, and $\hat{G}_j$ on q explicit. Define $\hat{G}_0 \equiv 0$, with $C_0 = \mathbb{R}^{nm}$ independently of q. If $\Theta(q, y, h, a_{-h}; c)$ has measure zero in $\mathbb{R}^{nm}$ for all c, then we let $K(q) = 0$, which completes the step. Otherwise, let $C^*(q) \subseteq \mathbb{R}^{nm}$ denote the open set of $c \in \mathbb{R}^{nm}$ such that $\Theta(q, y, h, a_{-h}; c)$ has positive measure. We then rewrite (19)–(21) as follows: for all $x_p \in X$, all $j \in N \setminus \{h\}$, and all $z \in Y(y, h, a_{-h})$,

$$\begin{aligned}
\theta_j \cdot (q - x_p) &\leq U(x_p; c_j) - U(q; c_j) \text{ if } a_{j,p} = 1 && (23)\\
\theta_j \cdot (x_p - q) &\leq U(q; c_j) - U(x_p; c_j) \text{ if } a_{j,p} = 0 && (24)\\
\theta_h \cdot (z - y) &\leq U(y; c_h) - U(z; c_h). && (25)
\end{aligned}$$

We can write these inequalities in the form $\alpha(q)\theta \leq \beta(c; q)$, where the first $(n-1)|X|$ rows of the matrix $\alpha(q)$ and the column vector $\beta(c; q)$ correspond to inequalities (23) or (24) (as appropriate), and the last $|Y(y, h, a_{-h})|$ rows correspond to inequalities (25). Note that the rows of $\alpha(q)$ are non-zero, as $q \notin X$. Of course, the constraint that θ belongs to Θ can also be formalized in terms of linear inequalities: for all $h = 1, \ldots, n$ and all $\ell = 1, \ldots, d$, $\theta_{h,\ell} \leq \overline{\theta}$ and $\theta_{h,\ell} \leq -\overline{\theta}$. Because these inequalities are fixed, we do not include them in the matrix representation of $\Theta(q, y, h, a_{-h}; c)$.

Given any $[(n-1)|X| + |Y(y, h, a_{-h})|] \times nd$ matrix A and any column vector b of dimension $(n-1)|X| + |Y(y, h, a_{-h})|$, define

$$\Theta(A, b) \quad = \quad \{\theta \in \Theta : A\theta \leq b\},$$

and note that the identity $\Theta(q, y, h, a_{-h}; c) \equiv \Theta(\alpha(q), \beta(c; q))$ holds on $C^*(q)$. Letting R index a subset of rows of A, we define $\Theta^R(A, b)$ as the polyhedral set of θ's satisfying the inequalities of R and the constraint that θ belong to Θ, so that

$$\Theta^R(A, b) \quad = \quad \{\theta \in \Theta : \text{ for all } r \in R,\ A_r\theta \leq b_r\},$$

where A_r is the rth row of the matrix A and b_r is the rth row of b. We say a set R is "minimal at (A, b)" if there is no set R' such that R' is a proper subset of R and $\Theta^R(A, b) = \Theta^{R'}(A, b)$. We let $\mathcal{R}(q)$ denote the collection of all sets of inequalities that are minimal at $(\alpha(q), \beta(c; q))$ for some $c \in C^*(q)$, i.e.,

$$\mathcal{R}(q) \quad = \quad \{R : \text{ there exists } c \in C^*(q) \text{ such that } R \text{ is minimal at } (\alpha(q), \beta(c; q))\},$$

and we set $K(q) = |\mathcal{R}(q)|$ and enumerate this collection as $R_1, \ldots, R_{K(q)}$. The domain of the function $\hat{G}_j(\cdot; q)$ to be defined will be

$$C_j(q) \quad = \quad \{c \in C^*(q) : \ R_j \text{ is minimal at } (\alpha(q), \beta(c; q))\},$$

an open set. For each $j = 1, \ldots, K(q)$, we define the functions Γ_j and Δ_j by

$$\Gamma_j(A,b) = \int_{\Theta^j(A,b)} f(\theta)d\theta \quad \text{and} \quad \Delta_j(A,b) = \int_{\Theta^j(A,b)} (\theta_i \cdot y) f(\theta) d\theta$$

for all (A,b), where we use $\Theta^j(A,b)$ for $\Theta^{R_j}(A,b)$. Finally, for each $j = 1, \ldots, K(q)$, we define the function $\hat{G}_j(\cdot;q)$ by

$$\hat{G}_j(c;q) = U(y;c_i)\Gamma_j(\alpha(q),\beta(c;q)) + \Delta_j(\alpha(q),\beta(c;q))$$

for all $c \in C_j(q)$.

Next, we claim that the functions $\hat{G}_j(\cdot;q)$, $j = 1,2,\ldots,K(q)$, are continuously differentiable at each $c \in C_j(q)$. Given $r \in R_j$, define

$$\Theta^j_r(A,b) = \{\theta \in \Theta^j(A,b) : A_r\theta = b_r\}$$

as the $(nd-1)$-dimensional face of $\Theta^j(A,b)$ determined by the hyperplane $A_r\theta = b_r$. We claim that when $\Theta^j(A,b)$ has positive measure in $\mathbb{R}^{nd}$ and R_j is minimal at (A,b), the solution set $\Theta^j_r(A,b)$ has positive volume in the $(nd-1)$-dimensional hyperplane spanned by $\Theta^j_r(A,b)$. To see this, note that since $\Theta^j(A,b)$ contains an open set, there is some $\tilde{\theta}$ that satisfies the inequalities corresponding to R_j strictly. Since R_j is minimal, the rth inequality is not redundant: there exists $\hat{\theta} \in \Theta^{R_j \setminus \{r\}}(A,b)$ such that $A_r\hat{\theta} > b_r$. Then there exists $\alpha \in (0,1)$ such that $A_r\theta' = b_r$, where $\theta' = (1-\alpha)\hat{\theta} + \alpha\tilde{\theta}$ satisfies the inequalities in $R_j \setminus \{r\}$ strictly. Letting $\tilde{\theta}$ vary while satisfying the inequalities of R_j, this implies that $\Theta^j_r(A,b)$ has positive $(nd-1)$-dimensional volume, as claimed. Therefore, Lasserre's (1998) Lemma 2.2 establishes that $\Gamma_j(A,b)$ and $\Delta_j(A,b)$ are continuously differentiable at such (A,b) with partial derivatives of the form

$$\frac{\partial \Gamma_j}{\partial b_r}(A,b) = \frac{1}{||A_r||}\int_{\Theta^j_r(A,b)} f(\theta)d\theta \quad \text{and} \quad \frac{\partial \Delta_j}{\partial b_r}(A,b) = \frac{1}{||A_r||}\int_{\Theta^j_r(A,b)} (\theta_i \cdot y)f(\theta)d\theta, \tag{26}$$

where $||\cdot||$ denotes the Euclidean norm and integrals are with respect to Lebesgue measure in $(nd-1)$-dimensional space.[23] Given $c \in C_j(q)$, note that $\Theta(\alpha(q),\beta(c;q))$ has positive measure in $\mathbb{R}^{nd}$ by construction, and therefore $\Theta^j(\alpha(q),\beta(c;q)) \supseteq \Theta(\alpha(q),\beta(c;q))$ does as well; furthermore, R_j is minimal at $(\alpha(q),\beta(c;q))$. Thus, Γ_j and Δ_j are continuously differentiable in b at $(\alpha(q),\beta(c;q))$, and the chain rule implies $\hat{G}_j$ is continuously differentiable in c with partials

$$\begin{aligned}\frac{\partial \hat{G}_j}{\partial c_{i,k}}(c;q) =& \frac{\partial U}{\partial c_{i,k}}(y;c_i)\Gamma_j(\alpha(q),\beta(c;q)) \\ &+ \sum_{r \in R_j}\left[U(y;c_i)\frac{\partial \Gamma_j}{\partial b_r}(\alpha(q),\beta(c;q))\frac{\partial \beta_r}{\partial c_{i,k}}(c;q) + \frac{\partial \Delta_j}{\partial b_r}(\alpha(q),\beta(c;q))\frac{\partial \beta_r}{\partial c_{i,k}}(c;q)\right],\end{aligned} \tag{27}$$

[23] Lasserre implicitly assumes that no rows of A are equal to zero. This is true for $\alpha(q)$ because we only consider $q \notin X$.

as claimed.

To show that $G_i(\cdot; q, y, h, a_{-h})$ is piecewise smooth with representation $\{\hat{G}_0, \hat{G}_1, \ldots, \hat{G}_K\}$, note that $G_i(c; q, y, h, a_{-h})$ takes values in the set $\{\hat{G}_0(c; q), \hat{G}_1(c; q), \ldots, \hat{G}_K(c; q)\}$. Indeed, if $c \notin C^*(q)$, then $G_i(c; q, y, h, a_{-h}) = 0 = G_0(c; q)$. Otherwise, consider any $c \in C^*(q)$, and let R_j be any set of inequalities that is minimal at $(\alpha(q), \beta(c; q))$ and such that $\Theta^j(\alpha(q), \beta(c; q)) = \Theta(q, y, h, a_{-h}; c)$; then $G_i(c; q, y, h, a_{-h}) = \hat{G}_j(c; q)$, as claimed. Finally, we must show that $G_i(c; q, y, h, a_{-h})$ is continuous in c. Indeed, consider a sequence $c^\ell \to c \in \mathbb{R}^{nm}$. For each ℓ, there exists R_{j_ℓ} such that $c^\ell \in C_{j_\ell}(q)$ and $\Theta^{j_\ell}(\alpha(q), \beta(c^\ell, q)) = \Theta(q, y, h, a_{-h}; c^\ell)$, which implies $G_i(c^\ell, q; y, h, a_{-h}) = \hat{G}_{j_\ell}(c^\ell, q)$. Without loss of generality (since the collection of subsets of inequalities is finite), we may suppose that j_ℓ is constant in ℓ, i.e., $j_\ell = j$. By continuity, we have $\Theta^j(\alpha(q), \beta(c; q)) = \Theta(q, y, h, a_{-h}; c)$, though R_j may not be minimal at $(\alpha(q), \beta(c; q))$. It may be that c lies outside $C_j(q)$, but we nevertheless have

$$G_i(c^\ell, q; y, h, a_{-h}) = \hat{G}_j(c^\ell, q) = \Gamma_j(\alpha(q), \beta(c^\ell, q)) \to \Gamma_j(\alpha(q), \beta(c; q)) = G_i(c; q, y, h, a_h),$$

where the limit follows from a straightforward dominated convergence argument. Indeed, we define $\phi^\ell(\theta)$ to be $[U(y; c_i^\ell) + \theta_i \cdot y] f(\theta)$ times the indicator function of $\Theta^j(\alpha(q), \beta(c^\ell, q))$, and we define $\phi(\theta)$ as $[U(y; c_i) + \theta_i \cdot y] f(\theta)$ times the indicator function of $\Theta^j(\alpha(q), \beta(c; q))$. The sequence $\{\phi^\ell\}$ is dominated by an integrable function and converges pointwise almost everywhere to ϕ, verifying the limit. This establishes continuity, and we conclude that $G_i(c; q, y, h, a_{-h})$ is piecewise smooth in c.

Step 2. *For all $i \in N$, all $q \in \tilde{X} \setminus X$, all $y \in X \cup \{q\}$, all $h \in N$, and all $a_{-h} \in A$, the function $G_i(\cdot; q, y, h, a_{-h})$ is directionally differentiable in c; furthermore, $G_i(\cdot; q, y, h, a_{-h})$ is locally Lipschitz continuous with uniform Lipschitz constant, i.e., for all $c \in \mathbb{R}^{nm}$, there is a constant M and an open set $\tilde{C}$ containing c such that for all $q \in \tilde{X} \setminus X$, all $y \in X \cup \{q\}$, all $h \in N$, all $a_{-h} \in A$, and all $c', c'' \in \tilde{C}$,*

$$|G_i(c''; q, y, h, a_{-h}) - G_i(c'; q, y, h, a_{-h})| \quad \leq \quad M||c'' - c'||.$$

Fix $i \in N$, $q \in \tilde{X} \setminus X$, $y \in X \cup \{q\}$, $h \in N$, and $a_{-h} \in A$. Since $G_i(\cdot; q, y, h, a_{-h})$ is piecewise smooth with representation $\{\hat{G}_0(\cdot; q), \ldots, \hat{G}_{K(q)}(\cdot; q)\}$, by Step 1, Proposition 2.1 of Kuntz and Scholtes (1994) establishes that $G_i(\cdot; q, y, h, a_{-h})$ has directional derivatives at all $c \in \mathbb{R}^{nm}$, completing the first part of the step. To prove local Lipschitz continuity with a uniform Lipschitz constant, consider any $c \in \mathbb{R}^{nm}$, and let $\tilde{C}$ be any open ball of finite radius containing c. We first claim that for every subset $R \subseteq \{1, \ldots, (n-1)|X| + |X| - 1\}$ (corresponding to possible inequalities in (23)–(25)), there exists a bound M_R such that for all $q \in \tilde{X} \setminus X$, all $y \in X \cup \{q\}$, all $h \in N$, all $a_{-h} \in A$, all j with $R_j = R$, and all $\tilde{c} \in \tilde{C} \cap C_j(q)$, we have

$$\left| \frac{\partial \hat{G}_j}{\partial c_{i,k}}(\tilde{c}; q) \right| \quad \leq \quad M_R.$$

Referring to (27), the term $\frac{\partial U}{\partial c_{i,k}}(y; \tilde{c}_i)\Gamma_j(\alpha(q), \beta(\tilde{c}; q))$ is bounded over such q, y, h, a_{-h}, j, and $\tilde{c}$. The summand includes the term $U(y; \tilde{c}_i)$, which is also bounded over such q, y, h,

a_{-h}, j, and $\tilde{c}$, and the terms

$$\frac{\partial \Gamma_j}{\partial b_r}(\alpha(q), \beta(\tilde{c}; q)) \frac{\partial \beta_r}{\partial c_{i,k}}(\tilde{c}; q) \quad \text{and} \quad \frac{\partial \Delta_j}{\partial b_r}(\alpha(q), \beta(\tilde{c}; q)) \frac{\partial \beta_r}{\partial c_{i,k}}(\tilde{c}; q).$$

We focus on the former, as the argument for the latter is similar, and rewrite it as

$$\frac{\partial \Gamma_j}{\partial b_r}(\alpha(q), \beta(\tilde{c}; q)) \frac{\partial \beta_r}{\partial c_{i,k}}(\tilde{c}; q) \;=\; \left(\int_{\Theta_r^j(\alpha(q), \beta(\tilde{c}; q))} f(\theta) d\theta \right) \left(\frac{1}{||\alpha_r(q)||} \frac{\partial \beta_r}{\partial c_{i,k}}(\tilde{c}; q) \right).$$

The first term in the product is clearly bounded across such q, y, h, a_{-h}, j, and $\tilde{c}$; but the analysis of the second term in the product is complicated by the fact that $||\alpha_r(q)||$ can approach zero for some status quos $q \in \tilde{X} \setminus X$.

In the expression $\frac{1}{||\alpha_r(q)||} \frac{\partial \beta_r}{\partial c_{i,k}}(\tilde{c}; q)$, row r corresponds to one of two cases: either to an inequality in (23) or (24) with alternative x_{p_r} and voter $j_r \neq h$, or to an inequality in (25) with alternative z_r and proposer $i_r = h$. In the first case, the norm of $\alpha_r(q) = q - x_{p_r}$ becomes arbitrarily small when $||q - x_{p_r}||$ is small. The partial derivative $\frac{\partial \beta_r}{\partial c_{j,k}}(\tilde{c}; q)$ is equal to zero if $j_r \neq j$, and it is equal to $\pm T_k(q) - T_k(x_{p_r})$ otherwise. When $j_r = j$, note that as q approaches x_{p_r}, we have

$$\lim_{q \to x_{p_r}} \left| \frac{1}{||\alpha_r(q)||} \frac{\partial \beta_r}{\partial c_{i,k}}(\tilde{c}; q) \right| = \lim_{q \to x_{p_r}} \frac{|T_k(q) - T_k(x_{p_r})|}{||q - x_{p_r}||} = |T'_k(x_p|s)| < \infty,$$

where $s = \frac{1}{||q - x_{p_r}||}(q - x_{p_r})$. This limit holds independently of $q \in \tilde{X}$, $y \in X \cup \{q\}$, and $\tilde{c} \in \tilde{C}$, and since h, a_{-h}, and j belong to finite sets, we conclude that the expression $\frac{1}{||\alpha_r(q)||} \frac{\partial \beta_r}{\partial c_{i,k}}(\tilde{c}; q)$ is uniformly bounded. In the second case, we have $\alpha_r(q) = z_r - y$, which is constant in q unless $y \in X$ and $z_r = q$. Then the norm of $\alpha_r(q) = q - y$ becomes small when $||q - y||$ is small, and an analogous argument shows that the expression is again uniformly bounded. We conclude that the partial derivative $\frac{\partial \hat{G}_j}{\partial c_{i,k}}(\tilde{c}; q)$ is bounded over $q \in \tilde{X} \setminus X$, $y \in X \cup \{q\}$, $h \in N$, $a_{-h} \in A$, j with $R_j = R$, and $\tilde{c} \in \tilde{C} \cap C_j(q)$. If there do not exist such q, y, h, a_{-h}, j, and $\tilde{c}$, we set $M_R = 0$, delivering the claim. Using the claim, we conclude that the directional derivatives $\hat{G}'_j(\cdot; q|s)$ are also bounded over directions s with $||s|| = 1$, $q \in \tilde{X} \setminus X$, $y \in X \cup \{q\}$, $h \in N$, $a_{-h} \in A$, j with $R_j = R$, and $\tilde{c} \in \tilde{C} \cap C_j(q)$. Let $\tilde{M}_R$ be such a bound.

Finally, we deduce that $G_i(q, y, h, a_{-h}; \cdot)$ is Lipschitz on $\tilde{C}$ with uniform constant $M = \sum_R \tilde{M}_R$, the sum being over all subsets of possible inequalities. Consider any $c', c'' \in \tilde{C}$, and let $q \in \tilde{X} \setminus X$, $y \in X \cup \{q\}$, $h \in N$, and $a_{-h} \in A$ be arbitrary. Let $[c', c'']$ be the convex hull of $\{c', c''\}$. For each $\tilde{c} \in [c', c'']$, let $J(\tilde{c}) = \{j : \tilde{c} \in C_j(q)\}$ index the functions $\hat{G}_j(\cdot; q)$ that contain $\tilde{c}$ in their domain, and let $B(\tilde{c})$ be any open ball in $\mathbb{R}^{nm}$ around $\tilde{c}$ such that for all $j \in J(\tilde{c})$, we have $B(\tilde{c}) \subseteq C_j(q)$. We can cover $[c', c'']$ with such open balls, and by compactness there is a finite subcover. We focus on one such ball, say $\tilde{B}$, and we show that $G_i(q, y, h, a_{-h}; \cdot)$ is Lipschitz on $[c', c''] \cap \tilde{B}$ with uniform constant M. To this end, consider any $\tilde{c}', \tilde{c}'' \in [c', c''] \cap \tilde{B}$, and let $\tilde{J}$ be the set of j such that $\tilde{B} \subseteq C_j(q)$ and such that for some $\tilde{c} \in [\tilde{c}', \tilde{c}'']$, we have $G_i(\tilde{c}; q, y, h, a_{-h}) = \hat{G}_j(\tilde{c}, q)$. For each $j \in \tilde{J}$, let $I_j = \hat{G}_j([\tilde{c}', \tilde{c}'']; q)$, and note that since $[\tilde{c}', \tilde{c}'']$ is compact and convex and $\hat{G}_j(\cdot; q)$ is continuous, the image of this set is a closed interval, say $I_j = [s_j, t_j]$. The image $G_i([\tilde{c}', \tilde{c}'']; q, y, h, a_{-h})$ is also a closed interval contained in the union $\bigcup_{j \in \tilde{J}} I_j$, and by construction we have $G_i([\tilde{c}', \tilde{c}'']; q, y, h, a_{-h}) \cap I_j \neq \emptyset$ for all $j \in \tilde{J}$. We therefore have

$$\begin{aligned} |G_i(\tilde{c}''; q, y, h, a_{-h}) - G_i(\tilde{c}'; q, y, h, a_{-h})| &\leq (\max_{j \in \tilde{J}} t_j) - (\min_{j \in \tilde{J}} s_j) \\ &\leq \sum_{j \in \tilde{J}} \tilde{M}_{R_j} ||\tilde{c}'' - \tilde{c}'|| \\ &\leq M||\tilde{c}'' - \tilde{c}'||. \end{aligned}$$

Since M is independent of the ball $\tilde{B}$, this completes the step.

Step 3. *The collocation function F is locally Lipschitz continuous and directionally differentiable.*

We consider each coordinate function $F_{i,k}$, the result then following by a straightforward argument. Referring to (22), the term of interest is the integral

$$\int_q \sum_{y \in X \cup \{q\}} G_i(c; q, y, h, a_{-h}) g(q|\hat{x}_k) dq. \tag{28}$$

For each q, we have shown that $G_i(c; q, y, h, a_{-h})$ is locally Lipschitz continuous and directionally differentiable, and that as a consequence, the integrand in (28) is locally Lipschitz continuous and directionally differentiable. In fact, we showed in Step 2 that for all $c \in \mathbb{R}^{nm}$, there is an open set $\tilde{C}$ containing c such that $G_i(\cdot; q, y, h, a_h)$ is Lipschitz continuous on $\tilde{C}$ with uniform constant. Then Proposition 1 of Qi, Shapiro, and Ling (2005) establishes that the expression in (28), and therefore $F_{i,k}$, is directionally differentiable and locally Lipschitz continuous.

∎

References

Acemoglu, D., G. Egorov, and K. Sonin. 2008. "Dynamics and Stability of Constitutions, Coalitions, and Clubs." *Unpublished manuscript.*

Armijo, L. 1966. "Minimization of functions having Lipschitz-continuous first patial derivatives." *Pacific Journal of Mathematics* 16: 1-3.

Banks, J., and J. Duggan. 2000. "A Bargaining Model of Collective Choice." *American Political Science Review* 94: 73–88.

Banks, J., and J. Duggan. 2006. "A General Bargaining Model of Legislative Policy-making." *Quarterly Journal of Political Science* 1: 49–85.

Baron, D. 1996. "A Dynamic Theory of Collective Goods Programs." *American Political Science Review* 90: 316–330.

Baron, D., and J. Ferejohn. 1989. "Bargaining in Legislatures." *American Political Science Review* 83: 1181–1206.

Baron, D., and M. Herron. 2003. "A Dynamic Model of Multidimensional Collective Choice." In *Computational Models in Political Economy*, ed. K. Kollman, J. Miller, and S. Page. Cambridge, MIT Press.

Battaglini, M., and S. Coate. 2007. "Inefficiency in Legislative Policy-Making: A Dynamic Analysis." *American Economic Review* 97: 118—149.

Binmore, K. 1987. "Perfect Equilibria in Bargaining Models." In *The Economics of Bargaining*, ed. Ken Binmore and Partha Dasgupta. New York, Basil Blackwell.

Broyden, C. G. 1965. "A class of methods for solving nonlinear simultaneous equations." *Mathematics of Computation* 19: 577-593.

Burkardt, John. 2007. "The Sobol Quasirandom Sequence." Website. http://people.sc.fsu.edu/~burkardt/m_src/sobol/sobol.html.

Cho, S.-J. 2005. "A Dynamic Model of Parliamentary Democracy." *Unpublished manuscript.*

Diermeier, D., and P. Fong. 2008. "Endogenous Limits on Proposal Power." *Unpublished manuscript.*

Duggan, J., and T. Kalandrakis. 2007. "Dynamic Legislative Bargaining." *W. Allen Wallis Institute Working Paper Series* 45.

Duggan, John, and Tasos Kalandrakis. 2008. "Dynamic Legislative Policy Making." *Unpublished manuscript.*

Duggan, John, Tasos Kalandrakis, and Vikram Manjunath. 2008. "Dynamics of the Presidential Veto: A Computational Analysis." *Mathematical and Computer Modeling* 48: 1570-1589.

Ip, Chi-Ming, and Jerzy Kyparisis. 1992. "Local Convergence of Quasi-newton Methods for B-differentiable Equations." *Mathematical Programming* 56: 71-89.

Judd, Kenneth L. 1998. *Numerical Methods in Economics.* First ed. The MIT Press.

Kalandrakis, T. 2004. "A Three-Player Dynamic Majoritarian Bargaining Game." *Journal of Economic Theory* 116: 294–322.

Kalandrakis, T. 2007. "Dynamics of Majority Rule with Endogenous Status Quo." *Wallis Working Paper no. 46.*

Kelley, Tim. 2003. *Solving nonlinear equations with Newton's method.* SIAM.

Kuntz, L., and S. Scholtes. 1994. "Structural Analysis of Nonsmooth Mappings, Inverse Functions, and Metric Projections." *Journal of Mathematical Analysis and Applications* 188: 346–386.

Kuntz, L., and S. Scholtes. 1995. "Qualitative Aspects of the Local Approximation of a Piecewise Differentiable Function." *Nonlinear Analysis, Theory, Methods and Applications* 25: 197–215.

Lasserre, Jean B. 1998. "Integration on a Convex Polytope." *Proceedings of the American Mathematical Society* 126 (8): 2433-2441.

Martinez, Jose Mario, and Liqun Qi. 1995. "Inexact Newton methods for solving nonsmooth equations." *Journal of Computational and Applied Mathematics* 60: 127-145.

Mas-Colell, A. 1985. *The Theory of General Economic Equilibrium: A Differentiable Approach*. Econometric Society Monographs, no. 9 Cambridge, Cambridge University Press.

Miranda, Mario J., and Paul L. Fackler. 2002. *Applied Computational Economics and Finance*. The MIT Press.

Pang, Jong-Shi. 1990. "Newton's Method for B-Differentiable Equations." *Mathematics of Operations Research* 15 (2): 311-341.

Pang, Jong-Shi. 1991. "A B-differentiable equation-based, globally and locally quadratically convergent algorithm for nonlinear programs, complementarity and variational inequality problems." *Mathematical Programming* 51: 101-131.

Penn, Elizabeth Maggie. 2009. "A Model of Farsighted Voting." *American Journal of Political Science* 53 (1): 36-54.

Plott, C. R. 1967. "A notion of equilibrium and its possibility under majority rule." *American Economic Review* 57: 787-806.

Press, William H., Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. 1992. *Numerical Recipes In C: The Art Of Scientific Computing*. Second ed. Cambridge University Press.

Qi, Liqun. 1993. "Convergence Analysis of Some Algorithms for Solving Nonsmooth Equations." *Mathematics of Operations Research* 18 (1): 227-244.

Qi, Liqun. 1997. "On superlinear convergence of quasi-Newton methods for nonsmooth equations." *Operations Research Letters* 20: 223-228.

Qi, Liqun, Alexander Shapiro, and Chen Ling. 2005. "Differentiability and semismoothness properties of integral functions and their applications." *Mathematical Programming, Series A* 102: 223-248.

Qi, Liqun, and J. Sun. 1993. "A Nonsmooth Version of Newton's Method." *Mathematical Programming* 58 (3): 353-367.

Rivlin, T. J. 1990. *Chebyshev Polynomials: From Approximation Theory to Algebra and Number Theory*. NewYork: Wiley Interscience.

Rubinstein, A. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50: 97–110.

Xu, H., and X. W. Chang. 1997. "Approximate Newton Methods for Nonsmooth Equations." *Journal of optimization theory and applications* 93 (2): 373-394.