Notes on Spatial Bargaining and Stochastic Games

John Duggan

February 24, 2016

Contents

The material in these notes on the theory of one-shot bargaining is mainly drawn from joint work of mine with Jeff Banks. It is somewhat specialized in that it is focussed on spatial models of bargaining, where we assume that utilities are strictly quasi-concave. This precludes applications with a transferable private good, but the results extend quite generally to such environments; I refer the reader to our papers for details. The main topics are existence of equilibrium, the possibility of delay in equilibrium, and connections to the core of the voting rule. Detailed characterizations are available for the one-dimensional model, so that special case receives emphasis. An advantage of the notes is that they give a synthetic treatment of the bad status quo model and the general status quo (with common discount factor) model, and key theorems on existence and delay rely on fundamental results for the setter model of Section 2. Another advantage is that I allow agents to use mixed voting strategies, rather than assume pure voting strategies, as in my work with Jeff. This flexibility is not needed for existence in the basic model, but it strengthens results precluding equilibria with delay, as we remove the possibility that there could exist equilibria in which some proposals are rejected with positive probability stemming from mixing by indifferent voters. Later material on dynamic bargaining with endogenous status quo, where today's choice of an alternative determines a status quo tomorrow, is drawn from work with Tasos Kalandrakis. An intervening section on stochastic games delves into issues with and results on existence of equilibria in stochastic games. It contains an existence proof for correlated stationary Markov perfect equilibria in general discounted stochastic games that relies on the theory of transition probabilities; to my knowledge, the proof is novel. To begin, however, we provide background on metric spaces, measure theory, transition probabilities, and the theory of correspondences.

1 Mathematical Background

Some familiarity with real analysis in finite-dimensional Euclidean space, $\mathbb{R}^n$, and a bit of linear algebra are assumed.

1.1 Metric Spaces

Metric space basics. Given a set X, we may define a concept of distance between any two elements in the form of a *metric,* which is a function $\rho\colon X \times X \to \mathbb{R}$ satisfying three properties:

(i) for all $x, y \in X$, $\rho(x, y) \geq 0$, where equality holds if and only if $x = y$,

(ii) for all $x, y \in X$, $\rho(x, y) = \rho(y, x)$,

(iii) for all $x, y, z \in X$, $\rho(x, z) \leq \rho(x, y) + \rho(y, z)$.

Property (iii) is the metric version of the "triangle inequality." A *metric space* is a set X with a metric ρ. A familiar metric space is Euclidean space with the *Euclidean metric,* defined by $\rho^e(x, y) = ||x - y||$. In fact, any subset $X \subseteq \mathbb{R}^n$ is a metric space in its own right when equipped with the Euclidean metric. For a different (but trivial) example, an arbitrary set X is a metric space with the *discrete metric,* which is defined as $\rho^d(x, y) = 0$ if $x = y$ and $\rho^d(x, y) = 1$ otherwise. For a more interesting example, we could consider a set of bounded functions defined on $\mathbb{R}^n$, in which case a common measure of distance is given by the the *supremum metric,* defined by $\rho^s(f, g) = \sup\{|f(x) - g(x)| \mid x \in \mathbb{R}^n\}$.

Though spare, the structure of a metric space permits us to make topological distinctions along the usual lines. The ball of radius r around $x \in X$ is $B_r(x) = \{y \in X \mid \rho(x, y) < r\}$, and then we define the boundary, interior, and closure of a set Y as usual, using the notation $\text{bd}(Y)$, $\text{int}(Y)$, and $\text{clos}(Y)$. For example, given $Y \subseteq X$, we say x is a *boundary point* of Y if for all $\epsilon > 0$, we have $B_\epsilon(x) \cap Y \neq \emptyset$ and $B_\epsilon(x) \setminus Y \neq \emptyset$. And we define open and closed sets following standard conventions: a set $Y \subseteq X$ is *open* if it is disjoint from its boundary, or equivalently, if for all $x \in Y$, there is some $\epsilon > 0$ such that $B_\epsilon(x) \subseteq Y$; and Y is *closed* if it contains its boundary. As usual, a set $Y \subseteq X$ is closed if and only if its complement $X \setminus Y$ is open; and Y is open if and only if $X \setminus Y$ is closed. Note that $\emptyset$ and X are both open; arbitrary unions of open sets are open; and finite intersections of open sets are open. Furthermore, $\emptyset$ and X are both closed; finite unions of closed sets are closed; and arbitrary intersections of closed sets are closed.

A sequence $\{x_m\}$ in a metric space X is a countably infinite list of elements indexed by the natural numbers $m \in \mathbb{N}$. A subsequence $\{x_{m_k}\}$ is a selection from this list, where $\{m_k\}$ is a strictly increasing sequence of natural numbers indexed by k. A sequence $\{x_n\}$ in X *converges* to $x \in X$ if for all $\epsilon > 0$, there exists n such that for all $m \geq n$, $x_m \in B_\epsilon(x)$; equivalently, if $\rho(x_n, x) \to 0$. Of course, a set $Y \subseteq X$ is closed if and only if every convergent sequence in Y has limit in Y; and as in Euclidean space, $\{x_m\}$ converges to x if and only if every subsequence converges to this same limit.

A set $Y \subseteq X$ is *bounded* if it is contained in a ball $B_r(x)$ for some $r > 0$ and some $x \in X$, in which case $\text{diam}(Y) = \sup\{\rho(x, y) \mid x, y \in Y\}$ is the *diameter* of Y. A set $Y \subseteq X$ is *compact* if every sequence in Y has a subsequence that converges to an element of Y; equivalently, Y is compact if and only if every open cover of Y has a finite subcover; and Y is compact if and only if every collection of closed subsets of Y with the finite intersection property has nonempty intersection. The empty set and all finite sets are compact; finite unions of compact sets are compact; and arbitrary intersections of compact sets are compact. Every compact set is closed and bounded, and the converse is true in $\mathbb{R}^n$; it does not hold for general metric spaces. A set $Y \subseteq X$ is *connected* if there do not exist open sets $U, V \subseteq X$ such that $U \cap V = \emptyset$, $Y \subseteq U \cup V$, $U \cap Y \neq \emptyset$, and $V \cap Y \neq \emptyset$.

In addition to compactness and connectedness, we can consider several more esoteric properties of metric spaces. A metric space X is *complete* if every Cauchy sequence in X converges to an element $x \in X$, where a sequence $\{x_n\}$ is *Cauchy* if for all $\epsilon > 0$, there exists m such that for all $k, \ell \geq m$, we have $\rho(x_k, x_\ell) < \epsilon$. The metric space X is *totally bounded* if for all $\epsilon > 0$, there exist n and elements $x_1, \ldots, x_n \in X$ such that $X \subseteq \bigcup_{i=1}^n B_\epsilon(x_i)$. The metric space X is *separable* if there is a countable subset $\{x_1, x_2, \ldots\} \subseteq X$ that is *dense,* i.e., for all $x \in X$ and all $\epsilon > 0$, there exists n satisfying $x_n \in B_\epsilon(x)$. And X is *locally compact* if for each x, there exists $r > 0$ such that the disc $D_r(x) = \{y \in X \mid \rho(x, y) \leq r\}$ is compact. Every totally bounded metric space is bounded and separable. Every locally compact metric space is complete. A metric space X is compact if and only if it is both complete and totally bounded. Every subset of a separable metric space is itself a separable metric space, and every closed subset of a complete metric space is itself a complete metric space. Euclidean space is complete, locally compact, and separable; the latter property follows, for example, because the subset of vectors with rational coordinates is countable and dense. In fact, if $X \subseteq \mathbb{R}^n$ is equipped with the Euclidean metric, then X is automatically separable, and it is complete as long as X is a closed subset of $\mathbb{R}^n$.

A real-valued function $f\colon X \to \mathbb{R}$ defined on a metric space X is *continuous* if for every sequence $\{x_n\}$ converging to x in X, we have $f(x_n) \to f(x)$; equivalently, if for every open $G \subseteq \mathbb{R}$, the preimage $f^{-1}(G)$ is open. More generally, given metric spaces X and Y, a function $f\colon X \to Y$ is *continuous* if for every sequence $\{x_n\}$ converging to x in X, we have $f(x_n) \to f(x)$ in Y; equivalently, f is continuous if for every open $G \subseteq Y$, the preimage $f^{-1}(G)$ is open. We can combine continuous functions in the usual way to produce new continuous functions, e.g., letting X, Y, and Z be any metric spaces, if $f\colon X \to Y$ and $g\colon Y \to Z$ are continuous with $f(X) \subseteq Y$, then $g \circ f$ is continuous; and if $f, g\colon X \to \mathbb{R}^n$ are continuous and $\alpha, \beta \in \mathbb{R}$, then $\alpha f + \beta g$ is continuous. By *Weierstrass' theorem,* the image of a compact set under a continuous function is compact; and by the *intermediate value theorem,* the image of a connected set under a continuous function is connected.

Given a metric space X and a function $f\colon X \to X$, an element x is a *fixed point* of f if $f(x) = x$. A function $f\colon X \to X$ is a *contraction mapping* on X if there exists $\beta \in [0, 1)$ such that for all $x, y \in X$, we have $\rho(f(x), f(y)) \leq \beta\rho(x, y)$, in which case β is the *modulus* of the contraction. By the *contraction mapping theorem,* if X is a nonempty, complete metric space and $f\colon X \to X$ is a contraction mapping, then f has a unique fixed point.

Metric mixture spaces. A *mixture space* is a set X for which a mapping $\xi\colon X \times X \times [0, 1] \to X$ is defined, where we write $\alpha x + (1-\alpha)y$ for $\xi(x, y, \alpha)$, and possesses certain intuitive properties; intuitively, the mapping ξ acts just like the usual idea of convex combination of vectors in Euclidean space. Since convex combinations are independent of order, we can write $\sum_{i=1}^{m} \alpha_i x_i$ without ambiguity.[1] Following the analysis of $\mathbb{R}^n$, we say x is a *convex combination* of m elements $x_1, \ldots, x_m$ if there exist non-negative coefficients $\alpha_1, \ldots, \alpha_m$ that sum to one and such that $x = \sum_{i=1}^{m} \alpha_i x_i$. A subset $Y \subseteq X$ of a metric mixture space is *convex* if for all $x, y \in Y$ and all $\alpha \in (0, 1)$, we have $\alpha x + (1 - \alpha)y$. And given any $A \subseteq X$, the *convex hull* of A is

$$\operatorname{conv}(Y) \quad = \quad \left\{ \sum_{i \in I} \alpha_i x_i \;\middle|\; \begin{array}{c} I \subseteq \mathbb{N} \text{ is finite, } x_i \in Y \text{ for} \\ \text{all } i \in I,\ \alpha_i \geq 0 \text{ for all } i \in I, \\ \text{and } \sum_{i \in I} \alpha_i = 1 \end{array} \right\},$$

which consists of all convex combinations of all finite subsets of Y. A set A is convex if and only if $A = \operatorname{conv}(A)$, so that the convex hull of a set is itself convex, and in fact it is the intersection of all convex supersets of A.

The set X is a *metric mixture space* if X is a metric space (with metric ρ), and the mapping ξ is continuous with respect to (x, y, α), i.e., for all (x, y, α)

[1] To be more precise, it must be that for all $x, y, z \in X$ and all $\alpha, \beta, \gamma \in [0, 1]$ with $\alpha + \beta + \gamma = 1$, we have (i) $1x + 0y = x$, (ii) $\alpha x + (1 - \alpha)y = (1 - \alpha)y + \alpha x$, and

$$\text{(iii) } \alpha x + (1 - \alpha)\left(\frac{\beta}{1 - \alpha} y + \frac{\gamma}{1 - \alpha} z\right) \;=\; \beta y + (1 - \beta)\left(\frac{\alpha}{1 - \beta} x + \frac{\gamma}{1 - \beta} z\right).$$

and all sequences $\{(x_m, y_m, \alpha_m)\}$ with $x_m \to x$, $y_m \to y$, and $\alpha_m \to \alpha$, we have $\rho(\alpha_m x_m + (1-\alpha_m)y_m, \alpha x + (1-\alpha)y) \to 0$. The metric ρ is *quasi-convex* if for all $x, y, z, w \in X$ and all $\alpha \in [0, 1]$,

$$\rho(\alpha x + (1-\alpha)z, \alpha y + (1-\alpha)w) \quad \leq \quad \max\{\rho(x, y), \rho(z, w)\}.$$

Obviously, a sufficient condition for quasi-convexity of ρ is that it is *convex,* i.e., for all $x, y, z, w \in X$ and all $\alpha \in [0, 1]$,

$$\rho(\alpha x + (1-\alpha)z, \alpha y + (1-\alpha)w) \quad \leq \quad \alpha\rho(x, y) + (1-\alpha)\rho(z, w).$$

In a metric mixture space, the convex hull of a finite set is compact; and assuming X is complete, the closure of the convex hull of each compact subset is compact. In contrast to $\mathbb{R}^n$, the latter result is not true if stated without taking the closure of the convex hull: in general metric spaces, there are examples of compact sets with non-compact convex hulls.

A metric version of *Schauder's fixed point theorem* states that if X is a nonempty, compact metric mixture space with quasi-convex metric, and if $f\colon X \to X$ is continuous, then f has at least one fixed point. In Subection 1.4, I state a metric version of Glicksberg's fixed point theorem, which generalizes Schauder's theorem to correspondences.

Countable Cartesian products. Let X_i be a metric space with metric ρ^i for $i = 1, \ldots, k$, and let $X = \prod_{i=1}^{k} X_i$ be the product of these sets. We can define the *product metric* ρ^π on the space X as follows: for all $x, y \in X$,

$$\rho^\pi(x, y) \quad = \quad \sum_{i=1}^{k} \rho^i(x_i, y_i),$$

where we use the convention that $x = (x_1, \ldots, x_k)$ and $y = (y_1, \ldots, y_k)$. With this metric, a sequence $\{x^m\}$ in X converges to $x \in X$ if and only if it converges in each component, i.e., we have $x_i^m \to x_i$ for each $i = 1, \ldots, k$. If each X_i is separable, then X is separable; if each X_i is complete, then X is complete; and by *Tychonoff's product theorem,* if each X_i is compact, then X is compact.

It is straightforward to extend this idea to the product of a countably infinite number of metric spaces X_i with metrics ρ^i, $i = 1, 2, \ldots$. Then the Cartesian product $X = \prod_{i=1}^{\infty} X_i$ consists of sequences $x = (x_1, x_2, \ldots)$, where each component x_i belongs to X_i. Assume for now that each X_i is bounded, so there is some r_i such that $\text{diam}(X_i) < r_i$. Define a metric on X as follows: for each $x = (x_1, x_2, \ldots)$ and $y = (y_1, y_2, \ldots)$ in X, let

$$\rho^\pi(x, y) = \sum_{k=1}^{\infty} \frac{1}{2^k r_k} \rho^k(x_k, y_k).$$

Note that given any k and any $x_k, y_k \in X_k$, we have $\rho^k(x_k, y_k) \leq r_k$, so $\rho^\pi(x, y) \leq \sum_{k=1}^{\infty} \frac{1}{2^k} < \infty$, so this metric is well-defined. With this metric, a sequence $\{x^m\}$ converges to x in X if and only if it converges in each component, i.e., letting $x^m = (x_1^m, x_2^m, \ldots)$ and $x = (x_1, x_2, \ldots)$, we have $x_i^m \to x_i$ for each $i = 1, 2, \ldots$. Again, if each X_i is separable, then X is separable; if each X_i is complete, then X is complete; and by *Tychonoff's product theorem,* if each X_i is compact, then X is compact.

If each X_i is a metric mixture space, then X is a metric mixture space with the product metric. Indeed, given any $x, y \in X$ and any $\alpha \in (0, 1)$, the convex combination $\alpha x + (1-\alpha)y = (\alpha x^1 + (1-\alpha)y^1, \alpha x^2 + (1-\alpha)y^2, \ldots)$ belongs to X; and given sequences $\{x^m\}$ converging to x in X, $\{y^m\}$ converging to y in X, and $\{\alpha_m\}$ converging to α in $[0, 1]$, we have $\alpha_m x^m + (1-\alpha_m)y^m \to \alpha x + (1-\alpha)y$ in the product metric. Moreover, if each X_i is bounded and each ρ^i is convex, then the product metric is convex: for all $x, y, z, w \in X$ and all $\alpha \in (0, 1)$, we have

$$\begin{aligned}
&\rho^\pi(\alpha x + (1-\alpha)y), \alpha z + (1-\alpha)w) \\
&= \sum_{k=1}^{\infty} \frac{1}{2^k r_k} \rho^k(\alpha x_k + (1-\alpha)y_k), \alpha z_k + (1-\alpha)w_k) \\
&\leq \sum_{k=1}^{\infty} \frac{1}{2^k r_k} [\alpha \rho_k(x_k, y_k) + (1-\alpha)\rho_k(z_k, w_k)] \\
&= \alpha \rho^\pi(x, z) + (1-\alpha)\rho^\pi(y, w),
\end{aligned}$$

as required.

The above choice of metric for the infinite product space relied on the assumption that all (or all but finitely many) component sets X_i are bounded. In general, when some X_i are not bounded, we can define a version of the product metric as follows:

$$\rho^\pi(x, y) = \sum_{k=1}^{\infty} \frac{1}{2^k} \cdot \frac{\rho^k(x_k, y_k)}{1 + \rho^k(x_k, y_k)}.$$

With this metric, a sequence $\{x_m\}$ again converges to x in X if and only if it converges in each component, i.e., we have $x_i^m \to x_i$ for each $i = 1, 2, \ldots$. Again, if each X_i is separable, then X is separable; if each X_i is complete, then X is complete; and by *Tychonoff's product theorem,* if each X_i is compact, then X is compact. In contrast to the initial definition of the product metric, however, the latter definition may not inherit convexity of its components.

1.2 Measure Theory

Throughout this subsection, we fix an arbitrary metric space X.

Borel measure. A *σ-algebra* is a collection $\mathcal{A}$ of subsets of X such that:

- the sets $\emptyset$ and X belong to $\mathcal{A}$,
- the collection $\mathcal{A}$ is closed with respect to complements, i.e., if $Y \in \mathcal{A}$, then $\overline{Y} \in \mathcal{A}$,
- the collection $\mathcal{A}$ is closed with respect to countable unions, i.e., if $\{Y_1, Y_2, \ldots\} \subseteq \mathcal{A}$, then $\bigcup_{i=1}^{\infty} Y_i \in \mathcal{A}$.

An implication, via De Morgan's law, is that a σ-algebra is closed with respect to countable intersection, i.e., if $\{Y_1, Y_2, \ldots\} \subseteq \mathcal{A}$, then $\bigcap_{i=1}^{\infty} Y_i \in \mathcal{A}$.

There is a unique smallest σ-algebra that contains the open subsets of X, and this is the *Borel σ-algebra,* denoted $\mathcal{B}$ (or when X is to be explicit, $\mathcal{B}_X$). Thus, if $\mathcal{A}$ is a σ-algebra that contains all of the open subsets of X, then $\mathcal{B} \subseteq \mathcal{A}$. A set $Y \in \mathcal{B}$ is called *Borel measurable.* Obviously, all open sets are Borel measurable, as are all closed sets. But many more sets are Borel measurable, e.g., every set obtained as the result of a countable number of intersections and/or unions of open and/or closed sets. Henceforth, "measurable" will mean "Borel measurable."

A function $\mu\colon \mathcal{B} \to \mathbb{R}_+^*$ mapping Borel measurable sets to the non-negative extended real numbers (including ∞) is a *Borel measure on* X if it satisfies the following:

(i) $\mu(\emptyset) = 0$,

(ii) for all pairwise disjoint collections $\{Y_1, Y_2, \ldots\}$ of measurable sets,

$$\mu\left(\bigcup_{i=1}^{\infty} Y_i\right) = \sum_{i=1}^{\infty} \mu(Y_i).$$

Condition (ii) is known as *countable additivity.* One example of a Borel measure is Lebesgue measure (restricted to the Borel σ-algebra), which is denoted λ and formalizes the intuitive idea of length or volume and will not be defined precisely here. Another simple example is the *counting measure*, γ, defined as $\gamma(Y) = |Y|$ if Y is finite and $\gamma(Y) = \infty$ otherwise. Henceforth, "measure" will refer to "Borel measure."

A particularly interesting measurable subset of the unit interval is the *Cantor set,* denoted C, defined as follows. Let $C_0 = [0, 1]$; then define $C_1 = C_0 \setminus (1/3, 2/3)$ by removing the "middle third" of C_0, leaving the union of two disjoint intervals; then define C_2 by removing the middle thirds of the intervals $[0, 1/3]$ and $[2/3, 1]$, leaving the union of four disjoint intervals; then define C_3 by removing their middle thirds, and so on. In general, for $n \geq 2$, define $C_n = C_{n-1} \setminus D_n$, where

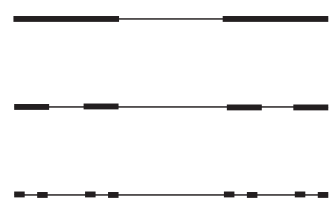

$D_n = \bigcup\{[\frac{k}{3^n}, \frac{k+1}{3^n}] \mid k = 1, \ldots, 3^n - 2, \text{ odd}\}$ is the union of alternating subintervals of length $1/3^n$. Then the Cantor set is $C = \bigcap_{i=1}^{\infty} C_i$. This set has the following properties:

(i) it is closed,

(ii) it is *nowhere dense,* i.e., its closure contains no open set,

(iii) it has Lebesgue measure zero,

(iv) it has the cardinality of the continuum.

In particular, there exist uncountable sets with Lebesgue measure zero. There are subsets of the unit interval that are not measurable, but the proof of this claim is non-constructive: the existence of such sets relies on the axiom of choice.

Given a measure μ, if $\mu(X) < \infty$, then μ is *finite.* If $\mu(X) = 1$, then μ is a *probability measure.* If there is a countable collection $\{Y_1, Y_2, \ldots\}$ of measurable sets such that $X = \bigcup_{i=1}^{\infty} Y_i$ and $\mu(Y_i) < \infty$ for all $i = 1, 2, \ldots$, then μ is *σ-finite.* Lebesgue measure is not finite, but it is σ-finite. If a measurable set Z satisfies $\mu(Z) = 0$, then it is *μ-measure zero,* or if μ is understood, simply *measure zero.* And if a property holds at all x outside a μ-measure zero set Z, then it holds "for μ-almost all x." Given measurable

$Y \subseteq X$, a property holds "for μ-almost every $x \in Y$" if it holds for all $x \in Y \setminus Z$.

The *Carathéodory extension theorem* establishes several facts about measures:

1. If Y and Z are measurable and $Y \subseteq Z$, then $\mu(Y) \subseteq \mu(Z)$.
2. If Y and Z are measurable and $\mu(Y) < \infty$, then $\mu(Y \setminus Z) = \mu(Y) - \mu(Y \cap Z)$.
3. If $Y_1, Y_2, \ldots$ are measurable, then

$$\mu\left(\bigcup_{i=1}^{\infty} Y_i\right) \leq \sum_{i=1}^{\infty} \mu(Y_i).$$

4. If $Y_1, Y_2, \ldots$ are measurable and $\mu(Y_j) < \infty$ for some $j = 1, 2, \ldots$, then

$$\mu\left(\bigcap_{i=1}^{\infty} Y_i\right) \geq \mu(Y_j) - \sum_{i=1}^{\infty} \mu(Y_i).$$

The property in condition 3 is known as *countable sub-additivity.*

For every measure μ, there is a unique closed set, called the *support* of μ and denoted $\mathrm{supp}(\mu)$, such that

(i) $\mu(X \setminus \mathrm{supp}(\mu)) = 0$

(ii) for all open $G \subseteq X$ with $G \cap \mathrm{supp}(\mu) \neq \emptyset$, we have $\mu(G \cap \mathrm{supp}(\mu)) > 0$.

If μ is finite, then the support of μ is the intersection of all closed sets F with $\mu(F) = \mu(X)$.

Given a measure μ, an element $x \in X$ is an *atom* of μ if $\mu(\{x\}) > 0$, and μ is *atomless* (or *non-atomic*) if it admits no atoms. If there is an atom $x \in X$ such that $\mu(X \setminus \{x\}) = 0$, then μ is *degenerate* on x. By *Lyapunov's theorem,* the range of a vector of atomless measures is convex, i.e., given any m and any atomless measures $\mu_1, \ldots, \mu_m$, the set

$$\{(\mu_1(Y), \ldots, \mu_m(Y)) \mid Y \text{ is measurable } \}$$

is convex. Of course, if $m = 1$ and μ is a probability measure, then the range is the unit interval $[0, 1]$. We can actually extend the idea of measure to a *vector measure,* which is a mapping μ which takes measurable sets Y to vectors $\mu(Y) = (\mu_1(Y), \ldots, \mu_m(Y))$ such that each μ_i is a measure, $i = 1, \ldots, m$. It is *atomless* if each μ_i is atomless. Then Lyapunov's theorem states that the range of an atomless vector measure is convex.

Measurable functions. A function $f\colon X \to \mathbb{R}$ is *Borel measurable* if for all $c \in \mathbb{R}$, $\{x \in X | f(x) \geq c\}$ is measurable. The following are equivalent definitions:

- $\{x \in X | f(x) > c\}$ is measurable for all $c \in \mathbb{R}$,
- $\{x \in X | f(x) \leq c\}$ is measurable for all $c \in \mathbb{R}$,
- $\{x \in X | f(x) < c\}$ is measurable for all $c \in \mathbb{R}$,
- $f^{-1}(G)$ is measurable for all open $G \subseteq \mathbb{R}$.

Implications are that level sets of a measurable function are measurable, and every continuous function is measurable. More generally, given a metric space Y and letting $\mathcal{B}_X$ and $\mathcal{B}_Y$ be the Borel σ-algebras, a function $f\colon X \to Y$ is *measurable* if for every open $G \subseteq Y$, we have $f^{-1}(G) \in \mathcal{B}_X$. Henceforth, "measurable function" will mean a Borel measurable function.

If $f\colon X \to Y$ is measurable, then in fact the pre-image of every measurable set is measurable, i.e., for all $B \in \mathcal{B}_Y$, we have $f^{-1}(B) \in \mathcal{B}_X$. An implication is that given metric spaces X, Y, and Z and two measurable mappings $f\colon X \to Y$ and $g\colon Y \to Z$ satisfying $f(X) \subseteq Y$, the composition $g \circ f\colon X \to Z$ is measurable.

Measurable functions can be combined in other ways to produce new measurable functions. Let $f, g\colon X \to \mathbb{R}$ be measurable, let $\{f_m\}$ be a sequence of measurable functions $f_m\colon X \to \mathbb{R}$, $m = 1, 2, \ldots$, and let $\alpha, \beta \in \mathbb{R}$. Then:

1. $\alpha f + \beta g$ is measurable,
2. fg is measurable,
3. $\max\{f(x), g(x)\}$ is measurable as a function of x,
4. if $h\colon X \to \mathbb{R}$ satisfies $f_m(x) \to h(x)$ for all $x \in X$, then h is measurable.

In particular, defining $f^+\colon X \to \mathbb{R}$ by $f^+(x) = \max\{f(x), 0\}$ and $f^-\colon X \to \mathbb{R}$ by $f^-(x) = \min\{f(x), 0\}$, both f^+ and f^- are measurable. In words, property 4 means that pointwise limits of measurable functions are measurable. This continues to hold if f is the pointwise limit of $\{f_m\}$ almost everywhere, i.e., if $f_m(x) \to f(x)$ for μ-almost all x, then f is measurable. Furthermore, this implies that if $\{f_m(x)\}$ is increasing and bounded outside a measurable set $Z \subseteq X$ with μ-measure zero, then the function $h\colon X \to \mathbb{R}$ defined by $h(x) = \sup f_m(x)$ for $x \in X \setminus Z$, and $h(x) = 0$ otherwise, is measurable.

Given measure μ and a non-negative, measurable function $f\colon X \to \mathbb{R}$, *integral of f with respect to μ,* denoted $\int f(x)\mu(dx)$, is

$$\lim_{m\to\infty} \sum_{k=1}^{m^2} \frac{k-1}{m} \mu\left(\left\{x \in X \,\middle|\, \frac{k-1}{m} \le f(x) < \frac{k}{m}\right\}\right),$$

which may be infinite. If $\int f(x)\mu(dx) < \infty$, then f is *integrable with respect to μ* (or *μ-integrable*). See Figure 1. For general measurable $f\colon X \to \mathbb{R}$, if either f^+ or f^- are integrable with respect to μ, then the integral of f with respect to μ is

$$\int f(x)\mu(dx) \;=\; \int f^+(x)\mu(dx) - \int f^-(x)\mu(dx).$$

If both f^+ or f^- are μ-integrable, then f is *μ-integrable.* Say f is *essentially bounded* with respect to μ (or *μ-bounded*) if there exists $c > 0$ such that $\mu(\{x \in X \mid f(x) > c\}) = 0$. If μ is finite and f is μ-bounded, then f is μ-integrable. We write $\int_Y f(x)\mu(dx)$ for $\int f(x) I_{X\cap Y}(x)\mu(dx)$. If μ is a probability measure, then the integral $\int f(x)\mu(dx)$ is called the *expected value* of f, and the vector of integrals $(\int x_1\mu(dx), \ldots, \int x_n\mu(dx))$ is the *expected value* (or *mean*) of x.

Let μ be a measure, and let $f, g\colon X \to \mathbb{R}$ be measurable. Then:

1. if f is non-negative, then $\int f(x)\mu(dx) = 0$ if and only if $f(x) = 0$ for μ-almost all x,

2. if either f or g is μ-integrable, then $\int [f(x)+g(x)]\mu(dx) = \int f(x)\mu(dx) + \int g(x)\mu(dx)$,

3. for all $\alpha \in \mathbb{R}$, $\int \alpha f(x)\mu(dx) = \alpha \int f(x)\mu(dx)$,

4. if $f(x) \le g(x)$ for μ-almost all $x \in X$, then $\int f(x)\mu(dx) \le \int g(x)\mu(dx)$.

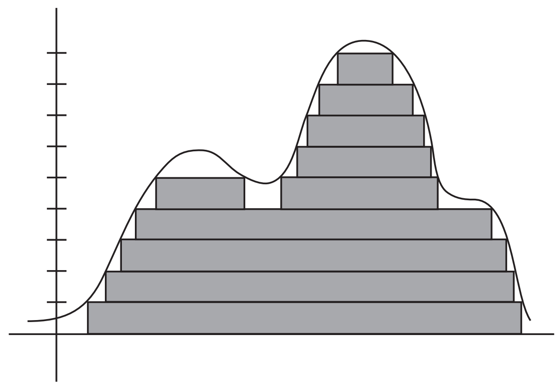

Figure 1: Lebesgue integral

Note that property 2 implies that if A and B are disjoint measurable sets, then

$$\int_{A\cup B} f(x)\mu(dx) \quad = \quad \int_A f(x)\mu(dx) + \int_B f(x)\mu(dx).$$

And properties 3 and 4 imply that $|\int f(x)\mu(dx)| \leq \int |f(x)|\mu(dx)$.

Properties 2 and 3 imply that the integral is linear in the integrand f, i.e., the integral of a linear combination of functions is the linear combination of integrals. We can also state a form of linearity in terms of the integrating measure. Given measures κ and μ on X and $\alpha, \beta \in \mathbb{R}_+$, we can define the new measure ν on X by the formula $\nu(Y) = \alpha\kappa(Y) + \beta\mu(Y)$ for all measurable $Y \subseteq$. Thus, ν is a non-negative linear combination of the first two measures. For all measurable $f\colon X \to \mathbb{R}$, if f is κ- and μ-integrable, then it is ν-integrable, and

$$\int_X f(x)\nu(dx) \quad = \quad \alpha\int_X f(x)\kappa(dx) + \beta\int_X f(x)\mu(dx),$$

so the integral with respect to a non-negative linear combination of measures is equal to the same linear combination of integrals.

Let $\{f_m\}$ be a sequence of μ-integrable functions $f_m\colon X \to \mathbb{R}$, $m = 1, 2, \ldots$, and suppose the functions are *dominated* by a μ-integrable function $g\colon X \to \mathbb{R}$, in the sense that $g(x) \geq |f_m(x)|$ for all m and μ-almost all $x \in X$. Then by *Fatou's lemma,* we have

$$\int \limsup_{m\to\infty} f_m(x)\mu(dx) \quad \geq \quad \limsup_{m\to\infty} \int f_m(x)\mu(dx),$$

and ignoring the subset of the domain on which it takes infinite values, the integrand on the left-hand side, $\limsup f_m(x)$, is a μ-integrable function of x. More precisely, if we define $f(x) = 0$ when $\limsup_{m\to\infty} f_m(x) = \infty$ and $f(x) = \limsup_{m\to\infty} f_m(x)$ otherwise, then the function f is integrable. The assumption that $\{f_m\}$ is dominated by integrable g can be weakened somewhat: it suffices to assume that each f_m is dominated by a μ-integrable $g_m\colon X \to \mathbb{R}$, and that there is a μ-integrable function $g\colon X \to \mathbb{R}$ satisfying $g_m(x) \to g(x)$ for μ-almost all $x \in X$ and $\int g_m(x)\mu(x) \to \int g(x)\mu(dx)$.

An easy consequence of Fatou's lemma is the following. Let $\{f_m\}$ be a sequence of μ-integrable functions $f_m\colon X \to \mathbb{R}$, $m = 1, 2, \ldots$. Assume that for μ-almost all $x \in X$, the sequence $\{f_m(x)\}$ is increasing, and assume that $\lim_{m\to\infty} \int f_m(x)\mu(dx) < \infty$. Then *Levi's monotone convergence theorem,* states that there exists a μ-integrable function $f\colon X \to \mathbb{R}$ such that $f_m(x) \uparrow f(x)$ for μ-almost all $x \in X$ and $\int f_m(x)\mu(dx) \to \int f(x)\mu(dx)$.

Again let $\{f_m\}$ be a sequence of μ-integrable functions $f_m\colon X \to \mathbb{R}$, and suppose that they are dominated by a μ-integrable function $g\colon X \to \mathbb{R}$. If $f\colon X \to \mathbb{R}$ satisfies $f_m(x) \to f(x)$ for μ-almost all $x \in X$, then *Lebesgue's dominated convergence theorem* states that f is μ-integrable and

$$\int f(x)\mu(dx) \quad = \quad \lim_{m\to\infty} \int f_m(x)\mu(dx).$$

Thus, integrals are (modulo domination by a μ-integrable function) continuous with respect to pointwise limits of μ-integrable functions.

Given measures μ and κ, a *density function* for μ with respect to κ is any measurable function $f\colon X \to \mathbb{R}$ with non-negative values such that for all measurable $Y \subseteq X$, we have $\mu(Y) = \int_Y f(x)\kappa(dx)$. The measure μ is *absolutely continuous* with respect to κ, written $\mu \ll \kappa$, if for all measurable Y, $\kappa(Y) = 0$ implies $\mu(Y) = 0$. For example, any measure defined by a density function on $\mathbb{R}^n$ is absolutely continuous with respect to Lebesgue measure; for a counterexample, the counting measure is not absolutely continuous with respect to Lebesgue measure. The *Radon-Nikodym theorem* states that if a measure μ is finite and absolutely continuous with respect to

κ, and if κ is σ-finite, then μ has a density with respect to κ that is unique up to sets of κ-measure zero, i.e., if f and g are both densities for μ with respect to κ, then $f(x) = g(x)$ for κ-almost all x. Conversely, starting with an arbitrary measurable mapping $f \colon X \to \mathbb{R}$ with non-negative values and any measure κ, we can define a mapping from measurable subsets $Y \subseteq X$ to the extended real numbers (possibly including infinity) by $\mu(Y) = \int_Y f(x)\kappa(dx)$. Then μ, so-defined, is a measure; it is absolutely continuous with respect to κ; and f is a density for μ with respect to κ.

Given a probability measure μ, the expected value of a concave function f is less than or equal to the value of the function at the mean of x. To be more precise, let $X \subseteq \mathbb{R}$ be convex, let μ be a probability measure on $\mathbb{R}$ with $\mu(X) = 1$, and let $f \colon X \to \mathbb{R}$ be concave; then f is measurable (in fact, it is directionally differentiable at μ-almost all $x \in \text{int}(X)$), and *Jensen's inequality* states that

$$f\left(\int x\mu(dx)\right) \quad \geq \quad \int f(x)\mu(dx).$$

Moreover, strict inequality holds if f is strictly concave and μ is not degenerate on some x. More generally, if $X \subseteq \mathbb{R}^n$ is measurable and convex, if μ is a probability measure on $\mathbb{R}^n$, and if $f \colon X \to \mathbb{R}$ is concave, then f is measurable, and

$$f\left(\int x_1\mu(dx), \ldots, \int x_n\mu(dx)\right) \quad \geq \quad \int f(x)\mu(dx),$$

again with strict inequality if f is strictly concave and μ is not degenerate.

We can define integration for vector-valued functions as well. Given a measure μ and a mapping $f \colon X \to \mathbb{R}^m$ with values $f(x) = (f_1(x), \ldots, f_m(x))$, the *integral* of the vector-valued function is just the vector of integrals of its components,

$$\int f(x)\mu(dx) \quad = \quad \left(\int f_1(x)\mu(dx), \ldots, \int f_m(x)\mu(dx)\right),$$

and f is μ*-integrable* if each component f_i is μ-integrable, $i = 1, \ldots, m$. When μ is a probability measure, the *expected value* of f is just $\int f(x)\mu(dx)$, permitting a more compact statement of Jensen's inequality: given measurable and convex $X \subseteq \mathbb{R}^n$ and $f \colon X \to \mathbb{R}$ concave, we have $f(\int x\mu(dx)) \geq \int f(x)\mu(dx)$.

Convergence of probability measures. There are several compelling notions of convergence for a sequence $\{\mu_m\}$ of measures. We say the sequence $\{\mu_m\}$ converges to μ...

- *uniformly set-wise* if for every sequence $\{Y_m\}$ of measurable sets of X, we have $|\mu_m(Y_m) - \mu(Y_m)| \to 0$,
- *set-wise* if for every measurable set $Y \subseteq X$, we have $\mu_m(Y) \to \mu(Y)$,
- *weakly* if for every measurable set $Y \subseteq X$ with $\mu(\mathrm{bd}(Y)) = 0$, we have $\mu_m(Y) \to \mu(Y)$.

Uniform set-wise convergence is usually called convergence *in total variation* (or sometimes *strong* convergence). Set-wise convergence is sometimes (at the risk of confusion) called *strong* convergence. Sometimes weak convergence is referred to as *weak** convergence.

These convergence concepts have equivalent formulations. First, $\mu_m \to \mu$ uniformly set-wise if and only if for all sequences $\{f_m\}$ of measurable functions $f_m\colon X \to \mathbb{R}$ with $\sup\{|f_m(x)| \mid x \in X\} \leq 1$ for all m, we have $|\int f_m(x)\mu_m(dx) - \int f_m(x)\mu(dx)| \to 0$. Second, $\mu_m \to \mu$ set-wise if and only if for all bounded, measurable functions $f\colon X \to \mathbb{R}$, we have $\int f(x)\mu_m(dx) \to \int f(x)\mu(dx)$. And third, $\mu_m \to \mu$ weakly if and only if any of the following conditions hold:

(i) $\int f d\mu_n \to \int f d\mu$ for all bounded, continuous functions $f\colon X \to \mathbb{R}$,

(ii) $\limsup_n \mu_n(F) \leq \mu(F)$ for each closed $F \subseteq X$,

(iii) $\liminf_n \mu_n(G) \geq \mu(G)$ for each open $G \subseteq X$.

In fact, weak convergence is often defined using condition (i).

The above convergence notions are listed in decreasing strength: uniform set-wise convergence implies set-wise convergence, which implies weak convergence. To see that weak convergence does not generally imply set-wise, let $n = 1$ and μ_m be the unit mass on $\frac{1}{m}$, i.e., $\mu_m(\{\frac{1}{m}\}) = 1$ and $\mu_m(\mathbb{R} \setminus \{\frac{1}{m}\}) = 0$. Then $\{\mu_m\}$ converges weakly to the measure μ defined as the unit mass on $\{0\}$, but $\mu_m(\{0\}) = 0$ for all m, so the sequence does not converge set-wise to μ. To see that set-wise convergence does not generally imply uniform set-wise convergence, let $f_1 = 2I_{[0,1/2]}$ be two times the

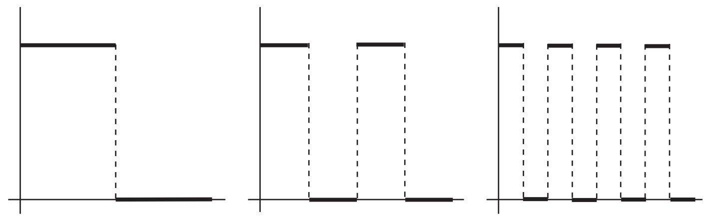

Figure 2: Sawtooth sequence

indicator function of $[0, 1/2]$, let f_2 be two times the indicator function of $[0, 1/4] \cup [2/4, 3/4]$, let f_3 be two times the indicator function of

$$[0, 1/8] \cup [2/8, 3/8] \cup [4/8, 5/8] \cup [6/8, 7, 8],$$

and so on. In general, given m, define intervals $I_m^1, \ldots, I_m^{2^m}$ so that $I_m^k = [(k-1)/2^m, k/2^m]$, and let $f_m \colon [0,1] \to \mathbb{R}$ be two times the indicator function of $J_m = \bigcup\{I_m^k \mid k = 1, \ldots, 2^m - 1, \text{ odd }\}$. Note that the Lebesgue measure of each J_m is one half. As depicted in Figure 2, these functions have an increasingly jagged, "sawtooth" appearance (I will refer to this useful example several times in the sequel). Defining the measure μ_m by $\mu_m(Y) = \int f_m(x)dx$ for all measurable $Y \subseteq \mathbb{R}$, it can be shown that the sequence $\{\mu_m\}$ converges set-wise to the uniform distribution, i.e., the measure μ given by the density $f = I_{[0,1]}$. But then

$$\mu_m(J_m) - \mu(J_m) = \int [f_m(x) - f(x)]dx = \int_{J_m} [2-1]dx = \frac{1}{2}$$

for all m, so the sequence does not converge uniformly set-wise.

Next, we state a *generalized version of Lebesgue's dominated convergence theorem* that allows the integrating probability measure to vary in a weakly continuous way. Let $\{\mu_m\}$ be a sequence of probability measures that converge weakly to the probability measure μ, let F be a set of measurable functions $f \colon X \to \mathbb{R}$ that is *uniformly bounded,* in the sense that $\sup\{f(x) \mid x \in X, f \in F\} < \infty$. Furthermore, assume that for every x, the function $s_x \colon X \to \mathbb{R}$ defined by

$$s_x(y) \quad = \quad \sup\{|f(x) - f(y)| \mid f \in F\}$$

is continuous. This implies that $\int f(x)\mu_m(dx) \to \int f(x)\mu(dx)$ uniformly in f, i.e., for all $\epsilon > 0$, there exists k such that for all $m \geq k$ and all $f \in F$, we have $|\int f(x)\mu_m(dx) - \int f(x)\mu(dx)| < \epsilon$. Now let $\{f_m\}$ be a uniformly bounded sequence of functions $f_m \colon X \to \mathbb{R}$ converging pointwise to the continuous function $f \colon X \to \mathbb{R}$. By the above observation, we then have $\int f_m(x)\mu_m(dx) \to \int f(x)\mu(dx)$.

The preceding formulation of the dominated convergence theorem has the following implication for parameterized integrals that is quite useful in equilibrium existence arguments. Let P be a metric space, and consider any bounded function $f \colon X \times P \to \mathbb{R}$. Assume that for all $x \in X$, the function $f_x \colon P \to \mathbb{R}$ defined by $f_x(p) = f(x,p)$ is continuous; and assume that for all $p \in P$, the function $f_p \colon X \to \mathbb{R}$ defined by $f_p(x) = f(x,p)$ is measurable. Let $\{\mu_m\}$ be a sequence of probability measures converging weakly to μ, and let $\{p_m\}$ be a sequence of parameters converging to p in P such that $f(x,p)$ is continuous in x. Then the parameterized integrals converge: $\int f(x,p_m)\mu_m(dx) \to \int f(x,p)\mu(dx)$.

Product measures. Given probability measures κ and μ on metric spaces Y and Z, respectively, we sometimes want to consider the joint distribution on $Y \times Z$ induced by the separate mixing on the two factors. We give $Y \times Z$ the product metric, and as usual $\mathcal{B}_{Y\times Z}$ is the σ-algebra of Borel sets. In fact, this collection contains all sets of the form $A \times B$, where $A \in \mathcal{B}_Y$ and $B \in \mathcal{B}_Z$, and it is the smallest σ-algebra on $Y \times Z$ that contains these sets. That is, if $\mathcal{A}$ is a σ-algebra on $Y \times Z$ and

$$\{A \times B \mid A \in \mathcal{B}_Y, B \in \mathcal{B}_Z\} \quad \subseteq \quad \mathcal{A},$$

then $\mathcal{B}_{Y\times Z} \subseteq \mathcal{A}$. Thus, we may write $\mathcal{B}_{Y\times Z} = \mathcal{B}_Y \otimes \mathcal{B}_Z$. Given a set $A \times B$, where A is a measurable subset of Y and B is a measurable subset of Z, the probability of $A \times B$ induced by κ and μ should clearly be $\kappa(A)\mu(B)$. In fact, there is a unique Borel probability measure on $Y \times Z$ for which the latter relationship holds, and it is terms the *product* of κ and μ and is denoted $\kappa \otimes \mu$. All of the above extends directly to arbitrary finite products of metric spaces and finite measures.

If a function $f \colon Y \times Z \to \mathbb{R}$ defined on the product of metric spaces Y and Z is measurable with the product metric on $Y \times Z$, then it is measurable in each argument: for each $y \in Y$, the function $f_y \colon Z \to \mathbb{R}$ defined by $f_y(z) = f(y,z)$ is measurable, and for each $z \in Z$, the function $f_z \colon Y \to \mathbb{R}$ defined by $f_z(y) = f(y,z)$ is measurable. We say $f \colon Y \times Z \to \mathbb{R}$ is a *Carathéodory function* if (i) for each $y \in Y$, f_y is measurable, and (ii) for each $z \in Z$, f_z is continuous. In fact, although a Carathéodory function $f(y, z)$ is defined to be measurable in one argument and continuous in the other, it is actually (jointly) measurable.

Given a measurable subset $A \subseteq Y \times Z$ and $y \in Y$, let $A_y = \{z \in Z \mid (y, z) \in A\}$ be the *section* of A at y; similarly, for $z \in Z$, let $A_z = \{y \in Y \mid (y, z) \in A\}$ be the section at z. Then for all $y \in Y$ outside a κ-measure zero exceptional set, the set A_y is measurable, and the function $\alpha\colon Y \to \mathbb{R}$ defined by $\alpha(y) = \mu(A_y)$ outside the exceptional set (and equal to zero on the exceptional set) is measurable; similarly, for all z outside a μ-measure zero exceptional set, A_z is measurable, and the function $\beta\colon Z \to \mathbb{R}$ defined by $\beta(z) = \kappa(A_z)$ outside the exceptional set (and equal to zero on the exceptional set) is measurable. Furthermore, the product measure of A can be computed as

$$(\kappa \otimes \mu)(A) \quad = \quad \int_Y \alpha(y)\kappa(dy) \quad = \quad \int_Z \beta(z)\mu(dz),$$

integrating across sections of the set.

Now consider a function $f\colon Y \times Z \to \mathbb{R}$. An implication of *Fubini's theorem* is that if f is $(\kappa \otimes \mu)$-integrable, then for almost all $y \in Y$, f_y is μ-integrable; and for almost all $z \in Z$, f_z is κ-integrable. In fact, the functions $g\colon Y \to \mathbb{R}$ and $h\colon Z \to \mathbb{R}$ defined by

$$g(y) \quad = \quad \int_Y f_y(z)\mu(dz) \qquad \text{and} \qquad h(z) \quad = \quad \int_Z f_z(y)\kappa(dy)$$

outside the corresponding exceptional sets (and equal to zero on the exceptional sets) are μ- and κ-integrable. Furthermore, the integral of f is independent of the order of integration:

$$\int_{Y \times Z} f(y, z)(\kappa \otimes \mu)(d(y, z)) \quad = \quad \int_Y g(y)\kappa(dy) \quad = \quad \int_Z h(z)\mu(dz).$$

Whereas Fubini's theorem assumes that f is $(\kappa \otimes \mu)$-integrable, *Tonelli's theorem* gives conditions for integrability: if $f\colon Y \times Z \to \mathbb{R}$ is measurable, then for almost all y, $f_y\colon Z \to \mathbb{R}$ is measurable; if f_y is μ-integrable for κ-almost all y, then the function g defined above is measurable; and if g is κ-integrable, then f is $(\kappa \otimes \mu)$-integrable, and Fubini's theorem then implies that the order of integration is irrelevant.

We can build on Fubini's theorem to state results on parameterized integrals. Letting X and P be metric spaces and endowing $X \times P$ with the product metric, consider a measurable function $f\colon X \times P \to \mathbb{R}$. Let κ be a measure on X, and let μ be a measure on P. By Fubini's theorem, if f is $(\kappa \otimes \mu)$-integrable, then it follows that the integral over x, $\int f(x,p)\kappa(dx)$, is a measurable function of p; here, we set the integral equal to zero on the measure zero set of p such that $f(\cdot, p)$ is not κ-integrable. Now assume that the parameterization is continuous: for each $x \in X$, the function $f_x\colon P \to \mathbb{R}$ defined by $f_x(p) = f(x,p)$ is continuous, i.e., f is a Carathéodory function. Suppose further that there is a κ-integrable function $g\colon X \to \mathbb{R}$ such that for κ-almost all x and for all p, $|f(x,p)| \leq g(x)$. By Lebesgue's dominated convergence theorem, the integral, $\int f(x,p)\kappa(dx)$, is a continuous function of p.

Not all measures on products of metric spaces are product measures. But any measure ν on $X = Y \times Z$ determines a measure on each factor. The *marginal measure* of ν on Y is the measure ν_Y defined by $\nu_Y(A) = \nu(A \times Z)$ for all measurable $A \subseteq Y$. Similarly, the marginal on Z is defined by $\nu_Z(B) = \nu(Y \times B)$ for all measurable $B \subseteq Z$. Assuming they are σ-finite, these marginal measures determine a product measure, $\nu_Y \otimes \nu_Z$, on X, but it will not be the case that $\nu = \nu_Y \otimes \nu_Z$ unless the initial measure ν was indeed a product measure. In that case, if $\nu = \kappa \otimes \mu$, then the marginals will coincide with the factors in the product: $\nu_Y = \kappa$ and $\nu_Z = \mu$.

We can characterize weak convergence of product measures in terms of convergence of marginals. Consider a sequence $\{\nu^k\}$ of measures on $X = Y \times Z$ converging weakly to the measure ν. Then the sequence $\{\nu_Y^k\}$ of marginals on Y converges weakly to the marginal ν_Y, and similarly for the marginals on Z. Indeed, consider any $A \subseteq Y$ with ν_Y-measure zero boundary in Y, and note that $\mathrm{bd}(A \times Z) = \mathrm{bd}(A) \times Z$, which has ν-measure zero in X. Then

$$\nu_Y^k(A) \;=\; \nu^k(A \times Z) \;\to\; \nu(A \times Z) \;=\; \nu_Y(A).$$

The converse holds for product measures: letting $\nu^k = \kappa^k \otimes \mu^k$ for all k and $\nu = \kappa \otimes \mu$, if $\kappa^k \to \kappa$ weakly and $\mu^k \to \mu$ weakly, then $\nu^k \to \nu$ weakly. These observations extend directly to arbitrary finite products of measures.

Special metric spaces.

Real-valued functions. Let $\mathcal{F}_b(X)$ consist of all bounded, real-valued functions $f\colon X \to \mathbb{R}$, equipped with the *supremum metric,* denoted ρ^s and defined by $\rho^s(f,g) = \sup\{|f(x) - g(x)| \mid x \in X\}$. The metric space $\mathcal{F}_b(X)$

with the supremum metric is complete. Note that a sequence $\{f_m\}$ converges to f in $\mathcal{F}_b(X)$ if and only if $\sup\{|f_m(x) - f(x)| \mid x \in X\} \to 0$. That is, $\rho^s(f_m, f) \to 0$ if and only if $\{f_m\}$ converges uniformly to f, showing that uniform convergence is metrizable.

Measurable functions. Let $\mathcal{M}_b(X)$ consist of all bounded and measurable mappings $f\colon X \to \mathbb{R}$, equipped with the supremum metric, i.e., $\rho^s(f,g) = \sup\{|f(x) - g(x)| \mid x \in X\}$. Then the metric space $\mathcal{M}_b(X)$ is complete.

Continuous functions. Let $\mathcal{C}_b(X)$ consist of all mappings $f\colon X \to \mathbb{R}$ that are bounded and continuous, again equipped with the supremum metric. Then $C_b(X)$ is a complete metric space, and if X is compact, then $\mathcal{C}_b(X)$ is in fact separable. A set $K \subseteq \mathcal{C}_b(X)$ is *equicontinuous* if for all $\epsilon > 0$, there exists $\delta > 0$ such that for all $f \in K$ and all $x, y \in X$, $||x - y|| < \delta$ implies $|f(x) - f(y)| < \epsilon$. Assuming X is compact, the *Arzela-Ascoli theorem* states that a subset of $\mathcal{C}_b(X)$ is compact if and only if it is bounded, closed, and equicontinuous. Thus, a bounded, closed, equicontinuous set $K \subseteq \mathcal{C}_b(X)$ is itself a compact metric space with the supremum metric. A subset $K \subseteq \mathcal{C}_b(X)$ is *convex* if for all $f, g \in K$ and all $\alpha \in (0,1)$, the convex combination of functions, $\alpha f + (1-\alpha)g$, belongs to K. Equipping any convex $K \subseteq \mathcal{C}_b(X)$ with the supremum metric, it becomes a metric mixture space: given sequences $\{f_m\}$ converging to f in K, $\{g_m\}$ converging to g in K, and $\{\alpha_m\}$ converging to α in $[0,1]$, we have $\alpha_m f_m + (1-\alpha_m)g_m \to \alpha f + (1-\alpha)g$ uniformly. Moreover, the supremum metric on K is convex: for all $f, g, \phi, \gamma \in \mathcal{C}_b(X)$ and all $\alpha \in (0,1)$, we have

$$
\begin{aligned}
&\rho^s(\alpha f + (1-\alpha)g, \alpha\phi + (1-\alpha)\gamma) \\
&\quad = \sup\{|\alpha(f(x) - \phi(x)) + (1-\alpha)(g(x) - \gamma(x)) \mid x \in X\} \\
&\quad \leq \alpha \sup\{|f(x) - \phi(x)| \mid x \in X\} + (1-\alpha)\sup\{|g(x) - \gamma(x)| \mid x \in X\} \\
&\quad = \alpha\rho^s(f,\phi) + (1-\alpha)\rho^s(g,\gamma),
\end{aligned}
$$

as required.

Probability measures. The set of Borel probability measures on X is denoted by

$$\Delta(X) \quad = \quad \{\mu \mid \mu \text{ is a Borel probability measure on } X\}.$$

We define a particular metric of interest, known as the *Prohorov metric* and denoted ρ^r, as follows: for $\mu, \nu \in \Delta(X)$, let

$$\rho^r(\mu,\nu) \quad = \quad \inf\left\{\epsilon > 0 \;\middle|\; \begin{array}{c} \text{for all mble } Y \subseteq X,\ \mu(Y) \leq \nu(Y^\epsilon) + \epsilon \\ \text{and } \nu(Y) \leq \mu(Y^\epsilon) + \epsilon \end{array}\right\},$$

where $Y^\epsilon = \{x \in X |$ there exists $y \in Y$ with $x \in B_\epsilon(y)\}$. A sequence $\{\mu_m\}$ of probability measures in $\Delta(X)$ converges to $\mu \in \Delta(X)$ if and only if for all measurable $Y \subseteq X$ with $\mu(\text{bd}(Y)) = 0$, we have $\mu_m(Y) \to \mu(Y)$. That is, $\rho^r(\mu_m, \mu) \to 0$ if and only if $\mu_m \to \mu$ weakly, showing that weak convergence of probability measures is metrizable. Interesting properties of the space $\Delta(X)$ equipped with the Prohorov metric are that if X is separable, then so is $\Delta(X)$; if X is complete and separable, then so is $\Delta(X)$; and if X is compact, then so is $\Delta(X)$. Note that $\Delta(X)$ is closed with respect to convex combinations: for all $\mu, \nu \in \Delta(X)$ and all $\alpha \in (0,1)$, we can define the measure $\alpha\mu + (1-\alpha)\nu$ by $(\alpha\mu + (1-\alpha)\nu)(Y) = \alpha\mu(Y) + (1-\alpha)\nu(Y)$ for all measurable $Y \subseteq X$. In fact, $\Delta(X)$ is a metric mixture space, and the Prohorov metric is quasi-convex.

Another metric of interest on $\Delta(X)$ is the *total variation metric,* denoted ρ^v and defined as:

$$\rho^v(\mu,\nu) \quad = \quad \sup\left\{\sum_{i=1}^{k} |\mu(Y_i) - \nu(Y_i)| \;\middle|\; \begin{array}{c} \{Y_1, \ldots, Y_k\} \text{ is a finite,} \\ \text{mble partition of } X \end{array}\right\},$$

where of course a collection $\mathcal{Y}$ is a *measurable partition* of X if it is a partition of X and each $Y \in \mathcal{Y}$ is measurable. In fact, we can define the sup metric on $\Delta(X)$ by $\rho^s(\mu,\nu) = \sup\{|\mu(Y) - \nu(Y)| \mid Y \in \mathcal{B}\}$, and it turns out that for all $\mu, \nu \in \Delta(X)$, we have $\rho^v(\mu,\nu) = 2\rho^s(\mu,\nu)$; thus, convergence in the total variation metric is equivalent to convergence in the sup metric. Given a sequence $\{\mu_m\}$ in $\Delta(X)$ and $\mu \in \Delta(X)$, we therefore have $\mu_m \to \mu$ uniformly set-wise if and only if $\rho^v(\mu_m, \mu) \to 0$, so that uniform set-wise convergence of probability measures is metrizable. The space $\Delta(X)$ with the total variation metric is complete, whether or not X is complete.

1.3 Transition Probabilities

Given metric spaces Y and Z, a mapping $\mu\colon \mathcal{B}_Y \times Z \to \mathbb{R}$ with values denoted $\mu(B|z)$ is a *transition probability* (or *stochastic kernel* or *Markov kernel*) if (i) for all $z \in Z$, $\mu(\cdot|z)$ is a probability measure on Y, and (ii) for all measurable $B \subseteq Y$, $\mu(B|z)$ is measurable as a function of z. More generally, $\mu(\cdot|z)$ may be any finite measure (not necessarily a probability measure), in which case the mapping μ is a *Young measure.* We focus here on transition probabilities. A useful fact is that given a transition probability μ and any bounded, continuous function $f\colon Y \to \mathbb{R}$, the function $g\colon Z \to \mathbb{R}$ defined by $g(z) = \int f(y)\mu(dy|z)$ is measurable.

Alternatively, we can view a transition probability μ as a mapping from the metric space Z to probability measures on the metric space Y, i.e., $\mu\colon Z \to \Delta(Y)$. In fact, giving $\Delta(Y)$ the Prohorov metric, this mapping μ is measurable if and only if conditions (i) and (ii) above are satisfied, i.e., μ is a transition probability. Given measurable $A \subseteq X$, define the *evaluation functional* $e_A\colon \Delta(X) \to \mathbb{R}$ by $e_A(\nu) = \nu(A)$, so the mapping simply evaluates ν at A. An implication of the foregoing observation is that e_A is measurable. Indeed, let $Z = \Delta(Y)$, and let μ be defined so that for all $\nu \in \Delta(Y)$, $\mu(\cdot|\nu) = \nu$; viewed as a mapping $\mu\colon Z \to \Delta(Y)$, it is the identity function, so it is trivially measurable. By condition (ii), for all measurable $A \subseteq Y$, $\mu(A|\nu)$ is measurable in ν, and the claim follows by noting that $\mu(A|\nu) = e_A(\nu)$.

The transition probability μ satisfies the *Feller property* if for every bounded, continuous function $f\colon Y \to \mathbb{R}$, the function $g\colon Z \to \mathbb{R}$ defined by $g(z) = \int f(y)\mu(dy|z)$ is continuous, rather than only measurable. Given μ, we can define the mapping $P_\mu\colon Z \to \Delta(Y)$ by $P_\mu(z) = \mu(\cdot|z)$. Then, in fact, μ satisfies the Feller property if and only if P_μ is continuous with the Prohorov metric on $\Delta(Y)$. Thus, if $z_k \to z$ in Z, then $\mu(\cdot|z_k) \to \mu(\cdot|z)$ weakly. Therefore, by the generalized version of Lebesgue's dominated convergence theorem, if $\{f^k\}$ is a uniformly bounded sequence of functions $f^k\colon Y \to \mathbb{R}$ converging pointwise to the continuous function $f\colon Y \to \mathbb{R}$, then $z_k \to z$ in Z implies $\int f^k(y)\mu(dy|z_k) \to \int f(y)\mu(dy|z)$.

Because the Feller property is quite useful, we note the following sufficient condition. Assume there is a probability measure ν on Y such that for all $z \in Z$, $\mu(\cdot|z) \ll \nu$. Let $h\colon Y \times Z \to \mathbb{R}$ be a density for μ, i.e., for all measurable $B \subseteq Y$ and all $z \in Z$, we have $\mu(B|z) = \int_B h(y|z)\nu(dy)$. Assume h is Caratheodory, i.e., the function $h_z\colon Y \to \mathbb{R}$ defined by $h_z(y) = h(y,z)$ is continuous for all z, and the function $h_y\colon Z \to \mathbb{R}$ defined by $h_y(z) = h(y,z)$ is measurable for all y. And assume that the set $\{h_z\colon z \in Z\}$ is dominated by the ν-integrable function $g\colon Y \to \mathbb{R}$. Let $z_k \to z$ in Z, and consider any bounded, continuous function $f\colon Y \to \mathbb{R}$. By Lebesgue's dominated convergence theorem, we have

$$\int f(y)\mu(dy|z_k) \;=\; \int f(y)h(y,z_k)dy \;\rightarrow\; \int f(y)h(y,z)dy \;=\; \int f(y)\mu(dy|z),$$

which establishes the Feller property.

A transition probability $\mu\colon \mathcal{B}_Y \times Z \rightarrow \mathbb{R}$ and a probability measure κ on Z determine a probability measure on $X = Y \times Z$, endowed with the product metric. This probability measure is denoted $\mu(\cdot|z)\otimes\kappa$ and is uniquely pinned down by products of measurable sets. That is, given any product $R = P \times Q$ such that $P \subseteq Y$ and $Q \subseteq Z$ are measurable, we specify that

$$(\mu(\cdot|z) \otimes \kappa)(R) \;=\; \int_Q \mu(P|z)\kappa(dz),$$

and this extends uniquely to the probability measure $\mu(\cdot|z) \otimes \kappa$ on X. Note that if the transition probability $\mu(\cdot|z)$ is constant in z, so we can write it simply as a probability measure μ on Y, then $\mu(\cdot|z) \otimes \kappa = \mu \otimes \kappa$, so the operation we have defined generalizes the concept of product measure.

Assume μ is a transition probability, as above. Given a measurable subset $A \subseteq X = Y \times Z$ and $z \in Z$, recall that $A_z = \{y \in Y \mid (y,z) \in A\}$ is the section of A at z. For all $z \in Z$ outside a κ-measure zero exceptional set, the set A_z is measurable, and the function $\alpha\colon Z \rightarrow \mathbb{R}$ defined by $\alpha(z) = \mu(A_z|z)$ outside the exceptional set (and equal to zero on the exceptional set) is measurable. Furthermore, the measure of A is

$$(\mu(\cdot|z) \otimes \kappa)(A) \;=\; \int \alpha(z)\kappa(dz) \;=\; \int \mu(A_z|z)\kappa(dz),$$

integrating across sections of the set; in contrast to the case of product measures, we now measure sections using the transition probability, which may itself vary with z. In fact, the above construction is quite general: it holds when κ is any σ-finite measure and the collection $\{\mu(\cdot|z) \mid z \in Z\}$ is *uniformly σ-finite,* i.e., there is a countable collection $\{Y_1, Y_2, \ldots\}$ of measurable subsets of Y such that $Y = \bigcup_{i=1}^{\infty} Y_i$ and for all $i = 1, 2, \ldots$, $\sup\{\mu(Y_i|z) \mid z \in Z\} < \infty$.

Maintaining that μ is a transition probability, let $f\colon Y \times Z \rightarrow \mathbb{R}$ be measurable. Then a general form of *Fubini's theorem* states that if f is $\mu(\cdot|z) \otimes \kappa$-integrable, then for κ-almost all z, the function $f_z\colon Y \rightarrow \mathbb{R}$ defined by

$f_z(y) = f(y,z)$ is $\mu(\cdot|z)$-integrable outside an exceptional set of κ-measure zero; the function defined by

$$g(z) \;=\; \int f_z(y)\mu(dy|z)$$

outside the exceptional set (and equal to zero on the exceptional set) is κ-integrable; and

$$\int f(y,z)(\mu(\cdot|z)\otimes\kappa)(d(y,z)) \;=\; \int g(z)\kappa(dz) \;=\; \int_z \left(\int_y f_z(y)\mu(dy|z)\right)\kappa(dy).$$

Conversely, by a general form of *Tonelli's theorem,* if f is measurable and for κ-almost all z, f_z is $\mu(\cdot|z)$-integrable, then the function g defined above is measurable; and if g is κ-integrable, then f is $\mu(\cdot|z)\otimes\kappa$-integrable.

Given that $\mu(\cdot|z)$ is absolutely continuous with respect to the probability measure ν for all $z \in Z$, it follows from the Radon-Nikodym theorem that for each z, there is a density $h_z\colon Y \to \mathbb{R}$ for $\mu(\cdot|z)$ with respect to ν. Because these densities are selected independently for each z, there is no guarantee that they vary in a measurable way across z. In fact, however, there is a measurable function $h\colon Y \times Z \to \mathbb{R}$ such that for all z, $h(\cdot, z)$ is a density for $\mu(\cdot|z)$. That is, the densities may be specified to be jointly measurable. Then we can write the integral of f with respect to $\mu(\cdot|z)\otimes\kappa$ as

$$\begin{aligned}\int f(y,z)(\mu(\cdot|z)\otimes\kappa)(d(y,z)) \;&=\; \int_z\int_y f(y,z)\mu(dy|z)\kappa(dz)\\ &=\; \int_z\int_y f(y,z)h(y,z)\nu(dy)\kappa(dz)\\ &=\; \int f(y,z)h(y,z)(\nu\otimes\kappa)(d(y,z)),\end{aligned}$$

and we can use the standard version of Fubini's theorem to conclude that the order of integration is irrelevant.

We have seen that a collection $\{\mu(\cdot|z) \mid z \in Z\}$ of probability measures on Y satisfying the measurability conditions of a transition probability, given a probability measure κ on Z, induces a probability measure on $X = Y \times Z$. Conversely, a probability measure ν on X induces a probability measure on Z via the marginal probability ν_Z, denoted κ in the present context, defined by $\kappa(C) = \nu(Y \times C)$ for all measurable $C \subseteq Z$. To construct a corresponding collection of probability measures on Y, first consider any measurable $A \subseteq X$. A function $P(A|\cdot)\colon Z \to \mathbb{R}$ is a *conditional probability of* A if it takes values between zero and one, it is measurable, and

$$\nu(A \cap (Y \times C)) \quad = \quad \int_C P(A|z)\kappa(dz)$$

for every measurable $C \subseteq Z$. In this context, we refer to A and C as "events," and $P(A|z)$ is interpreted as the conditional probability that the event A occurs given the information that z is the value of the conditioning variable.

A *system of conditional probabilities* is a collection $\{P(A|\cdot) \mid A \in \mathcal{B}_X\}$ of conditional probabilities for each measurable set. It is known that if $\{A_1, A_2, \ldots\}$ is any countable, pairwise disjoint collection of events, then for all z outside a κ-measure zero set, the conditional probabilities are countably additive:

$$P\left(\bigcup_{i=1}^{\infty} A_i|z\right) \quad = \quad \sum_{i=1}^{\infty} P(A_i|z).$$

A weakness of this countable additivity property is that the exceptional set can depend on the collection $\{A_1, A_2, \ldots\}$. In other words, there need not be a single κ-measure zero set such that for all z outside this exceptional set, $P(\cdot|z)$ is a probability measure.

In fact, any given probability measure ν on $X = Y \times Z$ admits a *regular conditional probability,* a mapping $\mu\colon \mathcal{B}_Y \times Z \to \mathbb{R}$ such that for κ-almost all z, $\mu(\cdot|z)$ is a probability measure on Y and $\{P(A|\cdot) \mid A \in \mathcal{B}_X\}$ is a system of conditional probabilities, where $P(A|z) = \mu(A_z|z)$ for all measurable $A \subseteq Z$ and all $z \in Z$. An implication is that for each measurable $B \subseteq Y$, $\mu(B|z)$ is measurable as a function of z, and therefore μ is a transition probability. Furthermore, we have

$$\nu(A) \quad = \quad \int P(A|z)\kappa(dz) \quad = \quad \int \mu(A_z|z)\kappa(dz)$$

for all measurable $A \subseteq Z$, and therefore $\nu = \mu(\cdot|z) \otimes \kappa$. That is, the probability measure ν induces the marginal κ and the transition probability μ (a regular conditional probability), which return the initial probability measure.

Transition probabilities are often used to model discrete time Markov chains, in which case we equate $Y = Z$ and view an element $z \in Z$ as the state of a system; then $\mu(\cdot|z)$ specifies the probability distribution over next period's state, conditional on the current state being z. Given a transition probability $\mu\colon \mathcal{B}_Z \times Z \to \mathbb{R}$, a distribution $\kappa \in \Delta(Z)$ over the current state induces a distribution $T^*_\kappa(\kappa)$ over next period's state as follows: for all measurable subsets $A \subseteq \mathbb{R}^m$,

$$T^*_\mu(\kappa)(A) \quad = \quad \int \mu(A|z)\kappa(dz).$$

The mapping T^*_μ is referred to as the *adjoint* of the transition probability.

A probability measure κ is an *invariant distribution* (or *stationary distribution*) if $T^*_\mu(\kappa) = \kappa$, i.e., it is a fixed point of T^*_μ. Note that if μ satisfies the Feller property, then the mapping $T^*_\mu\colon \Delta(Z) \to \Delta(Z)$ is continuous with the Prohorov metric on $\Delta(Z)$. Indeed, let $\kappa_k \to \kappa$ weakly, and let $f\colon Z \to \mathbb{R}$ be any bounded, continuous function. Then

$$\begin{aligned}
\int f(z)T^*_\mu(\kappa_k)(dz) &= \int_z f(z)\left(\int_{z'} \mu(dz|z')\kappa_k(dz')\right) \\
&= \int_{z'} \left(\int_z f(z)\mu(dz|z')\right)\kappa_k(dz') \\
&\to \int_{z'} \left(\int_z f(z)\mu(dz|z')\right)\kappa(dz') \\
&= \int f(z)T^*_\mu(\kappa)(dz),
\end{aligned}$$

where the second and third equalities follows from linearity of the integral in the integrating measure, and the limit follows from the Feller property and weak convergence. In fact, the converse holds as well.

To explore further the dynamic interpretation of transition probabilities, we consider a transition probability $\mu\colon \mathcal{B}_Z \times Z \to \mathbb{R}$ satisfying *Doeblin's condition:* there exist a finite measure ξ on $\mathbb{R}^m$ and $\epsilon > 0$ such that for every measurable set $C \subseteq Z$ and every $z \in Z$, $\xi(C) \leq \epsilon$ implies $\mu(C|z) \leq 1 - \epsilon$. In words, this requires that the transition probability not be too concentrated on small sets, where the meaning of "small" is given by ξ and ϵ. (Note that the measure ξ in Doeblin's condition must satisfy $\xi(Z) > 0$.) Because Doeblin's condition is quite useful, we note that it is satisfied if $\mu(\cdot|z) \ll \nu$ for all $z \in Z$, and μ has a density $h\colon Z \times Z \to \mathbb{R}$ with respect to ν that is bounded. Indeed, if $b \geq 1$ is an upper bound for h, then we can set $\xi = \nu$ and $\epsilon = 1/2b$; then given measurable C with $\xi(C) \leq \epsilon$, we have $\nu(C) \leq 1/2b$, and therefore

$$\mu(C|z) \;=\; \int_C h(y,z)\nu(dy) \;\leq\; \nu(C)b \;\leq\; \frac{1}{2} \;\leq\; 1 - \frac{1}{2b} \;=\; 1 - \epsilon$$

for all z, as required.

Then a measurable set $C \subseteq Z$ is *invariant* (or *absorbing* or *self-supporting*) if for all $z \in C$, we have $\mu(C|z) = 1$. And an invariant set E is *ergodic* if it is nonempty and contains no invariant set C with smaller ξ-measure, i.e., there is no invariant set $C \subseteq E$ with $\xi(C) < \xi(E)$. Every invariant set, and therefore every ergodic set, has ξ-measure of at least ϵ. Furthermore, if E and E' are ergodic sets, then it must be either that E and E' are equivalent with respect to ξ, i.e., $\xi(E \setminus E') = \xi(E' \setminus E) = 0$, or that they are disjoint with respect to ξ, i.e., $\xi(E \cap E') = 0$. As a consequence, the number of meaningfully distinct ergodic sets is bounded above by $\xi(Z)/\epsilon$. Furthermore, there is at least one ergodic set. An invariant distribution κ is an *ergodic distribution* if there is an ergodic set E such that $\kappa(E) = 1$.

Say a set F is *transient* if for all $z \in Z$, we have $\mu(F|z) < 1$. Then Z can be partitioned into a transient set F and a finite number k of ergodic sets, $E^1, \ldots, E^k$, such that every ergodic set E is equivalent to some E^i, $i = 1, \ldots, k$, with respect to ξ. For each ergodic set E^i, there is a unique ergodic distribution κ_i satisfying $\kappa_i(E^i) = 1$. Moreover, every invariant distribution is a convex combination of the ergodic distributions $\kappa_1, \ldots, \kappa_k$ associated with these ergodic sets: if κ is an invariant distribution, then there exist non-negative weights $\alpha_1, \ldots, \alpha_k$ with $\sum_{i=1}^k \alpha_i = 1$ such that $\kappa = \sum_{i=1}^k \alpha_i \kappa_i$.

To complete the analysis of dynamics, define the transition probabilities μ^r, $r = 1, 2, \ldots$, as follows: $\mu^1 = \mu$, and for $r \geq 2$, for all measurable $C \subseteq Z$ and all $z \in Z$, we specify $\mu^r(C|z) = \int \mu^{r-1}(C|z')\mu(dz'|z)$. That is, $\mu^r(C|z)$ is the probability that given initial state z, the state belongs to the set C after r steps of the process. Then, beginning from any state, the probability that the state reaches an ergodic set (and remains there) approaches one over time and at a geometric rate: there exist constants $c > 0$ and $d \in [0, 1)$

such that for all r and all $z \in Z$, we have

$$\mu^r \left(\bigcup_{i=1}^{k} E^i \middle| z \right) \geq 1 - cd^r.$$

Let $T_\mu^{*r} = T_{\mu^r}^*$, so that given an initial distribution κ over states, the distribution in r steps is $T_\mu^{*r}(\kappa)$. Beginning with any distribution over states, the average distribution over future states converges (in a strong sense) to an invariant distribution: for every probability measure κ with $\kappa(Z) = 1$, there is an invariant distribution κ^* such that the average distribution over r steps, i.e., $\frac{1}{r}\sum_{t=1}^{r} T_\mu^{*t}(\kappa)$, converges to κ^* as r goes to infinity; or equivalently,

$$\rho^v \left(\frac{1}{r} \sum_{t=1}^{r} T_\mu^{*t}(\kappa), \kappa^* \right) \rightarrow 0,$$

where ρ^v is the total variation metric on $\Delta(Z)$. In fact, if the initial distribution puts probability one on some ergodic set, say $\kappa(E^i) = 1$, then the limit distribution will be $\kappa^* = \kappa_i$, the ergodic distribution corresponding to the ergodic set.

The reason the above limit results are stated for the average distributions is that the Markov chain can "cycle," even within an ergodic set: for example, if $Z = \{z_1, z_2\}$ and $\mu(\{z_2\}|z_1) = \mu(\{z_1\}|z_2) = 1$, then starting from z_1, the chain alternates endlessly between the two states. To preclude such cycles, it is sufficient to add the Feller property and the following "mixing" condition: for every ergodic set E^i, there exists $z_i^* \in E^i$ such that for every open set $G \subseteq Z$ with $z_i^* \in G$ and for every $z \in E^i$, we have $\mu(G|z) > 0$. Under the latter condition, in combination with the Doeblin and Feller properties, the above limit results can be stated not for the sequence of average distributions, but for the sequence of distributions over states in each period, as in

$$\rho^v \left(T_\mu^{*r}(\kappa), \kappa^* \right) \rightarrow 0,$$

where κ is an initial distribution with $\kappa(Z) = 1$ and κ^* is invariant. And if we strengthen the mixing condition so that there exists $z^* \in Z$ such that for every open set G containing z^* and every $z \in Z$, we have $\mu(G|z) > 0$, then there is a unique ergodic set (up to ξ-equivalences) and a unique invariant distribution, say κ^*; then starting from any initial distribution over states, the induced distribution over states converges in the total variation metric to κ^*, providing the strongest possible ergodicity properties.

Note that these strong ergodicity properties hold if $\mu(\cdot|z) \ll \nu$ for all $z \in z$ and μ has a density $h\colon Z \times Z \to \mathbb{R}$ with respect to ν such that (i) h is Caratheodory, (ii) $\{h_z \mid z \in Z\}$ is dominated by ν-integrable $g\colon Z \to \mathbb{R}$, (iii) h is bounded, (iv) h is everywhere positive, i.e., $h(y,z) > 0$ for all $y, z \in Z$, and (v) for every open subset $G \subseteq Z$, we have $\nu(G) > 0$. Of course, (i) and (ii) together imply the Feller property, and (iii) yields Doeblin's condition. To see that the strong mixing condition obtains, choose any $z^* \in Z$, any open set G containing z^*, and any $z \in Z$. Note that $G = \bigcup_{k=1}^{\infty}\{y \in G \mid h(y,z) \geq \frac{1}{k}\}$ and $\nu(G) = \lim_{k\to\infty} \nu(\{y \in G \mid h(y,z) \geq \frac{1}{k}\}) > 0$. Thus, we have

$$\mu(G|z) \;\geq\; \int_G h(y,z)\nu(dy) \;\geq\; \frac{1}{k}\nu(\{y \in G \mid h(y,z) \geq \frac{1}{k}\}) \;>\; 0$$

for some k, as required.

To consider the metric properties of transition probabilities, let κ be a probability measure on Z. A mapping $f\colon Y \times Z \to \mathbb{R}$ is a *Carathéodory integrand* if it is a Carathéodory function, so (i) for all $y \in Y$, the function $f_y\colon Z \to \mathbb{R}$ defined by $f_y(z) = f(y,z)$ is measurable, and (ii) for all $z \in Z$, the function $f_z\colon Y \to \mathbb{R}$ defined by $f_z(y) = f(y,z)$ is continuous, and if in addition, (iii) there is a κ-integrable mapping $g\colon Z \to \mathbb{R}$ such that for all $z \in Z$, $\sup\{|f(y,z)| \mid y \in Y\} \leq g(z)$. An implication is that for all $z \in Z$, the function f_z is bounded and continuous. Let $\mu_k\colon \mathcal{B}_Y \times Z \to \mathbb{R}$, $k = 1, 2, \ldots$, and $\mu\colon \mathcal{B}_Y \times A \to \mathbb{R}$ be transition probabilities. Generalizing the notion of weak convergence of measures, we say the sequence $\{\mu_k\}$ of transition probabilities *converges weakly* to μ if for every Carathéodory integrand f, we have

$$\int_z \left(\int_y f(y,z)\mu_k(dy|z)\right)\kappa(dz) \;\;\to\;\; \int_z \left(\int_y f(y,z)\mu(dy|z)\right)\kappa(dy).$$

Weak convergence of the sequence $\{\mu_k\}$ of transition probabilities to μ is equivalent to the following condition: for every measurable $C \subseteq Z$ and every bounded, continuous function $f\colon Y \to \mathbb{R}$, we have

$$\int_C \left(\int_y f(y)\mu_k(dy|z)\right)\kappa(dz) \;\;\to\;\; \int_C \left(\int_y f(y)\mu(dy|z)\right)\kappa(dz).$$

Other possible terminology for this form of convergence is *narrow convergence* or *weak-strong convergence.*

To gain some insight into weak convergence of transition probabilities, reconsider the sawtooth sequence in Figure 2, but transform each element f_k of the sequence into a transition probability $\mu_k \colon \mathcal{B}_{\mathbb{R}} \times [0,1] \to \mathbb{R}$ all follows: for all $z \in [0,1]$, $\mu_k(\cdot|z)$ places probability one on $f_k(z)$, i.e., $\mu_k(\{f_k(z)\}|z) = 1$. Whereas the sequence of functions $\{f_k\}$ converges weakly to the function that takes a constant value equal to one, this sequence of transition probabilities converges weakly to the transition probability μ that places probability one half on 0 and 2, i.e., for all $z \in [0,1]$, $\mu(\{0\}|z) = \mu(\{2\}|z) = \frac{1}{2}$.

Let $\mathcal{R}(Y, Z, \kappa)$ denote the set of equivalence classes of transition probabilities $\mu \colon \mathcal{B}_Y \times Z \to \mathbb{R}$, where we identify any two transition probabilities μ and μ' that differ only on a κ-measure zero set, i.e., $\kappa(\{z \in Z \mid \mu(\cdot|z) \neq \mu'(\cdot|z)\}) = 0$. Then weak convergence of transition probabilities is metrizable in the sense that there is a metric, ρ^r, on $\mathcal{R}(Y, Z, \kappa)$ such that $\mu_k \to \mu$ weakly if and only if $\rho^r(\mu_k, \mu) \to 0$. We refer to this as the *narrow convergence metric,* but in case the transition probabilities are constant in z, so they are simply probability measures, the metric reduces to the Prohorov metric, justifying the similar notation. Indeed, there is a countable collection $\{f_i\}$ of bounded, continuous functions $f_i \colon Y \to \mathbb{R}$ such that we can set

$$\begin{aligned}\rho^r(\mu, \mu') \quad &= \quad \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \frac{1}{2^{i+j}\kappa(B_j)} \Bigg| \int_{z \in B_j} \left[\int_y f_i(y)\mu(dy|z) \right] \kappa(dz) \\ &\qquad - \int_{z \in B_j} \left[\int_y f_i(y)\mu'(dy|z) \right] \kappa(dz) \Bigg|,\end{aligned}$$

where $B_1, B_2, \ldots$ are the relatively open balls in Z defined by $B_j = Z \cap B_r(x)$ for rational radius r and vectors x with rational coordinates.

As with the case of probability measures, if Y is compact, then the set $\mathcal{R}(Y, Z, \kappa)$ endowed with the metric ρ^r is compact. In fact, the set $\mathcal{R}(Y, Z, \kappa)$ of transition probabilities equipped with the metric ρ^r is a metric mixture space, and the metric ρ^r is convex. Furthermore, if $\{\mu_k\}$ is a sequence of transition probabilities converging narrowly to μ in $\mathcal{R}(Y, Z, \kappa)$, then for κ-almost all z, we have:

$$\operatorname{supp}(\mu(\cdot|z)) \quad \subseteq \quad \bigcap_{k=1}^{\infty} \operatorname{clos}\left(\bigcup_{m=k}^{\infty} \operatorname{supp}(\mu_k(\cdot|z)) \right).$$

Although narrow convergence is not a pointwise notion, the above result does establish pointwise implications of the convergence concept, in terms of supports of the probability measures determined in each z.

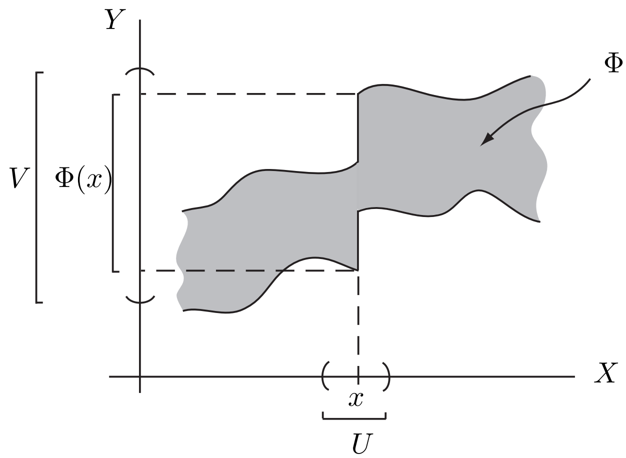

Figure 3: Upper hemicontinuity

1.4 Correspondences

Continuous correspondences. Given metric spaces X and Y, a *correspondence* from X to Y, denoted $\Phi\colon X \rightrightarrows Y$, is a mapping from X to subsets of Y, i.e., $\Phi(x) \subseteq Y$ for all $x \in X$. As a special case, we may view a function $f\colon X \to Y$ as a correspondence that takes only singleton sets as values; it is associated to the correspondence $\Phi\colon X \rightrightarrows Y$ defined by $\Phi(x) = \{f(x)\}$ for all $x \in X$. A correspondence $\Phi\colon X \rightrightarrows Y$ has *nonempty values* (resp. *closed values*, *compact values*) if for all $x \in X$, $\Phi(x) \neq \emptyset$ (resp. $\Phi(x)$ is closed, $\Phi(x)$ is compact). In contrast to the case of functions, there are two main notions of continuity of correspondences, upper and lower hemicontinuity, each generalizing the usual notion for functions.

The correspondence $\Phi\colon X \rightrightarrows Y$ is *upper hemicontinuous at* $x \in X$ if for every open set $V \subseteq Y$ with $\Phi(x) \subseteq V$, there is an open set $U \subseteq X$ such that $x \in U$ and for all $z \in U$, we have $\Phi(z) \subseteq V$. It is *upper hemicontinuous* if it is upper hemicontinuous at every $x \in X$. See Figure 3. Equivalently, the correspondence is upper hemi-continuous if for every closed set $F \subseteq Y$, the set $\{x \in X \mid \Phi(x) \cap F \neq \emptyset\}$ is closed. A correspondence Φ with closed values is upper hemicontinuous only if for all $x \in X$, all $y \in Y$, and all sequences $\{x_m\}$ converging to x in X and $\{y_m\}$ converging to y in Y such that $y_m \in \Phi(x_m)$ for all m, we have $y \in \Phi(x)$. Assuming Y is compact, the converse direction holds as well. Recall that the product space $X \times Y$ can be endowed with the product metric, making it a metric space, and that a sequence $\{(x_m, y_m)\}$ converges to (x, y) in the product space if and only if $x_m \to x$ in X and $y_m \to y$ in Y. Defining the *graph* of Φ as

$$\text{graph}(\Phi) \quad = \quad \{(x, y) \in X \times Y \mid y \in \Phi(x)\},$$

the correspondence has *closed graph* if, fittingly, its graph is closed in the product space $X \times Y$. Assuming Y is compact, a correspondence Φ with closed values is upper hemicontinuous if and only if it has closed graph.

The correspondence $\Phi\colon X \rightrightarrows Y$ is *lower hemicontinuous at* $x \in X$ if for every open set $V \subseteq Y$ with $\Phi(x) \cap V \neq \emptyset$, there is an open set $U \subseteq X$ such that $x \in U$ and for all $z \in U$, $\Phi(z) \cap V \neq \emptyset$. It is *lower hemicontinuous* if it is lower hemicontinuous at every $x \in X$. See Figure 4. Equivalently, the correspondence is lower hemi-continuous if for every closed set $F \subseteq Y$, the set $\{x \in X \mid \Phi(x) \subseteq F\}$ is closed. For yet another equivalence, the correspondence Φ is lower hemi-continuous if and only if for all $x \in X$, all $y \in \Phi(x)$, and all sequences $\{x_m\}$ converging to x in X, there exist a subsequence $\{x_{m_k}\}$ of $\{x_m\}$ and a corresponding sequence $\{y_k\}$ in Y such that $y_k \to y$ and for all k, $y_k \in \Phi(x_{m_k})$. The correspondence Φ has *open graph* if, fittingly, its graph is open in the product space $X \times Y$. Every correspondence with open graph is lower hemi-continuous. In fact, every correspondence with *open lower sections,* i.e., for all $y \in Y$,

$$\{x \in X \mid y \in \Phi(x)\}$$

is open, is lower hemi-continuous.

A correspondence $\Phi\colon X \rightrightarrows Y$ is *continuous at* $x \in X$ if it is both upper and lower hemi-continuous at x. It is *continuous* if it is continuous at all $x \in X$, i.e., it is both upper and lower hemi-continuous. If Φ has singleton values, so there is a function $f\colon X \to Y$ such that $\Phi(x) = \{f(x)\}$ for all $x \in X$, then upper hemi-continuity, lower hemi-continuity, and continuity of the correspondence Φ are equivalent conditions, and they in turn are equivalent to continuity of the function f.

Upper hemi-continuity is preserved by finite unions, and under weak conditions by arbitrary intersections and products: letting X and Y be metric spaces, and letting $\{\Phi_i \mid i \in I\}$ be a collection of upper hemi-continuous correspondences $\Phi_i\colon X \rightrightarrows Y$ indexed by elements of the set I,

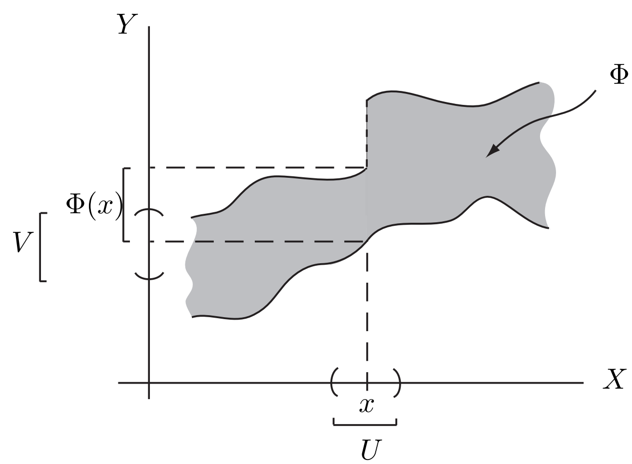

Figure 4: Lower hemicontinuity

- if $I = \{1, \ldots, m\}$ is finite, then the correspondence $\Psi\colon X \rightrightarrows Y$ defined by

$$\Psi(x) \quad = \quad \bigcup_{i=1}^{m} \Phi_i(x)$$

  is upper hemi-continuous,

- if each Φ_i has closed graph, or if each Φ_i has closed values and at least one is compact-valued, then the correspondence $\Psi\colon X \rightrightarrows Y$ defined by

$$\Psi(x) \quad = \quad \bigcap_{i \in I} \Phi_i(x)$$

  is upper hemi-continuous,

- if I is countable and each Φ_i has compact values, then the correspondence $\Psi\colon X \rightrightarrows Y^\infty$ defined by

$$\Psi(x) \quad = \quad \prod_{i \in I} \Phi_i(x)$$

  is upper hemi-continuous (with the product metric on Y^∞).

In fact, the last result extends to arbitrary products, but details are omitted here. If $\Phi\colon X \rightrightarrows Y$ is upper hemi-continuous, then the pointwise closure of

Φ is upper hemi-continuous; that is, the correspondence $\Psi\colon X \rightrightarrows Y$ defined by $\Psi(x) = \text{clos}(\Phi(x))$ is upper hemi-continuous.

Lower hemi-continuity is preserved by arbitrary unions and finite products: letting $\{\Phi_i \mid i \in I\}$ be a collection of lower hemi-continuous correspondences $\Phi_i\colon X \rightrightarrows Y$ indexed by elements of the set I,

- the correspondence $\Psi\colon X \rightrightarrows Y$ defined by

$$\Psi(x) \;=\; \bigcup_{i \in I} \Phi_i(x)$$

  is lower hemi-continuous,

- if $I = \{1, \ldots, k\}$ is finite, then the correspondence $\Psi\colon X \rightrightarrows Y^k$ defined by

$$\Psi(x) \;=\; \prod_{i=1}^{k} \Phi_i(x)$$

  is lower hemi-continuous (with the product metric on Y^k).

Intersections, even finite ones, of lower hemi-continuous correspondences need not be lower hemi-continuous. Given a finite collection of correspondences Φ_i, $i = 1, \ldots, n$, with open lower sections, however, the intersection of correspondences $\Psi = \bigcap_{i=1}^{n} \Phi_i$ will have open lower sections, and therefore it is lower hemi-continuous. If the correspondence $\Phi\colon X \rightrightarrows Y$ is lower hemi-continuous, then the pointwise closure of Φ is lower hemi-continuous; that is, the correspondence $\Psi\colon X \rightrightarrows Y$ defined by $\Psi(x) = \text{clos}(\Phi(x))$ is lower hemi-continuous. Conversely, if the pointwise closure of a correspondence Φ is lower hemi-continuous, then so is Φ.

When $Y = \mathbb{R}^m$ and $\Phi\colon X \rightrightarrows \mathbb{R}^m$ is upper hemi-continuous and has compact values, the pointwise convex hull of Φ is upper hemi-continuous; that is, the correspondence $\Psi\colon X \rightrightarrows Y$ defined by $\Psi(x) = \text{conv}(\Phi(x))$ is upper hemi-continuous. This uses the fact that the convex hull of a compact subset of $\mathbb{R}^n$ is compact, a property that does not hold in infinite-dimensional spaces.

When $Y = \mathbb{R}^m$ and $\Phi\colon X \rightrightarrows \mathbb{R}^m$ is lower hemi-continuous, the pointwise convex hull of Φ is lower hemi-continuous; that is, the correspondence $\Psi\colon X \rightrightarrows Y$ defined by $\Psi(x) = \text{conv}(\Phi(x))$ is lower hemi-continuous. In fact, this result remains true when Y is a metric mixture space and $\Phi\colon X \rightrightarrows Y$ is lower hemi-continuous.

For an example of an upper hemi-continuous correspondence, let X and P be metric spaces, and let $f\colon X \times P \to \mathbb{R}^m$ be a continuous function. Define the correspondence $\Phi\colon P \rightrightarrows X$ by

$$\Phi(p) \quad = \quad \{x \in X \mid f(x,p) = 0\}.$$

Assuming X is compact, the correspondence Φ is upper hemi-continuous. In words, the zeroes of a continuous function vary upper hemi-continuously with parameters. This construction does not generally yield a lower hemi-continuous correspondence, but assume that $X \subseteq \mathbb{R}^m$ is open, that f is not only continuous but for all $p \in P$, the function $f_p(x) = f(x,p)$ is a C^1 function of x, and that for all $p \in P$, zero is a regular value of f. Then the correspondence Φ defined above is lower hemi-continuous. Indeed, consider any $p \in P$ and any open set $V \subseteq \mathbb{R}^m$ with $V \cap \Phi(p) \neq \emptyset$, so there exists $x \in X$ with $f_p(x) = 0$. Since zero is a regular value, the derivative $Df_p(x)$ is non-singular, and then the metric version of the implicit function theorem yields open sets $V \subseteq \mathbb{R}^m$ and $Q \subseteq P$ with $(x,p) \in V \times Q$ and a continuous mapping $g\colon Q \to V$ such that $g(p) = x$ and for all $(x',p') \in V \times Q$, $f(x',p') = 0$ if and only if $g(p') = x'$. Setting $U = g^{-1}(V) \cap Q$, we have an open set such that for all $p' \in U$, $f(g(p'),p') = 0$, so $g(p') \in V \cap \Phi(p')$, as required for lower hemi-continuity.

For another example, let X and Y be metric spaces with Y compact, and let $\Phi\colon X \rightrightarrows Y$ be upper hemi-continuous with nonempty, closed values. Then the correspondence $\Psi\colon X \rightrightarrows \Delta(Y)$ defined by

$$\Psi(x) \quad = \quad \{\mu \in \Delta(Y) \mid \mu(\Phi(x)) = 1\} \quad = \quad \Delta(\Phi(x))$$

is upper hemi-continuous with nonempty, compact, convex values with the Prohorov metric on $\Delta(Y)$. Adding the assumption that X is a separable metric space, a stronger result is possible: then Φ is upper hemi-continuous if and only if Ψ is upper hemi-continuous; and Φ is lower hemi-continuous if and only if Ψ is lower hemi-continuous. For another special case, let X be separable, and define the support correspondence $\sigma\colon \Delta(X) \rightrightarrows X$ by

$$\sigma(\mu) \quad = \quad \mathrm{supp}(\mu).$$

Then σ is lower hemi-continuous and has closed values with the Prohorov metric on $\Delta(X)$.

Let X and P be metric spaces, let $f\colon X \times P \to \mathbb{R}$ be a continuous function, and let $\Phi\colon P \rightrightarrows X$ be a continuous correspondence with nonempty, compact values. Then *Berge's theorem of the maximum* (or simply the *maximum theorem*) states that the correspondence $\Psi\colon P \rightrightarrows X$ defined by

$$\Psi(p) \quad = \quad \arg\max_{x \in \Phi(p)} f(x,p)$$

is upper hemi-continuous and has nonempty, compact values; furthermore, the maximized value function $v\colon P \to \mathbb{R}$ defined by

$$v(p) \quad = \quad \max_{x \in \Phi(p)} f(x,p)$$

is continuous.

Extending the concept of fixed point to a correspondence $\Phi\colon X \rightrightarrows X$, say $x \in X$ is a *fixed point* of Φ if $x \in \Phi(x)$. *Kakutani's fixed point theorem* states that if $X \subseteq \mathbb{R}^n$ is nonempty, compact, and convex, and if $\Phi\colon X \rightrightarrows X$ is an upper hemi-continuous correspondence with nonempty, closed, convex values, then Φ has at least one fixed point. This can be extended to more general spaces. A metric version of *Glicksberg's fixed point theorem* states that if X is a nonempty, compact metric mixture space with quasi-convex metric, and if $\Phi\colon X \rightrightarrows X$ is an upper hemi-continuous correspondence with nonempty, closed, convex values, then Φ has at least one fixed point.

Measurable correspondences. Given metric spaces X and Y, a correspondence $\Phi\colon X \rightrightarrows Y$ is *lower measurable* (or *weakly measurable*) if for every open set $V \subseteq Y$, the set

$$\{x \in X \mid \Phi(x) \cap V \neq \emptyset\}$$

is measurable. If Φ has singleton values, so there is a function $f\colon X \to Y$ such that $\Phi(x) = \{f(x)\}$ for all $x \in X$, then lower measurability of Φ is equivalent to measurability of the function f. Clearly, a lower hemi-continuous correspondence is lower measurable. There are other notions of measurability for correspondences, but they will not be defined here.

Lower measurability is preserved by countable unions, often countable products, and sometimes countable intersections: letting $\{\Phi_i \mid i \in I\}$ be a countable collection of lower measurable correspondences $\Phi_i\colon X \rightrightarrows Y$ indexed by elements of I,

- the correspondence $\Psi\colon X \rightrightarrows Y$ defined by

$$\Psi(x) \quad = \quad \bigcup_{m=1}^{\infty} \Phi_m(x)$$

is lower measurable,

- if X is separable, then the correspondence $\Psi\colon X \rightrightarrows Y^I$ defined by

$$\Psi(x) \quad = \quad \prod_{i\in I} \Phi_i(x)$$

is lower measurable (with the product metric on Y^I),

- if X is separable, if Φ_i has closed values for each $i \in I$, and if for each $x \in X$, there is some $i \in I$ such that $\Phi_i(x)$ is compact, then the correspondence $\Psi\colon X \rightrightarrows Y^I$ defined by

$$\Psi(x) \quad = \quad \bigcap_{i\in I} \Phi_i(x)$$

is lower measurable (with the product metric on Y^I).

A correspondence $\Phi\colon X \rightrightarrows Y$ is lower measurable if and only if the correspondence $\Psi\colon X \rightrightarrows Y$ defined by $\Psi(x) = \operatorname{clos}(\Phi(x))$, the pointwise closure of Φ, is lower measurable.

When $Y = \mathbb{R}^m$ and the correspondence $\Phi\colon X \rightrightarrows \mathbb{R}^m$ is lower measurable, the pointwise convex hull of Φ is lower measurable; that is, the correspondence $\Psi\colon X \rightrightarrows \mathbb{R}^m$ defined by $\Psi(x) = \operatorname{conv}(\Phi(x))$ is lower measurable. And for a separable metric space Y, when $\Phi\colon X \rightrightarrows Y$ has nonempty, compact values, the correspondence Φ is lower measurable if and only if the correspondence $\Psi\colon X \rightrightarrows \Delta(Y)$ defined by

$$\Psi(x) \quad = \quad \{\mu \in \Delta(Y) \mid \mu(\Phi(x)) = 1\} \quad = \quad \Delta(\Phi(x))$$

is lower measurable with the Prohorov metric on $\Delta(Y)$.

Assuming Y is a complete, separable metric space, a special case of the *Kuratowski-Ryll-Nardzewski selection theorem* is that every lower measurable correspondence $\Phi\colon X \rightrightarrows Y$ with nonempty, closed values admits a *measurable selection,* i.e., there is a measurable function $f\colon X \to Y$ such that for all $x \in X$, we have $f(x) \in \Phi(x)$. In fact, more can be said. Again, let Φ have nonempty, closed values. *Castaing's theorem* states that if there is a countable set $\{f_m\}$ of measurable mappings such that for all $x \in X$,

$$\Phi(x) \quad = \quad \text{clos}(\{f_1(x), f_2(x), \ldots, \}),$$

then Φ is lower measurable. The converse holds as well if either Y is complete and separable or if Y is separable and Φ has compact values.

Next, we extend the notion of Carathéodory function to accommodate a metric space of parameters. Given metric spaces X, Y, and P, recall that a function $f\colon X \times P \to Y$ is a *Carathéodory function* if (i) for each $x \in X$, the mapping $f_x\colon P \to Y$ defined by $f_x(p) = f(x,p)$ is measurable, and (ii) for each $p \in P$, the mapping $f_p\colon X \to Y$ defined by $f_p(x) = f(x,p)$ is continuous. Given a Carathéodory function $f\colon X \times P \to Y$, let $\Phi\colon P \rightrightarrows X$ be a lower measurable correspondence with nonempty, compact values. Let $g\colon P \to Y$ be a measurable function such that for each $p \in P$, there exists $x \in X$ with $g(p) = f(x,p)$. Assuming X and Y are separable, a special case of *Filippov's implicit function theorem* establishes that the element x solving the latter equation can be chosen as a measurable function of p. More precisely, the correspondence $\Psi\colon P \rightrightarrows X$ defined by

$$\Psi(p) \quad = \quad \{x \in \Phi(p) \mid f(x,p) = g(p)\}$$

is lower measurable and admits a measurable selection, i.e., there is a measurable function $h\colon P \to X$ such that for all $p \in P$, we have $h(p) \in \Phi(p)$ and $g(p) = f(h(p),p)$.

Again, let $f\colon X \times P \to \mathbb{R}$ be a Carathéodory function, and let $\Phi\colon P \rightrightarrows X$ be a lower measurable correspondence. Assuming X is separable, the *measurable maximum theorem* implies that the correspondence $\Psi\colon P \rightrightarrows X$ defined by

$$\Psi(p) \quad = \quad \arg\max_{x \in \Phi(p)} f(x,p)$$

is lower measurable, admits a measurable selection, and has nonempty, compact values; furthermore, the maximized value function $v\colon P \to \mathbb{R}$ defined by $v(p) = \max_{x \in \Phi(p)} f(x,p)$ is measurable.

If a lower measurable correspondence $\Phi\colon X \rightrightarrows Y$ has nonempty, closed values in a separable metric space Y, then its graph is a measurable subset of $X \times Y$ with the product metric, i.e., $\text{graph}(\Phi) \in \mathcal{B}_{X \times Y}$. An implication is that if we let κ be a probability measure on X, then the subset of transition probabilities from X to Y that place probability one on the values of Φ, i.e.,

$$\{\mu \in \mathcal{R}(Y, X, \kappa) \mid \text{for } \kappa\text{-almost all } x,\ \mu(\Phi(x)|x) = 1\}$$

is a compact subset of $\mathcal{R}(Y, X, \kappa)$ with the narrow convergence metric.

2 Agenda Setting

Let $N = \{1, \ldots, n\}$ be a set of voters, let $j \in N$ be the proposer, and let X be a set of alternatives. Each voter i's preferences over alternatives are represented by the utility function $u_i \colon X \to \mathbb{R}$, and we extend these preferences to lotteries via expected utility. Let $r = (r_1, \ldots, r_n)$ be a vector of *reservation payoffs,* one for each voter. Assume that $X \subseteq \mathbb{R}^d$ is nonempty, compact, and convex, and that for all $i \in N$, u_i is continuous and strictly quasi-concave. In particular, each u_i has a unique maximizer, the *ideal point* of individual i, which we denote by $\hat{x}^i$. Without loss of generality, we assume $r \notin X$.

We consider the following two-period game, known as the *setter model,* or more explicitly as the *take-it-or-leave-it offer game.* (A special case is the well-known ultimatum game.) The proposer proposes an element $x \in X \cup \{r\}$, where $x = r$ means that no proposal is offered, and the game ends. If $x \in X$, then the voters observe this and vote simultaneously on the proposal, with each voter i casting ballot $b_i \in \{0, 1\}$, where $b_i = 1$ is interpreted as a vote to accept and $b_i = 0$ as a vote to reject. The *voting rule* is given by a collection $\mathcal{D}$ of groups such that a proposal passes if and only if the set of voters who vote to accept belongs to this collection, i.e.,

$$\{i \in N \mid b_i = 1\} \in \mathcal{D}.$$

Regarding the voting rule, we assume only that $\mathcal{D}$ is nonempty and *monotonic,* in the sense that for all $G \in \mathcal{D}$ and all $G' \supseteq G$, we have $G' \in \mathcal{D}$. For future reference, we say i is a *dummy voter* if there does not exist a group G such that $G \notin \mathcal{D}$ and $G \cup \{i\} \in \mathcal{D}$, and we let D be the set of dummy voters.

Payoffs are such that if the proposer proposes r, then the game ends with payoffs r_i for each voter i. If $x \in X$ is proposed and passes, then payoffs are $u_i(x)$ for each player; and if $x \in X$ is proposed and fails, then payoffs are r_i for each voter.

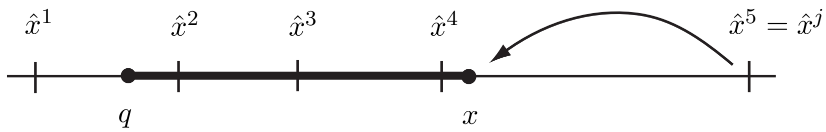

Figure 5: Setter model

In addition, we at times assume:

(a) there exists $x^r \in X$ such that for all $i \in N$, we have $u_i(x^r) \geq r_i$ and such that...

(b) there exists $y^r \in X$ and $G^r \in \mathcal{D}$ such that $u_j(y^r) > u_j(x^r)$ and for all $i \in G^r$, $u_i(y^r) \geq r_i$.

For example, consider Figure 5, where five individuals have Euclidean preferences given by $u_i(x) = -|x - \hat{x}^i|$ and the reservation payoff $r_i = u_i(q)$ is just the utility from a given status quo alternative. Then we can set $x^r = q$, $y^r = x$, and $G^r = \{3, 4, 5\}$ to fulfill (a) and (b). Note that assuming voters accept a proposal closer to their ideal points than the status quo, every proposal in the interior of the darkened interval will pass in a majority rule vote. Alternative x need not pass a priori, because voter 3 is indifferent between q and x, but to resolve the proposer's best response problem, we must assume that 3 votes in favor of x, so it passes and is the optimal proposal for the proposer.

Note, however, that if we move q to the opposite end of the darkened interval in Figure 5, then we cannot satisfy (b).

Given the reservation payoffs r, define the *acceptance set of voter i,* denoted $A_i(r)$, as $A_i(r) = \{x \in X \mid u_i(x) \geq r_i\}$. For a group G, the *acceptance set of G* and the *social acceptance set* are defined by

$$A_G(r) \;=\; \bigcap_{i \in G} A_i(r) \quad \text{and} \quad A(r) \;=\; \bigcup_{G \in \mathcal{D}} A_G(r),$$

respectively. In Figure 5, for example, the social acceptance set is just the darkened interval of alternatives. Note that (a) implies $x^r \in A_i(r)$ for every voter, so $A(r)$ is nonempty, and that (b) is equivalent to

$$u_j(x^r) \;\; < \;\; \sup\{u_j(y) \mid y \in A(r)\}.$$

Moreover, note that by continuity of u_i, each individual acceptance set $A_i(r)$ is closed, and so the social acceptance set is a closed subset of the compact set X, so $A(r)$ is itself compact; thus, the supremum above is actually achieved and can be written as a maximum.

A pure proposal strategy for the proposer j is an element $p_j \in X \cup \{r\}$, where we interpret $p_j = r$ as the option of not making a proposal, in which case payoffs are at the reservation level. A mixed strategy for the proposer is a probability measure π_j on $X \cup \{r\}$, i.e., a distribution on the set from which a proposal is realized, where $\pi_j(Y)$ is the probability that j proposes an alternative in (Borel) Y. A pure strategy for voter i is a (measurable) mapping $a_i\colon X \cup \{r\} \to \{0,1\}$, where we assume by convention that $a_i(r) = 0$. A mixed strategy for voter i is a (measurable) mapping $\alpha_i\colon X \cup \{r\} \to [0,1]$ which maps from the realized proposal to the probability of accepting it. Again, we impose $\alpha_i(r) = 0$. Note that voting subgames typically have multiple equilibria, many of them unreasonable, due to the fact no single voter may be pivotal for the outcome of the vote; for example, if voting is by majority rule and there are three or more voters, then everyone voting to reject x is a Nash equilibrium, regardless of the proposal or the voters' preferences. For this reason, we select subgame perfect equilibria in which no voter uses a dominated strategy.

Given voting strategies $\alpha = (\alpha_1, \ldots, \alpha_n)$, let $\overline{\alpha}(x)$ denote the probability that proposal x passes, i.e., for all $x \in X$,

$$\overline{\alpha}(x) \;\; = \;\; \sum_{G \in \mathcal{D}} \left(\prod_{i \in G} \alpha_i(x) \right) \left(\prod_{i \notin G} (1 - \alpha_i(x)) \right),$$

and $\overline{\alpha}(r) = 0$. Then (π_j, α) is a *subgame perfect Nash equilibrium (in undominated voting strategies)* if and only if

1. π_j puts probability one on optimal proposals, i.e.,

$$\pi_j \left(\arg \max_{x \in X \cup \{r\}} \overline{\alpha}(x) u_j(x) + (1 - \overline{\alpha}(x)) r_j \right) \;\; = \;\; 1,$$

2. for all $x \in X$ and all $i \in N \setminus D$, $u_i(x) > r_i$ implies $\alpha_i(x) = 1$; and $u_i(x) < r_i$ implies $\alpha_i(x) = 0$.

The next proposition establishes existence and gives a characterization of equilibria in the setter model; in particular, the proposer in equilibrium solves a relatively simple constrained optimization problem, and all equilibria are no-delay.

Proposition 2.1. *Assume (a). For every mixed proposal strategy π_j, if π_j puts probability one on optimal proposals in the social acceptance set, i.e.,*

(i) $\pi_j(\arg\max\{u_j(x) \mid x \in A(r)\}) = 1$,

then there exists α such that equilibrium proposals are accepted with probability one, i.e.,

(ii) $\int_X \overline{\alpha}(x)\pi_j(dx) = 1$

and such that (π_j, α) is a pure-strategy subgame perfect equilibrium (in undominated voting strategies). Assuming (b) in addition, if (π_j, α) is a subgame perfect Nash equilibrium (in undominated strategies), then (i) and (ii) hold.

Proof. For existence, assume (a), consider any π_j satisfying (i), and define pure voting strategies for each voter so that for all $x \in X$,

$$a_i(x) \quad = \quad \begin{cases} 1 & \text{if } x \in A_i(r) \\ 0 & \text{else.} \end{cases}$$

Given these voting strategies, note that a proposal x is accepted with probability one if $x \in A(r)$ and otherwise is rejected with probability one. By (a), there exists $x^r \in A(r)$ such that $u_j(x^r) \geq r_j$, which implies that $\max\{u_j(x) \mid x \in A(r)\} \geq r_i$, so π_j is indeed optimal. Since voting strategies are undominated, this completes the existence construction. For the characterization, assume (a) and (b), and let (π_j, α) be any subgame perfect equilibrium (in undominated voting strategies). Suppose that either (i) or (ii) is violated. Then I claim that the proposer's payoff is

$$\begin{aligned} &(1 - \pi_j(A(r)))r_j + \int_{A(r)} [\overline{\alpha}(x)u_j(x) + (1 - \overline{\alpha}(x))r_j]\, \pi_j(dx) \\ &\quad < \max_{x \in A(r)} u_j(x). \end{aligned} \tag{1}$$

The expression for the proposer's payoff on the left-hand side above uses the fact that for all $x \in X \setminus A(r)$, we have $\overline{\alpha}(x) = 0$, so the proposer's payoff is r_j

for such alternatives. To deduce the strict inequality, we consider two cases. If (ii) is violated, then the probability that the proposer receives payoff r_j is positive, i.e.,

$$1 - \pi_j(A(r)) + \int_{A(r)} (1 - \overline{\alpha}(x))\pi_j(dx) \;=\; 1 - \int_X \overline{\alpha}(x)\pi_j(dx) \;>\; 0.$$

Since we have $u_j(x) \leq \max_{x \in A(r)} u_j(x)$ for all $x \in A(r)$, and since $r_j < \max_{x \in A(r)} u_j(x)$ by assumption (b), this gives us the strict inequality claimed in (1). If condition (ii) holds but (i) is violated, then the left-hand side of (1) reduces to $\int_{A(r)} u_j(x)\pi_j(dx)$, and if π_j puts positive probability on proposals y such that $u_j(y) < \max_{x \in A(r)} u_j(x)$, then we again obtain the inequality in (1). This establishes the claim. Now, let x^* be any solution to the constrained maximization problem on the right-hand side of (1), so there exists $G \in \mathcal{D}$ such that for all $i \in G$, we have $u_i(x^*) \geq r_i$. By assumption (a), we have for all $i \in G$ the inequality $u_i(x^r) \geq r_i$, and note that by assumption (b), we have $x^* \neq x^r$. Then for all $\lambda \in (0,1)$, strict quasi-concavity implies that for all $i \in G$, we have $u_i(\lambda x^r + (1-\lambda)x^*) > r_i$, and in particular $\alpha_i(\lambda x^r + (1-\lambda)x^*) = 1$. Thus, by continuity, the proposer's payoff from proposing $\lambda x^r + (1-\lambda)x^*$ for λ approaching zero is

$$\lim_{\lambda \downarrow 0} u_j(\lambda x^r + (1-\lambda)x^*) \;=\; u_j(x^*),$$

which exceeds the payoff from using π_j, a contradiction. □

To see that the no-delay result of Proposition 2.1 relies on assumption (b), return to Figure 5 and move q to the right-hand endpoint of the social acceptance set. Then (b) is violated, and the proposition does not hold: there is an equilibrium in which the proposer proposes $x = q$ and voter 3 votes to reject, so the proposal fails. For that matter there are many equilibria of the model in which the proposer's proposals do not solve the constrained maximization problem: individual 5 can propose any alternative outside the social acceptance set in equilibrium, knowing that it will be rejected.

Proposition 2.1 establishes that under (a) and (b), in every subgame perfect Nash equilibrium (in undominated strategies), the proposer solves a constrained maximization problem and the proposals pass with probability one. This means that we can gain an understanding of equilibria in the setter

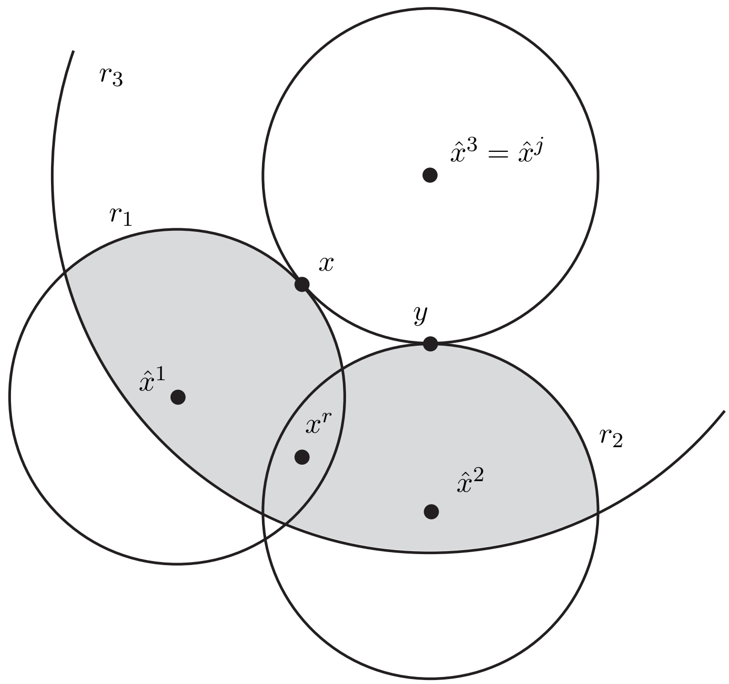

Figure 6: Multiple solutions

model by studying the solutions to the proposer's constrained maximization problem. Let

$$\Psi(r) \quad = \quad \arg\max\{u_j(x) \mid x \in A(r)\}$$

denote the set of solutions to this problem. That there is a maximizer follows from the facts that $A(r)$ is nonempty and compact and that u_j is continuous. Compactness of $\Psi(r)$ follows from compactness of $A(r)$ and continuity of u_j. In general, the set of maximizers need not be singleton or even convex, as depicted in Figure 6. Here, the set of alternatives is two-dimensional, voters have utility $u_i(x) = -||x - \hat{x}^i||$, and voting is by majority rule; individual acceptance sets are indicated by reservation payoffs r_i; and voting is by majority rule, with social acceptance set equal to the gray region. Clearly, the proposer's constrained maximization problem has two solutions, x and y.

When X is one-dimensional, however, the acceptance set $A(r)$ is an interval, and there is a unique solution to the proposer's problem.

Proposition 2.2. *Assume (a) and $X \subseteq \mathbb{R}$. Then $A(r)$ is nonempty and convex, and $\Psi(r)$ is a singleton.*

Proof. By assumption (a), $x^r \in A(r)$, and so the social acceptance set is nonempty. Since it is compact as well, we can define $\underline{x} = \min A(r)$ and $\overline{x} = \max A(r)$. Then there exist $\underline{G}, \overline{G} \in \mathcal{D}$ such that $\underline{x} \in A_{\underline{G}}(r)$ and $\overline{x} \in A_{\overline{G}}(r)$. Consider any $x \in [\underline{x}, \overline{x}]$, and assume without loss of generality that $\underline{x} \leq x \leq x^r$. Then for all $i \in \underline{G}$, we have $\min\{u_i(\underline{x}), u_i(x^r)\} \geq r_i$, so quasi-concavity implies that $u_i(x) \geq r_i$, which implies that $x \in A_{\underline{G}}(r) \subseteq A(r)$. Therefore, $A(r) = [\underline{x}, \overline{x}]$. Since u_j is strictly quasi-concave, it follows that there is a unique solution to $\max_{x \in A(r)} u_j(x)$, as required. □

Next, we examine the continuity properties of solutions to the constrained maximization problem as we vary the reservation payoffs. Let $\mathcal{R}^a$ consist of all vectors $r = (r_1, \ldots, r_n)$ such that (a) is satisfied. Then our definition of Ψ implies a correspondence $\Psi \colon \mathcal{R}^a \rightrightarrows X$, and we are interested in continuity properties of this corresopndence; in particular, our goal is to establish upper hemi-continuity of Ψ. Given that Ψ takes closed values in a compact set, recall that the correspondence Ψ satisfies *upper hemi-continuity* if for all $r \in \mathcal{R}^a$, for all sequences $\{r^m\}$ in $\mathcal{R}^a$ with limit r, and for every convergent sequence $\{x^m\}$ with limit x in X, if $x^m \in \Psi(r^m)$ for all m, then $x \in \Psi(r)$.

The theorem of the maximum provides general conditions for upper hemi-continuity of the solutions to a constrained maximization problem. Here, the objective function is u_j, and the constraint set is given by the social acceptance set, which we can view as a correspondence $A \colon \mathcal{R}^a \rightrightarrows X$. The key sufficient conditions are continuity of the objective function (which is satisfied by assumption) and full continuity of the correspondence A. This means that A is upper hemi-continuous in the sense defined above (replacing Ψ with A) and *lower hemi-continuous,* i.e., for all $r \in \mathcal{R}^a$, for all sequences $\{r^m\}$ in $\mathcal{R}^a$ with limit r, and for all $x \in A(r)$, there is a sequence $\{x^k\}$ in X and subsequence $\{r^{m_k}\}$ such that $x^k \in A(r^{m_k})$ for all k and $x^k \to x$.

Figure 7 illustrates upper hemi-continuity of the set of solutions to the proposer's problem. Here, we let $j = 3$ be proposer, and we increase voter 1's reservation payoffs, given by the sequence $r_1^m \uparrow r_1$. For each m, the unique solution to the proposer's optimization problem is to propose x^m to voter 1.

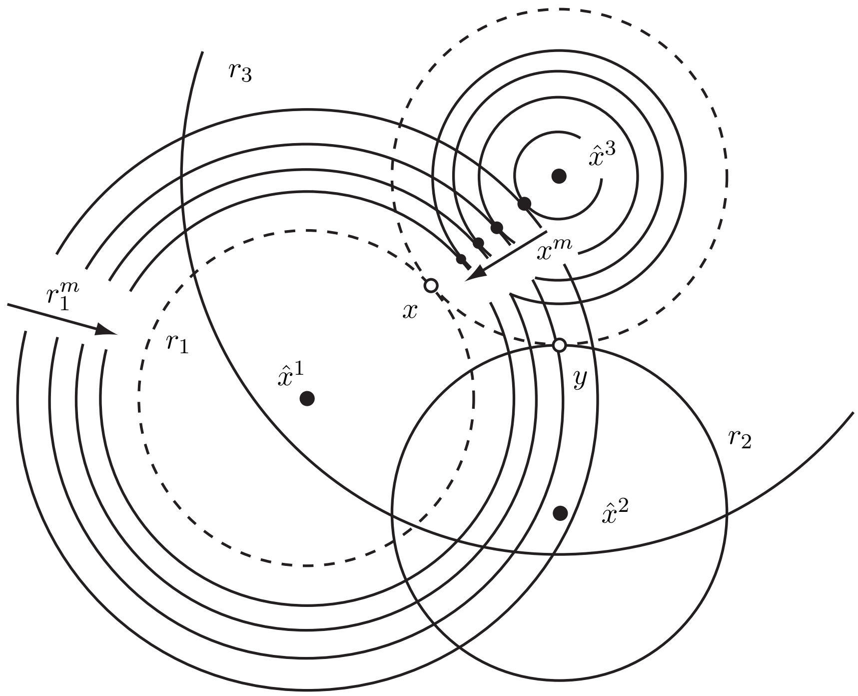

Figure 7: Upper hemi-continuity

In the limit, when 1's reservation payoff reaches r_1 and the optimal proposals reach x, the limiting proposal x is still optimal. Note that it is not uniquely optimal, because y now "pops in" to the set of maximizers, but while that violates lower hemi-continuity (and shows that we cannot generally expect that property), it is consistent with upper hemi-continuity.

To see how upper hemi-continuity of Ψ can be violated, consider Figure 8, where we now decrease the reservation payoff of voter 2, given by the sequence $r_2^m \downarrow r_2$. Now, assume that the level set of 2's utility function at the value r_2 is "thick," the gray ring in the figure. For each m, the unique solution to the proposer 3's maximization problem is $x^m = x$, but in the limit, when 2's reservation payoff reaches r_2, the social acceptance set expands to include the gray area, and y is now socially acceptable. This alternative yields a strictly higher utility to the proposer than x, and therefore x "pops

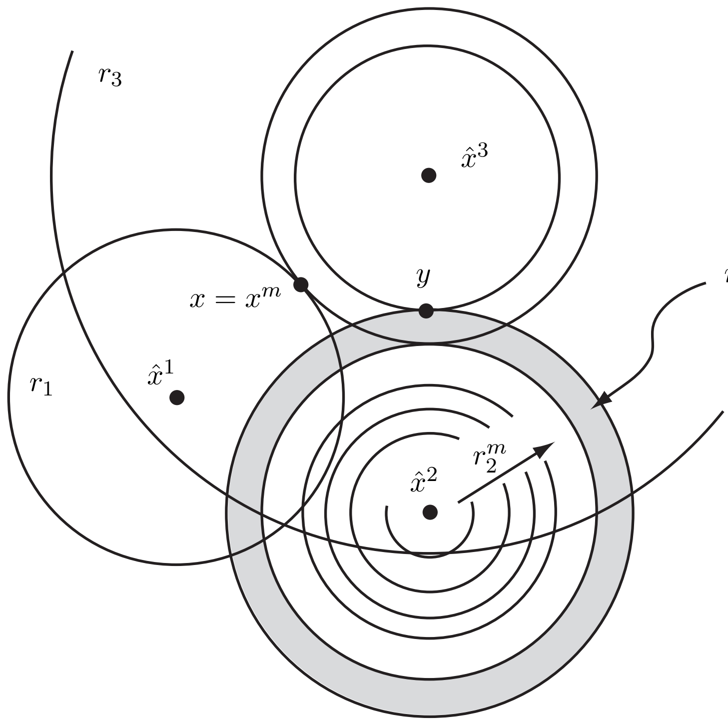

Figure 8: Upper hemi-continuity violation

out" of the solution set, a violation of upper hemi-continuity. Generating this pathology is the thick indifference curve of individual 2, which is consistent with quasi-concavity of u_2 but ruled out by strict quasi-concavity. Thus, the upper hemi-continuity result stated next must use strict quasi-concavity to avoid these kinds of discontinuities.

Proposition 2.3. *The correspondence* $\Psi \colon \mathcal{R}^a \rightrightarrows X$ *is upper hemi-continuous.*

Proof. The result follows by verifying the assumptions of the theorem of the maximum, which ensures that Ψ is upper hemi-continuous if the constraint correspondence A is both upper hemi-continuous and lower hemi-continuous. To verify upper hemi-continuity of the social acceptance set, consider any $r \in \mathcal{R}^a$, any sequence $\{r^m\}$ in $\mathcal{R}^a$ with limit r, and any sequence $\{x^m\}$ with limit x in X, and assume that $x^m \in A(r^m)$ for all m. Then for each m, there is a group $G_m \in \mathcal{D}$ such that $x^m \in A_{G_m}(r^m)$. Since the collection $\mathcal{D}$ is finite, there must be a group G that appears an infinite number of times in the list $G_1, G_2, G_3, \ldots$. We select such a group G and go to the subsequence such that $G = G_m$ (for simplicity, we continue to index the subsequence by m). Consider any $i \in G$. For each m, $x^m \in A_G(r^m)$ implies that $u_i(x^m) \geq r_i^m$, and by continuity this implies $u_i(x) \geq r_i$. Letting $G' = \{i \in N \mid u_i(x) \geq r_i\}$, this argument implies $G \subseteq G'$. By monotonicity, we have $G' \in \mathcal{D}$, and therefore $x \in A(r)$. This establishes upper hemi-continuity of A.

To verify lower hemi-continuity of A, consider any $r \in \mathcal{R}^a$, any sequence $\{r^m\}$ in $\mathcal{R}^a$ with limit r, and any $x \in A(r)$. Then there exists $G \in \mathcal{D}$ such that $x \in A_G(r)$, i.e., for all $i \in G$, we have $u_i(x) \geq r_i$. Using assumption (a), there are two cases. The first case is $x \neq x^r$. For each k, convexity of X implies $x^k = \frac{1}{k}x^r + \frac{k-1}{k}x \in X$, and strict quasi-concavity of u_i implies that for all $i \in G$, we have $u_i(x^k) > r_i$. It follows that for each $i \in G$, there exists M_k^i such that for all $m \geq M_k^i$, we have $u_i(x^k) > r_i^m$. Letting $M_k = \max\{M_k^i \mid i \in G\}$, we set $m_1 = M_1$ and for all $k \geq 2$, we set $m_k = \max\{M_k, m_{k-1} + 1\}$. For each k and for all $i \in G$, we have $u_i(x^k) > r_i^{m_k}$, so $x^k \in A(r^{m_k})$. Thus, we have a subsequence $\{r^{m_k}\}$ such that $x^k \in A(r^{m_k})$ for all k and $x^k \to x$, as required.

The second case is $x = x^r$. For each m, note that $x^{r^m} \in A(r^m)$. For notational clarity, define $y^m = x^{r^m}$. If there is a subsequence of $\{y^m\}$, say $\{y^{m_k}\}$, that converges to x^r, then we set $x^k = y^{m_k}$ to obtain a sequence $\{x^k\}$ and subsequence $\{r^{m_k}\}$ such that $x^k \in A(r^{m_k})$ for all k and $x^k \to x^r = x$. Otherwise, there is no subsequence of $\{y^m\}$ that converges to x^r. Since $\{y^m\}$ lies in X, a compact set, the sequence contains a convergent subsequence with limit $y \neq x^r$. By upper hemi-continuity, we have $y \in A(r)$, so there is a group $G \in \mathcal{D}$ such that $y \in A_G(r)$. Following a line of argument similar to that of the preceding paragraph, we set $x^k = \frac{1}{k}y + \frac{k-1}{k}x^r \in X$, so that for all $i \in G$, we have $u_i(x^k) > r_i$. For each $i \in G$, there exists M_k^i such that for all $m \geq M_k^i$, we have $u_i(x^k) > r_i^m$. Letting $M_k = \max\{M_k^i \mid i \in G\}$, we define m_k as in the preceding paragraph to obtain a sequence $\{x^k\}$ and subsequence $\{r^{m_k}\}$ such that $x^k \in A(r^{m_k})$ for all k and $x^k \to x^r = x$. □

Note that the proof of Proposition 2.3 uses condition (a), which assumes an alternative $x^r \in X$ such that for all $i \in N$, we have $u_i(x^r) \geq r_i$, but it does not use condition (b). This observation will be useful in the analysis of bargaining with a general status quo, later.

go deeper: mention order restriction, need $u_i(x) - r_i$ monotone in i, e.g., quadratic

3 One-shot Legislative Bargaining

In the agenda setting model discussed in the previous section, the reservation payoffs are exogenously given. A natural extension is to consider endogenizing the status quo policy. In this section, we use the insight from bargaining problems to study this issue. We consider the protocol of Baron and Ferejohn (1989) according to which: if the first proposal is accepted, then the game ends; if it is rejected, then we continue to the second period and randomly draw someone to make another proposal; if the second proposal is accepted, then the game ends; if it is rejected, then we continue to the next period and repeat the above process, ad infinitum. This extends the model of the previous section in the sense that, now, the reservation payoff r_i is the discounted expected payoff of individual i determined by the strategies of the players in the continuation game following rejection. Because the game ends once an agreement is reached, the model (albeit infinite-horizon) considers bargaining over a single decision, and so I refer to it as "one-shot" bargaining.

3.1 One-shot Bargaining Framework

To be more precise, the elements of the model are mostly unchanged: we assume $X \subseteq \mathbb{R}^d$ is nonempty, compact, and convex, and with no loss of applicability, we assume $|X| \geq 2$; we assume each u_i is continuous and strictly quasi-concave, and we add the assumption of concavity; we remove the exogenously specified reservation payoff vector r; we add a rate of time discounting $\delta_i \in [0, 1)$ for each individual; and we specify that at the beginning of each period, individual i is recognized as proposer with probability $\rho_i \in [0, 1]$, where of course $\sum_{i=1}^{n} \rho_i = 1$. In particular, each u_i has a unique maximizer, the ideal point of individual i, which we again denote by $\hat{x}^i$. We continue to assume a monotonic voting rule $\mathcal{D}$, and we let D be the set of dummy voters.

Let $C(\mathcal{D})$ be the *core* of the voting rule $\mathcal{D}$, which consists of each alternative x such that there there do not exist a group $G \in \mathcal{D}$ and alternative y strictly preferred to x by all members of G, i.e.,

$$C(\mathcal{D}) \quad = \quad \left\{ x \in X \mid \begin{array}{c} \text{for all } G \in \mathcal{D} \text{ and all } y \in X, \text{ there} \\ \text{exists } i \in G \text{ such that } u_i(x) \geq u_i(y) \end{array} \right\}.$$

We say $\mathcal{D}$ is *proper* if for all $G, G' \in \mathcal{D}$, we have $G \cap G' \neq \emptyset$. It is *strong* if for all $G \subseteq N$, $G \notin \mathcal{D}$ implies $N \setminus G \in \mathcal{D}$. When the voting rule is proper, the *median voter theorem* states that the core is nonempty and consists of the median ideal point if it is unique, or otherwise the interval of alternatives between the two median ideal points. When $\mathcal{D}$ is in addition strong, there is a unique median ideal point, which is in turn the unique element of the core. We say $\mathcal{D}$ is *collegial* if there is some individual belonging to every group in the voting rule, i.e., $\bigcap \mathcal{D} \neq \emptyset$. A typical example of a proper and strong voting rule is majority rule with n odd, while unanimity rule is collegial.

The timing of the model is as follows.

- A proposer i is drawn from the distribution $\rho = (\rho_1, \ldots, \rho_n)$,
- i makes a proposal $x \in X$,
- voters simultaneously decide either to accept or reject x,
- if x is accepted, the game ends; else, move to the next period and repeat.

In general, payoffs are as follows. Let $\overline{u}_i$ be a *status quo payoff* for i. If x is accepted and the game ends in period t, then i's payoff is $(1 - \delta_i^{t-1})\overline{u}_i + \delta_i^{t-1} u_i(x)$. We can interpret this as the discounted sum of payoffs from the stream

$$\underbrace{\overline{u}_i, \ldots, \overline{u}_i}_{t-1 \text{ periods}}, \underbrace{u_i(x), u_i(x), \ldots}_{\text{infinite periods}},$$

normalized by a factor of $1 - \delta_i$. If no proposal is ever accepted, then individual i's payoff is $\overline{u}_i$. Note that the status quo payoff of an individual i is different in principle than his or her reservation payoff, which is taken as primitive in the setter model.

Since the structure of the bargaining game is stationary, we focus on subgame prefect equilibria in stationary strategies, now assuming that players eliminate weakly dominated strategies in voting stages, holding future play fixed. A pure stationary proposal strategy for individual i is $p_i \in X$. A mixed stationary proposal strategy for i is a probability measure $\pi_i \in \Delta(X)$, where in general $\Delta(Y)$ denotes the set of probability measure on Y. A pure stationary voting strategy for i is a measurable mapping $a_i \colon X \to \{0, 1\}$, and a mixed stationary voting strategy for i is $\alpha_i \colon X \to [0, 1]$, where $\alpha_i(x)$ is the probability that i accepts x. In sum, a *mixed stationary strategy* for i

is $\sigma_i = (\pi_i, \alpha_i)$. A strategy profile is denoted by $\sigma = (\sigma_1, \ldots, \sigma_n)$. Let $\overline{\alpha}(x)$ be the probability that x is accepted.

In fact, Baron and Ferejohn (1989) consider a distributive model that is outside the above framework, in which $X = \{(x_1, \ldots, x_n) \in \mathbb{R}^n_+ | \sum_i x_i = 1\}$, $u_i(x) = x_i$, and $\delta_i = \delta$ for all i. At issue is the assumption of strict quasi-concavity of utilities: the latter specification is quasi-concave but not strictly so (when there are at least three agents). They consider stationary equilibria in a symmetric model where each agent has the same recognition probability, i.e., $\rho = (\frac{1}{n}, \ldots, \frac{1}{n})$, and k or more votes are required for a proposal to be accepted. In this case, the ex ante expected payoff of each agent is $\frac{1}{n}$, and the equilibrium proposal by agent i is to propose the minimum possible to the $k-1$ agents indexed below consistent with their acceptance and to keep the remainder, e.g.,

$$p_k = \Big(\underbrace{\frac{\delta}{n}, \ldots, \frac{\delta}{n}}_{1,\ldots,k-1}, \underbrace{1-(k-1)\frac{\delta}{n}}_{k}, 0, \ldots, 0\Big).$$

Eraslan (2002) further shows that the stationary equilibrium payoffs are unique in the symmetric model, and Eraslan and McLennan (2013) extend the uniqueness result to the general distributional model.

3.2 Bad Status Quo

We begin with the model in which delay is "bad," in the sense that for all $i \in N$ and all $x \in X$, we have $u_i(x) \geq \overline{u}_i$. Under that assumption, it is without loss of generality to assume that $\overline{u}_i = 0$ and that for all i and all x, $u_i(x) \geq 0$. We refer to this as the *bad status quo* model. Thus, if x is accepted in period t, then i's payoff is simply $\delta_i^{t-1} u_i(x)$. Finally, note that the assumptions that u_i is strictly quasi-concave and that $|X| \geq 2$ imply $u_i(\hat{x}^i) > 0$. The assumption of a bad status quo may not describe many bargaining environments, but this version of the bargaining model is relatively tractable analytically, and it is closest to the earlier literature on distributive bargaining, so we begin with it. Results on the bad status quo model in the following subsections are taken from Banks and Duggan (2000).

Given σ, individual i's *continuation value,* denoted $V_i(\sigma)$, is the expected discounted sum of payoffs accrued from the beginning of each period, prior to a proposer being selected; it uniquely satisfies the recursion

$$V_i(\sigma) \quad = \quad \sum_{j=1}^{n} \rho_j \int_X \left[\overline{\alpha}(x) u_i(x) + (1 - \overline{\alpha}(x))\delta_i V_i(\sigma)\right] \pi_j(dx).$$

Note that by stationarity, this continuation value is the same in each period. Solving, we obtain

$$V_i(\sigma) \quad = \quad \frac{\sum_j \rho_j \int_X \overline{\alpha}(x) u_i(x) \pi_j(dx)}{1 - \delta_i \sum_j \rho_j \int_X (1 - \overline{\alpha}(x))\pi_j(dx)}.$$

We say the strategy profile σ is *no-delay* if for all i, we have

$$\int_X \overline{\alpha}(x)\pi_i(dx) \quad = \quad 1.$$

In this case, continuation values simplify to

$$V_i(\sigma) \quad = \quad \sum_{j=1}^{n} \rho_j \int_X u_i(x)\pi_j(dx).$$

A *stationary bargaining equilibrium* is a stationary strategy profile σ such that σ is a subgame perfect equilibrium and voting strategies are stage-undominated, i.e.,

1. for all i, π_i puts probability one on optimal proposals, i.e.,

$$\pi_i\left(\arg\max_{x \in X} \overline{\alpha}(x) u_i(x) + (1 - \overline{\alpha}(x))\delta_i V_i(\sigma)\right) \quad = \quad 1,$$

2. for all $x \in X$ and all $i \in N \setminus D$, $u_i(x) > \delta_i V_i(\sigma)$ implies $\alpha_i(x) = 1$; and $u_i(x) < \delta_i V_i(\sigma)$ implies $\alpha_i(x) = 0$.

Therefore, when individual i is selected to propose, the equilibrium conditions (1) and (2) are analogous to the setter model, where we set $r_i = \delta_i V_i(\sigma)$. That is, the reservation payoff of a voter is her discounted continuation value. We define acceptance sets

$$\begin{aligned} A_i(\sigma) &= \{x \in X \mid u_i(x) \geq \delta_i V_i(\sigma)\} \\ A_G(\sigma) &= \bigcap_{i \in G} A_i(\sigma) \\ A(\sigma) &= \bigcup_{G \in \mathcal{D}} A_G(\sigma), \end{aligned}$$

as in the setter model with $r_i = \delta_i V_i(\sigma)$.

The next lemma uses concavity of utility functions to confirm conditions (a) and (b), substituting an arbitrary individual j and the social acceptance set $A(\sigma)$ in condition (b).

Lemma 3.1. *For every profile σ of stationary strategies, we have:*

(a) there exists $x^\sigma \in X$ such that for all $i \in N$, we have $u_i(x^\sigma) \geq \delta_i V_i(\sigma)$, and such that...

(b) for all $j \in N$,

$$\max_{x \in A(\sigma)} u_j(x) \quad > \quad u_j(x^\sigma).$$

Proof. Given any σ, I claim there exist $x^\sigma \in X$ and $\epsilon > 0$ such that x^σ is the ideal point of no individual, i.e., $x^\sigma \neq \hat{x}^i$ for all i, and such that for all $x \in B_\epsilon(x^\sigma) \cap X$ and all $i \in N$, we have $u_i(x) \geq \delta_i V_i(\sigma)$. Indeed, in case there is no possibility that a proposal is accepted, meaning that $\rho_j \int_X \overline{\alpha}(x)\pi_j(dx) = 0$ for all $j \in N$, then every individual's continuation value equals zero. Then $A(\sigma) = X$, and we can choose $x \in X \setminus \{\hat{x}^i \mid i \in N\}$ and $\epsilon > 0$ arbitrarily.

The remaining case is that there is some possibility that a proposal is accepted, i.e., we have $\rho_j \int_X \overline{\alpha}(x)\pi_j(dx) > 0$ for some $j \in N$. Note that

$$V_i(\sigma) \;\leq\; \frac{\sum_j \rho_j \int_X \overline{\alpha}(x) u_i(x) \pi_j(dx)}{1 - \sum_j \rho_j \int_X (1 - \overline{\alpha}(x))\pi_j(dx)} \;=\; \int_X u_i(x)\gamma(dx) \;\leq\; u_i(x^\gamma),$$

where γ is the *continuation distribution* generated by σ and is given by

$$\gamma(Y) \quad = \quad \frac{\sum_j \int_Y \rho_j \bar{\alpha}(x)\pi_j(dx)}{\sum_j \int_X \rho_j \bar{\alpha}(x)\pi_j(dx)},$$

for measurable $Y \subseteq X$; x^γ is the mean of the continuation distribution γ; and the last inequality follows from Jensen's inequality and concavity of u_i. Let $G = \{i \in N \mid V_i(\sigma) > 0\}$ denote the group of individuals with a positive continuation value, and note that for all $i \in N \setminus G$, we have $V_i(\sigma) = 0$, so $A_i(\sigma) = X$. For all $i \in G$, we in fact have $u_i(x^\gamma) \geq V_i(\sigma) > \delta_i V_i(\sigma)$, so by continuity of u_i, there exists $\eta_i > 0$ such that for all $x \in X$ with

$||x - x^\gamma|| < \eta_i$, we have $u_i(x) > \delta_i V_i(\sigma)$. Setting $\eta = \min\{\eta_i \mid i \in G\}$, we can choose $x^\sigma \in B_\eta(x^\gamma) \setminus \{\hat{x}^i \mid i \in N\}$ and $\epsilon > 0$ such that $B_\epsilon(x^\sigma) \subseteq B_\eta(x^\gamma)$ to satisfy the claim.

We use x^σ to fulfill (b). Now consider any proposer j with ideal point $\hat{x}^j$. Convexity of X implies that for all $\lambda \in (0,1)$, we have $\lambda\hat{x}^j + (1-\lambda)x^\sigma \in X$, and strict quasi-concavity of u_j and $x^\sigma \neq \hat{x}^j$ implies that $u_j(\lambda\hat{x}^j + (1-\lambda)x^\sigma) > u_j(x^\sigma)$. Moreover, by the above claim, we conclude that for all $i \in N$, there exists $\lambda_i \in (0,1)$ such that for all $\lambda \in (0, \lambda_i)$, we have $u_i(\lambda\hat{x}^j + (1-\lambda)x^\sigma) \geq \delta_i V_i(\sigma)$. Letting $\lambda = \min\{\lambda_i \mid i \in G\}$, it follows that for all $i \in N$, we have $\lambda\hat{x}^j + (1-\lambda)x^\sigma \in A_i(\sigma)$, and in particular, $\lambda\hat{x}^j + (1-\lambda)x^\sigma \in A(\sigma)$. Finally, this implies that $\max\{u_j(x) \mid x \in A(\sigma)\} > u_j(x^\sigma)$, fulfilling (b). □

We have proved conditions (a) and (b) for an arbitrary strategy profile, and in particular they hold for equilibrium profiles. By Proposition 2.1, we then have the following result: in any equilibrium, π_i puts probability one on the set of solutions to the problem

$$\max_{x \in A(\sigma)} u_i(x),$$

and proposals are accepted with probability one. Thus, there is no delay in equilibrium. An implication is that the continuation distribution has a very simple form: $\gamma = \sum_j \rho_j \pi_j$.

Proposition 3.2. *In every stationary bargaining equilibrium σ, we have for all $j \in N$,*

(i) $\pi_j(\arg\max\{u_j(x) \mid x \in A(\sigma)\}) = 1$, *and*

(ii) $\int_X \overline{\alpha}(x)\pi_j(dx) = 1$.

In one dimension, Propositions 2.2 and 3.2 yield the following necessary condition for equilibria: in any equilibrium, the acceptance set is convex and all proposal strategies are pure. Note that, while there is no delay, individuals may still use non-degenerate voting strategies off the path of play, so technically there may be equilibria that are not in pure strategies; each such equilibrium will, however, be equivalent to a pure strategy equilibrium.

Proposition 3.3. *Assume $d = 1$. In every stationary bargaining equilibrium σ, the acceptance set $A(\sigma)$ is a closed interval, and for all $i \in N$, the proposal strategy π_i places probability one on the unique solution to*

$$\max_{x \in A(\sigma)} u_i(x).$$

A further property of interest is whether it is possible that all individuals propose the same alternative with probability one, i.e., there exists $x^* \in X$ such that for all $i \in N$, we have $\pi_i(\{x^*\}) = 1$. In this case, i's continuation value equals the utility from x^*, i.e., $V_i(\sigma) = u_i(x^*)$. As in the argument for Proposition 3.2, there is an open ball $B_\epsilon(x^*)$ around x^* such that $B_\epsilon(x^*) \cap X \subseteq A(\sigma)$. By strict quasi-concavity of u_i, it is then suboptimal for i to propose x^* unless x^* is the individual's ideal policy. Thus, a necessary condition for "policy coincidence" is that all individuals have the same ideal policy, an obviously restrictive condition.

The next proposition extends this simple observation by deducing a necessary condition for asymptotic policy coincidence as the discount factors of the individuals approach one. We show that if the equilibrium proposal strategies of all agents approach the unit mass on a given alternative as the agents become patient, then that alternative must belong to the core. The result is general with respect to the dimensionality of the set of alternatives and the form of utilities, but of course it takes as an assumption that condition that proposal strategies approach a single alternative; this will be useful in the later analysis of equilibria in the one-dimensional model.

Proposition 3.4. *For each $i \in N$, let $\{\delta_i^m\}$ be a sequence in $[0, 1)$ with $\delta_i^m \to 1$, and for each m, let σ^m be a stationary bargaining equilibrium. If there exists $x^* \in X$ such that for all $i, j \in N$, we have*

$$\int_X u_i(x) \pi_j^m(dx) \to u_i(x^*),$$

then $x^* \in C(\mathcal{D})$.

Proof. Let $\sigma^m = (\pi^m, \alpha^m)$ be an equilibrium for discount factors δ_i^m, $i = 1, \ldots, n$, as in the proposition. Let $V_i^m(\sigma^m)$ and $A_i^m(\sigma^m)$ denote the continuation value and acceptance set of individual i from σ^m in the model with discount factor δ_i^m, and let $A^m(\sigma^m)$ be the corresponding social acceptance set. Suppose, in order to deduce a contradiction, that $x^* \notin C(\mathcal{D})$.

Then there exists $y \in X$ and $G \in \mathcal{D}$ such that for all $j \in G$, we have $u_j(y) > u_j(x^*)$. Consider any $i \in G$, and note that

$$V_i^m(\sigma^m) \;=\; \sum_{j=1}^{n} \rho_j \int_X u_i(x)\pi_j^m(dx) \;\to\; u_i(x^*) \;<\; u_i(y).$$

Thus, for high enough m, we have $y \in A_i^m(\sigma^m)$, and therefore $y \in A^m(\sigma^m)$. But then

$$\begin{aligned} \int_X u_i(x)\pi_i^m(dx) &= \max_{x \in A^m(\sigma^m)} u_i(x) \;\geq\; u_i(y) \\ &> u_i(x^*) \;=\; \lim_{m\to\infty} \int_X u_i(x)\pi_i^m(dx), \end{aligned}$$

where the first equality follows from optimality of equilibrium proposals. Taking the limit of the left-hand side of the above inequality, we obtain a contradiction, as required. □

The antecedent of Proposition 3.4 is that the expected payoff to every agent i from the equilibrium proposal strategy of every agent j converges to the utility from x^*. Because u_i is continuous, this antecedent condition is satisfied if, for example, each π_j^m is degenerate on an alternative x_j^m and $x_j^m \to x^*$. But it is satisfied much more generally, extending the applicability of the proposition. Recall that a sequence $\{\mu^m\}$ of probability measures on X *converges weakly* to a probability measure μ if for every (bounded) continuous function $f\colon X \to \mathbb{R}$, the expected values converge, i.e.,

$$\int_X f(x)\mu^m(dx) \to \int_X f(x)\mu(dx).$$

The case of interest is when μ is degenerate on some alternative x^*, in which case this notion of convergence is permissive: it is satisfied if each μ^m is degenerate on some x^m and $x^m \to x$, but it also allows for the possibility that μ^m is a continuous distribution, as long as the probability mass of μ^m is "piling up" near x^*. Since each u_i is continuous, the antecedent of Proposition 3.4 is satisfied if each π_j^m converges weakly to the unit mass on x^*.

Of course, if the core is empty, i.e., $C(\mathcal{D}) = \emptyset$, then Proposition 3.4 implies that as the agents become patient, their proposal strategies cannot converge weakly to a single alternative.

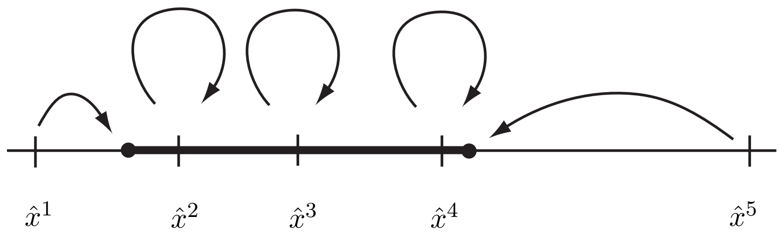

Figure 9: Equilibrium in one dimension

For an example of the form of stationary bargaining equilibria in the one-dimensional model, consider Figure 9, where the social acceptance set is the darkened interval. Here, individuals whose ideal policies are acceptable simply propose those, while other individuals propose the acceptable policies closest to their ideal policies. In multiple dimensions with an empty core, equilibria can still exist and may be of the form in Figure 10, where the social acceptance set is the shaded region, and each individual mixes over his or her (two) optimal proposals.

3.3 Existence of Equilibrium

To this point, we have not established that stationary bargaining equilibria exist generally. The next result fills that lacuna. In fact, Proposition 3.2 establishes that if there is an equilibrium, then it is no delay, so continuation values have the form

$$V_i(\sigma) \quad = \quad \sum_j \rho_j \int_X u_i(x)\pi_j(dx),$$

an insight that guides the construction of the fixed point correspondence used in the proof. Of note, the fixed point argument takes the set of profiles of proposal strategies as its domain and has the following form.

$$\pi = (\pi_1, \ldots, \pi_n) \longrightarrow (V_1(\pi), \ldots, V_n(\pi)) \longrightarrow A(\pi) \longrightarrow \{(\pi'_1, \ldots, \pi'_n)\}$$

Given a profile $\pi = (\pi_1, \ldots, \pi_n)$, we calculate continuation values $V_i(\pi)$ under the assumption that all proposal are accepted; these continuation values determine a social acceptance set $A(\pi)$ assuming that $V_i(\pi)$ is individual

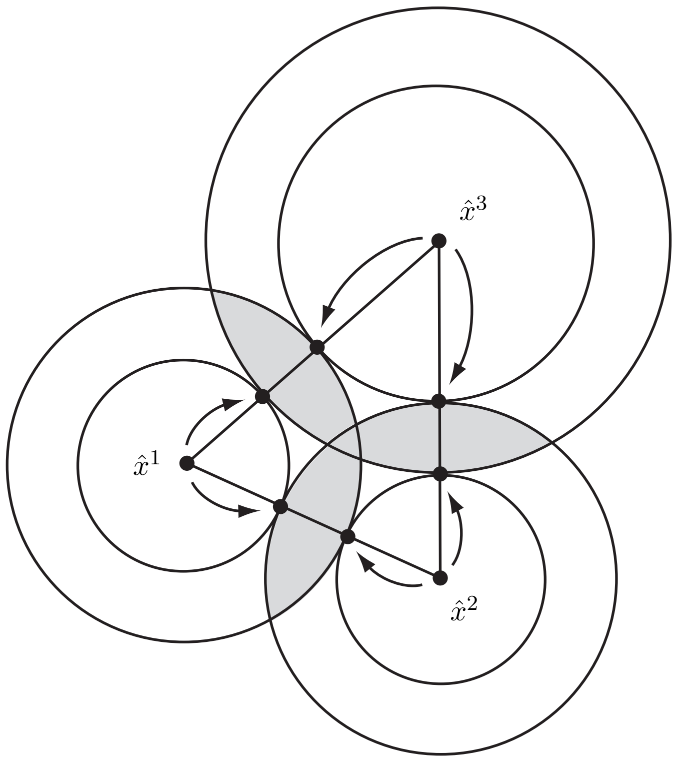

Figure 10: Equilibrium in multiple dimensions

i's continuation value; and we then consider optimal proposal strategies π_i' for each individual, given this social acceptance set. The proof consists of simply of verifying the conditions needed for the application of Glicksberg's fixed point theorem.

Proposition 3.5. *There exists a no-delay stationary bargaining equilibrium.*

Proof. Let $\Delta(X)$ be the set of probability measures on X with the Prohorov metric, and give $\Delta(X)^n$ the product metric. Then $\Delta(X)^n$ is nonempty, convex, and compact. Letting $\pi = (\pi_1, \ldots, \pi_n)$ denote a profile of mixed proposal strategies, we define $V_i \colon \Delta(X)^n \to \mathbb{R}$ for each individual i by

$$V_i(\pi) \;=\; \sum_{j=1}^{n} \rho_j \int_X u_i(x)\pi_j(dx).$$

Note that a sequence $\{\pi^m\}$ converges to π in $\Delta(X)^n$ if and only if for each i, $\pi_i^m \to \pi_i$ weakly; and because u_i is bounded and continuous, we have $\int_X u_i(x)\pi_j^m(dx) \to \int_X u_i(x)\pi_j(dx)$. Therefore, V_i is continuous. Define the mapping $r \colon \Delta(X)^n \to \mathbb{R}^n$ by

$$r(\pi) \;=\; (\delta_1 V_1(\pi), \ldots, \delta_n V_n(\pi)),$$

which is also continuous. To see that condition (a) from the setter model holds when we specify reservation payoffs $r(\pi)$, note by convexity of X that the mean

$$x(\pi) \;=\; \left(\sum_{j=1}^{n} \rho_j \int_X x_1 \pi_j(dx), \ldots, \sum_{j=1}^{n} \rho_j \int_X x_d \pi_j(dx) \right)$$

belongs to X. Moreover, by concavity of u_i, we have $u_i(x(\pi)) \geq V_i(\pi)$ for all i, and therefore $u_i(x(\pi)) \geq r_i(\pi)$. Setting $x^r = x(\pi)$, we fulfill condition (a). Thus, $r(\Delta(X)^n) \subseteq \mathcal{R}^a$, and Proposition 2.3 implies that Ψ_i is upper hemi-continuous on $r(\Delta(X)^n)$.

Define the correspondences $\hat{\Psi}_i \colon \Delta(X)^n \rightrightarrows X$ and $\Psi_i^* \colon \Delta(X)^n \rightrightarrows \Delta(X)$ by

$$\begin{aligned} \hat{\Psi}_i(\pi) &= \Psi_i(r(\pi)) \\ \Psi_i^*(\pi) &= \Delta(\hat{\Psi}_i(\pi)). \end{aligned}$$

To see that $\hat{\Psi}_i$ is upper hemi-continuous, consider any sequence $\{\pi^m\}$ converging to π in $\Delta(X)^n$ and any sequence $\{x^m\}$ converging to x in X such that for all m, $x^m \in \hat{\Psi}_i(\pi^m)$. By continuity of r, we have $r(\pi^m) \to r(\pi)$. Since $x^m \in \Psi_i(r(\pi^m))$ for all m, Proposition 2.3 implies $x \in \Psi_i(r(\pi))$, i.e., $x \in \hat{\Psi}_i(\pi)$, as required. Thus, $\hat{\Psi}_i$ has nonempty, closed values and is upper hemi-continuous. These properties carry over to Ψ_i^*, which is in addition convex-valued. Finally, we define a correspondence $\Phi \colon \Delta(X)^n \rightrightarrows \Delta(X)^n$ by

$$\Phi(\pi) \;=\; \prod_{i=1}^{n} \Psi_i^*(\pi) \;=\; \prod_{i=1}^{n} \Delta(\hat{\Psi}_i(r(\pi))).$$

As the product of correspondences that have nonempty, closed, and convex values and that are upper hemi-continuous, Φ inherits these properties.

Thus, Glicksberg's theorem yields a fixed point $\pi \in \Phi(\pi)$. Then for each i, we specify the voting strategy α_i so that i accepts a proposal x if and only if $u_i(x) \geq \delta_i V_i(\pi)$, and we set $\sigma_i = (\pi_i, \alpha_i)$ and $\sigma = (\sigma_1, \ldots, \sigma_n)$. Given these voting strategies, all proposals are accepted, and therefore $V_i(\sigma) = V_i(\pi)$ and $A(\sigma) = A(\pi)$. It follows that proposal strategies are optimal, and voting strategies are not stage-dominated, so σ is a no-delay stationary bargaining equilibrium. □

It is now straightforward to show that in one dimension, there is an equilibrium in pure strategies. Proposition 3.6 establishes general existence of equilibrium in mixed strategies, and Proposition 3.3 tells us that in such an equilibrium, each individual i's proposal strategy places probability one on the unique solution to $\max_{x \in A(\sigma)} u_i(x)$. Although mixed voting could occur off the path of play, given any such equilibrium, we can redefine voting strategies so that for all $i \in N$ and all $x \in X$, we have

$$\alpha_i(x) \quad = \quad \begin{cases} 1 & \text{if } u_i(x) \geq \delta_i V_i(\sigma), \\ 0 & \text{else.} \end{cases}$$

This modification leaves the social acceptance set unchanged and gives us an equilibrium in pure strategies. That said, I provide a simple, direct proof of existence that uses Brouwer's fixed point theorem and may be more transparent than the general proof.

Proposition 3.6. *Assume $d = 1$. There exists a pure-strategy, no-delay stationary bargaining equilibrium.*

Proof. We give the n-fold product X^n the product metric, making it nonempty, convex, and compact. We define a mapping $\phi\colon X^n \to X^n$ as follows. Given $(p_1, \ldots, p_n) \in X^n$, we define the implied reservation payoff for individual i as

$$r_i(p_1, \ldots, p_n) \quad = \quad \delta_i \sum_j \rho_j u_i(p_j),$$

where we implicitly assume that all proposals are accepted. To see that condition (a) from the setter model is satisfied given these reservation payoffs, define $x(p_1, \ldots, p_n) = \sum_j \rho_j p_j$, and note that by convexity of X, we have $x(p_1, \ldots, p_n) \in X$. By concavity and non-negativity of u_i, we have

$$u_i(x(p_1, \ldots, p_n)) \;\geq\; \sum_j \rho_j u_i(p_j) \;\geq\; r_i(p_1, \ldots, p_n)$$

for all $i \in N$. Setting $x^r = x(p_1, \ldots, p_n)$, we fulfill (a). Defining acceptance sets as in the setter model, we apply Proposition 2.2 to conclude that the social acceptance set

$$A(r_1(p_1, \ldots, p_n), \ldots, r_n(p_1, \ldots, p_n))$$

is a nonempty, closed interval. Therefore, we can define $\phi_i(p_1, \ldots, p_n)$ as the unique solution to

$$\max_{x \in A(r_1(p_1,\ldots,p_n),\ldots,r_n(p_1,\ldots,p_n))} u_i(x).$$

Clearly, $\phi(p_1, \ldots, p_n) = (\phi_1(p_1, \ldots, p_n), \ldots, \phi_n(p_1, \ldots, p_n)) \in X^n$. To prove continuity of each ϕ_i, we note that the reservation payoffs $u_j(p_1, \ldots, p_n)$ are continuous as a function of proposals, and by Proposition 2.3, the correspondence of maximizers is upper hemicontinuous as a function of reservation payoffs; in the present setting, because individual i has a unique maximizer, it follows that this maximizer is a continuous function of reservation payoffs. Thus, $\phi_i(p_1, \ldots, p_n)$ is a composition of two continuous functions, so it is continuous. Finally, Brouwer's theorem implies that ϕ admits a fixed point, $(p_1^*, \ldots, p_n^*) = \phi(p_1^*, \ldots, p_n^*)$. We then define a pure-strategy equilibrium such that each individual i proposes p_i^* with probability one, and i accepts proposal x with probability one if $u_i(x) \geq r_i(p_1^*, \ldots, p_n^*)$ and rejects with probability one otherwise. It is straightforward to verify that these proposal strategies are optimal, and that these voting strategies are not stage-dominated. □

3.4 Core Equivalence

The next result is an asymptotic median voter theorem in the one-dimensional model: as the individuals become arbitrarily patient, equilibrium proposals must converge to a single element of the core. Note that the statement of the next proposition assumes that each individual's equilibrium proposals converge; this is essentially without loss of generality, because they belong to the compact set X, every sequence $\{(p_1^m, \ldots, p_n^m)\}$ of proposal vectors will have a convergent subsequence, and the result can then be applied.

Proposition 3.7. *Assume that $d = 1$; that for all $i \in N$, we have $\rho_i > 0$; and that $\mathcal{D}$ is proper. For each $i \in N$, let $\{\delta_i^m\}$ be a sequence in $[0, 1)$ with $\delta_i^m \to 1$, and for each m, let σ^m be a stationary bargaining equilibrium such that individual i proposes $p_i^m \in X$ and $\{p_i^m\}$ has limit $p_i \in X$. Then there exists $x^* \in C(\mathcal{D})$ such that $p_1 = \cdots = p_n = x^*$.*

Proof. For each m, let $V_i^m(\sigma^m)$ and $A_i^m(\sigma^m)$ be the continuation value and acceptance set of individual i in equilibrium σ^m for the model with discount factors δ_i^m, $i = 1, \ldots, n$, and let $A^m(\sigma^m) = [\underline{x}^m, \overline{x}^m]$ be the social acceptance set. I claim that $\overline{x}^m - \underline{x}^m \to 0$, for suppose otherwise. Going to a subsequence (still indexed by m for simplicity), we can then assume that there exist $\underline{x}, \overline{x} \in X$ such that

$$\underline{x} \;=\; \lim_{m\to\infty} \underline{x}^m \;<\; \lim_{m\to\infty} \overline{x}^m \;=\; \overline{x}.$$

For all m, $\underline{x}^m, \overline{x}^m \in A^m(\sigma^m)$ implies that there exist $\underline{G}^m, \overline{G}^m \in \mathcal{D}$ such that $\underline{x}^m \in A^m_{\underline{G}^m}(\sigma^m)$ and $\overline{x}^m \in A^m_{\overline{G}^m}(\sigma^m)$. Since $\mathcal{D}$ is proper, there exists $i_m \in \underline{G}^m \cap \overline{G}^m$. Since n is finite, there is some individual i who appears infinitely many times in the list $i_1, i_2, i_3, \ldots$. We go to a further subsequence consisting of m such that $i_m = i$. Continuing to index this subsequence by m, it follows that for all m, we have

$$\min\{u_i(\underline{x}^m), u_i(\overline{x}^m)\} \;\;\geq\;\; \delta_i^m V_i^m(\sigma^m) \;\;=\;\; \delta_i^m \sum_j \rho_j u_i(p_j^m).$$

Taking limits, the above inequality and continuity of u_i imply that

$$\min\{u_i(\underline{x}), u_i(\overline{x})\} \;\;\geq\;\; \sum_j \rho_j u_i(p_j).$$

But, again using continuity of u_i, we have

$$\begin{aligned} u_i(p_i) \;&=\; \lim_{m\to\infty} u_i(p_i^m) \;=\; \lim_{m\to\infty} \max_{x\in[\underline{x}^m,\overline{x}^m]} u_i(x) \\ &=\; \max_{x\in[\underline{x},\overline{x}]} u_i(x) \;>\; \min_{x\in[\underline{x},\overline{x}]} u_i(x), \end{aligned}$$

where the third equality uses the theorem of the maximum, and the inequality follows from $\underline{x} < \overline{x}$ and strict quasi-concavity. But, using $\rho_i > 0$ and again using strict quasi-concavity, we then have

$$\sum_j \rho_j u_i(p_j) \;>\; \min_{x\in[\underline{x},\overline{x}]} u_i(x) \;=\; \min\{u_i(\underline{x}), u_i(\overline{x})\},$$

a contradiction. We conclude that, as claimed, $\overline{x}^m - \underline{x}^m \to 0$. Thus, there exists x^* such that $p_1 = \cdots = p_n = x^*$. Finally, Proposition 3.4 implies that $x^* \in C(\mathcal{D})$, as required. □

Whereas Proposition 3.4 shows that a necessary condition for equilibrium proposal strategies to converge to alternative x is that it belongs to the core, the next corollary assumes $\mathcal{D}$ is strong and shows that belonging to the core is actually sufficient for the convergence result. The result follows directly from the preceding proposition after noting that when $\mathcal{D}$ is proper and strong, the core $C(\mathcal{D})$ consists of a single alternative.

Corollary 3.8. *Assume that $d = 1$; that for all $i \in N$, we have $\rho_i > 0$; and that $\mathcal{D}$ is proper and strong. For each $i \in N$, let $\{\delta_i^m\}$ be a sequence in $[0, 1)$ with $\delta_i^m \to 1$, and for each m, let σ^m be a stationary bargaining equilibrium such that individual i proposes $p_i^m \in X$. Then for all $i \in N$, we have $p_i^m \to x^*$, where $C(\mathcal{D}) = \{x^*\}$.*

The proof of Proposition 3.7 relies on the fact that if the diameter of the social acceptance set did not converge to zero as individuals became patient, then some individual would have to vote to accept the worst possible proposal. As long as the individual's recognition probability is positive, however, this is not possible in equilibrium, for the individual would strictly prefer to reject: the cost is that there may be delay for one or more periods, but this cost goes to zero as the agent becomes patient; the benefit is that with positive probability, she will be recognized as proposer and be able to propose the best alternative in the social acceptance set.

The result does not extend directly to multiple dimensions, but it does hold if the voting rule $\mathcal{D}$ is collegial, in the sense that $\bigcap \mathcal{D} \neq \emptyset$. The next proposition states this formally, now in terms of mixed proposal strategies converging to the unit mass on a single element of the core.

Proposition 3.9. *Assume that for all $i \in N$, we have $\rho_i > 0$; and that $\mathcal{D}$ is collegial. For each $i \in N$, let $\{\delta_i^m\}$ be a sequence in $[0, 1)$ with $\delta_i^m \to 1$, and for each m, let σ^m be a stationary bargaining equilibrium such that $\{\pi_i^m\}$ has weak limit $\pi_i \in \Delta(X)$. Then there exists $x^* \in C(\mathcal{D})$ such that $\pi_1(\{x^*\}) = \cdots = \pi_n(\{x^*\}) = 1$.*

In fact, Propositions 3.7 and 3.9 have a unified proof that exploits the deep structure of the voting rule. Define the *Nakamura number* of $\mathcal{D}$ as follows:

if $\mathcal{D}$ is not collegial, then

$$N(\mathcal{D}) \;\; = \;\; \left| \left\{ |\mathcal{G}| \mid \;\; \mathcal{G} \subseteq \mathcal{D} \text{ and } \bigcap \mathcal{G} = \emptyset \;\; \right\} \right|,$$

i.e., it is the size of the smallest collection of groups in the voting rule having empty intersection; otherwise, if $\mathcal{D}$ is collegial, set $N(\mathcal{D}) = \infty$. The key step in the core convergence proof is to show that equilibrium proposal strategies of all individuals converge to the unit mass on one alternative x^*. (The argument that this alternative belongs to the core is relatively straightforward.) Note that $A_G(\sigma^m)$ is a convex set, and so the maximization problem $\max_{x \in A_G^m(\sigma^m)} u_i(x)$ has a unique solution. It follows that

$$\max_{x \in A^m(\sigma^m)} u_i(x)$$

has at most $|\mathcal{D}|$ solutions, and so π_i^m has finite support, denoted $\text{supp}(\pi_i^m)$, and puts positive probability on no more than $|\mathcal{D}|$ alternatives. Suppose that the diameter of these supports sets does not converge to zero, i.e.,

$$\text{diam}\left(\bigcup_{i \in N} \text{supp}(\pi_i^m) \right) \not\to 0.$$

I claim that for sufficiently high m, it must be that the union of support sets, $\bigcup_{i \in N} \text{supp}(\pi_i^m)$, has at least $N(\mathcal{D})$ extreme points. Indeed, if this were not the case, then for arbitrarily high m, the set E^m of extreme points would have fewer than $N(\mathcal{D})$ elements. For each $z \in E^m$, there is a group $G_z \in \mathcal{D}$ such that $z \in A_{G_z}^m(\sigma^m)$, and $|\{G_z \mid z \in E^m\}| < N(\mathcal{D})$. Therefore, there is an individual $i \in \bigcap_{z \in E^m} G_z$ who votes to accept every extreme point. But because u_i is strictly quasi-concave, the worst alternative for i that is proposed with positive probability is an element of E^m, so i accepts the worst possible alternative. As in the proof of Proposition 3.7, this leads to a contradiction.

Applied to the one-dimensional model with $\mathcal{D}$ proper, it follows that if the diameter of $\{p_i^m \mid i \in N\}$ does not go to zero, then $\bigcup_{i \in N} \text{supp}(\pi_i^m)$ must have at least $N(\mathcal{D})$ extreme points; but in one dimension, a set can have at most two extreme points, and if $\mathcal{D}$ is proper, then $N(\mathcal{D}) \geq 3$, an impossibility. In the multidimensional model with $\mathcal{D}$ collegial, the set $\bigcup_{i \in N} \text{supp}(\pi_i^m)$ has at most $|\mathcal{D}|$ extreme points, but the Nakamura number is infinite, again an impossibility. Thus, a key step in the proof of core convergence relies on a unified analysis of the number of extreme proposals relative to the Nakamura number of the voting rule.

Example: To illustrate the core convergence results for the one-dimensional model, assume utilities are quadratic (plus a common constant to ensure non-negativity), i.e., $u_i(x) = -(x - \hat{x}^i)^2 + c$, and $\mathcal{D}$ is proper and strong. Let individual k be such that $C(\mathcal{D}) = \{\hat{x}^k\}$, and let δ be a common discount factor. Then individual k is a representative voter, and in every equilibrium σ, the social acceptance set coincides with the acceptance set of k, i.e., $A(\sigma) = A_k(\sigma)$. For general $\delta \in [0, 1)$, Proposition 3.6 yields a pure strategy equilibrium, which has the form in Figure 9 with social acceptance set $A(\sigma) = [\hat{x}^k - \Delta, \hat{x}^k + \Delta]$ and Δ satisfying

$$\begin{aligned} & -\Delta^2 + c \\ & = \delta \Bigg[\sum_{j:\hat{x}^j < \hat{x}^k - \Delta} \rho_j(-\Delta^2 + c) + \sum_{j:\hat{x}^k - \Delta \leq \hat{x}^j \leq \hat{x}^k + \Delta} \rho_j(-(\hat{x}^j - \hat{x}^k)^2 + c) \\ & \qquad + \sum_{j:\hat{x}^k + \Delta < \hat{x}^j} \rho_j(-\Delta^2 + c) \Bigg]. \end{aligned}$$

In the limiting case with $\delta = 1$, the above equation becomes

$$\Delta^2 \sum_{j:\hat{x}^k - \Delta \leq \hat{x}^j \leq \hat{x}^k + \Delta} \rho_j = \sum_{j:\hat{x}^k - \Delta \leq \hat{x}^j \leq \hat{x}^k + \Delta} \rho_j(\hat{x}^j - \hat{x}^k)^2,$$

and since $\rho_k > 0$, this can only be satisfied if $\Delta = 0$.

Consistent with Proposition 3.8, the radius Δ of the equilibrium social acceptance set therefore goes to zero as δ goes to one; assuming no other individual has ideal point at $\hat{x}^k$, then individual k is the only individual whose ideal point belongs to the acceptance set. The condition defining Δ simplifies to

$$\Delta = \sqrt{\frac{(1-\delta)c}{1 - \delta(1 - \rho_k)}},$$

and it is simple to do comparative statics. In particular, the social acceptance set shrinks as we increase the recognition probability of the core individual or the discount factor, and clearly $\Delta \to 0$ as $\delta \to 1$. □

Cho and Duggan (2003) show that if utilities are quadratic (plus a common constant to ensure non-negativity), if $\mathcal{D}$ is proper and strong, and if individuals have a common discount factor, then there is exactly one stationary bargaining equilibrium. For general continuous, concave, strictly quasi-concave utilities and for general voting rules, there may be multiple equilibria, but the authors show that the social acceptance sets of different equilibria are nested: if σ and σ' are equilibria, then either $A(\sigma) \subseteq A(\sigma')$ or $A(\sigma') \subseteq A(\sigma)$. The uniqueness result is extended by Cardona and Ponsati (2011) to the case of common curvature of utilities, i.e., there exists $v\colon \mathbb{R} \to \mathbb{R}$ such that $u_i(x) = v(x - \hat{x}^i)$ for all i, and a quota rule, i.e., $\mathcal{D} = \{G \mid |G| \geq q\}$ and $\frac{n}{2} < q \leq n$.

While the dynamic median voter theorem is stated above only for stationary bargaining equilibria, it is known that in repeated games (and in a class of stochastic games), arbitrary payoff vectors can be supported in subgame perfect equilibrium when players are patient. Cho and Duggan (2009) prove an "anti-folk theorem" for the one-dimensional model: if the number of individuals is odd, if voting is sequential (and certain orders of voting have positive probability) and by majority rule, if the median voter's recognition probability is positive, and if individuals have a common discount factor, then the set of alternatives that are supportable as outcomes of subgame perfect equilibria in the bargaining model converges to the median ideal point as the individuals become patient.

The result carries over to simultaneous voting (in stage-undominated strategies), as long as we restrict attention to equilibria in which players do not condition their actions on the votes of particular voters (e.g., when voting is by secret ballot). Cho and Duggan (2013) show that when voting is simultaneous and players are allowed to condition on votes of particular individuals, the elimination of stage-game dominated votes loses its bite, and a strong folk theorem holds: in the model with a bad status quo, arbitrary alternatives can be supported as outcomes of subgame perfect equilibria in undominated voting strategies for *arbitrary positive discount factors.*

3.5 General Status Quo Alternative

We now allow for any status quo alternative $q \in X$ and assume that for all $i \in N$, the status quo payoff is just the utility from q, i.e., $u_i(q) = \overline{u}_i$. We no longer impose any sign restrictions on utilities, and there may be individuals who prefer q to other alternatives, i.e., we can have $u_i(q) > u_i(x)$ for some x. In fact, it could be that the status quo is the ideal point of some (or every!) individual. In the spatial setting, where alternatives are often viewed as choices of public policy, the *general status quo* model may have significantly greater applicability than the bad status quo model. An important restriction that we now add is that there is a common discount factor: there exists $\delta \in [0, 1)$ such that for all i, $\delta_i = \delta$. Results on the general status quo model in the following subsections are taken from Banks and Duggan (2006).

The assumption of a common discount factor aside, the general status quo model does not, strictly speaking, generalize the bad status quo model: given an instance of the bad status quo model, it may be that there does not exist an alternative $q \in X$ such that $u_i(q) = 0$ for all $i \in N$, and it could in fact be that every alternative is Pareto optimal, precluding the possibility of such a q. Given such a model, however, we can capture it indirectly in the current framework by expanding the set of alternatives and extending utilities.

To provide some detail, suppose that in the bad status quo model, X is an interval, and all utilities are non-negative, i.e., $u_i(x) \geq 0$. Now construct a general status quo model in which we imbed X in a two-dimensional set X' of alternatives: let $\tilde{X} = \{\tilde{x} \in \mathbb{R}^2_+ \mid \tilde{x}_1 + \tilde{x}_2 \leq 1\}$ be a new set of alternatives, where we let the face $\{\tilde{x} \in \tilde{X} \mid \tilde{x}_1 + \tilde{x}_2 = 1\}$ correspond to the original set of alternatives. Define utility functions $\tilde{u}_i$ in the general status quo model so that on this face, $\tilde{u}_i$ is equivalent to u_i, and extend $\tilde{u}_i$ to the rest of $\tilde{X}$ in the following manner: given $\tilde{x} \in \tilde{X}$ such that $0 < \tilde{x}_1 + \tilde{x}_2 < 1$, define

$$\tilde{u}_i(\tilde{x}) \quad = \quad (\tilde{x}_1 + \tilde{x}_2)\tilde{u}_i\left(\frac{1}{\tilde{x}_1 + \tilde{x}_2}\tilde{x}\right),$$

and define $\tilde{u}_i(0) = 0$. That is, we scale $\tilde{x}$ up to obtain an alternative in the original model, calculate the utility from that scaled alternative, then scale the utility down by a factor of $\tilde{x}_1 + \tilde{x}_2$. The new utility function $\tilde{u}_i$ is continuous, strictly quasi-concave, and concave. We then set $\tilde{q} = 0$ in the transformed model. This gives us a general status quo model with an expanded set of alternatives and with $\tilde{q} \in \tilde{X}$ such that $\tilde{u}_i(\tilde{q}) = 0$ for all $i \in N$. Of course, we have added alternatives, but every alternative $\tilde{x} \in \tilde{X}$ with $\tilde{x}_1 + \tilde{x}_2 < 1$ is Pareto dominated, and they will not affect the equilibrium outcomes of bargaining.

An interesting feature of the bad status quo model is that delay cannot occur in equilibrium. In contrast, it is easy to see that delay can occur in the general model: if $q = \hat{x}^i$ is the ideal point of every individual, then there is trivially an equilibrium in which every individual proposes $x \neq q$, and all proposals are rejected. Of course, delay is good in this example, and what is more, there is no initial period in which the status quo obtains before eventual acceptance of a different alternative; rather, proposals are always rejected, and the status quo remains in place by default. It is a simple matter to specify that every individual propose the status quo, and then to specify mixed voting strategies so that proposals are initially rejected with positive probability, but eventually the status quo is proposed and accepted. Again, however, the status quo is in effect in every period. Thus, delay in this example may not correspond to the intuitive meaning. This points us to a distinction that becomes important in the general model.

We say a stationary strategy profile σ is *static* if

$$\sum_j \rho_j \int_{X\setminus\{q\}} \overline{\alpha}(x)\pi_j(dx) \quad = \quad 0,$$

or in words, with probability one, no individual proposes an alternative that is different from the status quo and is accepted. Otherwise, the profile is *non-static.* This concept is distinct from delay, giving us a fourfold classification of strategy profiles.

	static	non-static
delay	$qq\|qq\cdots$	$qq\|xx\cdots$
no-delay	$\|qqq\cdots$	$\|xxx\cdots$

In the above table, the vertical bar | indicates that a proposal is accepted, followed by the sequence of that policy ad infinitum. In the bad status quo model, all equilibria belonged to the bottom right cell; we have seen that in the general status quo model, we can have equilibria in the upper and lower left cells.

Given σ, individual i's continuation value, $V_i(\sigma)$, now uniquely satisfies the recursion

$$V_i(\sigma) \quad = \quad \sum_{j=1}^{n} \rho_j \int_X \left[\overline{\alpha}(x)u_i(x) + (1-\overline{\alpha}(x))((1-\delta)u_i(q) + \delta V_i(\sigma))\right]\pi_j(dx).$$

Solving, we have

$$V_i(\sigma) = \frac{\sum_j \rho_j \int_X \overline{\alpha}(x)u_i(x)\pi_j(dx) + (1-\delta)u_i(q)\sum_j \rho_j \int_X (1-\overline{\alpha}(x))\pi_j(dx)}{1-\delta\sum_j \rho_j \int_X (1-\overline{\alpha}(x))\pi_j(dx)},$$

where the expression $\sum_j \rho_j(1-\overline{\alpha}(x))\pi_j(dx)$ is the ex ante (prior to a proposer being recognized) probability that a proposal is rejected. In case σ is no-delay, the latter probability is zero, and the continuation value simplifies to

$$V_i(\sigma) = \sum_j \rho_j \int_X u_i(x)\pi_j(dx),$$

just as in the bad status quo model.

A *stationary bargaining equilibrium* is a stationary strategy profile σ such that σ is a subgame perfect equilibrium and voting strategies are stage-undominated, i.e.,

1. for all i, π_i puts probability one on optimal proposals, i.e.,

$$\pi_i\left(\arg\max_{x\in X} \overline{\alpha}(x)u_i(x) + (1-\overline{\alpha}(x))((1-\delta)u_i(q)+\delta V_i(\sigma))\right) = 1,$$

2. for all $x \in X$ and all $i \in N \setminus D$, $u_i(x) > (1-\delta)u_i(q) + \delta V_i(\sigma)$ implies $\alpha_i(x) = 1$; and $u_i(x) < (1-\delta)u_i(q) + \delta V_i(\sigma)$ implies $\alpha_i(x) = 0$.

This is analogous to the bad status quo model and the setter model, with $r_i = (1-\delta)u_i(q) + \delta V_i(\sigma)$. We define acceptance sets in the usual way, subject to that modification:

$$\begin{aligned} A_i(\sigma) &= \{x \in X \mid u_i(x) \geq (1-\delta)u_i(q) + \delta V_i(\sigma)\} \\ A_G(\sigma) &= \bigcap_{i\in G} A_i(\sigma) \\ A(\sigma) &= \bigcup_{G\in\mathcal{D}} A_G(\sigma). \end{aligned}$$

The next lemma uses concavity of utility functions and the assumption of a common discount factor to confirm condition (a) from the setter model, suitably translated to the bargaining model.

Lemma 3.10. *For every profile σ of stationary strategies, we have:*

(a) there exists $x^\sigma \in X$ such that for all $i \in N$, we have

$$u_i(x^\sigma) \geq (1-\delta)u_i(q) + \delta V_i(\sigma).$$

Proof. If there is no possibility that a proposal is accepted, meaning that $\rho_j \int_X \overline{\alpha}(x)\pi_j(dx) = 0$ for all $j \in N$, then $V_i(\sigma) = u_i(q)$ for all $i \in N$, and we set $x^\sigma = q$ to satisfy (a). Next, suppose there is some possibility that a proposal is accepted, i.e., for some $j \in N$, we have $\rho_j \int_X \overline{\alpha}(x)\pi_j(dx) > 0$. Then we can write $V_i(\sigma) = \int_X u_i(x)\nu(dx)$ for a probability distribution ν defined as follows: for all measurable $Y \subseteq X$, if $q \notin Y$, then

$$\nu(Y) \quad = \quad \frac{\sum_j \rho_j \int_Y \overline{\alpha}(x)\pi_j(dx)}{1 - \delta \sum_j \rho_j \int_X (1 - \overline{\alpha}(x))\pi_j(dx)},$$

and if $q \in Y$, then

$$\nu(Y) \quad = \quad \frac{\sum_j \rho_j \int_Y \overline{\alpha}(x)\pi_j(dx) + (1-\delta)\sum_j \rho_j (1 - \overline{\alpha}(x))\pi_j(dx)}{1 - \delta \sum_j \rho_j \int_X (1 - \overline{\alpha}(x))\pi_j(dx)}.$$

Letting μ denote the probability distribution that places probability one on q, we then define the *continuation distribution* generated by σ as the convex combination $\gamma = (1-\delta)\mu + \delta\nu$. Then for all $i \in N$, we have

$$(1-\delta)u_i(q) + \delta V_i(\sigma) \quad = \quad \int_X u_i(x)\gamma(dx),$$

so we can write the reservation payoff of individual i as an expected payoff with respect to a probability distribution that is common to all individuals. Letting x^γ denote the mean of the continuation distribution γ, we can set $x^\sigma = x^\gamma$, and then concavity of u_i implies

$$u_i(x^\sigma) \;\; \geq \;\; \int_X u_i(x)\gamma(dx) \;\; = \;\; (1-\delta)u_i(q) + \delta V_i(\sigma)$$

for all $i \in N$, fulfilling (a). □

3.6 Possibility of Delay

Absent from the previous result is a statement that condition (b) from the setter model (or the bad status quo model) holds, and indeed it may not: in the trivial example above, where the status quo is the ideal point of every individual, the equilibrium is static, so $V_i(\sigma) = u_i(q)$ for every individual, and we must set $x^\sigma = q$, and there is no way to give a proposer greater utility than from x^σ, violating (b).

There exist less trivial equilibria exhibiting delay. The next example shows that there are no-delay equilibria that are non-static, occupying the upper right cell of the above table. Note that individuals are risk neutral in the example.

Example: Let $X = [0,1]$, $n = 3$, $\mathcal{D}$ equal to majority rule, $q = 0$, $\rho_i = \frac{1}{3}$ for all i, and $u_1(x) = 1 - x$ and $u_2(x) = u_3(x) = x$. Consider the strategy profile σ such that individual 1 proposes $p_1 = 0$ and individuals 2 and 3 propose $p_2 = p_3 = 1$, and individual 1 accepts a proposal if and only if it belongs to the interval $[0, \frac{2\delta}{3-\delta}]$ and 2 and 3 accept a proposal if and only if it belongs to $[\frac{2\delta}{3-\delta}, 1]$. Note that the social acceptance set given these strategies is $A(\sigma) = [\frac{2\delta}{3-\delta}, 1]$, so individual 1's proposals are rejected, while individuals 2 and 3 propose their ideal point, which passes. Using the above specifications of continuation values, we have

$$\begin{aligned} V_1(\sigma) &= \frac{0 + (1-\delta)(1)(\frac{1}{3})}{1 - \delta\frac{1}{3}} = \frac{1-\delta}{3-\delta} \\ V_i(\sigma) &= \frac{(\frac{2}{3})(1) + 0}{1 - \delta\frac{1}{3}} = \frac{2}{3-\delta}, \end{aligned}$$

for $i = 2, 3$. Individual 1 is indifferent between a proposal x and rejection when

$$1 - x = (1-\delta)(1) + \delta\left(\frac{1-\delta}{3-\delta}\right),$$

or equivalently, $x = \frac{2\delta}{3-\delta}$, so the individual's voting strategy is not stage-dominated. Individual $i = 2, 3$ is indifferent when

$$x = (1-\delta)(0) + \delta\left(\frac{2}{3-\delta}\right) = \frac{2}{3-\delta},$$

so individual 2's and 3's voting strategies are not stage-dominated. Clearly, the proposals of 2 and 3 are optimal. And for individual 1, there is no alternative in the social acceptance set better than $\frac{2\delta}{3-\delta}$, and we have shown that the individual is indifferent between accepting this alternative and proposing $p_1 = 0$, which is rejected. Thus, σ is a stationary bargaining equilibrium, and there is a positive probability of delay before alternative $x = 1$ is eventually accepted. □

Note that in the above example of a non-static equilibrium exhibiting delay, the individuals are risk neutral. In fact, under weak background conditions, strict concavity of utility functions precludes this phenomenon.

Lemma 3.11. *Assume that $\delta > 0$, and that for all $i \in N$, u_i is strictly concave. For every non-static profile σ of stationary strategies, we have:*

(a) there exists $x^\sigma \in X$ such that for all $i \in N$, we have

$$u_i(x^\sigma) \geq (1-\delta)u_i(q) + \delta V_i(\sigma),$$

(b) for all $j \in N$,

$$\max_{x \in A(\sigma)} u_j(x) \quad > \quad u_j(x^\sigma).$$

Proof. Define x^γ as in the proof of Lemma 3.10. Because $\delta > 0$ and σ is non-static, it follows that the continuation distribution γ is non-degenerate, and therefore strict concavity implies that for all $i \in N$, $u_i(x^\gamma) > (1-\delta)u_i(q) + \delta V_i(\sigma)$. By continuity of u_i, there exists $\epsilon_i > 0$ such that for all $x \in X$ with $||x - x^\gamma|| < \epsilon_i$, we have $u_i(x) > (1-\delta)u_i(q) + \delta V_i(\sigma)$. Setting $\epsilon = \min\{\epsilon_i \mid i \in N\}$, and using $|X| \geq 2$, we can set $x^\sigma \in B_\epsilon(x) \cap X$ so that for all $i \in N$,

$$u_i(x^\sigma) \quad > \quad (1-\delta)u_i(q) + \delta V_i(\sigma), \tag{2}$$

and to also satisfy $x^\sigma \neq \hat{x}^i$ for all $i \in N$.

Now consider any proposer j with ideal point $\hat{x}^j$. For all $i \in N$, inequality (2) and continuity of u_i imply that there exists $\lambda_i \in (0,1)$ such that for all $\lambda \in (0, \lambda_i)$, we have $u_i(\lambda\hat{x}^j + (1-\lambda)x^\sigma) > (1-\delta)u_i(q) + \delta V_i(\sigma)$. Letting $\lambda = \min\{\lambda_i \mid i \in N\}$, it follows that $\lambda\hat{x}^j + (1-\lambda)x^\sigma \in A(\sigma)$, and by strict quasi-concavity, we have $u_j(\lambda\hat{x}^j + (1-\lambda)x^\sigma) > u_j(x^\sigma)$, which implies $u_j(x^\sigma) < \max_{x \in A(\sigma)} u_j(x)$, fulfilling (b). ☐

Thus, when strict concavity is assumed to rule out the previous example, it follows that the only equilibria with delay are the static sort described in the trivial example above. Proposition 2.1, on the setter model, then yields the following result; a direct implication, assuming $\delta > 0$, is that if an equilibrium exhibits delay, then it is static.

Proposition 3.12. *Assume that $\delta > 0$ and that for all $i \in N$, u_i is strictly concave. In every stationary bargaining equilibrium σ that is non-static, we have for all $j \in N$,* *(i)* $\pi_j(\arg\max\{u_j(x) \mid x \in A(\sigma)\}) = 1$, *and*

(ii) $\int_X \overline{\alpha}(x)\pi_i(dx) = 1$.

The following example provides an illustration of equilibrium with a two-dimensional set of alternatives.

Example: Assume $X \subseteq \mathbb{R}^2$, $n = 5$, and each $\rho_i = 1/5$. Let $\mathcal{D}$ be the quota rule such that four or more votes are required for passage: $\mathcal{D} = \{G \subseteq N \mid |G| \geq 4\}$. Let utility functions u_i have ideal points and circular indifference curves and let the status quo be located as in Figure 11. We specify a strategy profile σ so that individual 1 mixes over x and y with equal probabilities, and we let each legislator $i = 2, 3, 4, 5$ propose p_i, as in the figure. We further specify utilities as follows. Define

$$v_i \quad = \quad \frac{u_i(x) + u_i(y)}{10} + \frac{1}{5}\sum_{j=2}^{5} u_i(p_j)$$

for each legislator, we require

$$\begin{aligned}
u_1(p_3) &= u_1(p_4) \;=\; (1-\delta)u_1(q) + \delta v_1 \\
u_2(x) &= (1-\delta)u_2(q) + \delta v_2 \\
u_3(y) &= u_3(p_5) \;=\; u_3(p_4) \;=\; (1-\delta)u_3(q) + \delta v_3 \\
u_4(x) &= u_4(p_2) \;=\; u_4(p_3) \;=\; (1-\delta)u_4(q) + \delta v_4 \\
u_5(y) &= (1-\delta)u_5(q) + \delta v_5.
\end{aligned}$$

It is clear that this is consistent with our assumption that utilities be concave. We then specify that each individual accept a proposal x if and only if $u_i(x) \geq (1-\delta)u_i(q) + \delta v_i$. Then all proposals are accepted, $v_i = V_i(\sigma)$ is the continuation determined by σ, and voting strategies are not weakly dominated. The core consists of the shaded polygon in the figure. It contains the status quo, yet legislators 1, 2, and 5 would vote against the status quo if it were proposed, and so it would be rejected. □

We also obtain the result that in one dimension, the social acceptance set is an interval, meaning that no-delay equilibrium proposal strategies are pure.

Proposition 3.13. *Assume $d = 1$. In every stationary bargaining equilibrium σ, the acceptance set $A(\sigma)$ is a closed interval; and if σ is no-delay,*

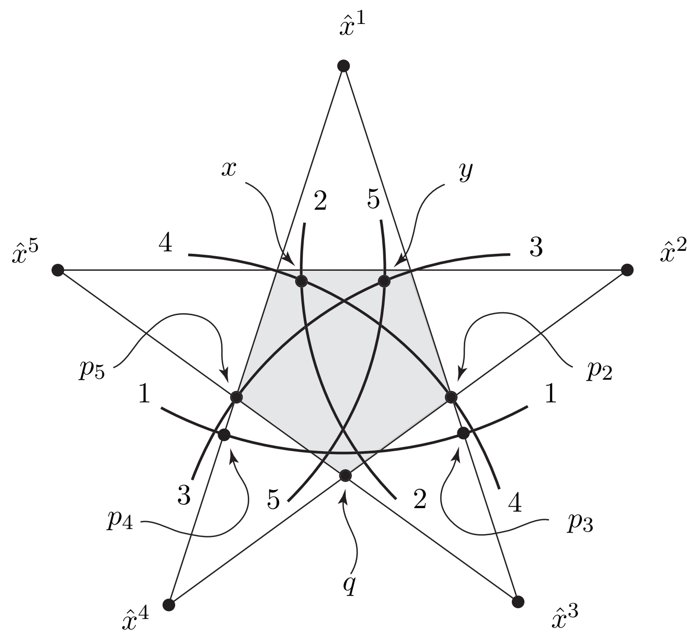

Figure 11: Equilibrium in multiple dimensions

then for all $i \in N$, the proposal strategy π_i places probability one on the unique solution to

$$\max_{x \in A(\sigma)} u_i(x).$$

3.7 Existence and the Core

We have seen examples of equilibria with delay, but they also possess no-delay equilibria. In fact, existence of such an equilibrium relies only on condition (a), which holds generally, permitting the following existence result; the proof parallels that of Proposition 3.6 and is omitted.

Proposition 3.14. *There exists a no-delay stationary bargaining equilibrium.*

In one dimension, we obtain an equilibrium in pure strategies.

Proposition 3.15. *Assume $d = 1$. There exists a pure-strategy, no-delay stationary bargaining equilibrium.*

We have seen that under weak background conditions (including strictly concave utilities), the only possible equilibria with delay are static. The next proposition gives an exact characterization of when such equilibria are possible: it is necessary and sufficient that the status quo belong to the core.

Proposition 3.16. *Assume $\rho_i > 0$ for all $i \in N$. Then there is a static stationary bargaining equilibrium if and only if $q \in C(\mathcal{D})$.*

Proof. First, assume σ is a static equilibrium, so that for all $i \in N$, we have $V_i(\sigma) = u_i(q)$, and therefore $(1-\delta)u_i(q) + \delta V_i(\sigma) = u_i(q)$. Suppose that $q \notin C(\mathcal{D})$, so there exists $G \in \mathcal{D}$ and $y \in X$ such that for all $i \in G$, $u_i(y) > u_i(q)$. Since we eliminate stage-dominated voting strategies, it follows that for all $i \in G$, $\alpha_i(y) = 1$, so $\overline{\alpha}(y) = 1$. Consider any $i \in G$, and note that the expected payoff from proposing y is $u_i(y) > u_i(q)$. This implies, since $\rho_i > 0$, that σ is not static, a contradiction. Now assume $q \in C(\mathcal{D})$. Define proposal strategies so that each π_i puts probability one on q, and define

$$\alpha_i(x) \quad = \quad \begin{cases} 1 & \text{if } u_i(x) \geq u_i(q), \\ 0 & \text{else.} \end{cases}$$

Given these strategies, we have $V_i(\sigma) = u_i(q)$ for all $i \in N$, and therefore the above voting strategies are not stage-dominated. To see that proposal strategies are are optimal, suppose that some individual can increase her expected payoff by proposing $x \neq q$. Then it must be that x is accepted with positive probability, so $x \in A(\sigma)$. Thus, there exists $G \in \mathcal{D}$ such that for all $j \in G$, we have $u_i(x) \geq u_i(q)$. Since X is convex, we have $\frac{1}{2}x + \frac{1}{2}q \in X$, and by strict quasi-concavity, we have $u_j(\frac{1}{2}x + \frac{1}{2}q) > u_j(q)$ for all $j \in G$, contradicting the assumption that q belongs to the core. □

The following corollary of Propositions 3.12 and 3.16.

Corollary 3.17. *Assume that $\delta > 0$ and that for all $i \in N$, $\rho_i > 0$ and u_i is strictly concave. If $q \notin C(\mathcal{D})$, then every stationary bargaining equilibrium is no-delay.*

The example following Proposition 3.12 and depicted in Figure 11 shows that when the set of alternatives is multi-dimensional, there may exist no-delay equilibria alongside static ones. We obtain a more precise characterization of no-delay equilibria when the set of alternatives is one-dimensional. It strengthens Proposition 3.16 by showing that when the status quo belongs to the core, not only does there exist a static equilibrium, but actually *all* equilibria are static. Moreover, it shows that if the status quo is socially acceptable, then it is in fact the only acceptable alternative, so that no-delay equilibrium proposal strategies are unique: every individual proposes the status quo. Note that the result depends on the assumption of one dimension, for the previous example shows that in two dimensions, the status quo may belong to the core, yet alternatives distinct from the status quo are socially acceptable; indeed, in that example, the status quo belongs to the core but is not itself socially acceptable.

Proposition 3.18. *Assume $d = 1$, $\mathcal{D}$ is proper, $\delta > 0$, and $\rho_i > 0$ for all $i \in N$. For every no-delay stationary bargaining equilibrium σ, we have*

$$A(\sigma) = \{q\} \;\Leftrightarrow\; q \in C(\mathcal{D}) \;\Leftrightarrow\; q \in A(\sigma).$$

Proof. We prove only that $q \in A(\sigma)$ implies $A(\sigma) = \{q\}$. To this end, assume $q \in A(\sigma)$. It suffices to show that $|A(\sigma)| = 1$. Otherwise, by Proposition 3.13, the social acceptance set is $A(\sigma) = [\underline{x}, \overline{x}]$, with $\underline{x} < \overline{x}$. Then there exist $\underline{G}, \overline{G} \in \mathcal{D}$ such that $\underline{x} \in A_{\underline{G}}(\sigma)$ and $\overline{x} \in A_{\overline{G}}(\sigma)$. Since $\mathcal{D}$ is proper, there exists $i \in \underline{G} \cap \overline{G}$, and it follows that

$$\min\{u_i(\underline{x}), u_i(\overline{x})\} \quad \geq \quad (1-\delta)u_i(q) + \delta V_i(\sigma). \tag{3}$$

Note that by strict quasi-concavity, we have

$$\max_{x \in A(\sigma)} u_i(x) \quad > \quad \min_{x \in A(\sigma)} u_i(x) \quad = \quad \min\{u_i(\underline{x}), u_i(\overline{x})\}.$$

Moreover, because σ is no-delay and $\rho_i > 0$, we have

$$V_i(\sigma) \quad \geq \quad (1-\rho_i) \min_{x \in A(\sigma)} u_i(x) + \rho_i \max_{x \in A(\sigma)} u_i(x) \quad > \quad \min\{u_i(\underline{x}), u_i(\overline{x})\}.$$

And since $q \in A(\sigma)$, we have $u_i(q) \geq \min\{u_i(\underline{x}), u_i(\overline{x})\}$. Because $\delta > 0$, however, we then have

$$\min\{u_i(\underline{x}), u_i(\overline{x})\} \quad < \quad (1-\delta)u_i(q) + \delta V_i(\sigma), \tag{4}$$

but (3) and (4) cannot simultaneously hold. □

Finally, we return to the relationship between equilibria of the bargaining model and the core in one dimension. The next result establishes an asymptotic median voter theorem analogous to Proposition 3.7: as individuals become arbitrarily patient, equilibrium proposals converge to a core alternative. As for the bad status quo model, a version of the result holds in multiple dimensions when the voting rule is collegial; I omit the formal statement of the result.

Proposition 3.19. *Assume that $d = 1$; that for all $i \in N$, we have $\rho_i > 0$; and that $\mathcal{D}$ is proper. For each $i \in N$, let $\{\delta_i^m\}$ be a sequence in $[0, 1)$ with $\delta_i^m \to 1$, and for each m, let σ^m be a stationary bargaining equilibrium such that individual i proposes $p_i^m \in X$ and $\{p_i^m\}$ has limit $p_i \in X$. Then there exists $x^* \in C(\mathcal{D})$ such that $p_1 = \cdots = p_n = x^*$.*

Of course, when $\mathcal{D}$ is strong, the core consists of a single alternative, and we obtain the following analogue to Corollary 3.8.

Corollary 3.20. *Assume that $d = 1$; that for all $i \in N$, we have $\rho_i > 0$; and that $\mathcal{D}$ is proper and strong, with $C(\mathcal{D}) = \{x^*\}$. For each $i \in N$, let $\{\delta_i^m\}$ be a sequence in $[0, 1)$ with $\delta_i^m \to 1$, and for each m, let σ^m be a stationary bargaining equilibrium such that individual i proposes $p_i^m \in X$. Then for all $i \in N$, we have $p_i^m \to x^*$.*

In the general status quo model, we have an additional parameter of interest, the status quo. By Proposition 3.18, the unique equilibrium proposal strategies when the status quo belongs to the core is that every individual proposes the status quo. The next proposition establishes a continuity result along these lines: as the status quo approaches the core, then equilibrium proposal strategies must approach the status quo (and therefore the core). Note that in the statement of the following proposition, we do not need to assume that equilibrium proposals are convergent (implicitly going to a convergent subsequence), because they must converge to the status quo.

Proposition 3.21. *Assume that $d = 1$; that for all $i \in N$, we have $\rho_i > 0$; and that $\mathcal{D}$ is proper. For each $i \in N$, let $\{q^m\}$ be a sequence in X with $q^m \to q \in C(\mathcal{D})$, and for each m, let σ^m be a stationary bargaining*

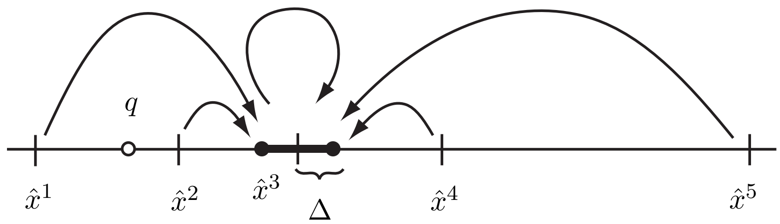

Figure 12: Equilibrium in one dimension

equilibrium such that individual i proposes $p_i^m \in X$. Then for all $i \in N$, we have $p_i^m \to q$.

The next example illustrates the above core convergence results for the one-dimensional model.

Example: Assume utilities are quadratic, and $\mathcal{D}$ is proper and strong. Let individual k be such that $C(\mathcal{D}) = \{\hat{x}^k\}$, assume $q \neq \hat{x}^k$, and let δ be a common discount factor. Then individual k is a representative voter, and in every equilibrium σ, the social acceptance set coincides with the acceptance set of k, i.e., $A(\sigma) = A_k(\sigma)$. For general $\delta \in [0, 1)$, Proposition 3.15 yields a pure strategy equilibrium, and by choosing δ sufficiently close to one, we have the situation in Figure 12, where $k = 3$ is the only individual whose ideal point belongs to the social acceptance set $A(\sigma) = [\hat{x}^k - \Delta, \hat{x}^k + \Delta]$. Then Δ is defined by the indifference condition

$$\Delta^2 \quad = \quad (1-\delta)(q-\hat{x}^k)^2 + \delta[(1-\rho_k)\Delta^2 + \rho_k(0)].$$

Solving, we obtain

$$\Delta \quad = \quad \sqrt{\frac{(1-\delta)(q-\hat{x}^k)^2}{1-\delta(1-\rho_k)}},$$

and it is simple to do comparative statics. In particular, the social acceptance set shrinks as we increase the recognition probability of the core individual or the discount factor or as we move the status quo toward the core. Consistent with Propositions 3.19 and 3.21, $\Delta \to 0$ as $\delta \to 1$ or $q \to \hat{x}^k$. □

Finally, the uniqueness result of Cho and Duggan (2003) holds in the general status quo model when utilities are quadratic and the voting rule is proper and strong. The anti-folk theorem of Cho and Duggan (2009) for the model with sequential voting also holds for the general status quo model, and the folk theorem result of Cho and Duggan (2013) for the model with simultaneous voting (eliminating stage-dominated voting strategies) also holds for the general status quo model; in contrast to the bad status quo model, the result for the general model requires discount factors close to one.

4 Stochastic Games

Stochastic games form a general class of discrete-time, infinite-horizon dynamic games with symmetric information.

Elements:

N	set of players (finite)
X_i	set of conceivable actions for i
S	set of states
$A_i(s)$	set of feasible actions for i in s (nonempty)
p or μ	transition probability (see below)
$u_i(s,a)$	stage payoff for i from a in s
δ_i	discount factor for i

Let $X = \prod_{i \in N} X_i$ be the set of conceivable action profiles. We denote an action profile by $a = (a_1, \ldots, a_n)$, and a_{-i} may denote the actions chosen by players other than i. We assume that each X_i is a metric space, with metric $\rho^{x,i}$, and we give X the product metric, denoted ρ^x. When X_i is finite, we always give it the discrete metric, so every set is open (and closed and measurable). Typically, X_i would be a subset of Euclidean space $\mathbb{R}^{d_i}$, and we would use the usual Euclidean metric.

We assume S is a metric space, with metric ρ^s, and when needed, we give the set $S \times X$ the product metric. If S is finite or countably infinite, then we assume that ρ^s is the discrete metric, so every set is open (and closed and measurable). In this case, we let $p(s'|s,a)$ be the probability that the next period's state is s' given that the current state is s and players chose actions $a = (a_1, \ldots, a_n)$. If S is uncountably infinite, then we endow it with the Borel σ-algebra, and we assume that for all (s,a), $\mu(\cdot|s,a)$ is a probability measure on S. That is, given a measurable set S' of states, $\mu(S'|s,a)$ is the probability that tomorrow's state belongs to S' given (s, a) today. Further measurable structure will be imposed, as needed.

Each player discounts streams of stage payoffs in the usual fashion, i.e., given a sequence $(s^1, a^1), (s^2, a^2), \ldots$ of state-action pairs, the payoff to player i is

$$\sum_{t=1}^{\infty} \delta_i^{t-1} u_i(s^t, a^t).$$

In the general case of infinite action sets or uncountable states, we assume that $u_i(s, a)$ is measurable in s and continuous in a, and we assume that the correspondence $A_i \colon S \rightrightarrows X_i$ is lower measurable; the latter property is satisfied if A_i is lower hemi-continuous or, of course, if it is constant. It is very general.

Special cases:

- infinitely repeated games
- finitely repeated games (use terminal state)
- finite games of perfect information
- spatial bargaining games

When the spatial bargaining model is mapped into the stochastic game framework, we let the set of states be $\tilde{S} = N \cup X \cup \{\emptyset\}$, where $s = i$ indicates that i is proposer, $s = x$ indicates that the proposer has proposed x, and $\emptyset$ is the terminal state. (For simplicity, assume $N \cap X = \emptyset$.) Periods alternate between "proposer stages" and "voting stages." In a proposer stage, if the state is i, then the feasible actions for i are $\tilde{A}_i(i) = X$, and for all $j \neq i$, we specify the trivial action set $\tilde{A}_j(i) = \{\emptyset\}$. After x is proposed, we move to a voting stage, and the state is $s' = x$. In particular, the *spatial bargaining model* is a particular type of stochastic game in which the set of states is uncountably infinite, and feasible action sets are uncountably infinite in proposer states. Moreover, the transition probability from proposer states to voting states is deterministic: in state i, if player i proposes x, then this becomes the state in the next voting stage. This class of stochastic games is exceedingly difficult, and equilibrium existence results are not available "off the shelf."

4.1 Finite-state/action Games

We begin with the analysis of finite-state/action games, i.e., we assume S and X are finite sets, and we give them the discrete metrics. A *strategy for* i is any mapping from histories to feasible actions. A strategy is *Markovian* if for each period t and for all $(t-1)$-period histories h and h' and all states s, the strategy specifies the same action at histories (h, s) and (h', s). That is, given any period t, the action specified is a function of the current state only. A Markovian strategy is *stationary* if the function mapping states to actions is the same in each period. We denote a stationary Markovian strategy by σ_i, and a profile is $\sigma = (\sigma_1, \ldots, \sigma_n)$. We may write σ_{-i} for the profile of strategies of all players other than i.

Technically, we view the set of probability measures on X_i as the unit simplex in $\mathbb{R}^{|X_i|}$, i.e., $\Delta(X_i) = \{z \in \mathbb{R}_+^{|X_i|} \mid \sum_{k=1}^{|X_i|} z_k = 1\}$, and a stationary Markovian strategy for i is a vector with $|S|\,|X_i|$ coordinates telling us the probability of every conceivable action in every state. Then the set of stationary Markov strategies for player i is

$$\Sigma_i \;=\; \Delta(X_i)^{|S|} \subseteq \mathbb{R}^{|S|\,|X_i|},$$

and the set of stationary Markovian profiles is

$$\Sigma \;=\; \prod_i \Sigma_i \;=\; \prod_i \left(\Delta(X_i)^{|S|}\right) \subseteq \mathbb{R}^{|S| \sum_i |X_i|},$$

which is a subset of finite-dimensional Euclidean space. Because $\Delta(X_i)$ is compact, it follows that Σ is a nonempty, convex, and compact (i.e., it is closed and bounded) subset of Euclidean space with the usual Euclidean metric, ρ^e. We view Σ as a metric space in its own right by endowing it with the metric ρ^e. Of course, a sequence $\{\sigma^m\}$ converges to σ in this metric if and only if $\rho^e(\sigma^m, \sigma) \to 0$, i.e., it converges in every coordinate.

We use the convention that the coordinate of σ_i corresponding to the probability of a_i in s has value $\sigma_i(a_i|s)$, and we write $\sigma_i(\cdot|s) \in \mathbb{R}^{|X_i|}$ for the vector of probabilities of actions in state s. Using this notation, note that $\sigma^m \to \sigma$ if and only if for all $i \in N$, all $s \in S$, and all $a_i \in X_i$, $\sigma_i^m(a_i|s) \to \sigma_i(a_i|s)$. We let $\sigma(a|s) = \prod_i \sigma_i(a_i|s)$ denote the probability that action profile a is chosen in state s. Given a subset $Y \subseteq X_i$ of actions, we use the shorthand that $\sigma_i(Y|s)$ is the probability that i chooses an action belonging to Y in state s. Note that by definition of a strategy, we have $\sigma_i(A_i(s)|s) = 1$ for all i and all s. Of course, $\sigma_{-i}(a_{-i})$ is the probability that players other than i choose a_{-i}. A subgame perfect equilibrium in which each player uses a stationary Markov strategy is a *stationary Markov perfect equilibrium.*

Given strategy profile σ, we let $V_i(s|\sigma)$ denote i's expected discounted payoff calculated at the beginning of a period in state s. This satisfies the recursion: for all $s \in S$,

$$V_i(s|\sigma) \quad = \quad \sum_a \left[u_i(s,a) + \delta_i \sum_{s'} V_i(s'|\sigma) p(s'|s,a) \right] \sigma(a|s).$$

Note that the right-hand side of the above equality is a function of s alone (since a and s' are "integrated out"). We could insert any mapping $f\colon S \to \mathbb{R}$ in place of $V_i(\cdot|\sigma)$, and the right-hand side would give us a new mapping, say f', from states to real numbers. Thus, we can view the right-hand side itself as a mapping (from functions like f to new functions f'), and then $V_i(\cdot|\sigma)$ is a fixed point of this mapping. In fact, the right-hand side defines a contraction mapping, so it has a unique fixed point, so $V_i(\cdot|\sigma)$ is the unique function satisfying the recursion.

A more formal argument to this effect is contained in the proof of the following lemma, which establishes an important continuity result for continuation values in stochastic games. The proof technique, which consists of an application of the contraction mapping theorem, is a standard method for deducing properties of continuation values in dynamic decision problems and games. We give $S \times \Sigma$ the product metric, so that a sequence $\{(s^m, \sigma^m)\}$ converges to (s, σ) if and only if $s^m \to s$ and $\sigma^m \to \sigma$. Since S has the discrete metric, note that the former condition implies that $s^m = s$ for sufficiently high m.

Lemma 4.1. *In every finite-state/action stochastic game, each continuation value mapping* $V_i\colon S \times \Sigma \to \mathbb{R}$ *is continuous.*

Proof. Let $\mathcal{C}_b(S \times \Sigma)$ denote the set of bounded, continuous mappings $f\colon S \times \Sigma \to \mathbb{R}$. In fact, with the discrete metric on S, $f \in \mathcal{C}_b(S \times \Sigma)$ if and only if each $s \in S$, $f(s,\sigma)$ is continuous in σ. It is known that $\mathcal{C}_b(S \times \Sigma)$ is a complete metric space when equipped with the sup metric, denoted ρ^s. Define the operator $\Psi\colon \mathcal{C}_b(S \times \Sigma) \to \mathcal{C}_b(S \times \Sigma)$ so that for all $f \in \mathcal{C}_b(S \times \Sigma)$,

$$\Psi(f)(s,\sigma) \quad = \quad \sum_a \left[u_i(s,a) + \delta_i \sum_{s'} f(s',\sigma) p(s'|s,a) \right] \sigma(a|s).$$

To verify that $\Psi(f)$ is indeed continuous, consider any state s and any convergent sequence $\{\sigma^m\}$ with limit σ in Σ. Then for all a, we have $\sigma^m(a|s) \to \sigma(a|s)$, and since f is continuous, we have $f(s', \sigma^m) \to f(s', \sigma)$. Thus, we have

$$\begin{aligned}
\Psi(f)(s, \sigma^m) &= \sum_a \left[u_i(s,a) + \delta_i \sum_{s'} f(s', \sigma^m) p(s'|s,a) \right] \sigma^m(a|s) \\
&\to \sum_a \left[u_i(s,a) + \delta_i \sum_{s'} f(s', \sigma) p(s'|s,a) \right] \sigma(a|s) \\
&= \Psi(f)(s, \sigma),
\end{aligned}$$

as required. For every $f, g \in \mathcal{C}_b(S \times \Sigma)$, note that

$$\begin{aligned}
&\rho^s(\Psi(f), \Psi(g)) \\
&= \sup_{(s,\sigma)} \delta_i \left| \sum_a \sum_{s'} f(s', \sigma) p(s'|s,a) \sigma(a|s) - \sum_a \sum_{s'} g(s', \sigma) p(s'|s,a) \sigma(a|s) \right| \\
&\leq \delta_i \sup_{(s,\sigma)} \sum_a \sum_{s'} |f(s', \sigma) - g(s', \sigma)| p(s'|s,a) \sigma(a|s) \\
&\leq \delta_i \rho^s(f, g),
\end{aligned}$$

and therefore Ψ is a contraction mapping with modulus $\delta_i < 1$. It follows that Ψ has a unique fixed point, and that this belongs to $\mathcal{C}_b(S \times \Sigma)$. By construction, the fixed point is V_i, completing the proof of the lemma. □

Given σ and a state s, let us now define the strategic form game *induced by σ in state s* as the game with player set N, strategy set $A_i(s)$ for each player, and payoff functions $U_i^\sigma(\cdot|s)$ defined by

$$U_i^\sigma(a|s) \quad = \quad u_i(s,a) + \delta_i \sum_{s'} V_i(s'|\sigma) p(s'|s,a).$$

We denote this induced game by $\Gamma^\sigma(s)$. Let $N^\sigma(s)$ denote the set of mixed strategy Nash equilibria of the induced game, nonemptiness of which follows from Nash's theorem. Of course, Lemma 4.1 implies that the induced payoff $U_i^\sigma(a|s)$ is continuous in σ for each i, s, and a.

By the one-shot deviation principle, we have the following characterization of subgame perfect equilibria in stationary Markovian strategies.

Proposition 4.2. *In a finite-state/action stochastic game, a profile σ is a stationary Markov perfect equilibrium if and only if for all s, σ determines a mixed strategy Nash equilibrium of the game induced by σ in state s, i.e., $(\sigma_1(\cdot|s), \ldots, \sigma_n(\cdot|s)) \in N^\sigma(s)$.*

The latter proposition suggests an approach to existence of stationary Markov perfect equilibria. Given σ, we calculate the best responses of player i in the induced game at state s as

$$BR_i^\sigma(s) \;\; = \;\; \arg\max_{a_i \in A_i(s)} \sum_{a_{-i}} U_i^\sigma(a_i, a_{-i}|s)\sigma_{-i}(a_{-i}|s).$$

This set is nonempty and closed (indeed, finite), but not necessarily convex. Then we take mixtures over best responses to give us the set $\Delta(BR_i^\sigma(s))$, which we view as the face of $\Delta(X_i)$ such that $BR_i^\sigma(s)$ has probability one. We obtain the best response stationary Markov strategies for i as

$$\Phi_i(\sigma) \;\; = \;\; \prod_s \Delta(BR_i^\sigma(s)),$$

which gives us a correspondence $\Phi_i \colon \Sigma \rightrightarrows \Sigma_i$ with nonempty, closed, and convex values. Finally, we define the correspondence $\Phi \colon \Sigma \rightrightarrows \Sigma$ by

$$\Phi(\sigma) \;\; = \;\; \prod_i \Phi_i(\sigma) \;\; = \;\; \prod_i \left(\prod_s \Delta(BR_i^\sigma(s)) \right),$$

which also has nonempty, closed, and convex values. A fixed point of this correspondence is a stationary Markov strategy profile $\sigma \in \Phi(\sigma)$, so that for all i and all s, $\sigma_i(\cdot|s)$ puts probability one on $BR_i^\sigma(s)$, so $(\sigma_1(\cdot|s), \ldots, \sigma_n(\cdot|s))$ is a mixed strategy Nash equilibrium of the induced game in state s; by Proposition 4.2, σ is therefore a subgame perfect equilibrium.

Thus, the existence proof consists of verifying that Φ satisfies the conditions needed for Kakutani's theorem, i.e., it has non-empty, convex, closed values, and it is upper hemi-continuous. We have noted all but the last property, and so the crux of the proof consists of checking upper hemi-continuity. With the assumption of a finite set of actions, this property follows from a simply continuity argument using Lemma 4.1. The following result generalizes Shapley's (1953) analysis of two-player, zero-sum stochastic games and was proved independently by Fink (1964), Rogers (1969), and Sobel (1971);

Takahashi (1964) assumed finite states but allowed infinite action sets, a case covered in the next subsection.

Proposition 4.3. *Every finite-state/action stochastic game admits a stationary Markov perfect equilibrium.*

Proof. By the above line of argument, it suffices to confirm that Φ is upper hemi-continuous. As the product of upper hemi-continuous correspondences is upper hemi-continuous, it suffices to confirm that for all $i \in N$ and all $s \in S$, the correspondence $\sigma \to \Delta(BR_i^\sigma(s))$ is upper hemi-continuous. It is known that if $\sigma \to BR_i^\sigma(s)$ is upper hemi-continuous, then the correspondence from strategy profiles to mixtures over best responses is upper hemi-continuous. Thus, consider any state s, a sequence $\{\sigma^m\}$ converging to σ in Σ, and a sequence $\{a_i^m\}$ such that $a_i^m \in BR_i^{\sigma^m}(s)$ for all m and $a_i^m \to a_i$. In fact, since X_i is finite and endowed with the discrete topology, we have $a_i^m = a_i$ for sufficiently high m. We must show that $a_i \in BR_i^\sigma(s)$. Suppose not. Then there exists $a_i' \in A_i(s)$ such that

$$\sum_{a_{-i}} U_i^\sigma(a_i', a_{-i}|s)\sigma_{-i}(a_{-i}|s) \quad > \quad \sum_{a_{-i}} U_i^\sigma(a_i, a_{-i}|s)\sigma_{-i}(a_{-i}|s).$$

By Lemma 4.1 (and the fact that $a_i^m = a_i$ for sufficiently high m), we have

$$\sum_{a_{-i}} U_i^{\sigma^m}(a_i', a_{-i}|s)\sigma_{-i}^m(a_{-i}|s) \quad \to \quad \sum_{a_{-i}} U_i^\sigma(a_i', a_{-i}|s)\sigma_{-i}(a_{-i}|s)$$
$$\sum_{a_{-i}} U_i^{\sigma^m}(a_i^m, a_{-i}|s)\sigma_{-i}^m(a_{-i}|s) \quad \to \quad \sum_{a_{-i}} U_i^\sigma(a_i, a_{-i}|s)\sigma_{-i}(a_{-i}|s).$$

But then we have

$$\sum_{a_{-i}} U_i^{\sigma^m}(a_i', a_{-i}|s)\sigma_{-i}^m(a_{-i}|s) \quad > \quad \sum_{a_{-i}} U_i^{\sigma^m}(a_i^m, a_{-i}|s)\sigma_{-i}^m(a_{-i}|s)$$

for high enough m, contradicting optimality of a_i^m given σ_{-i}^m in state s. □

4.2 Countable-state, Infinite-action Games

When we generalize the framework to allow for a countable set of states, but maintain the assumption of finite actions, very little changes. One additional assumption we impose is that stage payoffs are bounded in absolute value for each player, a condition that is automatically satisfied in finite-state/action games.

(A0) there exists $b \in \mathbb{R}$ such that for all i, all s, and all a, $|u_i(s,a)| \leq b$.

Because S is countable, we equip S with the discrete metric, so every subset of states is measurable, skirting possible technical issues of measurability. A difference now is that because S may be infinite, $\mathbb{R}^{|S||X_i|}$ is not finite-dimensional, and the "usual" Euclidean metric is not even defined. But we can define Σ_i as the set of mappings from S to $\Delta(X_i)$, or equivalently as the product

$$\Sigma_i \;\; = \;\; \Delta(X_i)^S.$$

Here, $\Delta(X_i)$ is still a nonempty, convex, and compact subset of $\mathbb{R}^{|X_i|}$, and we give it the usual Euclidean metric, denoted $\rho^{e,i}$.

We then endow Σ_i with the product metric, i.e., indexing $S = \{s_1, s_2, \ldots\}$, the distance between two strategies $\sigma_i, \sigma_i' \in \Sigma_i$ is

$$\rho^{\sigma,i}(\sigma_i, \sigma_i') \;\; = \;\; \sum_{k=1}^{\infty} \frac{1}{2^k \sqrt{|X_i|}} \rho^{e,i}(\sigma_i(\cdot|s_k), \sigma_i'(\cdot|s_k)),$$

where we use the fact that the distance between any two elements of the unit simplex is bounded by $\sqrt{|X_i|}$. Then the distance between two profiles $\sigma, \sigma' \in \Sigma = \prod_i \Sigma_i$ is given by the product metric over players, i.e.,

$$\rho^{\sigma}(\sigma, \sigma') \;\; = \;\; \sum_{i=1}^{n} \rho^{\sigma,i}(\sigma_i, \sigma_i').$$

This gives us a metric on Σ such that $\sigma^m \to \sigma$ if and only if for all $i \in N$, all $s \in S$, and all $a_i \in X_i$, we have $\sigma_i^m(a_i|s) \to \sigma_i(a_i|s)$. The set Σ is nonempty, convex, and by the Tychonoff product theorem, it inherits compactness of the set $\Delta(X_i)$. Note that this metric also satisfies the convexity property needed for the application of Glicksberg's fixed point theorem.

Given σ, the continuation value function $V_i(\cdot|\sigma)$ is again the unique function satisfying the recursion

$$V_i(s|\sigma) \;\; = \;\; \sum_a \left[u_i(s,a) + \delta_i \sum_{s'} V_i(s'|\sigma) p(s'|s,a) \right] \sigma(a|s)$$

for all s. Lemma 4.1 continues to hold by the original argument, augmented with (A0), and payoffs $U_i^{\sigma}(a|s)$ in the induced game are still continuous in

σ for each i, s, and a. Proposition 4.2 does not rely on the assumption of finite states and continues to hold. We again define $BR_i^{\sigma}(s)$ as the best responses for player i in the induced game at state s given strategies $\sigma_{-i}(\cdot|s)$ of the other players. By the above arguments, the correspondence Φ again has nonempty, convex, closed values and is upper hemi-continuous. To obtain a fixed point, we simply use the more general fixed point theorem of Glicksberg. Thus, we easily obtain existence of an equilibrium.

Proposition 4.4. *Every stochastic game with a countable set of states and finite action sets satisfying (A0) admits a stationary Markov perfect equilibrium.*

More interesting—and difficult—technical issues arise when we allow for infinite action sets. In this case, in addition to bounded stage payoffs, we assume:

(A1) X_i is a compact metric space

(A2) $A_i(s)$ is closed (and therefore compact)

(A3) $p(s'|s,a)$ is continuous in a

(A4) $u_i(s,a)$ is continuous in a.

A familiar special case is that in which X_i is simply a closed, bounded subset of a finite-dimensional Euclidean space, say $\mathbb{R}^{d_i}$.

The *first* technical issue that arises is the definition of stationary Markov strategy, as we can no longer view $\sigma_i(\cdot|s)$ as a vector with a finite number of coordinates. And notationally, we cannot write $\sigma_i(a_i|s)$ for the probability of a_i in state s: if player i mixes according to a continuous distribution in state s, then the probability of any given action is zero. Instead, we let $\sigma_i(\cdot|s)$ be a probability measure on X, and we replace summation notation by integral notation. In the finite-action case, we defined $\sigma(a|s) = \prod_i \sigma_i(a_i|s)$ as the probability of action profile given state s. Again, σ determines a probability distribution over action profiles, but we cannot define this as a pointwise product as we did before. Instead, we define a product measure on X determined by $(\sigma_1(\cdot|s), \ldots, \sigma_n(\cdot|s))$ in the standard way. The usual notation for the product measure is $\otimes_i \sigma_i(\cdot|s)$, but we abuse notation and simply write

$\sigma(\cdot|s)$ for this measure. Then continuation values are again obtained as the unique solution to a system of equations, now of the form: for all s,

$$V_i(s|\sigma) \quad = \quad \int_X \left[u_i(s,a) + \delta_i \sum_{s'} V_i(s'|\sigma) p(s'|s,a) \right] \sigma(da|s).$$

Note that continuation values are bounded in absolute value by $\frac{b}{1-\delta_i}$.

The *second* technical issue to arise is that of compactness. We again let Σ_i denote the set of stationary Markov strategies for i, but now we view $\sigma_i(\cdot|s)$ as an element of the set $\Delta(X_i)$ of probability measures on X_i, and we define a stationary Markov strategy as a mapping from the set S of states to this set $\Delta(X_i)$ of mixtures over actions. We still view a stationary Markov strategy for i as a mapping from S to $\Delta(X_i)$, so $\Sigma_i = \Delta(X_i)^S$, but the interpretation of this is somewhat deeper now. We give $\Delta(X_i)$ the Prohorov metric, denoted $\rho^{r,i}$, and we then give Σ_i the product metric $\rho^{\sigma,i}$, i.e., given $\sigma_i, \sigma_i' \in \Sigma$, we let

$$\rho^{\sigma,i}(\sigma_i, \sigma_i') \quad = \quad \sum_{k=1}^{\infty} \frac{1}{2^k r_i} \rho^{r,i}(\sigma_i(\cdot|s_k), \sigma_i'(\cdot|s_k)),$$

where states are indexed $S = \{s_1, s_2, \ldots\}$ and r_i is a bound for the Prohorov metric on $\Delta(X_i)$. This gives us a metric on Σ_i such that $\sigma_i^m \to \sigma_i$ if and only if for all $s \in S$, we have $\rho^{r,i}(\sigma_i^m(\cdot|s), \sigma_i(\cdot|s)) \to 0$. Recall this means that $\sigma_i^m(\cdot|s) \to \sigma_i(\cdot|s)$ weakly, i.e., for every bounded, continuous function $f\colon X_i \to \mathbb{R}$, we have

$$\int_{X_i} f(a_i) \sigma_i^m(da_i|s) \quad \to \quad \int_{X_i} f(a_i) \sigma_i(da_i|s).$$

Moreover, since the Prohorov metric is convex, the product metric inherits this convexity condition, facilitating the application of Glicksberg's fixed point theorem.

Then we give $\Sigma = \prod_i \Sigma_i$ the product metric, denoted ρ^σ, i.e., given $\sigma, \sigma' \in \Sigma$, we specify

$$\rho^\sigma(\sigma, \sigma') \quad = \quad \sum_{i=1}^{n} \rho^{\sigma,i}(\sigma_i, \sigma_i').$$

This gives us a metric on Σ such that $\sigma^m \to \sigma$ if and only if for all i and all s, $\sigma_i^m(\cdot|s) \to \sigma_i(\cdot|s)$ weakly. The set Σ is nonempty, convex, and by the Tychonoff product theorem, it inherits compactness of the set $\Delta(X_i)$. Again, the metric satisfies the convexity condition needed for the application of Glicksberg's fixed point theorem. Additionally, we will use the fact that if $\sigma^m \to \sigma$ in the metric ρ^σ, then in fact for all s, the sequence $\{\otimes_i \sigma_i^m(\cdot|s)\}$ of product measures converges weakly to the product measure $\otimes_i \sigma_i(\cdot|s)$. This means that for every bounded, continuous function $f\colon X \to \mathbb{R}$, we have

$$\int_X f(a)\sigma^m(da|s) \;\to\; \int_X f(a)\sigma(da|s),$$

a fact that is useful in verifying continuity of continuation values.

A *third* potential issue is continuity, but this is already taken care of by our use of the Prohorov metric. Specifically, with this metric on mixtures over actions, the continuation value $V_i(s|\sigma)$ is continuous in σ for each state s. The argument for this uses the contraction mapping theorem as in the proof of Lemma 4.1, but we must re-check that the operator Ψ indeed maps bounded, continuous mappings $f\colon S \times \Sigma \to \mathbb{R}$ to bounded, continuous mappings. Under current assumptions, we specify

$$\Psi(f)(s,\sigma) \;=\; \int_X \left[u_i(s,a) + \delta_i \sum_{s'} f(s',\sigma)p(s'|s,a)\right]\sigma(da|s),$$

which is well-defined since f is bounded. Given a sequence $\{\sigma^m\}$ of strategy profiles converging to σ, note that the integrand above is continuous (and therefore bounded) in a, and that the product measures $\{\sigma^m(\cdot|s)\}$ converge weakly to $\sigma(\cdot|s)$. Although the integrand varies along the sequence with m, continuity of f and a generalized version of Lebesgue's dominated convergence theorem yield

$$\begin{aligned}\Psi(f)(s,\sigma^m) \;&=\; \int_X \left[u_i(s,a) + \delta_i \sum_{s'} f(s',\sigma^m)p(s'|s,a)\right]\sigma^m(da|s)\\ &\to\; \int_X \left[u_i(s,a) + \delta_i \sum_{s'} f(s',\sigma)p(s'|s,a)\right]\sigma(da|s)\\ &=\; \Psi(f)(s,\sigma),\end{aligned}$$

and we conclude that $\Psi(f)(s,\sigma)$ is continuous in σ, as required. Of course, $\Psi(f)$ is bounded as well.

Lemma 4.5. *In every countable-state, infinite-action stochastic game satisfying (A0)–(A4), each continuation value mapping $V_i\colon S \times \Sigma \to \mathbb{R}$ is continuous.*

We then define induced games as before, with the induced payoff of player i in the game induced by σ in state s given by

$$U_i^\sigma(a|s) \;=\; u_i(s,a) + \delta_i \sum_{s'} V_i(s'|\sigma) p(s'|s,a).$$

In finite-action games, we noted that $U_i^\sigma(a|s)$ was continuous in σ for fixed a, and this was sufficient because convergence of a sequence $\{a^m\}$ to a entailed $a^m = a$ for high enough m. In the current framework, however, it is crucial to check that induced payoffs are *jointly continuous* in (a,σ). To see that this holds, let $\sigma^m \to \sigma$ and $a^m \to a$. For each $\epsilon > 0$, there is a finite set S' such that $\sum_{s' \in S'} p(s'|s,a) > 1-\epsilon$. Then for m high enough, we have

$$\sum_{s' \in S'} p(s'|s,a^m) \;>\; 1-\epsilon \quad \text{and} \quad \sup_{s' \in S'} |V_i(s'|\sigma^m) - V_i(s'|\sigma)| \;<\; \epsilon,$$

which implies

$$\left| \sum_{s'} V_i(s'|\sigma^m) p(s'|s,a^m) - \sum_{s'} V_i(s'|\sigma) p(s'|s,a) \right| \;<\; \epsilon \frac{2b}{1-\delta_i} + (1-\epsilon)\epsilon,$$

which implies the continuity claim.

Lemma 4.6. *In every countable-state, infinite-action stochastic game satisfying (A0)–(A4), the induced payoff $U_i^\sigma(a|s)$ is jointly continuous in (a,σ) in the induced game at each state s.*

As for finite-state/action games, we use the one-shot deviation principle to state a characterization of stationary Markov perfect equilibria in terms of equilibria of induced games.

Proposition 4.7. *In a countable-state, infinite-action stochastic game satisfying (A0), a profile σ is a stationary Markov perfect equilibrium if and only if for all s, σ determines a mixed strategy Nash equilibrium of the game induced by σ in state s, i.e., $(\sigma_1(\cdot|s), \ldots, \sigma_n(\cdot|s)) \in N^\sigma(s)$.*

Given σ, we calculate the best responses of player i in the induced game at state s as

$$BR_i^\sigma(s) \;=\; \arg\max_{a_i \in A_i(s)} \int_{a_{-i}} U_i^\sigma(a_i, a_{-i}|s) \sigma_{-i}(da_{-i}|s).$$

In the above, continuity of $U_i^\sigma(a|s)$ in a_i and Lebesgue's dominated convergence theorem imply that the objective function

$$\int_{a_{-i}} U_i^\sigma(a_i, a_{-i}|s)\sigma_{-i}(da_{-i}|s)$$

is continuous in a_i. Thus, the set of best responses is nonempty and closed, but not necessarily convex. Then we take mixtures over best responses, as in $\Delta(BR_i^\sigma(s))$, and we obtain the best response stationary Markov strategies for i as

$$\Phi_i(\sigma) \quad = \quad \prod_s \Delta(BR_i^\sigma(s)),$$

which gives us a correspondence $\Phi_i \colon \Sigma \rightrightarrows \Sigma_i$ with nonempty, closed, and convex values. Finally, we define the correspondence $\Phi \colon \Sigma \rightrightarrows \Sigma$ by

$$\Phi(\sigma) \quad = \quad \prod_i \Phi_i(\sigma) \quad = \quad \prod_i \left(\prod_s \Delta(BR_i^\sigma(s)) \right).$$

As in finite-state/action games, a fixed point of this correspondence is a stationary Markov strategy profile $\sigma \in \Phi(\sigma)$, so that for all i and all s, $\sigma_i(\cdot|s)$ puts probability one on $BR_i^\sigma(s)$, so $(\sigma_1(\cdot|s), \ldots, \sigma_n(\cdot|s))$ is a mixed strategy Nash equilibrium of the induced game in state s; by Proposition 4.7, σ is therefore a subgame perfect equilibrium.

Accordingly, the existence proof consists of verifying that Φ satisfies the conditions needed for Glicksberg's theorem, i.e., it has non-empty, convex, closed values, and it is upper hemi-continuous. Because $\Delta(BR_i^\sigma(s))$ is nonempty, convex, and closed, so is $\Phi_i(\sigma)$, and therefore so is $\Phi(\sigma)$. Thus, the crux of the proof consists of checking upper hemi-continuity, and this boils down to an argument using joint continuity of $U_i^\sigma(a|s)$ in (a, σ) from Lemma 4.6. The continuity argument is rather routine step in standard existence arguments. The following generalizes a result of Takahashi (1964), who assumed a finite set of states, and is due to Federgruen (1978).

Proposition 4.8. *Every countable-state, infinite-action stochastic game satisfying (A0)–(A4) admits a stationary Markov perfect equilibrium.*

Proof. By the above line of argument, it suffices to confirm that Φ is upper hemi-continuous. As the product of upper hemi-continuous correspondences is upper hemi-continuous, it suffices to confirm that for all $i \in N$ and all $s \in S$, the correspondence $\sigma \to \Delta(BR_i^\sigma(s))$ is upper hemi-continuous. Moreover, if $\sigma \to BR_i^\sigma(s)$ is upper hemi-continuous, then the correspondence from strategy profiles to mixtures over best responses is upper hemi-continuous. Thus, consider any state s, a sequence $\{\sigma^m\}$ converging to σ in Σ, and a sequence $\{a_i^m\}$ such that $a_i^m \in BR_i^{\sigma^m}(s)$ for all m and $a_i^m \to a_i$. We must show that $a_i \in BR_i^\sigma(s)$. Suppose not. Then there exists $a_i' \in A_i(s)$ such that

$$\int_{X_{-i}} U_i^\sigma(a_i', a_{-i}|s)\sigma_{-i}(da_{-i}|s) \quad > \quad \int_{X_{-i}} U_i^\sigma(a_i, a_{-i}|s)\sigma_{-i}(da_{-i}|s).$$

Although $U_i^{\sigma^m}(a_i', a_{-i}|s)$ and $U_i^{\sigma^m}(a_i^m, a_{-i}|s)$ vary with m, Lemma 4.6 and a generalized version of Lebesgue's dominated convergence theorem yield

$$\begin{aligned} \int_{X_{-i}} U_i^{\sigma^m}(a_i', a_{-i}|s)\sigma_{-i}^m(da_{-i}|s) \quad &\to \quad \int_{X_{-i}} U_i^\sigma(a_i', a_{-i}|s)\sigma_{-i}(da_{-i}|s) \\ \int_{X_{-i}} U_i^{\sigma^m}(a_i^m, a_{-i}|s)\sigma_{-i}^m(da_{-i}|s) \quad &\to \quad \int_{X_{-i}} U_i^\sigma(a_i, a_{-i}|s)\sigma_{-i}(da_{-i}|s). \end{aligned}$$

But then we have

$$\int_{X_{-i}} U_i^{\sigma^m}(a_i', a_{-i}|s)\sigma_{-i}^m(da_{-i}|s) \quad > \quad \int_{X_{-i}} U_i^{\sigma^m}(a_i^m, a_{-i}|s)\sigma_{-i}^m(da_{-i}|s)$$

for high enough m, contradicting optimality of a_i^m given σ_{-i}^m in state s. □

4.3 Uncountable-State Games

We now consider the class of stochastic games with uncountable sets of states—which finally includes the spatial bargaining model—and we will note several technical dangers that arise in such games. In particular, we assume that S is a complete, separable metric space, and we give $S \times X$ the product metric, and we continue to allow for infinite action sets, assuming X is a compact metric space and imposing the continuity conditions of the previous subsection. Now, however, we must take care to address measurability issues. Thus, in addition to our continuity conditions, we assume:

(A5) S is a complete, separable metric space,

(A6) $A_i\colon S \rightrightarrows X_i$ is lower measurable,

(A7) for all (s, a), $\mu(\cdot|s, a)$ is a probability measure on states,

(A8) for all measurable S', $\mu(S'|s,a)$ is measurable in (s,a),

(A9) $u_i(s,a)$ is measurable in s.

Note that we now use the general notation for the transition probability μ, and we will replace summations over states by integrals.

As in the previous subsection, $\sigma_i(\cdot|s)$ is viewed as a probability measure on actions, but now we must impose a measurability restriction: for all measurable $Y \subseteq X_i$, $\sigma_i(Y|s)$ is measurable in s. Thus, we actually are assuming that σ_i is itself a transition probability, now from states to actions (rather than state-action pairs to states). Again, $\sigma(\cdot|s) = \otimes_i \sigma_i(\cdot|s)$ is the probability measure over action profiles determined by $(\sigma_1, \ldots, \sigma_n)$, and it is known that σ is itself a transition probability, now from states to profiles of actions. Continuation values are again obtained as the unique solution to a system of equations, now of the form: for all s,

$$V_i(s|\sigma) \quad = \quad \int_X \left[u_i(s,a) + \delta_i \int_S V_i(s'|\sigma)\mu(ds'|s,a) \right] \sigma(da|s).$$

Measurability follows from a contraction mapping argument using completeness of the set $\mathcal{M}_b(S)$ of bounded, measurable functions from states to real numbers. Later, we switch to a different type of continuation value and prove uniqueness and measurability.

To illustrate the potential difficulties of compactness and continuity, and to motivate a series of increasing restrictions on the transition probability μ, we consider a trivial, two-player, two-period (i.e., the state transitions to a terminal state after the second period) stochastic game.

Example: The set of states is $S = [0,1]$, but they are irrelevant in the first period, so we suppress the first-period state. In the first period, the only active player is player 1, who chooses an action a_1 in the feasible set $A_1 = [0,1]$. The state in the second period is s, which is drawn from $\mu(\cdot|a_1)$. Then the players play a two-by-two coordination game, depicted below.

$0,0$	s,s
$0,0$	$0,0$

Obviously, there are two pure-strategy equilibria, (U,R) and (D,L), and one mixed strategy equilibrium in the second-period subgames with $s > 0$. A strategy profile specifies an action for player 1 in the first period and, as a function of the state, an action for each player in the matrix game. One possibility is that the players play (U, R) for $s \in [0, 1/2)$ and (D, L) for $s \in [1/2, 1]$. This profile is depicted in the top panel of Figure 13.

Suppose that the transition is deterministic, so that $\mu(\{a_1\}|a_1) = 1$, i.e., player 1's choice in the first period carries over directly as the state in the second period. Given the second-period strategies described above, consider player 1's induced payoff in the first period:

$$U_1^\sigma(a_1) \quad = \quad \begin{cases} a_1 & \text{if } a_1 < \frac{1}{2}, \\ 0 & \text{if } a_1 \geq \frac{1}{2}\ . \end{cases}$$

Clearly, these strategies create a best response problem for player 1, as he or she would like to choose an action less than but arbitrarily close to one half: the player has no optimal action. This problem is rooted a discontinuity—not in the fundamentals of the game, but one that arises from the behavior of the players, which is endogenous. Although the state is continuous (in one sense) as a function of the first-period action, the players' responses to the state are not. And because this behavior is endogenous, we cannot rule out a priori such discontinuous responses. This points to a fundamental complexity of stochastic games and, in particular, a pathology of deterministic transitions. □

If the second-period state is drawn from a distribution that is independent of the first-period action, then the above anomaly cannot arise, but that is too restrictive. Instead, we assume that $\mu(\cdot|s, a)$ is *set-wise continuous* in a.

(A10) for all $s \in S$ and all measurable $S' \subseteq S$, $\mu(S'|s, a)$ is continuous in a.

To see that this precludes deterministic transitions (and indeed, transitions that have even a deterministic component), return to the example in which $a_1 = s$. Let a_1^m approach a_1, and note that $\mu(\{a_1\}|a_1^m) = 0$ for all m, but $\mu(\{a_1\}|a_1) = 1$ in the limit, violating set-wise continuity. A further implication of set-wise continuity is that given a strategy profile σ, the induced payoff in state s,

$$U_i^\sigma(a|s) \quad = \quad u_i(s, a) + \delta_i \int_S V_i(s'|\sigma)\mu(ds'|s, a),$$

is continuous in a, avoiding the discontinuity highlighted in the above example.

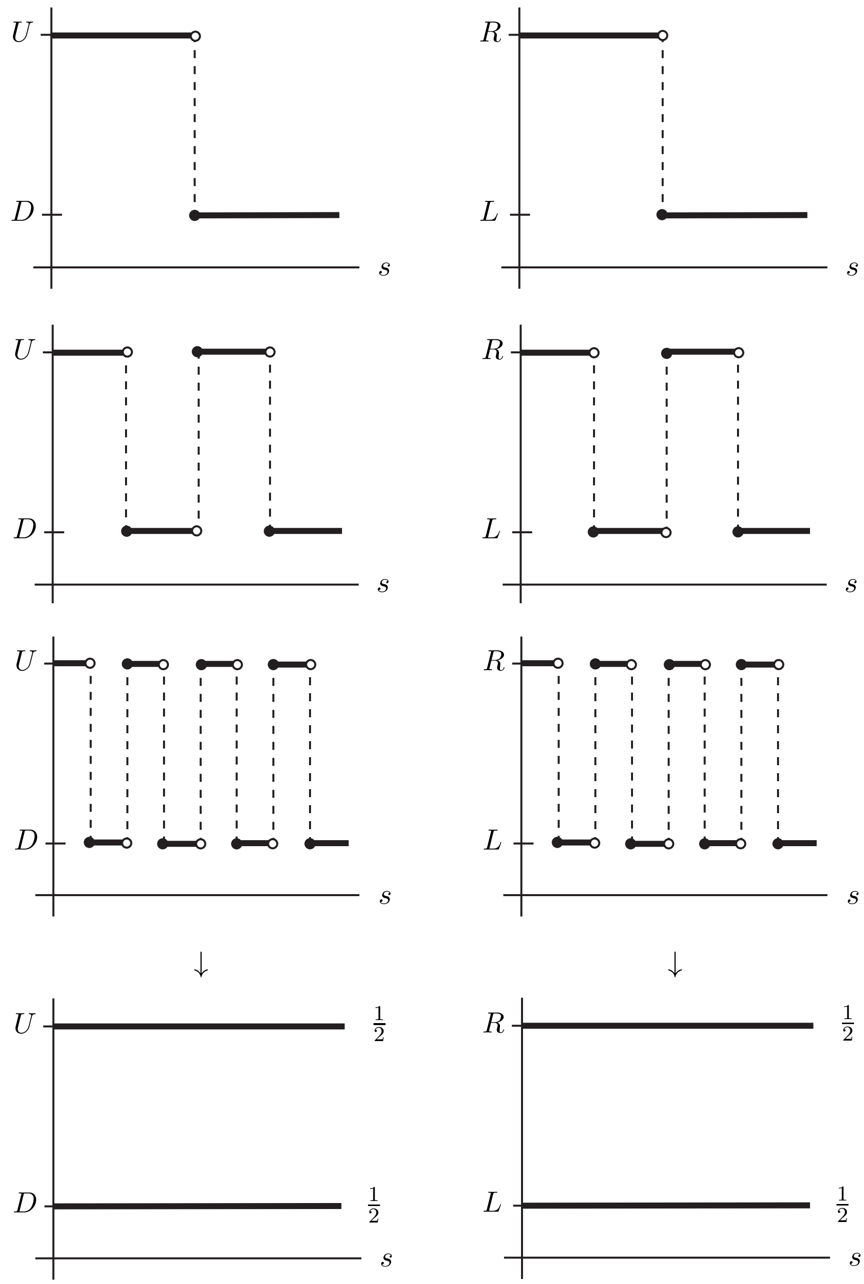

Figure 13: Continuity problems

But recall that in countable-state, infinite-action stochastic games, a key step (established in Lemma 4.6) for the existence argument is to show that the induced payoff be *jointly* continuous in (a, σ). Set-wise continuity is not sufficient for this; in fact, it would be if $V_i(s|\sigma)$ were guaranteed to be continuous in σ, but the next example shows that this property does not generally hold.

Example (cont.): Returning to the above example, we simplify the game by assuming the second-period payoffs are independent of the state,

$0,0$	$1,1$
$0,0$	$0,0$

and we assume that the second-period state s is drawn uniformly from $[0,1]$. Consider a sequence $\{\sigma^m\}$ of strategy profiles such that σ_i^1 is as described above: row player chooses U for $s \in [0,1/2)$ and D for $s \in [1/2,1]$, and column chooses R for $s \in [0,1/2)$ and L for $s \in [1/2,1]$. In σ_i^2, the players switch between strategies at twice the rate: row chooses U and column chooses R for $s \in [0,1/4) \cup [1/2,3/4)$, and otherwise the players choose D and L. In σ_i^3, the players again double the rate at which they switch actions, doubling again in σ_i^4, and so on. Graphing the second-period strategies of the players, each $\{\sigma_i^m\}$ has the form of the sawtooth sequence of functions (familiar from Figure 2), depicted in Figure 13.

For the existence argument to go through, we need compactness of the set of strategy profiles, and so the sequence $\{\sigma^m\}$ must have a convergent subsequence, but to what limit? There is no possible pointwise limit, as for almost all s, the corresponding sequence of actions for row player consists of alternating strings of U and D (and similarly for column). The only reasonable limit for $\{\sigma_i^m\}$ is a fifty-fifty randomization between the player's two actions for almost all states s, so $\sigma_1^m \to \sigma_1$ and $\sigma_2^m \to \sigma_2$, where for almost all s,

$$\sigma_1(\{U\}|s) = \sigma_1(\{D\}|s) = \frac{1}{2} \quad \text{and} \quad \sigma_2(\{L\}|s) = \sigma_2(\{R\}|s) = \frac{1}{2}.$$

Note that for almost all s, the second-period actions switch between (U,R) and (D,L) an infinite number of times, and so the players' payoffs switch between 0 and 1 an infinite number of times; but in the limit, the players' are mixing fifty-fifty, giving them an expected payoff of $\frac{1}{4}$ in almost all states. Thus, $V_i(s|\sigma)$ is discontinuous at almost all states. Moreover, the second-period strategies are not best responses, e.g., row player has an incentive to deviate to U in almost all states. Finally, the players' expected payoffs evaluated in the first period, before the second-period state is realized, jump down from $\frac{1}{2}$ to $\frac{1}{4}$. □

The preceding example raises the issues of compactness and continuity. In the countable state case, we simply give $\Delta(X_i)$ the Prohorov metric, and we view σ_i as an element of $\Delta(X_i)^S$ with the product metric, giving us compactness of Σ_i. We could continue to specify $\Sigma_i = \Delta(X_i)^S$ and give this set the "product topology" (a generalization of the product metric), and in fact the Tychonoff product theorem does imply compactness of this space, but it is not metrizable, and we have to deal with "nets" (or generalized sequences), and we lose the dominated convergence theorem. Moreover, that specification does not impose any measurability conditions, so we lose the ability to take integrals representing expected payoffs. Instead, we define Σ_i as the set of transition probabilities from states S to actions X_i, so a mapping $\sigma_i \colon \mathcal{B}_{X_i} \times S \to \mathbb{R}$ belongs to Σ_i if and only if for all $s \in S$, $\sigma_i(\cdot|s)$ is a probability measure on X_i with $\sigma_i(A_i(s)|s) = 1$, and for all measurable $Y \subseteq X_i$, $\sigma_i(Y|s)$ is measurable in s.

To define a useful metric on player i's stationary Markov strategies, we assume the transition probability μ is dominated by a fixed probability measure on states.

(A11) there is a probability measure κ on S such that for all $(s, a) \in S \times X$, we have $\mu(\cdot|s, a) \ll \kappa$.

Recall that a mapping $f \colon X \times S \to \mathbb{R}$ is a *Carathéodory integrand* if it is a Carathéodory function, so (i) for all $a \in X$, the function $f_a \colon S \to \mathbb{R}$ defined by $f_a(s) = f(a, s)$ is measurable, and (ii) for all $s \in S$, the function $f_s \colon X \to \mathbb{R}$ defined by $f_s(a) = f(a, s)$ is continuous, and if in addition, (iii) there is a κ-integrable mapping $g \colon S \to \mathbb{R}$ such that for all $s \in S$, $\sup\{|f(a, s)| \mid a \in X\} \leq g(s)$. Then we say $\{\sigma_i^m\}$ converges *narrowly* to σ_i if for every Carathéodory integrand f, we have

$$\int_S \int_X f(a, s)\sigma_i^m(da|s)\kappa(ds) \quad \to \quad \int_S \int_X f(a, s)\sigma_i(da|s)\kappa(ds).$$

This notion of convergence is metrizable by the narrow convergence metric, denoted ρ^r, so that $\rho^r(\sigma_i^m, \sigma_i) \to 0$ if and only if $\sigma_i^m \to \sigma_i$ narrowly. An advantage of using this metric is that with it, Σ_i is compact, so every sequence of stationary Markov strategies has a subsequence that converges narrowly.

In the above example, the sawtooth sequence $\{\sigma_i^m\}$ does converge to the strategy that mixes equally between U and D (for row) and L and R (for column). Thus, we endow Σ_i with the narrow convergence metric, ρ^r, and for now we give Σ the product metric, so $\sigma^m \to \sigma$ if and only if for all i, $\{\sigma_i^m\}$ converges to σ_i in the narrow convergence metric.

But solving the compactness problem alone is not enough: we noted the discontinuity of expected payoffs and the fact that the limiting strategies did not form equilibria in second-period subgames. To address this problem, instead of giving $\Sigma = \prod_i \Sigma_i$ the product metric, we let $\overline{\Sigma}$ be the set of transition probabilities from states S to action profiles X, so a mapping $\overline{\sigma}\colon \mathcal{B}_X \times S \to \mathbb{R}$ belongs to $\overline{\Sigma}$ if and only if for all $s \in S$, $\overline{\sigma}(\cdot|s)$ is a Borel probability measure on X with $\overline{\sigma}\big(\prod_i A_i(s)|s\big) = 1$, and for all measurable $Y \subseteq X$, $\overline{\sigma}(Y|s)$ is measurable in s. We then endow $\overline{\Sigma}$ directly with the narrow convergence metric. Given $\overline{\sigma} \in \overline{\Sigma}$, note that we now allow for the possibility that $\overline{\sigma}(\cdot|s)$ is correlated, i.e., not a product measure on X, so to avoid confusion we refer to $\overline{\sigma}$ as a *joint strategy.*

Note that X is compact and therefore separable. Moreover, because A_i is lower measurable with closed values, the correspondence $A\colon S \rightrightarrows X$ defined by $A(s) = \prod_i A_i(s)$ is lower measurable with closed values. Then the subset of transition probabilities from S to X that place probability one on the values of $A(s)$ for κ-almost all s, i.e.,

$$\overline{\Sigma} \;=\; \{\overline{\sigma} \in \mathcal{R}(X, S, \kappa) \mid \text{for } \kappa\text{-almost all } s,\, \overline{\sigma}(A(s)|s) = 1\}$$

is a compact in the narrow convergence metric, as desired.

Having addressed the compactness issue, we return to the example to see that there is hope for continuity as well. In taking this route, however, the example suggests that we must now admit equilibria in *correlated* strategies.

Example (cont.): We return to the sawtooth sequence $\{\overline{\sigma}^m\}$ of the previous example, depicted in Figure 14, now viewing $\overline{\sigma}^m$ as a joint strategy. By compactness of $\overline{\Sigma}$ in the narrow convergence metric, this sequence has a convergent subsequence (in fact, it is convergent) with limit depicted in the lower right-hand panel of the figure. That is, the sequence converges to $\overline{\sigma}$ such that for almost all s, the players choose (U, R) and (D, L) with equal probability. This again creates a discontinuity in the players' continuation values for almost all states: $V_i(s)$ alternates between strings of zeroes and

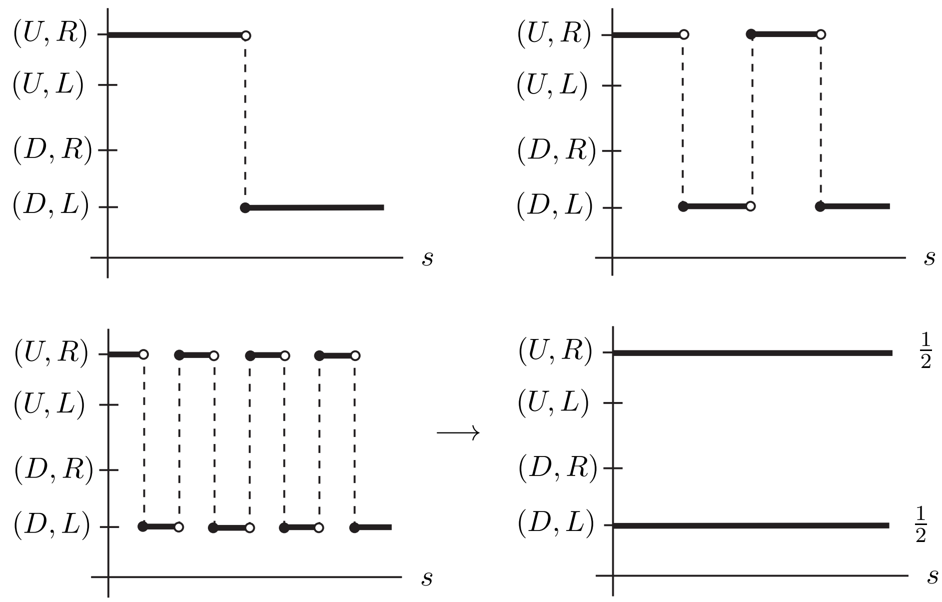

Figure 14: Continuity solutions

ones, but in the limit each player's expected payoff is one half. More importantly, however, the players' expected payoff calculated in period one *is* continuous: for each m, the ex ante (before the state is realized) expected payoff in period two is

$$\int_S V_i(s|\overline{\sigma}^m)ds \quad = \quad \frac{1}{2} \quad = \quad \int_S V_i(s|\overline{\sigma})ds.$$

The latter continuity property is key for existence: discontinuities in the ex post expected payoff given the realization of a state are not actually critical, for it is continuity of future payoffs as a function of today's action that is needed. Note that while this solves the most serious continuity problem, we have introduced a wrinkle in the equilibrium analysis, as the players' actions in the second period are now correlated. □

So, by giving $\overline{\Sigma}$ the narrow convergence metric, we obtain not only compactness, but continuity of the ex ante expected payoff in period two of the example. Although the distribution of the second-period state is assumed uniform in the example, we can let the state s be drawn from a jointly continuous distribution $F(s|a_1)$, with density $f(s|a_1)$, as a function of player 1's first-period action. For simplicity assume that the second-period strategies are such that the players play (U, R) for $s < c$ and (D, L) for $s \geq c$. Then player 1's expected payoff as a function of the first-period action is

$$U_1^c(a_1) \quad = \quad \int_{[0,c]} (1) f(s|a_1) ds + \int_{[c,1]} (0) f(s|a_1) ds \quad = \quad F(c|a_1),$$

which is jointly continuous in (a_1, c).

Motivated by this observation, we now formulate continuation values as functions of states and action profiles, where we endow $X \times \overline{\Sigma}$ with the product metric, so that a sequence $\{(a^m, \overline{\sigma}^m)\}$ converges to $(a, \overline{\sigma})$ in $X \times \overline{\Sigma}$ if and only if $a^m \to a$ and $\overline{\sigma}^m \to \overline{\sigma}$. We write $V_i(s, a)$ as the expected discounted payoff of player i at the beginning of next period, assuming action profile a is taken given state s in the current period. This function satisfies the following recursion: for all s and all a,

$$V_i(s, a|\overline{\sigma}) \quad = \quad \int_S \int_X \Big[u_i(s', a') + \delta_i V_i(s', a'|\overline{\sigma}) \Big] \overline{\sigma}(da'|s') \mu(ds'|s, a).$$

This recursion is uniquely satisfied by these continuation values, and a contraction mapping argument establishes that $V_i(\cdot|\overline{\sigma})$ is measurable; Lemma 4.10 contains details. Induced games are defined in the usual way, except that now we define induced payoffs using the new formulation of continuation values:

$$U_i^{\overline{\sigma}}(a|s) \quad = \quad u_i(s, a) + \delta_i V_i(s, a|\overline{\sigma}).$$

Set-wise continuity of μ, in (A10), implies that $U_i^{\overline{\sigma}}(a|s)$ is continuous in a, and it follows that every induced game admits at least one mixed strategy Nash equilibrium, but as in the previous subsection, the crucial step is to show that these induced payoffs are jointly continuous in $(a, \overline{\sigma})$.

To obtain joint continuity, we strengthen our continuity assumption on the transition probability to *norm continuity* of $\mu(\cdot|s, a)$ in a.

(A12) for all $s \in S$, all a, and all convergent sequences $\{a^m\}$ with limit a, we have $\mu(\cdot|s, a^m) \to \mu(\cdot|s, a)$ in total variation.

Although more restrictive than our previous assumption (A10), norm continuity is still arguably quite weak. Suppose, for example, that for all (s, a), the probability measure $\mu(\cdot|s, a)$ has a density $f(\cdot|s, a)$ with respect to κ, that these densities are continuous in a, and that they are uniformly bounded.

Then given a sequence $\{a^m\}$ converging to a in X, and given any sequence $\{Y_m\}$ of measurable subsets of X, we have

$$\begin{aligned}
&|\mu(Y_m|s,a^m) - \mu(Y_m|s,a)| \\
&\quad = \left| \int_S I_{Y_m}(s')f(s'|s,a^m)\kappa(ds') - \int_S I_{Y_m}(s')f(s'|s,a)\kappa(ds') \right| \\
&\quad \leq \int_S I_{Y_m}(s')|f(s'|s,a^m) - f(s'|s,a)|\kappa(ds') \\
&\quad \leq \int_S |f(s'|s,a^m) - f(s'|s,a)|\kappa(ds') \\
&\quad \to 0,
\end{aligned}$$

where the limit follows from continuity of f and Lebesgue's dominated convergence theorem. Thus, $\mu(\cdot|s,a^m) \to \mu(\cdot|s,a)$ in total variation, and (A12) is satisfied.

Before proceeding to the continuity argument, we first address a small technical issue, which is verifying that the set of real-valued mappings that are measurable in s and continuous in $(a,\overline{\sigma})$ is complete with the sup metric. The proof follows very standard lines. Let $\mathcal{V}(S \times X \times \overline{\Sigma})$ denote the set of mappings $f\colon S \times X \times \overline{\Sigma} \to \mathbb{R}$ such that $f(s,a,\overline{\sigma})$ is measurable in s, continuous in $(a,\overline{\sigma})$, and bounded in absolute value by $\frac{b}{1-\delta_i}$.

Lemma 4.9. *The space $\mathcal{V}(S \times X \times \overline{\Sigma})$ equipped with the sup metric is complete.*

Proof. Let $\{f^m\}$ be a Cauchy sequence in $\mathcal{V}(S\times X\times\overline{\Sigma})$ with the sup metric, ρ^s. Giving $S \times X \times \overline{\Sigma}$ the product metric, it is known that $\mathcal{V}(S \times X \times \overline{\Sigma})$ is a subset of the bounded, measurable mappings $\mathcal{M}_b(S \times X \times \overline{\Sigma})$, which is itself is complete. Therefore, there exists $f \in \mathcal{M}_b(S \times X \times \overline{\Sigma})$ such that $\rho^s(f^m, f) \to 0$. Clearly, f is bounded in absolute value by $\frac{b}{1-\delta_i}$. We must show that $f(s,a,\overline{\sigma})$ is continuous in $(a,\overline{\sigma})$. Consider any $(a,\overline{\sigma})$ and any sequence $\{(a^m,\overline{\sigma}^m)\}$ converging to $(a,\overline{\sigma}$ in $X\times\overline{\Sigma}$. Let $\epsilon > 0$. There is some k such that $\rho^s(f^k, f) < \frac{\epsilon}{3}$. Since f^k is continuous, there exists M such that for all $m \geq M$, we have $|f^k(s,a^m,\overline{\sigma}^m) - f^k(s,a,\overline{\sigma})| < \frac{\epsilon}{3}$. Therefore, for

$m \geq M$, we have

$$\begin{aligned}
&|f(s,a^m,\overline{\sigma}^m) - f(s,a,\overline{\sigma})| \\
&\quad \leq \; |f(s,a^m,\overline{\sigma}^m) - f^k(s,a^m,\overline{\sigma}^m)| + |f^k(s,a^m,\overline{\sigma}^m) - f^k(s,a,\overline{\sigma})| \\
&\qquad + |f^k(s,a,\overline{\sigma}) - f(s,a,\overline{\sigma})| \\
&\quad < \; \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} \\
&\quad = \; \epsilon,
\end{aligned}$$

and we conclude that $f(s,a,\overline{\sigma})$ is continuous in $(a,\overline{\sigma})$, so $f \in \mathcal{V}(S \times X \times \overline{\Sigma})$, as required. □

Joint continuity of continuation values, and therefore of induced payoffs, is established in the next lemma, the proof of which is considerably more complex than continuity arguments encountered in more structured games.

Lemma 4.10. *In every stochastic game satisfying (A0)–(A12), each continuation value mapping* $V_i \colon S \times X \times \overline{\Sigma} \to \mathbb{R}$ *is such that* $V_i(s,a|\overline{\sigma})$ *is measurable in* s *and continuous in* $(a,\overline{\sigma})$.

Proof. Define the operator $\Psi \colon \mathcal{V}(S \times X \times \overline{\Sigma}) \to \mathcal{V}(S \times X \times \overline{\Sigma})$ so that for all $f \in \mathcal{V}(S \times X \times \overline{\Sigma})$,

$$\Psi(f)(s,a,\overline{\sigma}) \;=\; \int_S \int_X \Big[u_i(s',a') + \delta_i f(s',a',\overline{\sigma}) \Big] \overline{\sigma}(da'|s') \mu(ds'|s,a).$$

To verify that $\Psi(f)(s,a,\overline{\sigma})$ is indeed continuous in $(a,\overline{\sigma})$, consider any state s and any convergent sequence $\{(a^m,\overline{\sigma}^m)\}$ with limit $(a,\overline{\sigma})$. Because $u_i(s,a)$ is a Carathéodory integrand, this term is unproblematic. We must show that

$$\begin{aligned}
&\int_S \int_X f(s',a',\overline{\sigma}^m)\overline{\sigma}^m(da'|s')\mu(ds'|s,a^m) \\
&\quad \to \; \int_S \int_X f(s',a',\overline{\sigma})\overline{\sigma}(da'|s')\mu(ds'|s,a).
\end{aligned}$$

Write

$$\begin{aligned}
&\int_S \int_X f(s', a', \overline{\sigma}^m)\overline{\sigma}^m(da'|s')\mu(ds'|s, a^m) \\
&= \left[\int_S \int_X f(s', a', \overline{\sigma}^m)\overline{\sigma}^m(da'|s')\mu(ds'|s, a^m)\right. \\
&\quad \left. - \int_S \int_X f(s', a', \overline{\sigma}^m)\overline{\sigma}^m(da'|s')\mu(ds'|s, a)\right] \\
&\quad + \int_S \int_X f(s', a', \overline{\sigma}^m)\overline{\sigma}^m(da'|s')\mu(ds'|s, a).
\end{aligned}$$

For each m, define the function $g^m(s') = \int_X f(s', a', \overline{\sigma}^m)\overline{\sigma}^m(da'|s')$. Then the difference between the first two terms on the right-hand side above is $\int_S g^m(s')\mu(ds'|s, a^m) - \int_S g^m(s')\mu(ds'|s, a)$, which goes to zero by norm continuity of μ and the fact that $\{g^m\}$ are uniformly bounded. Next, rewrite the third term on the right-hand side as

$$\begin{aligned}
&\int_S \int_X f(s', a', \overline{\sigma}^m)\overline{\sigma}^m(da'|s')\mu(ds'|s, a) \\
&= \int_S \int_X [f(s', a', \overline{\sigma}^m) - f(s', a', \overline{\sigma})]\overline{\sigma}^m(da'|s')\mu(ds'|s, a) \\
&\quad + \int_S \int_X f(s', a', \overline{\sigma})\overline{\sigma}^m(da'|s')\mu(ds'|s, a).
\end{aligned}$$

Note that for each s', we have

$$\int_X [f(s', a', \overline{\sigma}^m) - f(s', a', \overline{\sigma})]\overline{\sigma}(da'|s') \leq \max_{a' \in X} |f(s', a', \overline{\sigma}^m) - f(s', a', \overline{\sigma})| \to 0,$$

where the limit follows from continuity of $f(s', a', \overline{\sigma})$ in $\overline{\sigma}$ and the theorem of the maximum. Then by Lebesgue's dominated convergence theorem, we have

$$\int_S \int_X [f(s', a', \overline{\sigma}^m) - f(s', a', \overline{\sigma})]\overline{\sigma}^m(da'|s')\mu(ds'|s, a) \quad \to \quad 0.$$

Finally, using the assumption that $\mu(\cdot|s, a) \ll \kappa$, the Radon-Nikodym theorem yields a density $h(\cdot|s, a)$ for $\mu(\cdot|s, a)$ with respect to κ, and we write

$$\begin{aligned}
&\int_S \int_X f(s', a', \overline{\sigma})\overline{\sigma}^m(da'|s')\mu(ds'|s, a) \\
&= \int_S \int_X f(s', a', \overline{\sigma})h(s'|s, a)\overline{\sigma}^m(da'|s')\kappa(ds').
\end{aligned}$$

Letting f be bounded in absolute value by c, we can define the function $g(s') = c\,h(s'|s,a)$ to verify that $f(\cdot,\overline{\sigma})h(\cdot|s,a)$ is a Carathéodory integrand. Then narrow convergence implies

$$\begin{aligned}
&\lim_{m\to\infty}\int_S\int_X f(s',a',\overline{\sigma}^m)\overline{\sigma}^m(da'|s')\mu(ds'|s,a^m)\\
&= \lim_{m\to\infty}\int_S\int_X f(s',a',\overline{\sigma})\overline{\sigma}^m(da'|s')\mu(ds'|s,a)\\
&= \int_S\int_X f(s',a',\overline{\sigma})h(s'|s,a)\overline{\sigma}(da'|s)\kappa(ds')\\
&= \int_S\int_X f(s',a',\overline{\sigma})\overline{\sigma}(da'|s)\mu(ds'|s,a),
\end{aligned}$$

as required. For every $f,g\in\mathcal{V}(S\times X\times\overline{\Sigma})$, note that

$$\begin{aligned}
&\rho^s(\Psi(f),\Psi(g))\\
&= \sup_{(s,\overline{\sigma})}\delta_i\left|\int_{s'}\int_{a'} f(s',a',\overline{\sigma})\overline{\sigma}(da'|s')\mu(ds'|s,a)\right.\\
&\qquad\left. -\int_S\int_X g(s',a',\overline{\sigma})\overline{\sigma}(da'|s)\mu(ds'|s,a)\right|\\
&\le \delta_i\sup_{(s,\overline{\sigma})}\int_S\int_X |f(s',a',\overline{\sigma})-g(s',a',\overline{\sigma})|\overline{\sigma}(da'|s')\mu(ds'|s,a)\\
&\le \delta_i\rho^s(f,g),
\end{aligned}$$

and therefore Ψ is a contraction mapping with modulus $\delta_i<1$. Since $\mathcal{V}(S\times X\times\overline{\Sigma})$ is complete, by Lemma 4.9, the contraction mapping theorem implies that Ψ has a unique fixed point, and that this belongs to $\mathcal{V}(S\times X\times\overline{\Sigma})$. By construction, the fixed point is V_i, completing the proof of the lemma. □

Obviously, Lemma 4.10 implies that the induced payoff $U_i^{\overline{\sigma}}(a|\overline{\sigma})$ is also jointly continuous in $(a,\overline{\sigma})$. We can now define a fixed point correspondence using a construction that is rather different from that used in previous subsections; we do so in two steps. We have seen that correlation now plays an important role in maintaining upper hemi-continuity, so it becomes important to consider probability measures on mixed strategy profiles in induced games. Given $\overline{\sigma}$, continuity of induced payoffs implies that the set $N^{\overline{\sigma}}(s)$ of mixed strategy Nash equilibria is nonempty in each state. Let $\alpha_i\in\Delta(A_i(s))$ denote a mixed strategy in the induced game, and let $\alpha=(\alpha_1,\ldots,\alpha_n)$ denote a mixed strategy profile; following above conventions, we also let

$\alpha = \otimes_i \alpha_i$ denote the product probability measure determined by the mixed strategies of the players. Then we can view each $\alpha \in N^{\overline{\sigma}}(s)$ as an element of $\Delta(X)$. Endowing $\Delta(X)$ with the Prohorov metric, it is compact, and it is known that $N^{\overline{\sigma}}(s)$ is a closed (and therefore compact) subset of $\Delta(X)$, but it need not be convex. We can, however, endow $\Delta(X)$ with the Borel $\overline{\sigma}$-algebra $\mathcal{B}_{\Delta(X)}$, and then the set $\Delta(\Delta(X))$ of probability measures over probability measures on action profiles is compact. In particular, the set $\Delta(N^{\overline{\sigma}}(s))$ of mixtures of mixed strategy equilibria is nonempty, closed, and convex.

We make use of the following fundamental continuity property of mixed strategy Nash equilibria of induced games. The next lemma follows directly from joint continuity of $U_i^{\overline{\sigma}}(a|s)$ in $(a, \overline{\sigma})$, established in Lemma 4.10; the proof relies on standard continuity arguments and is omitted.

Lemma 4.11. *In every stochastic game satisfying (A0)–(A12), the correspondence $N^{(\cdot)}(s)\colon \overline{\Sigma} \rightrightarrows \Delta(X)$ is upper hemi-continuous with closed values for each state s.*

In fact, we will consider transition probabilities from the set S of states to the set $\Delta(X)$; denoting such a transition probability by ν, we interpret $\nu(\cdot|s)$ as a randomization over mixed strategy profiles α in the induced game at state s. Recall that the set of such transition probabilities is denoted $\mathcal{R}(\Delta(X), S, \kappa)$. The first step in our construction is to define a correspondence from joint strategies to this set of transition probabilities: define $\Psi\colon \overline{\Sigma} \rightrightarrows \mathcal{R}(\Delta(X), S, \kappa)$ so that

$$\Psi(\overline{\sigma}) \quad = \quad \left\{ \nu \in \mathcal{R}(\Delta(X), S, \kappa) \;\middle|\; \begin{array}{c} \text{for } \kappa\text{-almost all } s, \\ \nu\big(N^{\overline{\sigma}}(s)\big) = 1 \end{array} \right\}.$$

Thus, $\nu \in \Psi(\overline{\sigma})$ if and only if for κ-almost all states, $\nu(\cdot|s)$ is a mixture over mixed strategy equilibria in the induced game at state s.

It is not obvious that Ψ has nonempty values: nonemptiness of $N^{\overline{\sigma}}(s)$ for each s does not immediately imply the existence of a transition probability ν that with support on $N^{\overline{\sigma}}(s)$ for almost all states, because ν must be appropriately measurable. But the assumption that the correspondence $A_i\colon S \rightrightarrows X_i$ is lower measurable ensures that there is such a ν; the proof consists of verifying that the correspondence $N^{\overline{\sigma}}\colon S \rightrightarrows \Delta(X)$ is lower measurable, followed by an application of the Kuratowski-Ryll-Nardzewski selection theorem. In addition, Ψ has convex values. Moreover, for each $\overline{\sigma} \in \overline{\Sigma}$, the set $\Psi(\overline{\sigma})$ is a compact subset of $\overline{\Sigma}$ with the narrow convergence metric. The following lemma establishes these properties; I omit the nonemptiness proof for brevity.

Lemma 4.12. *In every stochastic game satisfying (A0)–(A12), the correspondence Ψ has nonempty, compact, convex values and is upper hemi-continuous.*

Proof. Consider any $\overline{\sigma} \in \overline{\Sigma}$. Note that $\Delta(X)$ is compact and therefore separable. Moreover, because $N^{\overline{\sigma}}$ is lower measurable with closed values, by Lemma 4.11, the subset of transition probabilities from S to $\Delta(X)$ that place probability one on the values of $N^{\overline{\sigma}}(s)$ for κ-almost all s, i.e.,

$$\Psi(\overline{\sigma}) \;=\; \{\nu \in \mathcal{R}(\Delta(X), S, \kappa) \mid \text{for } \kappa\text{-almost all } s,\ \nu(N^{\overline{\sigma}}(s)|s) = 1\}$$

is a compact in the narrow convergence metric, so Ψ has compact values. Given any $\nu, \nu' \in \Psi(\overline{\sigma})$ and any $\gamma \in (0,1)$, the transition probability $\gamma\nu + (1-\gamma)\nu'$ also places probability one on $N^{\overline{\sigma}}(s)$ for κ-almost all s. Indeed, since ν does so outside a κ-measure zero set $Z \subseteq S$ of states, and ν' does so outside a κ-measure zero set $Z' \subseteq S$ of states, it follows that the convex combination $\gamma\nu + (1-\gamma)\nu'$ does so outside $Z \cup Z'$, which is measure zero. Thus, Ψ has convex values.

Next, we consider upper hemi-continuity. We begin by recalling that, from Lemma 4.11, the correspondence defined by $N^{\overline{\sigma}}(s)$ is upper hemi-continuous in $\overline{\sigma}$ and has closed values for each s; thus it has closed graph. Let $\{\overline{\sigma}^m\}$ be a sequence converging narrowly to $\overline{\sigma}$ in $\overline{\Sigma}$, and consider $\{\nu^m\}$ such that $\nu^m \in \Psi(\overline{\sigma}^m)$ for all m and $\nu^m \to \nu$ narrowly. It is known that for κ-almost all s,

$$\operatorname{supp}(\nu(\cdot|s)) \;\subseteq\; \bigcap_{m=1}^{\infty} \operatorname{clos}\left(\bigcup_{k=m}^{\infty} \operatorname{supp}(\nu^k(\cdot|s))\right).$$

Let s be a state such that the above inclusion holds. Consider any $\alpha \in \operatorname{supp} \nu(\cdot|s)$, and note that for all m, we have $\alpha \in \operatorname{clos}\big(\bigcup_{k=m}^{\infty} \operatorname{supp} \nu^k(\cdot|s)\big)$. This means that there is a sequence $\{\alpha^{m,\ell}\}_{\ell=1}^{\infty}$ in $\bigcup_{k=m}^{\infty} \operatorname{supp} \nu^k(\cdot|s)$ such that $\alpha^{m,\ell} \to \alpha$ weakly as $\ell \to \infty$. For each m, we can choose ℓ_m high enough that α^{m,ℓ_m} is within a distance of $\frac{1}{m}$ to α. As well, there exists k_m such that $\alpha^{m,\ell_m} \in \text{supp}\, \nu^{k_m}(\cdot|s)$. Letting $\beta^m = \alpha^{m,\ell_m}$, we then have $\beta^m \in \text{supp}\, \nu^{k_m}(\cdot|s)$ for all m and $\beta^m \to \alpha$ weakly. By construction, $\beta^m \in N^{\overline{\sigma}^{k_m}}(s)$ for all m, and by closed graph, we conclude that $\alpha \in N^{\overline{\sigma}}(s)$. Therefore, we have shown that $\nu(N^{\overline{\sigma}}(s)|s) = 1$ for κ-almost all s, i.e., $\nu \in \Psi(\overline{\sigma})$, as required. ☐

The second step in the construction is to define a mapping $\phi\colon \mathcal{R}(\Delta(X), S, \kappa) \to \overline{\Sigma}$ as follows. First, given measurable $Y \subseteq X$, define the evaluation functional $e_Y\colon \Delta(X) \to \mathbb{R}$ by $e_Y(\alpha) = \alpha(Y)$. It is known that for a given Y, the function e_Y is measurable. Then define ϕ so that for κ-almost all s and all measurable $Y \subseteq X$,

$$\phi(\nu)(Y|s) \;=\; \int_{\Delta(X)} e_Y(\alpha)\nu(d\alpha|s).$$

That is, the probability that we realize an action profile in Y at state s is the expected probability derived from mixtures $\alpha(Y)$, as we integrate over α according to $\nu(\cdot|s)$.

Lemma 4.13. *In every stochastic game satisfying (A0)–(A12), the mapping ϕ is continuous.*

Proof. Consider a sequence $\{\nu^m\}$ converging narrowly to ν in $\mathcal{R}(\Delta(X), S, \kappa)$. For each m, let $\overline{\sigma}^m \in \overline{\Sigma}$ be defined by $\overline{\sigma}^m = \phi(\nu^m)$, and let $\overline{\sigma} = \phi(\nu)$. We must show that $\overline{\sigma}^m \to \overline{\sigma}$ narrowly. It suffices to show that for every bounded, continuous $f\colon X \to \mathbb{R}$ and for all measurable $S' \subseteq S$, we have

$$\int_{S'}\int_X f(a)\overline{\sigma}^m(da|s)\kappa(ds) \;\to\; \int_{S'}\int_X f(a)\overline{\sigma}(da|s)\kappa(ds).$$

For each m, we have

$$\int_{S'}\int_X f(a)\overline{\sigma}^m(da|s)\kappa(ds) \;=\; \int_{S'}\int_{\Delta(X)}\int_X f(a)\alpha(da)\nu^m(d\alpha|s)\kappa(ds).$$

and we have

$$\int_{S'}\int_X f(a)\overline{\sigma}(da|s)\kappa(ds) \;=\; \int_{S'}\int_{\Delta(X)}\int_X f(a)\alpha(da)\nu(d\alpha|s)\kappa(ds).$$

Defining $F\colon \Delta(X) \to \mathbb{R}$ by

$$F(\alpha) \;=\; \int_X f(a)\alpha(da),$$

we must show that

$$\int_{S'}\int_{\Delta(X)} F(\alpha)\nu^m(d\alpha|s)\kappa(ds) \;\to\; \int_{S'}\int_{\Delta(X)} F(\alpha)\nu(d\alpha|s)\kappa(ds).$$

By the assumption that f is bounded and continuous, it follows that F is bounded and continuous, and then the desired limit follows from narrow convergence of $\{\nu^m\}$ to ν. □

Finally, we define $\Phi\colon \overline{\Sigma} \rightrightarrows \overline{\Sigma}$ by the composition $\phi \circ \Psi$, i.e., for all $\overline{\sigma} \in \overline{\Sigma}$,

$$\Phi(\overline{\sigma}) \;=\; \phi(\Psi(\overline{\sigma})).$$

For a fixed point $\overline{\sigma}$ of the correspondence Φ, we have $\overline{\sigma} \in \Phi(\overline{\sigma})$, and so there exists $\nu \in \Psi(\overline{\sigma})$ such that $\overline{\sigma} = \phi(\nu)$. The transition probability ν gives us a mixture over mixed strategy Nash equilibria of induced games at κ-almost all states s that engenders the distribution $\overline{\sigma}(\cdot|s)$. To interpret this, imagine that we augment the stochastic game so that at the beginning of each period, the players observe the realization of a public randomization device, or "sunspot," allowing them to coordinate on a given equilibrium of the induced game depending on this realization. In the game augmented with this randomization device, the players are in fact best responding in all states after all realizations of the device, and $\overline{\sigma}$ represents a marginal distribution over action profiles (after integrating out the sunspot).

For this reason, we call ν a *stationary Markov perfect equilibrium with public randomization.* The following result is due to Nowak and Raghavan (1992), who used a fixed point argument in the space of continuation value functions; the proof here, which is in the space of transition probabilities, appears to be new. Naturally, the existence argument consists of verifying the conditions of Glicksberg's fixed point theorem to deduce a fixed point.

Proposition 4.14. *Every stochastic game satisfying (A0)–(A12) admits a stationary Markov perfect equilibrium with public randomization.*

Proof. We have already noted that $\overline{\Sigma}$ is nonempty, convex, and compact in the narrow convergence metric. Because Ψ has nonempty and compact values, continuity of ϕ immediately implies that $\Phi(\overline{\sigma})$ is nonempty and compact. Given any $\overline{\sigma} \in \overline{\Sigma}$, consider $\hat{\overline{\sigma}}, \hat{\overline{\sigma}}' \in \Phi(\overline{\sigma})$ and $\gamma \in (0,1)$. Then there exist $\nu, \nu' \in \Psi(\overline{\sigma})$ such that $\hat{\overline{\sigma}} = \phi(\nu)$ and $\hat{\overline{\sigma}}' = \phi(\nu')$. Define $\tilde{\nu} = \gamma\nu + (1-\gamma)\nu'$, and note since Ψ has convex values, we have $\tilde{\nu} \in \Psi(\overline{\sigma})$; moreover, we have $\phi(\tilde{\nu}) \in \Phi(\overline{\sigma})$. Now note that

$$\begin{aligned}
\phi(\tilde{\nu})(Y|s) &= \int_{\Delta(X)} e_Y(\alpha)\tilde{\nu}(d\alpha|s) \\
&= \int_{\Delta(X)} e_Y(\alpha)(\gamma\nu(\cdot|s) + (1-\gamma)\nu'(\cdot|s))(d\alpha) \\
&= \gamma \int_{\Delta(X)} e_Y(\alpha)\nu(d\alpha|s) + (1-\gamma)\int_{\Delta(X)} \nu'(d\alpha|s)) \\
&= \gamma\phi(\nu)(Y|s) + (1-\gamma)\phi(\nu'(Y|s)),
\end{aligned}$$

so $\phi(\tilde{\nu}) = \gamma\phi(\nu) + (1-\gamma)\phi(\nu') = \gamma\hat{\overline{\sigma}} + (1-\gamma)\hat{\overline{\sigma}}'$, and thus $\Phi(\overline{\sigma})$ is convex. Finally, I claim that because Ψ is upper hemi-continuous, continuity of ϕ implies upper hemi-continuity of the composition. Indeed, given any $\overline{\sigma} \in \overline{\Sigma}$ and any every open set $V \subseteq \overline{\Sigma}$ with $\Phi(\overline{\sigma}) \subseteq V$, it follows that $\phi^{-1}(V)$ is open, and therefore there exists open $U \subseteq \overline{\Sigma}$ with $\overline{\sigma} \in U$ such that for all $\overline{\sigma}' \in U$, we have $\Psi(\overline{\sigma}') \subseteq \phi^{-1}(V)$, which implies $\Phi(\overline{\sigma}') = \phi(\Psi(\overline{\sigma}')) \subseteq V$. This establishes the claim. Thus, Φ fulfills the conditions of Glicksberg's theorem, and we conclude that there is a fixed point $\overline{\sigma} \in \Phi(\overline{\sigma})$, which gives us a stationary Markov perfect equilibrium with public randomization. □

The proof of Nowak and Raghavan (1992) uses a fixed point argument in the space of continuation value functions, $v\colon S \to \mathbb{R}^n$. Given any v, they consider the induced games $\{\Gamma^v(s)\}$ and take selections from the convex hull of payoffs from mixed strategy Nash equilibria of the induced games. This determines a set of new continuation value functions $\hat{v}\colon S \to \mathbb{R}^n$, and this correspondence satisfies the conditions for Glicksberg's theorem; in particular, it is convex-valued. The mathematics is somewhat different than that used here.

We have given sufficient conditions for existence of a stationary Markov equilibrium with public randomization and pointed to technical dangers that suggest we should not expect a general existence theorem in the absence of correlation. A recent paper by Levy (2013) gives counterexamples to existence of stationary Markov perfect equilibria (without correlation) for the general class of stochastic games considered in this subsection.

Alternative routes to existence without public randomization:

- ϵ-equilibrium
- semi-stationary strategies
- noisy stochastic games
- singe-mover games

4.4 Back to One-shot Bargaining

The bargaining model of Section 3, when formulated as a stochastic game, does not satisfy assumptions (A0)–(A12) of the previous section. In particular, the transition probability from a proposal to the state in the following voting stage is deterministic: if i proposes x, then next period's state is $s' = x$. This violates set-wise continuity (A10) and the stronger condition of norm continuity (A12). In fact, it violates the absolute continuity condition (A11) as well. One possibility—though not desirable on modeling grounds—would be to introduce some noise on the proposal, so that if x is chosen by the proposer, then the actual proposal before the voters is y, which is realized with some noise as a function of x, e.g., $y = x + \epsilon$, where ϵ is a random variable with support close to zero. With appropriate adjustments to ensure that proposals are feasible, this addresses the continuity problem and allows Proposition 4.14 to be applied.

Further issues arise, however. First, the result of Nowak and Raghavan (1992) delivers a correlated equilibrium. In proposer stages, this potential problem is moot, because there is a single mover: the induced game in proposer state i is simply a decision problem, and an equilibrium of the induced game is simply an optimal proposal for i, so public randomization reduces to mixing over optimal choices for the proposer, which is consistent with the concept of stationary bargaining equilibrium. Of course, in voting stages, multiple individuals cast votes for or against a proposal, so public randomization introduces non-trivial correlation. In fact, if voting were stage-undominated, even this correlation could be dealt with: if there were equilibria in stage-undominated voting strategies such that x is accepted and rejected, then some individuals would have to be indifferent, and these voters are pivotal; this would allow us to specify independent mixed voting strategies for the indifferent voters to obtain the same probability of accepted determined by the correlated equilibrium.

But the second problem is that voting strategies are not guaranteed to be stage-undominated. In fact, assuming majority voting with three or more individuals, a trivial stationary Markov perfect equilibrium of the bargaining game (formulated as a stochastic game) is that every individual propose a fixed alternative x and all individuals vote to reject every proposal. Proposition 4.14 does not imply the existence of any other kind of equilibrium.

A possible approach to this problem is to formulate the bargaining model so that voting is sequential, rather than simultaneous, and to simply impose subgame perfection in voting stages. This leaves the stationary bargaining equilibria of the bargaining model essentially unchanged, and it relies only on subgame perfection—this, unlike the stage dominance refinement, is delivered by stationary Markov perfect equilibrium. But now the first problem arises again. Suppose that individuals vote in the order $1, 2, \ldots, n$. When i chooses x, suppose a proposal y is drawn, and we move to state $(y, 1)$, where individual 1 votes to accept or reject y. Following that vote, we move to a new state, in which individual 2 votes. But if that state is (y, b_1), indicating that 1 voted b_1 on proposal the same y, then the state transition is deterministic. Because action sets are finite, condition (A12) is actually satisfied, but the absolute continuity condition (A11) is not. In order to satisfy it, we would have to denote the initial proposal drawn as y^1, and after individual 1 votes b_1 to accept or reject y^1, we would need to draw another proposal y^2, and individual 2 would vote to accept or reject y^2, and so on. In sum, we have transformed the bargaining game so that the proposed alternatives is subject to a shock not only after the original proposal, but after each individual casts his or her vote. Although technically fulfilling the requirements of Proposition 4.14, these are not appealing modeling assumptions, and we conclude that existence results for general stochastic games do not appear to be well-suited to bargaining applications.

That said, we have provided separate arguments for the existence of stationary bargaining equilibria. Two features of the bargaining model facilitate the existence argument. First, in proposal stages, only a single individual has a non-trivial action set; this ameliorates a difficult convexity issue present in general stochastic games. Second, because the bargaining game ends once an agreement is reached, there is an absorbing state in the model. After eliminating stage-dominated votes, this drastically simplifies equilibrium behavior in voting stages as a function of the proposal. We have seen that condition (a) holds, for example, with the consequence, established in Proposition 2.3, that if we vary continuation values exogenously in the bargaining model, then each individual's optimal proposals vary upper hemi-continuously.

5 Bargaining with Endogenous Status Quo

Next, we extend the one-shot bargaining framework to allow for continual bargaining, in which the alternative chosen in one period endogenously determines the status quo in the next. This structure is analyzed by Baron (1996) in a one-dimensional spatial environment and by Kalandrakis (2004, 2010) in the distributive model. This class of model should be expected to present all of the challenges of the class of one-shot models, so that results for general stochastic games are not easily applied. In fact, although the endogenous status quo bargaining models possess one key simplifying feature in common with the one-shot models—namely, the fact that the proposer moves unilaterally—they do not possess an absorbing state, and the analysis is drastically more complex. Proposal and voting strategies must now be conditioned on the current status quo, creating a vast multiplicity of histories at which we must specify behavior. Of course, equilibrium strategies must specify behavior that is optimal in all subgames given future expectations, and these expectations must be consistent with the strategies used by the individuals. We elaborate on the complexity of these models, and especially the complications entailed by the absence of an absorbing state, at a later point when more formalism is in place.

5.1 Dynamic Bargaining Framework

To address continuity issues that are by now familiar, we augment the setting of Section 3 with two types of noise: noise on the status quo and on the preferences of individuals. Respectively, we assume that the outcome in one period determines the status quo in the next stochastically; this captures uncertainty about how a decision in the current period will carry over to the next. We further assume that while preferences in the current period are common knowledge, there is some uncertainty about preferences in the next period. Both types of uncertainty can be made arbitrarily small. Results on the dynamic bargaining model in the following subsections are taken from

Duggan and Kalandrakis (2012) and the antecedent working paper, Duggan and Kalandrakis (2007).

Elements of the model:

$N = \{1, \ldots, n\}$	individuals
$X \subseteq \mathbb{R}^d$	set of alternatives
$\Phi \colon \mathbb{R}^d \rightrightarrows X$	feasibility correspondence
$\rho \colon \mathbb{R}^d \to [0,1]^n$	recognition probability mapping
$\mathcal{D} \colon \mathbb{R}^d \rightrightarrows 2^N$	voting rule
$\Theta = \prod_i \Theta_i \subseteq \mathbb{R}^{nm}$	preference shocks
$u_i \colon X \times \Theta_i \to \mathbb{R}$	utility function
$\delta_i \in [0,1)$	discount factor
$F \colon \mathbb{R}^{nm} \to [0,1]$	distribution of preference shocks
$G \colon \mathbb{R}^d \times X \to [0,1]$	status quo transition

Here, the set X of alternatives can be multidimensional, and we now allow the feasible set of alternatives, $\Phi(q)$, and the recognition probabilities, $\rho(q)$, to depend on the status quo in any period. Of course, $\sum_i \rho_i(q) = 1$. A vector of preference shocks is denoted $\theta = (\theta_1, \ldots, \theta_n)$, where for all i, $\theta_i \in \mathbb{R}^m$ is a multidimensional shock. This enters i's utility function as in $u_i(x, \theta_i)$. The vector θ is drawn iid across periods from a distribution F, and the distribution of next period's status quo given outcome x in the current period is $G(\cdot|x)$. Note that the vector θ is publicly observed at the beginning of each period.

Timing: in each period, we begin with a status quo q and a vector θ of preference shocks. Then...

- a proposer i is drawn from the distribution $\rho(q)$,
- i proposes $y \in \Phi(q)$,
- all individuals simultaneously vote to accept or reject the proposal,
- if the set of individuals who vote to accept belongs to $\mathcal{D}(q)$, then the proposal passes, and the outcome in the period is $x = y$; otherwise, the outcome is $x = q$,
- each individual receives utility $u_j(x, \theta_j)$,

- a new vector θ' of preference shocks is drawn from F, and a new status quo q' is drawn from $G(\cdot|x)$,

and the process is repeated.

As usual, given the sequence $\theta^1, \theta^2, \ldots$ of shocks, each individual evaluates a sequence $x^1, x^2, \ldots,$ of alternatives according to (normalized) discounted utility,

$$(1-\delta_i)\sum_{t=1}^{\infty}\delta^{t-1}u_i(x^t,\theta_i^t),$$

and we extend preferences to lotteries over sequences via expected utility.

We impose a number of regularity assumptions on the primitives of the model; the following are not stated in full generality. Let $r = \max\{3, d+1\}$.

Alternatives Assume that X is compact, and that for each q, we have $\Phi(q) \subseteq X$. Moreover, $\Phi(q)$ is cut out by smooth constraints, i.e., there are functions $h_\ell \colon \mathbb{R}^d \times \mathbb{R}^d$, indexed by $K = \{1, \ldots, k\}$, such that for all q,

$$\Phi(q) \;=\; \{x \in \mathbb{R}^d \mid h_\ell(x,q) \geq 0, \ell \in K^{in}, h_\ell(x,q) = 0, \ell \in K^{eq}\},$$

where $K = K^{in} \cup K^{eq}$ and K^{in} and K^{eq} partition K into, respectively, inequality and equality constraints. Assume the status quo is feasible, i.e., $q \in \Phi(q)$. Assume each function h_ℓ is measurable in q and r-times continuously differentiable in x. Finally, we assume the weak linear independence condition that gradients of binding constraints are linearly independent: for all q and all $x \in \Phi(q) \setminus \{q\}$, $\{D_x h_\ell(x,q) \mid \ell \in K(x,q)\}$ is linearly independent, where $K(x,q)$ indexes the constraints such that $h_\ell(x,q) = 0$.

Protocol We assume that ρ is measurable, that for all q, $\mathcal{D}(q)$ is nonempty and monotonic, and that for all G, the set $\{q \in \mathbb{R}^d \mid G \in \mathcal{D}(q)\}$ is measurable. Let $D(q)$ be the set of dummy voters for $\mathcal{D}(q)$.

Preferences Each u_i is r-times continuously differentiable, that Θ is open, that u_i is bounded, and that for all i, all $\theta \in \Theta$, and all distinct $x, y \in X$, $D_{\theta_i}[u_i(x,\theta_i) - u_i(y,\theta_i)] \neq 0$. This holds, e.g., if $m = d$ and (abusing notation slightly) $u_i(x,\theta_i) = u_i(x) + x \cdot \theta_i$.

Transition probabilities Assume that F has density f with support in Θ, and that f is bounded. Assume that there is a fixed probability measure ν_q such that for all x, the probability measure determined by $G(\cdot|x)$ is absolutely continuous with respect to ν_q and has support in X, and let $g(\cdot|x)$ be a density for this distribution. Assume that g is bounded, measurable in q, and r-times continuously differentiable in x. Moreover, assume that derivatives of all orders $1, \ldots, r$ are uniformly bounded.

5.2 Special Cases

The structure imposed above may seem more restrictive than it is, as many special cases can be obtained by suitable (and sometimes creative) specification of the model. For simplicity, assume except in the first special case that utilities $u_i(x, \theta_i)$ have the linear form, i.e., $u_i(x) + \theta_i \cdot x$. I do not attempt to fill in all of the details, but the aim is rather to give an idea of the flexibility of the dynamic bargaining framework.

Finite set of alternatives. To obtain a finite set of m alternatives as a special case, we let $d = 1$, and we define one equality constraint and two inequality constraints, all independent of the status quo. The equality constraint is $h_1(x) = \sin(x)$, and the two inequality constraints are $h_2(x) = x + \frac{1}{2}$ and $h_3(x) = -x + (m-1)\pi + \frac{1}{2}$. Then the set of feasible alternatives is

$$X = \left\{ x \in \mathbb{R} \mid -\frac{1}{2} \leq x \leq (m-1)\pi + \frac{1}{2},\ \sin(x) = 0 \right\},$$

which consists of m elements, indexed $x_1, \ldots, x_m$. It can be checked that our linear independence condition is satisfied, giving us a finite set of feasible alternatives. The status quo transition can be any distribution $G(\cdot|x)$ as a function of $x \in X$, where we let ν_q be the uniform distribution on X to satisfy our absolute continuity condition. A plausible formulation of preference shocks in this model is to simply perturb payoffs from alternatives independently, so that $\theta_i = (\theta_{i,1}, \ldots, \theta_{i,m}) \in \mathbb{R}^m$, and $u_i(x_h, \theta_i) = x_h + \theta_{i,h}$ for all $h = 1, \ldots, m$.

One-shot bargaining. It is straightforward to incorporate an absorbing state in the model to produce noisy versions of the models of Section 3. Consider the a bad status quo model, for example, and assume the smoothness restrictions of the dynamic bargaining model. Assuming without loss of generality that $0 \notin X$ we add an alternative $q = 0$ for use as a status quo;

and choosing $a \notin X \cup \{q\}$, we use a as an absorbing state. Then we specify a new set of alternatives as $\tilde{X} = X \cup \{q\} \cup \{a\}$; this can be done by appropriate modification of feasibility constraints. We define utilities $\tilde{u}_i \colon \tilde{X} \to \mathbb{R}$ in the extended model so that

$$\tilde{u}_i(\tilde{x}) \;\; = \;\; \begin{cases} \frac{u_i(\tilde{x})}{1-\delta_i} & \text{if } \tilde{x} \in X, \\ 0 & \text{if } \tilde{x} = q \text{ or } \tilde{x} = a. \end{cases}$$

Recognition probabilities and the voting rule are specified as in the original model, and we define the feasibility correspondence so that $\Phi(\tilde{q}) = X$ if $\tilde{q} = q$, and $\Phi(\tilde{q}) = \{0\}$ if $\tilde{q} = 0$. The state transition $\tilde{G}(\cdot|\tilde{x})$ is defined so that the status quo is equal to q with probability one if $\tilde{x} = q$, and it is equal to 0 with probability one if $\tilde{x} \in X \cup \{0\}$. Note that this distribution is absolutely continuous with respect to ν_q defined as the even chance distribution over $\{0, q\}$, and since the status quo belongs to X with zero probability, it is immaterial how we define Φ on this set. Assuming that θ_i has mean zero, this means that if the status quo q remains in place, each individual receives utility $\tilde{u}_i(q) = 0 + \theta_i \cdot 0 = 0$; and if a proposal $x \in X$ is accepted, then individual i's expected (normalized) discounted payoff is

$$(1 - \delta_i)\frac{u_i(x)}{1 - \delta_i} + \delta_i \int_{\Theta} [0 + a \cdot \theta_i] f(\theta) d\theta \;\; = \;\; u_i(x),$$

as in the original model. Thus, we obtain a version of the one-shot bargaining model in which there is uncertainty about preferences next period but is otherwise unchanged.

Costly policy making. Consider a version of the model in which there is a fixed cost of implementing (proposing and accepting) an alternative, rather than leaving the status quo in place. To capture this as a special case, we create an extra copy of the set of alternatives, say X', that is disjoint from X. Let v be a vector such that $X' = X + v$. We use X' to "host" first-time choices of alternatives and X to host status quos that are left in place. We then define an extended model with set $\tilde{X} = X \cup X'$ of alternatives such that for all $\tilde{x} \in X'$, we have

$$\tilde{u}_i(\tilde{x}) \;\; = \;\; u_i(\tilde{x} - v) - c_i,$$

where $c_i > 0$ is a fixed cost of policy change for individual i. The status quo transition is such that given $\tilde{x} \in X'$, the status quo transitions to X with probability one, as in the original model: $\tilde{G}(\cdot|\tilde{x}) = G(\cdot|\tilde{x} - v)$. The feasible set is such that for all $\tilde{q} \in X$, $\tilde{\Phi}(\tilde{q}) = \Phi(\tilde{q}) + v$, so that a proposer's feasible set is as in the original model but translated to X', where the cost c_i is incurred. Utilities, recognition probabilities, and the voting rule are defined on X exactly as in the original model. Because the status quo is drawn from X with probability one, it is immaterial how the feasible set, recognition probabilities, and voting rule are extended to X'.

In the above formulation, the policy cost is fixed. An alternative modeling assumption is that the cost can depend on the change in the chosen alternative from the status quo, i.e., $x - q$. In such a model, the utility of individual i from accepting x given status quo q can be written generally as $u_i(x, q)$, where u_i is appropriately smooth. We capture this model as a special case by specifying the set of alternatives as $\tilde{X} = X^2$, so that an alternative $\tilde{x} = (x_1, x_2)$ consists of two components. We define utility functions $\tilde{u}_i(\tilde{x}) = u_i(x_1, x_2)$, so that if x_2 is the status quo and x_1 is the chosen alternative, i's utility is as in the original model. Given $\tilde{q} = (q_1, q_2)$, define

$$\tilde{\Phi}(\tilde{q}) = \Phi(q_2) \times \{q_2\},\ \tilde{\rho}(\tilde{q}) = \rho(q_2),\ \mathcal{D}(\tilde{q}) = \mathcal{D}(q_2),$$

so recognition probabilities and the voting rule are defined as in the original model with status quo q_2. The feasible set allows the proposer to propose any alternative in the original feasible set, now with the restriction that a feasible alternative carries the status quo in its second component. Finally, the status quo transition $\tilde{G}(\cdot|\tilde{x})$ is the same as $G(\cdot|x_1)$ but now extended to the diagonal of X^2, i.e.,

$$\tilde{g}(\tilde{q}|\tilde{x}) \quad = \quad \begin{cases} g(q_1|x_1) & \text{if } q_1 = q_2 \\ 0 & \text{else.} \end{cases}$$

Note that this satisfies our absolute continuity assumption by suitably extending ν_q to the a probability measure $\tilde{\mu}_q$ concentrated on the diagonal. Thus, given status quo (q, q), an individual can propose any (x, q) such that $x \in \Phi(q)$, and if accepted, utilities are $u_i(x, q)$ and the new status quo is determined by $G(\cdot|x)$.

Sunset provisions. In the basic bargaining model, an alternative that is proposed and accepted remains in place (with noise on the status quo) until a different alternative is proposed and accepted. A more nuanced model would allow a sunset provision, whereby the alternative, if left in place, persists for a predefined number of periods, say h, and then reverts to a default. We assume that the horizon, h, is bounded above by T. For simplicity, assume that the set of feasible alternatives is X, independent of the status quo. We capture the model with sunset provisions as a special case by specifying the set of alternatives as $\tilde{X} = X^{T+1}$, the $(T+1)$-fold product of alternatives in the original model. Given an alternative $\tilde{x} = (x_1, \ldots, x_{T+1}) \in \tilde{X}$, we specify utility $\tilde{u}_i(\tilde{x}) = u_i(x_1)$, so that utility (modulo preference shocks) is accrued from the first component of $\tilde{x}$ only. The status quo transition maps the current alternative $\tilde{x} = (x_1, \ldots, x_{T+1})$ to a status quo $\tilde{q} = (q_1, \ldots, q_{T+1})$ (plus noise), where the components of $\tilde{x}$ are shifted to the left, i.e.,

$$\tilde{q} \quad = \quad (x_2, x_3, \ldots, x_T, x_{T+1}, x_{T+1}).$$

Technically, we could let $\tilde{G}(\cdot|\tilde{x})$ be the distribution of $\tilde{q}$ plus a random variable with mean zero and appropriate smoothness properties. Thus, if the proposal $\tilde{x} = (x_1, \ldots, x_{T+1})$ is accepted and left in place, the outcome in the current period is essentially x_1 (as it is the first coordinate of $\tilde{x}$ that determines payoffs), and in the next period is x_2 (plus noise), and so on. After T periods, the alternative transitions to x_{T+1} and remains in place thereafter. This provides a noisy model of sunset provisions, in which proposals consist of a sequence of alternatives for up to T periods, after which the alternative moves to a default.

Finite state variable. Consider a version of the bargaining model with a finite set S of abstract states that enter the feasible set, recognition probabilities, the voting rule, payoffs, and the status quo transition, e.g.,

$$\Phi(s, q),\ \rho(s, q),\ \mathcal{D}(s, q),\ u_i(s, x),\ G(q|s, x),$$

where the state transition $p(s'|s, x)$ is a function of the current state and an r-times continuously differentiable function of the alternative chosen. We obtain such a model by creating $|S|$ copies of the dynamic bargaining model, indexing them by $m = 1, \ldots, |S|$, and specifying that the sets of alternatives $X^1, \ldots, X^{|S|}$ of the replicated model are disjoint. We let $X^1 = X$, making the first copy a baseline; note that for all copies m, there is a vector v^m such that $X^m = X + v^m$. We then define a new model such that the set of alternatives is $\tilde{X} = \bigcup_{m=1}^{|S|} X^m$. Also using m to index states, we specify utilities $\tilde{u}_i$ as follows: given any $\tilde{x} \in \tilde{X}$, there is some copy such that $\tilde{x} \in X^m$, and we set $\tilde{u}_i(\tilde{x}) = u_i(s_m, \tilde{x} - v^m)$. We specify $\tilde{\Phi}(\tilde{q}) = \Phi(s_m, \tilde{q} - v^m) + v^m$, $\tilde{\rho}(\tilde{q}) = \rho(s_m, \tilde{q} - v^m)$, and $\tilde{\mathcal{D}}(\tilde{q}) = \mathcal{D}(s_m, \tilde{q} - v^m)$, for $\tilde{q} \in X^m$. In particular, to define the feasible set $\tilde{\Phi}(\tilde{q})$, we locate the copy in which $\tilde{q}$ lies to identify the state, say s_m; we translate $\tilde{q}$ back to the baseline model, find the feasible set given s_m and the underlying status quo $\tilde{q} - v^m$; and we then translate this back to the mth copy of the model. Finally, the status quo transition $\tilde{G}(\tilde{q}|\tilde{x})$ is defined so that the density satisfies

$$\tilde{g}(\tilde{q}|\tilde{x}) \;=\; p(s_\ell|s_m, \tilde{x} - v^m)g(\tilde{q} - v^\ell|\tilde{x} - v^m),$$

for $\tilde{q} \in X^\ell$ and $\tilde{x} \in X^m$. That is, a given $\tilde{q} \in \tilde{X}$ belongs to some copy, say X^ℓ, which identifies the status quo with state s_ℓ. Then the probability of transitioning to such a status quo is equal to the probability of state s_ℓ given underlying alternative $\tilde{x} - v^m$ in state s_m.

Two-party competition. Consider a model in which two parties, A and B, with mixed motivations compete for a political office. Assume that each period, the incumbent party is locked into a position $x \in X$, while the challenging party chooses a policy platform $y \in X$. After doing so, an odd number n of voters cast their ballots in a majority rule election, and the winner implements her platform. In the next period, the incumbent party is again fixed at her position x' (either x or y), and the challenging party chooses a policy platform y', and so on. Given two platforms, each voter calculates the expected discounted payoff from electing each party and votes for the party offering the higher payoff. A stationary Markov perfect equilibrium determines a sequence $z^1, z^2, \ldots$ of policies and a sequence $w^1, w^2, \ldots$ of winners, and the payoff of party A is

$$\sum_{t=1}^{\infty} \delta_A^{t-1}[u_A(z^t) + \beta I_A(w^t)],$$

where $\beta \geq 0$ is an office benefit term, and $I_A(\cdot)$ is an indicator function with value 1 if $w^t = A$ and 0 otherwise, and B's payoff is analogous. Voters also receive the discounted utility from policies, where the utility function of voter i from policy z is $u_i(z)$. We can capture a version of this model with transient preference shocks and noise added to the platform of the incumbent. We specify the set of individuals as $\tilde{N} = \{1, \ldots, n, n+1, n+2\}$, where $n+1$ and $n+2$ correspond to parties A and B, respectively, and the set of alternatives is $\tilde{X} = X \times \{A, B\}$. The status quo $\tilde{q} = (q, A)$ indicates that A is the incumbent and has platform q, and we define $\tilde{\Phi}(\tilde{q}) = X \times \{B\}$, $\tilde{\rho}_{n+2}(\tilde{q}) = 1$, and $\tilde{\mathcal{D}}$ as majority rule among $\{1, \ldots, n\}$. Since B is the proposer with probability one, a proposal $(y, B) \in \tilde{\Phi}(\tilde{q})$ indicates that B is the challenging party and locates at y. When B is incumbent, feasible alternatives and recognition probabilities are defined symmetrically. An outcome $\tilde{x} = (x, A)$ indicates that party A holds office and implements its platform x, which generates utility $\tilde{u}_i(\tilde{x}) = u_i(x)$ for voters $i = 1, \ldots, n$, while parties receive

$\tilde{u}_{n+1}(\tilde{x}) = u_A(x) + \beta$ and $\tilde{u}_{n+2}(\tilde{x}) = u_B(x)$, with a symmetric specification when B holds office. Finally, the status quo transition probability is defined so that $G(\cdot|(x,j))$ puts probability one on $X \times \{j\}$, reflecting the fact that the office holder j becomes incumbent in the following period.

5.3 Equilibrium Existence and Characterization

A pure stationary strategy for individual i is now a pair $s_i = (p_i, a_i)$, where $p_i \colon X \times \Theta \to X$ such that for all $(q,\theta) \in X \times \Theta$, we have $p_i(q,\theta) \in \Phi(q)$, and $a_i \colon X \times X \times \Theta \to \{0,1\}$. Here, $p_i(q,\theta)$ is the alternative proposed by i with status quo q and preference shocks θ, and $a_i(y,q,\theta) = 1$ if i accepts proposal y with status quo q and shocks θ. A pure stationary strategy profile is denoted $s = (s_1, \ldots, s_n)$. We say s is *no-delay* if for all q, all θ, and all i, we have $\{j \in N \mid a_j(p_i(q,\theta), q, \theta) = 1\} \in \mathcal{D}(q)$. As in the bad status quo model, we can focus on no-delay strategy profiles generally. In contrast to one-shot bargaining and stochastic games, and as stated formally later, the noise in the dynamic bargaining model permits us to focus on pure strategies with no loss of generality.

Given a pure, no-delay stationary strategy profile s, the continuation value $V_i(x|s)$ represents i's expected discounted payoff beginning next period given the choice of x in the current period. This satisfies the following recursion: for all x,

$$\begin{aligned} &V_i(x|s) \\ &= \int_X \int_\Theta \sum_j \rho_j(q) \Big[(1-\delta_i)u_i(p_j(q,\theta),\theta_i) + \delta V_i(p_j(q,\theta)|s)\Big] f(\theta) g(q|x) d\theta dq. \end{aligned}$$

Then we define the *dynamic utility*

$$U_i(x,\theta_i|s) \;\; = \;\; (1-\delta_i)u_i(x,\theta_i) + \delta_i V_i(x|s).$$

Note that these quantities are analogous to the continuation values $V_i(s,a|\overline{\sigma})$ and induced payoff $U_i^{\overline{\sigma}}(a|s)$ in the general stochastic game framework from the previous section. For a strategy profile in which delay occurs with positive probability following some status quos and preference shocks, we can characterize continuation values using a similar, but more complicated, recursion.

We can define mixed strategies $\sigma_i = (\pi_i, \alpha_i)$ in the obvious way, where now $\pi_i \colon X \times \Theta \to \Delta(X)$ is a transition probability such that for all (q,θ),

$\pi_i(\Phi(q)|q,\theta) = 1$, and $\alpha_i \colon X \times X \times \Theta \to [0,1]$ is measurable. We let $\overline{\alpha}(x,q,\theta)$ denote the probability that the proposal x is accepted given status quo q and shocks θ. Given a profile $\sigma = (\sigma_1, \ldots, \sigma_n)$, continuation values $V_i(x|\sigma)$ satisfy

$$\begin{aligned} V_i(x|\sigma) &= \int_X \int_\Theta \sum_j \rho_j(q) \int_X \Big[\overline{\alpha}(x,q,\theta) U_i(x,\theta_i|\sigma) + \\ &\qquad (1 - \overline{\alpha}(x,q,\theta) U_i(q,\theta_i|\sigma)\Big] \pi_j(dx|q,\theta) f(\theta) g(q|x) d\theta dq, \end{aligned}$$

while dynamic utilities satisfy

$$U_i(x,\theta_i|\sigma) \;=\; (1-\delta_i)u_i(x,\theta_i) + \delta_i V_i(x|\sigma).$$

We say a mixed stationary strategy profile σ is *equivalent* to a pure, no-delay stationary strategy profile s if for all q and almost all θ, the alternatives determined by σ are the same as those determined by s with probability one, i.e., for all q, for almost all θ, and for all i, we have: (i) if $p_i(q,\theta) \neq q$, then $\pi_i(\{p_i(q,\theta)\}|q,\theta) = \overline{\alpha}(p_i(q,\theta),q,\theta|s) = 1$, and (ii) if $p_i(q,\theta) = q$, then no proposals other than the status quo pass with positive probability,

$$\int_{X \setminus \{q\}} \overline{\alpha}(y,q,\theta|s) \pi_i(dy|q,\theta) \;=\; 0.$$

Note that the definition of equivalence is defined in two parts, because there are two, payoff-equivalent ways the status quo can prevail during a given period—the status quo can be proposed and accepted, or a proposal can be rejected.

A *stationary bargaining equilibrium* is a stationary strategy profile σ such that σ is a subgame perfect equilibrium and voting strategies are stage-undominated, i.e.,

1. for all q, all θ, and all i,

$$\pi_i\Bigg(\arg\max_{y \in \Phi(q)} \overline{\alpha}(y,q,\theta) U_i(y,\theta_i|\sigma) + (1-\overline{\alpha}(y,q,\theta)) U_i(q,\theta_i|\sigma) \Bigg| q,\theta \Bigg),$$

2. for all q, all θ, all y, and all $i \in N \setminus D(q)$, $U_i(y,\theta_i|\sigma) > U_i(q,\theta_i|\sigma)$ implies $\alpha_i(y,q,\theta) = 1$, and $U_i(y,\theta_i|\sigma) < U_i(q,\theta_i|\sigma)$ implies $\alpha_i(y,q,\theta) = 0$.

Obviously, a pure stationary bargaining equilibrium is a profile s such that for all q, all θ, and all i, $p_i(q,\theta)$ solves the maximization problem in part 1, and for all y, a_i satisfies the implications in part 2.

Using by now familiar conventions, we can define acceptance sets

$$\begin{aligned} A_i(q,\theta_i|\sigma) &= \{x \in \Phi(q) \mid U_i(x,\theta_i|\sigma) \geq U_i(q,\theta_i|\sigma)\} \\ A_G(q,\theta|\sigma) &= \bigcap_{i \in G} A_i(q,\theta|\sigma) \\ A(q,\theta|\sigma) &= \bigcup_{G \in \mathcal{D}(q)} A_G(q,\theta|\sigma), \end{aligned}$$

with similar notation used for profiles of pure strategies. Note that $A_i(q,\theta_i|\sigma)$ depends on the preference shock of individual i alone.

The next proposition justifies our focus on stationary bargaining equilibria in pure, no-delay strategies.

Proposition 5.1. *Every stationary bargaining equilibrium σ is equivalent to a pure stationary bargaining equilibrium s such that for all $i \in N$, all q, and all θ,*

(i) $p_i(q,\theta)$ solves

$$\max_{x \in A(q,\theta|s)} U_i(x,\theta_i|s),$$

(ii) the proposal $p_i(q,\theta)$ is accepted, i.e., $\{j \in N \mid a_j(p_i(q,\theta),q,\theta) = 1\} \in \mathcal{D}(q)$.

The previous result does not touch on existence of equilibrium. This is established next, along with several regularity properties possessed by all pure, no-delay stationary bargaining equilibria.

Proposition 5.2. *There is a pure, no-delay stationary bargaining equilibrium, and every such equilibrium s satisfies the following conditions:*

1. *continuation values are smooth: for all i, $V_i(x|s)$ is r-times continuously differentiable in x,*

2. *proposals are almost always strictly best: for all i, all q, almost all θ, and all $y \in A(q,\theta|s)$ with $y \neq p_i(q,\theta)$, we have $U_i(p_i(q,\theta),\theta_i|s) > U_i(y,\theta_i|s)$,*

3. *proposal strategies are almost always smooth: for all i, all q, and almost all θ, p_i is continuously differentiable in an open set around (q,θ).*

The proofs of Propositions 5.1 and 5.2 rely on similar insights, although the existence result consists of a fixed point argument, while the characterization results essentially take a fixed point as given. We focus first on the necessary conditions in Proposition 5.2. Upon inspection of

$$\int_X \int_\Theta \sum_j \rho_j(q)\Big[(1-\delta_i)u_i(p_j(q,\theta),\theta_i) + \delta V_i(p_j(q,\theta)|s)\Big] f(\theta)g(q|x)d\theta dq,$$

differentiability of continuation values, in part 1, follows directly from differentiability of the status quo transition density: the integral depends on x through the density only, and so we have

$$\begin{aligned} DV_i(x|s) \quad &= \quad \int_X \int_\Theta \sum_j \rho_j(q)\Big[(1-\delta_i)u_i(p_j(q,\theta),\theta_i) \\ &\qquad + \delta V_i(p_j(q,\theta)|s)\Big] f(\theta)D_x g(q|x)d\theta dq \end{aligned}$$

by passing the derivative through the integral.

The claim that proposals are almost always strictly best, in part 2, relies on the definition of the set

$$A_{-i}(q,\theta_{-i}|s) \quad = \quad \bigcup_{G\in\mathcal{D}(q)} \bigcap_{j\in G\setminus\{i\}} A_j(q,\theta_j|s)$$

of alternatives that would be accepted by the members, other than i, of some decisive group; assuming i accepts his or her own proposal, these are the proposals by i that will pass. Note that

$$\max_{y\in A(q,\theta|s)} U_i(y,\theta_i|s) \quad \text{and} \quad \max_{y\in A_{-i}(q,\theta_{-i}|s)} U_i(y,\theta_i|s)$$

have the same solutions. Indeed, let x solve the first problem and y solve the second. Since $q \in A_{-i}(q,\theta_{-i}|s)$, we have $U_i(y,\theta_i|s) \geq U_i(q,\theta_i|s)$, which

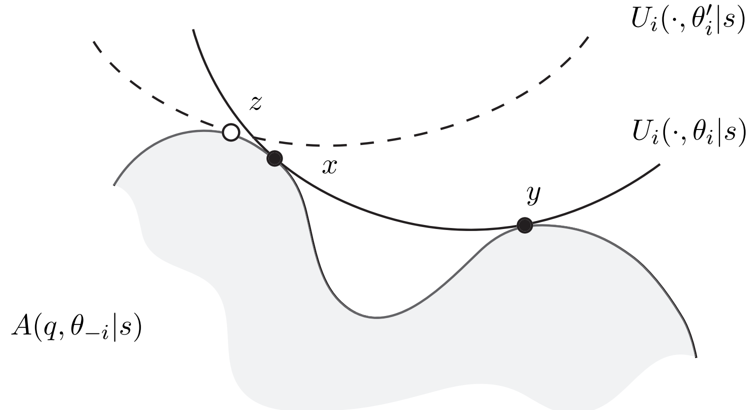

Figure 15: Genericity of unique maximizer

implies $y \in A(q, \theta|s)$. It follows that $U_i(x, \theta_i|s) \geq U_i(y, \theta_i|s)$, so x solves the second problem. Furthermore, since $A(q, \theta|s) \subseteq A_{-i}(q, \theta_{-i}|s)$, we conclude that y solves the first problem. Given q and θ_{-i}, it may be that for some preference shocks θ_i, the second maximization problem above has multiple solutions, but a perturbation to a shock such as θ_i', as in Figure 15, leads to a unique maximizer. By our assumption that $D_{\theta_i}[u_i(x, \theta_i) - u_i(y, \theta_i)] \neq 0$ for all distinct x and y, it is known that the set of θ_i for which the maximization problem has multiple solutions has Lebesgue measure zero. By Fubini's theorem, it follows that for all q and almost all θ, individual i's optimal proposal is unique.

To see part 3, we note that the problem of maximizing the dynamic utility of a proposer i subject to the acceptance of a given group G takes the standard form of maximization subject to a finite number of equality and inequality constraints:

$$\begin{aligned} &\max\nolimits_{y \in \mathbb{R}^d} U_i(y, \theta_i|s) \\ \text{s.t. } &U_j(y, \theta_j|s) \geq U_j(q, \theta_j|s), \ j \in G \\ &h_\ell(y, q) \geq 0, \ \ell \in K^{in} \\ &h_\ell(y, q) = 0, \ \ell \in K^{eq}. \end{aligned}$$

A critical step in the proof is to show that for all q and almost all θ, the constraint set of this restricted problem (and for all restricted problems as G varies across N) satisfies the linear independence constraint qualification (LICQ), i.e., at every $y \neq q$, the gradients of binding constraints are linearly

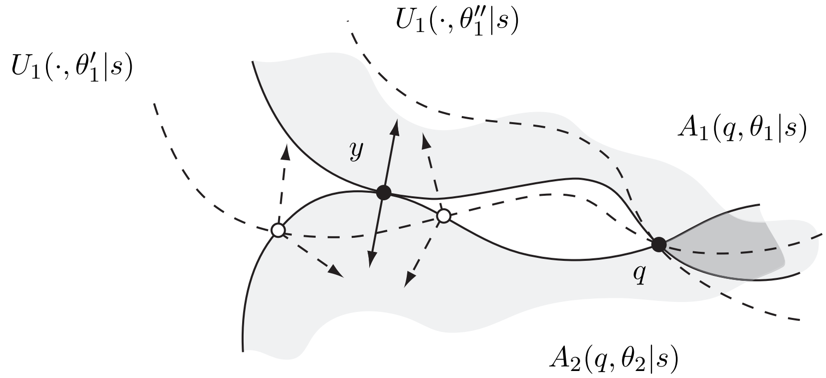

Figure 16: Genericity of LICQ

independent. To gain intuition for this result, consider Figure 16, where $G = \{1, 2\}$ and LICQ is violated at y given shocks θ, as the gradients of individual 1's and 2's gradients are collinear. But perturbing 1's shock to θ_1' or θ_1'', the indifference curves of the individuals intersect transversally, and LICQ is satisfied. An implication is that the optimal proposals of individual i to the group G will vary continuously—and in fact differentiably—with parameters of the problem. For almost all (q, θ), the above problem has a unique solution, say $p_i^G(q, \theta)$, and the unique optimal proposal $p_i(q, \theta)$ is clearly among these, i.e.,

$$U_i(p_i(q, \theta), \theta_i|s) \quad = \quad \max_{G \in \mathcal{D}(q)} U_i(p_i^G(q, \theta), \theta_i|s).$$

In fact, for almost all (q, θ), there is a unique group G such that $p_i(q, \theta) = p_i^G(q, \theta)$, and then the optimal proposal of individual i varies continuously—and in fact differentiably—with parameters of the proposer's problem.

To digress slightly, we note a fundamental difference between the dynamic bargaining model and the one-shot bargaining model of previous sections. In the latter framework, the objective function of the proposer is u_i, and the inequality constraints faced by the proposer have the form $u_j(x) \geq r_j$, where the reservation values r_j are endogenous, but the utility functions are primitives given in the specification of the model. We impose assumptions of concavity and strict quasi-concavity on utilities, and so the acceptance set $A_G(\sigma)$ for a group G is guaranteed to be not only nonempty and compact (properties that hold in the current framework), but convex as well. By strict quasi-concavity of utilities, the proposer has a unique optimal proposal to G, and this proposal will vary continuously in the parameters of the model. In the dynamic bargaining framework, the objective is $U_i(\cdot, \theta_i|s)$ and constraints have the form $U_j(x, \theta_i|s) \geq r_j$, where now dynamic utilities are themselves endogenous. Thus, even if we assume desirable properties of the underlying stage utilities, these need not carry over to the relevant objects. For this reason, we introduce the structure of differentiability and the preference shocks: pathologies can still arise, but they are avoided for all status quos and almost all shocks.

Existence of equilibrium follows, using the observations in parts 1–3, by a fixed point argument. The argument of Duggan and Kalandrakis (2012) takes place in the space of continuation value functions, specifying a nonempty, convex, compact set $\mathcal{V}$ of mappings $v \colon \mathbb{R}^d \to \mathbb{R}^n$ and a continuous function $\psi \colon \mathcal{V} \to \mathcal{V}$; this function takes each v and updates it to $\psi(v)$, and it is defined so that a fixed point corresponds to a pure, no-delay stationary bargaining equilibrium. In keeping with the analysis in previous sections, I sketch out a fixed point argument in the space of strategy profiles. Let Π consist of transition probabilities from (q, θ) pairs to $\Delta(X)$ that respect feasibility, i.e.,

$$\Pi \quad = \quad \{\pi \in \mathcal{R}(X, X \times \Theta, \nu_q \otimes \lambda) \mid \text{for a.e. } (q, \theta),\, \pi(\Phi(q)|q, \theta) = 1\},$$

where we identify transition probabilities that are equal for $(\nu_q \otimes \lambda)$-almost all (q, θ). Of course, Π is the set of possible (equivalence classes of) proposal strategies, and Π^n is the set of possible profiles of proposal strategies. We endow Π with the narrow convergence metric and Π with the product metric. The idea is to take a profile $\pi = (\pi_1, \ldots, \pi_n)$ of proposal strategies as given; assuming all proposals are accepted, these induce continuation values $V_i(x|\pi)$, which determined social acceptance sets $A(q, \theta|\pi)$, which determine essentially unique optimal proposal strategies $(\hat{\pi}_1, \ldots, \hat{\pi}_n)$, as below.

$$\pi = (\pi_1, \ldots, \pi_n) \longrightarrow \{V_i(x|\pi)\} \longrightarrow \{A(q, \theta|\pi)\} \longrightarrow (\hat{\pi}_1, \ldots, \hat{\pi}_n)$$

Call the mapping so-defined $\varphi \colon \Pi^n \to \Pi^n$.

Induced continuation values are characterized by the recursion

$$V_i(x|\pi) \quad = \quad \int_X \int_\Theta \sum_j \rho_j(q) \Big[(1 - \delta_i) u_i(y, \theta_i) + \delta V_i(y|\pi) \Big] \pi_j(dy|q, \theta) f(\theta) g(q|x) d\theta dq,$$

and the induced dynamic utility is

$$U_i(x,\theta_i|\pi) \;=\; (1-\delta_i)u_i(x,\theta_i) + \delta_i V_i(x|\pi).$$

By a contraction mapping argument, $V_i(x|\pi)$ will be jointly continuous in (x,π), and thus so will $U_i(x,\theta_i|x)$.

Now let $\pi^m \to \pi$ and for each m, let $\hat{\pi}^m = \varphi(\pi^m)$. Letting $\pi^* = \varphi(\pi)$, it suffices to argue that $\hat{\pi}^m \to \pi^*$. For all i and almost all (q,θ), the proposal strategy $\pi_i^*(\cdot|q,\theta)$ puts probability one on the solutions to

$$\max_{y\in A(q,\theta|\pi)} U_i(y,\theta_i|\pi).$$

By the argument for part 2, we know for almost all (q,θ) that there is in fact a unique solution, say $p_i^*(q,\theta)$, to this problem, and that the constraint set $A_G(q,\theta|\pi)$ satisfies LICQ for every group G. Therefore, the optimal solution $p_i^*(q,\theta)$ will vary continuously in the parameters of the problem. In particular, this is true when we vary the profile π (and with it the induced continuation values), and we conclude that for all i and almost all (q,θ), $\hat{\pi}_i^m(q,\theta)$ converges to $p_i^*(q,\theta)$. Therefore, $\hat{\pi}_i^m \to \pi_i^*$ narrowly for each individual, as required.

5.4 Ergodic Properties and the Core

In contrast to the one-shot bargaining model, stationary bargaining equilibria in the current framework entail non-trivial dynamics. Given measurable $Y \subseteq \mathbb{R}^d$, let I_Y denote the indicator function of Y. The transition probability on alternatives induced by a pure, no-delay stationary bargaining equilibrium s is defined by

$$\mu(Y|x) \;=\; \int_X \int_\Theta \sum_{i\in N} \rho_i(q) I_Y(p_i(q,\theta)) f(\theta) g(q|x)\, d\theta\, dq,$$

which is the probability, conditional on alternative x this period, that the chosen alternative belongs to the set Y next period. The associated Markov operator T is defined on the space of bounded, measurable functions $\phi\colon X \to \mathbb{R}$ by

$$T\phi(x) \;=\; \int \phi(z)\mu(dz|x).$$

The adjoint T^* operates on measures on X, denoted ξ, and is defined by

$$T^*\xi(Y) \;=\; \int \mu(Y|x)\xi(dx).$$

This describes the distribution of alternatives in the next period, given a distribution ξ over alternatives in the current period. The iterates of T^*, denoted T^{*m}, give the distribution of alternatives m periods hence and are therefore key in describing the long run outcomes of the equilibrium.

It is straightforward to show that μ satisfies the Feller property and Doeblin's condition, so from any initial distribution ξ on X, the sequence of long run average distributions, $\frac{1}{m}\sum_{t=1}^{m} T^{*t}\xi$, $m = 1, 2, \ldots$, converges to an invariant distribution in the total variation metric. While this provides a minimal characterization of long run dynamics, however, the claim is weak in several respects: it concerns the long run average distributions, rather than the distribution of alternatives in each period t; the limiting invariant distribution can depend on the initial distribution; and the rate of convergence is only known to be arithmetic. If there is a positive probability that the status quo determined by x remains close to x, however, then we can deduce geometric convergence of the distribution of alternatives in each period. Moreover, if the supports of status quo distributions satisfy an overlapping condition, then each equilibrium induces a unique ergodic distribution over alternatives.

Proposition 5.3. *Let s be a pure, no-delay stationary bargaining equilibrium and T^* the adjoint of the associated Markov operator.*

1. *The transition probability μ satisfies the Feller property and Doeblin's condition, and it admits at least one invariant distribution.*

2. *Given any initial distribution ξ, the sequence of long run average distributions, $\frac{1}{m}\sum_{t=1}^{m} T^{*t}\xi$, converges arithmetically to an invariant distribution ξ^* of T^* in total variation: there is a constant $M > 0$ such that for all m, we have*

$$\rho^v\left(\frac{1}{m}\sum_{t=1}^{m} T^{*t}\xi, \xi^*\right) \;\leq\; \frac{M}{m}.$$

3. *If for every policy $x \in X$, we have $g(x|x) > 0$, if $g(q|x)$ is jointly continuous in (q, x), and if there is an individual i such that $\rho_i(q) > 0$*

for all q, then μ is aperiodic and given any initial distribution ξ, the sequence of per period distributions of alternatives converges geometrically to an invariant distribution ξ^* in total variation: there are constants M and β, with $M > 0$ and $\beta \in (0,1)$, such that for all m, we have $\rho^v(T^{*m}\xi, \xi^*) \leq M\beta^m$.

4. *If for all $x, y \in X$ there exists $q \in X$ such that $g(q|x)g(q|y) > 0$, then μ admits a unique invariant distribution, say ξ^*. Given any initial distribution ξ, the sequence of iterates, $T^{*m}\xi$, converges geometrically to ξ^* in total variation: there are constants M and β, with $M > 0$ and $\beta \in (0,1)$, such that for all m, we have $\rho^v(T^{*m}\xi, \xi^*) \leq M\beta^m$.*

The proof of part 3 relies on the regularity properties of stationary bargaining equilibria established in Proposition 5.2. Given any ergodic set Y and an individual i whose recognition probability is always positive, there is a preference shock θ_i such that i's dynamic utility is maximized over the closure of Y at a unique alternative, say x. If the alternative chosen in a given period is near x, then there is a positive probability that next period's status quo is near x, i's preference shock is near θ_i, and the selected proposer is again i, who implements an alternative near x. This precludes the possibility of cyclic subsets and implies aperiodicity.

Next, we take the analysis of dynamics to the model in which the core is nonempty and the core individual is a representative voter. To simplify matters, assume the voting rule $\mathcal{D}$ is fixed independently of the status quo, and assume it is a *weighted majority rule,* i.e., there exist weights $w_i \geq 0$ with $\sum_i w_i = 1$ such that $G \in \mathcal{D}$ if and only if $\sum_{i \in G} w_i > \frac{1}{2}$, and assume that $\mathcal{D}$ is strong, so there is no group G such that $\sum_{i \in G} w_i = \frac{1}{2}$. Assume utilities are quadratic, i.e., $u_i(x, \theta_i) = -\|x - \hat{x}^i\|^2 + x \cdot \theta_i$, and assume the core $C(\mathcal{D})$ for the underlying Euclidean preferences (i.e., $\theta = 0$) is nonempty; this is automatically true in the one-dimensional model. Let individual k be such that $\hat{x}^k$ is the unique element of the core, so k is the "core voter." Assume that recognition probabilities ρ are fixed independently of the status quo, and that $\rho_k > 0$. Finally, assume a common discount factor δ. In this context, we say the *canonical model* is such that the distribution of shocks θ is degenerate with probability one on $\theta = 0$, and the distribution $G(\cdot|x)$ of the status quo is degenerate with probability one on x. This model assumes that there is no noise on preferences or the status quo, and it is outside the framework of this section.

We say a model is ϵ-*canonical* if it satisfied the above assumptions, along with our maintained assumptions of boundedness of the densities f and g and differentiability of g, and if furthermore, we have

- $\mathrm{supp}(f) \subseteq B_\epsilon(0)$,
- for all $x \in X$, $\mathrm{supp}(g(\cdot|x)) \subseteq B_\epsilon(x)$,

where $\mathrm{supp}(\cdot)$ indicates the subset of the domain on which a function takes strictly positive values. That is, the distribution of preference shocks is concentrated around zero, and noise on the status quo is restricted to the ϵ-ball around the alternative chosen in any period. A sequence $\{\gamma^m\}$ is *approximately canonical* if there is a sequence $\{\epsilon^m\}$ such that γ^m is ϵ^m-canonical for all m and $\epsilon^m \to 0$. Such a sequence approximates the canonical model within the dynamic bargaining framework of this section.

Given a status quo q and proposal x, individual i's vote depends on a comparison of dynamic utilities: $U_i(x, \theta_i|s)$ vs. $U_i(q, \theta_i|s)$. Note that x determines a distribution over alternatives in the current period (namely, the probability measure μ_x that is degenerate on x) that is weighted by $1-\delta$, and it determines a distribution $T^*\mu_x = \mu(\cdot|x)$ over alternatives in the next period weighted by $(1-\delta)\delta$, and a distribution $T^{*2}\mu_x$ two periods hence weighted by $(1-\delta)\delta^2$, and so on. If we sum these distributions, we obtain a continuation distribution

$$\gamma_x \quad = \quad (1-\delta)\left[\mu_x + \delta T^*\mu_x + \delta^2 T^{*2}\mu_x + \cdots\right]$$

that represents the distribution of current and future alternatives if x is accepted. A continuation distribution γ_q can similarly be constructed for the status quo. It is known that when utilities are quadratic and the voting rule is a strong, weighted majority rule, the core voter is "representative," in the sense that the voter prefers one lottery to another if and only if all members of a decisive group $G \in \mathcal{D}$ have the same preference. In the absence of preference shocks in the model, it follows that the fate of a proposal would be determined simply by the preferences of individual k.

In the dynamic bargaining framework, however, we do impose preference shocks, and future alternatives will depend on the realization of those shocks. That is, future realizations of θ and x will be correlated, so individuals cannot evaluate the flow of future alternatives by simply using quadratic utility with their current ideal point. Nevertheless, when the preference shocks are concentrated near zero, so that uncertainty about future preferences is small, the core voter k will be "nearly representative," in the sense that if k's expected utility from one lottery exceeds that of another by a given, positive increment, then the members of a decisive group will have the same preference; and this increment can be made arbitrarily small as preference shocks become negligible.

We are interested in the connection between the equilibrium dynamics of the bargaining model and the core of the voting rule. In the canonical model, it is reasonable to expect that chosen alternatives will gravitate to the core, eventually reach it, and then stay there. But in the model with noise, there can be no such expectation: noise on the status quo implies that chosen alternatives can drift away from the core, and even if the core individual k is selected as proposer, the preference shock θ_k will perturb the ideal point of the individual, so the core may never be reached in the first place. Nevertheless, the following result establishes that as noise in the model goes to zero, the long run distribution of chosen alternatives does indeed pile up at the core. Note that this result is fundamentally different in nature than the dynamic median voter theorems (Propositions 3.7 and 3.19) in the one-shot bargaining model: there, we required individuals to be arbitrarily patient, and we concluded that bargaining ends in the first period with an alternative close to the core; here, we do not impose any restrictions on patience, and bargaining continues over time, with chosen alternatives diverging from the core for long intervals of time with positive probability, but spending most time very close to the core.

Proposition 5.4. *Let $\{\gamma^m\}$ be approximately canonical. For each m, let s^m be a pure, no-delay stationary bargaining equilibrium in γ^m, and let ξ^m be an invariant distribution corresponding to s^m. Then $\{\xi^m\}$ converges weakly to the unit mass on $\hat{x}^k$.*

The proof, which is omitted, is complicated by the inability to work in the limit, because we do not have a general existence theorem for the canonical model or an applicable continuity result. Rather, we work along the sequence by supposing that there is an open ball B around $\hat{x}^k$ and that a subsequence of invariant distributions places probability outside B with a positive lower bound, $\beta > 0$. Thus, $\xi^m(X \setminus B) \geq \beta$ for infinitely many m. A contradiction is deduced from these suppositions to infer the weak convergence claimed in the theorem. A sketch of the argument is roughly as follows. As we have

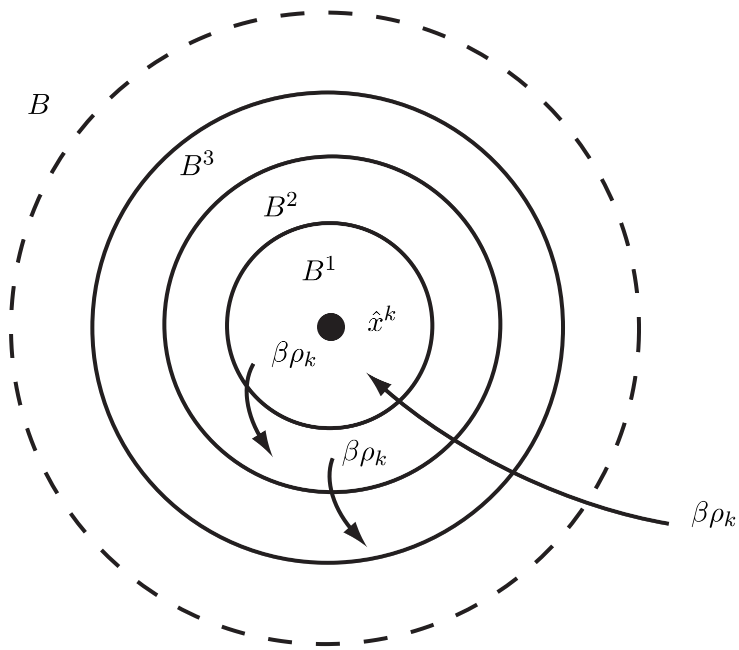

Figure 17: Probability flow under invariant distribution

noted, the core individual k becomes nearly representative as the model approximates the canonical model, and as a consequence, the core individual is able to implement policies close to $\hat{x}^k$ whenever selected to propose. In particular, we may specify an arbitrarily small open ball B^1 around $\hat{x}^k$, contained in B, and we obtain that for high enough m, whenever the status quo in the current period lies outside B and k is selected to propose, next period's policy will belong to B^1, as in Figure 17. Thus, evaluated using the invariant distribution, there is a flow of probability at least equal to $\beta\rho_k$ into B^1. Because the core individual is nearly representative, we conclude that for high enough m, given any status quo in B^1 in one period, the next period's chosen alternative cannot be too far away; we can bound it by an open ball B^2 containing and arbitrarily close to B^1. We can repeat this construction, generating a number $\ell > 1/\rho_k\beta$ of such nested balls, denoted $B^1, B^2, \ldots, B^\ell$, all contained in B. Obviously, a formal proof must account for details missing here, including the fact that once an alternative is chosen in the current period, next period's status quo is realized with small noise.

We have established that, in intuitive terms, there is a net flow of probability at least equal to $\beta\rho_k$ into B^1 from outside B, giving us a lower bound on the probability of B^1 under the invariant distribution: for high enough m, we have $\xi^m(B^1) \geq \beta\rho_k$. Therefore, because ξ^m is invariant, there must be a net flow of probability of at least $\beta\rho_k$ out of B^1 and into B^2, yielding $\xi^m(B^2 \setminus B^1) \geq \beta\rho_k$. Continuing this logic, there is a net flow of probability at least equal to $\beta\rho_k$ from B^2 to B^3, and so on, and we conclude that each of the sets $B^1, B^2 \setminus B^1, B^3 \setminus B^2, \ldots, B^\ell \setminus B^{\ell-1}$ has probability at least $\beta\rho_k > 0$ under the invariant distribution. But recall that $\ell > 1/\rho_k\beta$, so we have

$$\xi^m(X) \;\; \geq \;\; \xi^m(B^1) + \xi^m(B^2 \setminus B^1) + \cdots + \xi^m(B^\ell \setminus B^{\ell-1}) \;\; > \;\; 1,$$

a contradiction.

References

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

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

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

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

[5] D. Cardona and C. Ponsati (2011) "Uniqueness of Stationary Equilibria in Bargaining One-dimensional Policies under (Super) Majority Rules," *Games and Economic Behavior,* 73: 65–75.

[6] S.-J. Cho and J. Duggan (2003) "Uniqueness of Stationary Equilibria in a One-dimensional Model of Bargaining," *Journal of Economic Theory,* 113: 118–130.

[7] S.-J. Cho and J. Duggan (2009) "Bargaining Foundations of the Median Voter Theorem," *Journal of Economic Theory,* 144: 851–868.

[8] S.-J. Cho and J. Duggan (2013) "A Folk Theorem for the One-dimensional Spatial Bargaining Model," unpublished paper.

[9] J. Duggan and T. Kalandrakis (2007) "Dynamic Legislative Policy Making," W. Allen Wallis Institute of Political Economy working paper No. 45, University of Rochester.

[10] J. Duggan and T. Kalandrakis (2012) "Dynamic Legislative Policy Making," *Journal of Economic Theory,* 147: 1653–1688.

[11] H. Eraslan (2002) "Uniqueness of Stationary Equilibrium Payoffs in the Baron-Ferejohn Model," *Journal of Economic Theory,* 103: 11–30.

[12] H. Eraslan and A. McLennan (2013) "Uniqueness of Stationary Equilibrium Payoffs in Coalitional Bargaining," *Journal of Economic Theory,* forthcoming.

[13] A. Federgruen (1978) "On n-Person Stochastic Games with Denumerable State Space," *Advances in Applied Probability,* 10: 452–471.

[14] A. Fink (1964) "Equilibrium in a Stochastic n-Person Game," *Journal of Science of the Hiroshima University,* 28: 89–93.

[15] T. Kalandrakis (2004) "A Three-player Dynamic Majoritarian Bargaining Game," *Journal of Economic Theory,* 116: 294–322.

[16] T. Kalandrakis (2010) "Minimum Winning Coalitions and Endogenous Status Quo," *International Journal of Game Theory,* 39: 617–643.

[17] Y. Levy (2013) "Discounted Stochastic Games with No Stationary Nash Equilibrium: Two Examples," *Econometrica,* 81: 1973–2007.

[18] A. Nowak and T. Raghavan (1992) "Existence of Stationary Correlated Equilibria with Symmetric Information for Discounted Stochastic Games," *Mathematics of Operations Research,* 17: 519–526.

[19] P. Rogers (1969) "Non-zero Sum Stochastic Games," Operations Research Center Report 69-8, University of California, Berkeley.

[20] L. Shapley (1953) "Stochastic Games," *Proceedings of the National Academy of Sciences,* 39: 1095–1100.

[21] M. Sobel (1971) "Non-cooperative Stochastic Games," *Annals of Mathematical Statistics,* 42: 1930–1935.

[22] M. Takahashi (1964) "Equilibrium Points of Stochastic Non-cooperative n-Person Games," *Journal of Science of the Hiroshima University,* 28: 95–99.