Purification of Bayes Nash Equilibrium with Correlated Types and Interdependent Payoffs

Paulo Barelli and John Duggan

October 26, 2012

Outline

Motivation

- We develop a technique to obtain existence and purification results in a variety of models exhibiting product structure.
- Goal is to impose structure that is plausible in applications, but to conduct analysis at a general level.
- The main restrictions are finite actions and one component being atomless. Payoff is general correlation/externalities across actors.
- We do not assume monotonicity or strategic complementarity or restrict correlation/externalities to be finite-dimensional.
- The general idea is to exploit Lyapunov's theorem to purify component by component, and a theorem of Artstein to combine these in a mble way.

Overview

- **Bayesian games:**
  - Assume types have product structure (general component, private component).
  - Private component is atomless but may be payoff-relevant.
  - General types may be arbitrarily correlated, subject to absolute continuity of information condition. Type sets may be infinite.
  - Payoffs may be interdependent, as general component can contain information that is payoff relevant for other players.
  - For every BNE, there is an equivalent pure-strategy BNE.
  - Result extends to undominated strategies.

Overview (cont.)

- **Large games:**
  - Assume player set has product structure (zip code, address).
  - Allow infinite-dimensional externalities across "zip codes."
  - There is a mixed strategy Nash equilibrium, and for every such equilibrium, there is an equivalent equilibrium in pure strategies.
- **Stochastic games:**
  - Assume states have product structure (general component, noise component).
  - Continuum of actions and states. No correlation used.
  - There is a stationary MPE, and for every such equilibrium, there is an "extremal" stationary MPE.
  - In sequential move games, every equilibrium can be purified.

Literature review: Bayesian games

- Existence of mixed strategy BNE is known from Milgrom and Weber (1985) and Balder (2002). Balder (2002) gives pure strategy BNE existence result. The set up is general in other ways but assumes finite-dimensional payoff interdependence.
- Early purification results are given by Radner and Rosenthal (1982) and Milgrom and Weber (1985). These results assume finite actions, finite state t_0, and conditionally independent, atomless, private value types t_i.
- Balder (2008) gives more general purification theorem but assumes finite actions and finite-dimensional interdependence.
- Existence of pure strategy BNE via monotonicity or strategic complementarity are due to Athey (2001), McAdams (2003), Reny (2011), and De Castro (2012).

Literature review: Large and stochastic games

- Existence of mixed strategy equilibria is established by Mas-Colell (1984) and Balder (2002). Schmeidler (1973) and Balder (2002) (and others) establish existence of pure strategy equilibria.
- Balder (2008) gives purification result assuming finite actions and finite-dimensional externalities.
- Existence of stationary MPE in noisy stochastic games is proved in Duggan (2012) using similar methods.
- Earlier work uses correlation (Nowak and Raghavan (1992), Duffie et al. (1994)), or allows history dependence (Mertens and Parthasarathy (1987,2003)), or assumes stronger conditions on states and actions.

Example: Multi-unit auctions

- Finite set N of bidders, each bids $b_i = (b_i^1, \ldots, b_i^d) \in B$ (a finite grid on $[0,1]^d$) over finite set of objects O.
- Outcomes are determined by $O_i = g_i(b_1, \ldots, b_n)$ and prices $p_i = p_i(b_1, \ldots, b_n)$.
- Each i receives signal t_i about continuous state variable s, and preference parameter u_i is drawn independently conditional on t_i from continuous distribution.
- Bidder i's payoff is mble function $\pi_i(O_i, p_i, s, u_i)$.
- An undominated BNE exists, and for each such BNE, there is an equivalent undominated, pure strategy BNE.

Example: Environmental externalities

- A continuum of consumers indexed by U at a continuum of locales T make consumption decisions that generate environmental externalities.
- Consumer (t, u) chooses $\sigma(t, u) \in A(t, u)$ from a finite-dimensional simplex of feasible actions.
- Payoffs are $\pi_i(a, t, u; \alpha)$, where α is the externality function

$$\alpha(t) = \int_u \sigma(t, u)\nu(du|t).$$

- Assume $A(t, u)$ mble, $\pi(a, t, u; \alpha)$ mble and continuous in (a, α), multilinear in a, and each $\nu(\cdot|t)$ nonatomic.
- A Nash equilibrium exists, and for every Nash equilibrium, there is an equivalent pure strategy Nash equilibrium.

Example: Industry dynamics

- A finite set N of firms make entry/exit decisions and choose production plans over time.
- Each period begins with a vector $q = (z, k)$, recording active firms and capital stock levels, and with a vector r of shocks to profit functions.
- Firms choose to enter or exit, e_i, and a production plan $p_i \in \phi_i(z, k, r)$, and profits are $\pi_i((z, k, r), (e, p))$.

Example: Industry dynamics (cont.)

- Active firms in the next period, z', are determined by e, and new capital stocks k' are drawn from $g(\cdot|(z,k),(e,p))$. New shocks r' are drawn iid from a continuous distribution.
- Assume $\phi_i(z,k,r)$ compact, $\pi_i((z,k,r),(e,p))$ mble and continuous in p, $g(r'|(z,k),(e,p))$ mble and continuous in p.
- A stationary MPE exists, and for every equilibrium, there is an equivalent "extremal" equilibrium.
- In the sequential move version of this game, extremal equilibria are in pure strategies.

Outline

Elements of the Model

- players $i = 1, \ldots, n$
- types (t_i, u_i), where
  - $t_i \in T_i$ is i's general type
  - $u_i \in U_i$ is i's private type
- feasible actions $A_i(t_i, u_i) \subseteq \Delta_d$
  - a_i is mixture of d pure actions
- payoffs $\pi_i(a, t, u_i)$, where $a = (a_i)$ and $t = (t_i)$
- prior beliefs μ on $T \times U = (\bigtimes T_i) \times (\bigtimes U_i)$

Assumptions

- T_i and U_i are complete, separable metric spaces.
- $A_i(t_i, u_i)$ is a nonempty face of Δ_d and is lower measurable.
- $\pi_i(a, t, u_i)$ is multilinear in a, bounded, and measurable.
- μ can be decomposed as

  $$\mu \quad = \quad \left(\bigotimes \nu_i(\cdot|t_i)\right) \otimes \kappa,$$

  where κ is the prior on t.
- κ is absolutely continuous with respect to the product of the marginals, $\bigotimes \kappa_i$.

Strategies

- A **strategy** for i is a measurable mapping $\sigma_i \colon T_i \times U_i \to \Delta_d$ such that $\sigma_i(t_i, u_i) \in A_i(t_i, u_i)$ for almost all (t_i, u_i).
- The expected action corresponding to σ_i is $\alpha_i \colon T_i \to \Delta_d$ defined by

$$\alpha_i(t_i) \quad = \quad \int_{u_i} \sigma_i(t_i, u_i)\nu_i(du_i|t_i).$$

- Say σ_i and σ_i' are **equivalent** if $\alpha_i = \alpha_i'$.
- Say σ_i is a **pure strategy** if $\sigma_i(t_i, u_i) \in \text{ext}\Delta_d$ for almost all (t_i, u_i).

Payoffs

- Interim payoffs are

$$\Pi_i(a_i, t_i, u_i; \alpha_{-i}) \;=\; \int_{t_{-i}} \pi_i(a_i, \alpha_{-i}(t_{-i}), t, u_i)\kappa(dt_{-i}|t_i),$$

  where $\kappa(\cdot|t_i)$ is a conditional probability on t_{-i}.

- Action a_i is **dominated** at (t_i, u_i) if there exists $a_i' \in A_i(t_i, u_i)$ such that (i) $\Pi_i(a_i', t_i, u_i; \alpha_{-i}) \geq \Pi_i(a_i, t_i, u_i; \alpha_{-i})$ for all α_{-i}, and (ii) strict inequality for some α_{-i}.

- Strategy σ_i is **undominated** if it almost always specifies an undominated action.

Outline

Bayes Nash equilibrium

- Strategy profile σ is a **Bayes Nash equilibrium** if for all i and all σ_i',

$$\int_t \int_{u_i} \pi_i(\sigma_i(t_i, u_i), \alpha_{-i}(t_{-i}), t, u_i)\nu_i(du_i|t_i)\kappa(dt)$$
$$\geq \int_t \int_{u_i} \pi_i(\sigma_i'(t_i, u_i), \alpha_{-i}(t_{-i}), t, u_i)\nu_i(du_i|t_i)\kappa(dt).$$

Existence and purification

Theorem

A BNE exists, and if $\{\nu_i(\cdot|t_i) : t_i \in T_i\}$ are nonatomic, then for every BNE, there is an equivalent BNE.

- The result improves literature by allowing arbitrary correlation of general types subject to absolute continuity of information.
- We allow a complete, separable metric space T_0 of environmental states by adding a dummy player 0.
- The approach can accommodate finite-dimensional correlation of private types u_i, as in the literature.

Undominated strategies

Theorem

Assume $\{\nu_i(\cdot|t_i) : t_i \in T_i\}$ are nonatomic. A pure-strategy, undominated BNE exists, and for every undominated BNE, there is an equivalent pure-strategy, undominated BNE.

- Let $E_i(t_i, u_i)$ denote the set of undominated pure actions at (t_i, u_i). This is nonempty and lower measurable.
- Define an associated game with $A_i'(t_i, u_i) = \text{conv} E_i(t_i, u_i)$ and apply the first theorem.
- A pure-strategy BNE in the associated game is a pure-strategy, undominated BNE in the original game.

Infinite games

- Let X be a compact metric space, and let $Y_i(t_i, u_i) \subseteq X$ be a nonempty, compact set of feasible actions.
- Assume $A_i(t_i, u_i) = \mathcal{P}(Y_i(t_i, u_i))$ for all (t_i, u_i).
- Assume $\pi_i(a, t, u_i)$ is continuous in a.

Theorem

In the infinite-action product Bayesian game, a BNE exists. If $\{\nu_i(\cdot|t_i) : t_i \in T_i\}$ are nonatomic, then for every $\epsilon > 0$, there is a pure-strategy ϵ-BNE.

- The proof follows by applying the first theorem to a sequence of finite approximations of the infinite-action game.

Outline

Topologizing average actions

- Place average actions $\alpha_i \colon T_i \to \Delta_d$ in $L^d_\infty(T_i)$ with the weak* topology.
- So $\{\alpha_i^\nu\} \to \alpha_i$ if for all κ_i-integrable $g \colon T_i \to \mathbb{R}^d$,

$$\int_{t_i} (\alpha_i^\nu(t_i) - \alpha_i(t_i)) \cdot g(t_i)\kappa_i(dt_i) \quad \to \quad 0.$$

- This makes the space of average actions compact (and convex).
- By absolutely continuous information, interim payoffs $\Pi_i(a_i, t_i, u_i; \alpha_{-i})$ are jointly continuous in $(a_i, (\alpha_j)_{j\neq i})$.

Best response actions

- Define the best response action correspondence by

$$M_i(t_i, u_i; \alpha) \quad = \quad \arg \max_{a_i \in A_i(t_i, u_i)} \Pi_i(a_i, t_i, u_i; \alpha_{-i}).$$

- Properties of this correspondence: for almost all (t_i, u_i),
  (i) $M_i(t_i, u_i; \alpha)$ is a face of Δ_d
  (ii) $M_i(\cdot; \alpha)$ is lower measurable.
  (iii) $M_i(t_i, u_i; \cdot)$ has closed graph.

Average best response actions

- Define the average best response action correspondence by

$$\begin{aligned} M_i^*(t_i;\alpha) &= \int_{u_i} M_i(t_i,u_i;\alpha)\nu_i(du_i|t_i) \\ &= \left\{ \int_{u_i} \gamma_{t_i}(u_i)\nu_i(du_i|t_i) : \begin{array}{c} \gamma_{t_i} \text{ is a mble selection} \\ \text{from } M_i(t_i,\cdot|\alpha) \end{array} \right\}. \end{aligned}$$

- Properties of this correspondence: for almost all t_i,
  (i) $M_i^*(t_i;\alpha)$ is convex and compact
  (ii) $M_i^*(t_i;\alpha)$ is nonempty

Average best response actions (cont.)

- Moreover, for every measurable selection from $M_i^*(\cdot;\alpha)$, there is a strategy σ_i such that

$$\sigma_i(t_i, u_i) \in M_i(t_i, u_i; \alpha)$$

  for almost all (t_i, u_i).

- Indeed, if β_i is a mble selection from $M_i^*(\cdot;\alpha)$, then $\beta_i(t_i) \in M_i^*(t_i)$ means there is a mble selection γ_{t_i} from $M_i(t_i, \cdot; \alpha)$ with

$$\beta_i(t_i) \quad = \quad \int_{u_i} \gamma_{t_i}(u_i)\nu_i(du_i|t_i).$$

- These selections are made independently across t_i.

Average best response actions (cont.)

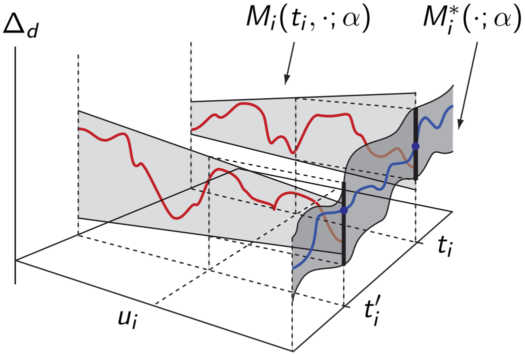

- The problem is to sew the mappings γ_{t_i} together across t_i in a measurable way.
- We use a theorem of Artstein (1989), or the "mble mble choice" theorem of Mertens (2003).
- In particular, every mble selection from $M_i^*(\cdot; \alpha)$ is an average action.

Fixed point argument

- We consider fixed points of the correspondence $\mathscr{S}$ from average action profiles to average action profiles defined by:

$$\alpha = (\alpha_1, \ldots, \alpha_n) \quad \rightarrow \quad \bigtimes_{i=1}^{n} \begin{pmatrix} \text{selections} \\ \text{from } M_i^*(\cdot; \alpha) \end{pmatrix}$$

- This correspondence has nonempty, convex, closed values and is sequentially upper hemi-continuous.
- By a version of Glicksberg's theorem, it has a fixed point $\alpha^* \in \mathscr{S}(\alpha^*)$.
- We back out a BNE σ^* using Artstein (1989).

Purification

- Consider a BNE σ^* with average actions α^*. Under non-atomicity,

$$
\begin{aligned}
M_i(t_i; \alpha^*) &= \int_{u_i} M_i(t_i, u_i; \alpha)\nu_i(du_i|t_i) \\
&= \int_{u_i} \text{ext} M_i(t_i, u_i; \alpha^*)\nu_i(du_i|t_i)
\end{aligned}
$$

- Note that α_i^* is a mble selection from $M_i^*(\cdot; \alpha^*)$.
- Then $\alpha_i^*(t_i) \in M_i^*(t_i; \alpha^*)$ means there is a mble selection γ_{t_i} from $\text{ext} M_i(t_i, \cdot; \alpha^*)$ with $\alpha_i^*(t_i) = \int_{u_i} \gamma_{t_i}(u_i)\nu_i(du_i|t_i)$.

Purification (cont.)

- We use Artstein (1989) to sew the mappings γ_{t_i} together.
- This gives us a strategy σ_i' for each player such that for almost all t_i,

$$\alpha_i'(t_i) \;=\; \int_{u_i} \gamma_{t_i}(u_i)\nu_i(du_i|t_i) \;=\; \alpha_i^*(t_i)$$

  and for almost all (t_i, u_i),

$$\sigma_i'(t_i, u_i) \in \text{ext} M_i(t_i, u_i; \alpha^*) \;=\; \text{ext} M_i(t_i, u_i; \alpha').$$

- Therefore, σ_i' is a pure-strategy BNE and is equivalent to σ^*.

Outline

Generalizing the approach

- We can use these tools by discarding payoff functions and the polyhedral structure of action sets.
- Instead, we require conditions on $A_i(t_i, u_i)$ and $M_i(t_i, u_i; \alpha)$: we need...
  - $A_i(t_i, u_i)$ to be a nonempty and compact subset of $\mathbb{R}^d$, lower measurable, and μ-integrably bounded.
  - $M_i(t_i, u_i; \alpha)$ to be a nonempty compact subset of $A_i(t_i, u_i)$, lower measurable in (t_i, u_i), and to have closed graph in α.

Product large games: Elements

- Players $T \times U$ (set $n = 1$)
  - $t \in T$ is "zip code" of player (t, u).
  - $u \in U$ is her "street address".
- $A(t, u)$ is the feasible actions for player (t, u) satisfying the above assumptions.
- $P(a, t, u; \alpha)$ is the set of actions strictly preferred to a.
- μ is the distribution of players.

Product large games: Strategies

- A **strategy profile** is a mble mapping $\sigma\colon\, T \times U \to \mathbb{R}^d$ such that for almost all (t, u), $\sigma(t, u) \in A(t, u)$.
- We allow infinite-dimensional externalities:

$$\alpha(t) \;\; = \;\; \int_u \sigma(t, u)\nu(du|t).$$

- The maximal feasible actions of player (t, u), given α, are

$$M_i(t, u; \alpha) \;\; = \;\; \{a \in A(t, u) : P(a, t, u; \alpha) = \emptyset\}.$$

- Assume $M(t, u; \alpha)$ is nonempty and closed, $M(\cdot; \alpha)$ is lower measurable, and $M(t, u; \cdot)$ has closed graph.

Product large games: Result

Theorem

Assume $\{\nu(\cdot|t) : t \in T\}$ are nonatomic. A Nash equilibrium exists, and for every Nash equilibrium, there is an equivalent extremal Nash equilibrium.

- We do not need $M(t, u; \alpha)$ to be convex – we obtain convexity of $M^*(t; \alpha) = \int_u M(t, u; \alpha)\nu(du|t)$ via Lyapunov's theorem.
- Purification relies on a lemma on lower measurability of $\overline{\text{ext}} M(t, u)$.

Noisy stochastic games: Elements

- players $i = 1, \ldots, n$
- states $s = (t, u)$
  - t is a general component
  - u is a noise component
- feasible actions $A_i(s)$
- stage payoffs $\pi_i(s, a)$
- discount factor δ_i
- law of motion $t' \sim \mu(\cdot|s, a)$ and $u' \sim \nu(\cdot|t')$.

Noisy stochastic games: Assumptions

- $A_i(s) \subseteq X_i$ is a nonempty, closed subset of a fixed, compact metric space and is lower measurable.
- $\pi_i(s, a)$ is continuous in a and mble in s.
- $\mu(\cdot|s, a)$ is norm-continuous in a and mble in s.
- distribution of noise $\nu(\cdot|t)$ is absolutely continuous with respect to fixed, non-atomic λ.

Strategies and equilibrium

- A **stationary Markov strategy** for i is a measurable mapping $\sigma_i \colon S \to \mathcal{P}(X_i)$ such that for all s, $\sigma_i(s)(A_i(s)) = 1$.
- Let $v_i(s;\sigma)$ be player i's expected discounted payoff beginning in state s.
- Then $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a **stationary Markov perfect equilibrium** if $\sigma_i(s)$ puts probability one on

$$\arg\max_{a_i \in A_i(s)} \int_{a_{-i}} \left[u_i(s,a) + \delta_i \int_S v_i(s';\sigma)\mu(ds'|s,a) \right] \sigma_{-i}(da_{-i}|s)$$

for all i and all s.

Noisy stochastic games: Result

Theorem

A stationary MPE exists, and for every stationary MPE, there is an equivalent extremal stationary MPE.

- A stochastic game is **sequential move** if for all states s, there is a player i such that for all $j \neq i$, $A_j(s)$ is singleton.

Corollary

In a sequential move stochastic game, every stationary MPE is equivalent to a pure-stratgy stationary MPE.

Noisy stochastic games: Idea of proof

- Given σ, let $v_i(t;\sigma)$ be the discounted expected payoff conditional on general component t at the beginning of a period.
- Then v_i satisfies:

$$
\begin{aligned}
& v_i(t;\sigma) \\
& \quad = \int_u \left[\int_a \left[\pi_i(s,a) + \delta_i \int_{t'} v_i(t';\sigma)\mu(du|s,a) \right] \sigma(da|s) \right] \nu(du|t),
\end{aligned}
$$

where $s = (t,u)$.

Noisy stochastic games: Idea of proof

- Given functions $v = (v_i)$, define the induced game $\Gamma_v(s)$ with actions $A_i(s)$ and payoffs

$$\pi_i(s, a) + \delta_i \int_{t'} v_i(t')\mu(dt'|s, a).$$

- Let $M(t, u; v)$ denote the mixed strategy equilibrium payoffs in the induced game $\Gamma_v(s)$.
- Then $M(t, u; v)$ need not be convex, but it is nonempty and compact, $M(\cdot; v)$ is lower measurable, and $M(t, u; \cdot)$ has closed graph.
- So the above approach applies.

Outline

Concluding remarks

- In Bayesian games with a private type component (e.g., preference shocks), purification of BNE is possible at a general level.
- No need to assume finite-dimensional correlation or for monotonicity or strategic complementarity assumptions.
- Approach extends to large games with product structure on the set of players with infinite-dimensional externalities...
- ...and to stochastic games in which states have a noise component, without the need for correlation.