A Bayesian Nonatomic Game and Its Applicability to Finite-player Situations

Jian Yang
Department of Management Science and Information Systems
Business School, Rutgers University
Newark, NJ 07102

September 2021

Research Outline

- We formulate a *nonatomic* game (NG) with *Bayesian* features
- This NG allows players to have *correlated* signals and has very *general* finite-player counterparts
- After establishing *equilibrium* existence, we show how any of this NG's equilibria could be used by its *randomly generated* finite counterparts to achieve *approximate equilibrium*
- Mixed NG equilibria could yield approximate *pure* equilibria for large finite games randomly generated in NG's neighborhood
- When *anonymous*—joint external player-action distribution influences a player only through *marginal* action distribution, NG itself can be shown to have *pure* equilibria

Literature Overview

- Normal-form NGs were studied by Schmeidler (1973), Mas-Colell (1984), Balder (1995, 2002), Khan, Rath, and Sun (1997), Loeb and Sun (2006), Podczeck (2009), and Khan et al. (2013), etc.
- Finite n-player Bayesian games were treated by Harsanyi (1967-8), Radner and Rosenthal (1982), Milgrom and Weber (1985), Balder (1988), He and Sun (2019), and so on
- Kalai (2006) showed *ex-post stability* of *large* Bayesian games; extensions and generalizations were made by Carmona (2008), Carrtwright and Wooders (2009), Gradwohl and Reingold (2010), Carmona and Podczeck (2012), and Deb and Kalai (2015), etc.
- Our NG–finite-game connections convey a different but still *robustness*-themed message: players would *not* be much bothered by their opponents' realized characteristics in a large game

Infinite Players and Incomplete Information

- Khan and Rustichini (1991), Balder (1991), and Balder and Rustichini (1994) all studied games involving *infinite* numbers of players possessing *incomplete* information
- Kim and Yannelis (1997) used *sub-sigma-fields* to model players' differentiated knowledge about true *state of world* and demonstrated existence of pure equilibria for case involving *concave* payoff functions over action spaces with *linear* structures
- Carmona and Podczeck (2020) let players have *independent* types and demonstrated any NG equilibrium would be *limiting point* of equilibria for a converging sequence of *finite* games
- We study a somewhat *complementary* situation where a substantial portion of information could be held by any one single player while her and others' received signals could be very much *correlated*

One Observation

- Let $|I|$ be player $\#$, $|\Sigma|$ be signal $\#$, and $|A|$ be action $\#$
- The greater an extent to which $|I| \times |\Sigma| \gg |A|$, the easier it would be for *deterministic* (i, σ)-dependent actions a to weave out desired externalities and hence for *pure* equilibria to emerge
- For normal-form NG where $|\Sigma| = 1$, most works, e.g., Schmeidler (1973) and Khan et al. (2013), effectively required $|I| \gg |A|$
- For finite Bayesian games where $|I| = n$, Milgrom and Weber (1982) and He and Sun (2019) demanded $|\Sigma| \gg |A|$
- However, relative predominance concerning cardinalities seems relaxable when only *approximate purification* is sought; see, e.g., Carmona (2008) and our *mixed*-to-*pure* result

Essential Elements

- Our Bayesian NG has a few spaces:
  - ω) some Ω is for *states of world*
  - i) $I \equiv [0,1]$ is set of players or player *characteristics*
  - σ) a finite $\Sigma \equiv \left\{\bar{\sigma}_1, \bar{\sigma}_2, ..., \bar{\sigma}_{|\Sigma|}\right\}$ is space of *signals*
  - a) a finite $A \equiv \left\{\bar{a}_1, \bar{a}_2, ..., \bar{a}_{|A|}\right\}$ is set of *actions*
- Let $\mathcal{D} \equiv \mathscr{P}(I \times A)$ be space of joint player-action distributions
- A *player* $i \in I$ would receive *signal* $\tilde{s}(\omega, i) \in \Sigma$ in a *state* $\omega \in \Omega$
- Under a worldly *state* $\omega \in \Omega$, a *player* $i \in I$ would receive a $[0,1]$-valued *payoff* $\tilde{u}(\omega, i, a, \delta)$ when she takes an *action* $a \in A$ while facing an *external environment* $\delta \in \mathcal{D}$

Game Formulation

- For some $\tilde{\gamma} \in \mathscr{P}(\Omega)$ as a *prior distribution* of worldly states and some *atomless* $\tilde{\lambda} \in \mathscr{P}(I)$ as a *player distribution*, we can identify our NG as $(\Omega, I, \Sigma, A, \tilde{s}, \tilde{u}, \tilde{\gamma}, \tilde{\lambda})$
- For emphasis on *player distribution*, let us name this NG $\Gamma(\tilde{\lambda})$
- Let $\mathcal{I}(\sigma|\omega) \equiv [\tilde{s}(\omega, \cdot)]^{-1}(\{\sigma\})$ (think *horizontal*) be set of players who would receive signal σ under state ω——$(\mathcal{I}(\sigma|\omega))_{\sigma \in \Sigma}$ would form a *partition* of player space I for every state ω
- Let $\mathcal{W}(\sigma|i) \equiv [\tilde{s}(\cdot, i)]^{-1}(\{\sigma\})$ (think *vertical*) be set of states that would let player i receive signal σ——$(\mathcal{W}(\sigma|i))_{\sigma \in \Sigma}$ would form a *partition* of state space Ω for every player i

A Depiction with $|\Sigma| = 3$

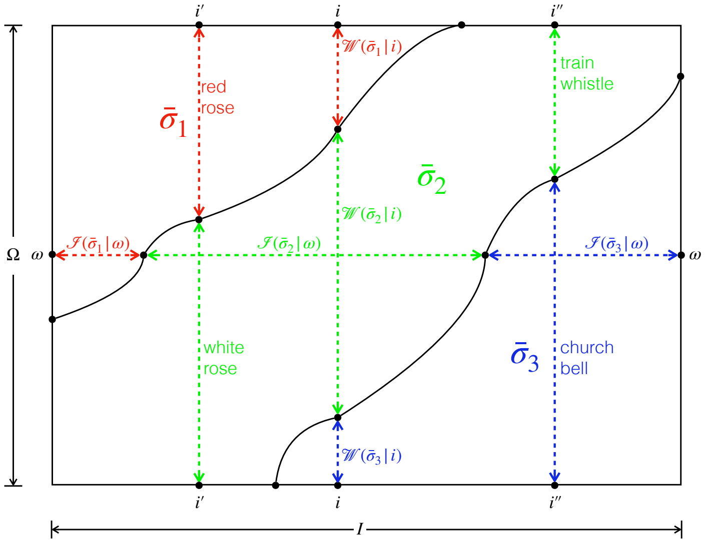

Mixed Strategy Profile for NG

- Let $\Theta_{|A|}$ be probability *simplex* embedded in $\Re^{|A|}$

  It is topologically equivalent to $\mathscr{P}(A)$ and its set of extreme points $\Theta^0_{|A|}$ contains vectors $\bar{\theta}^0_{|A|,j}$ that correspond to actions $\bar{a}_j$

- By a *mixed strategy profile* for $\Gamma(\tilde{\lambda})$, we mean

$$\mu \equiv (\mu(i,\sigma))_{i\in I,\sigma\in\Sigma} \equiv (\mu(a|i,\sigma))_{i\in I,\sigma\in\Sigma,a\in A},$$

  which is an element of $\mathcal{M} \equiv \mathscr{M}(I \times \Sigma, \Theta_{|A|})$, so that

  each $\mu(a|i,\sigma)$ is chance for an (i,σ)-player's action to be a

- Under a state $\omega \in \Omega$, any player $i \in I$ would receive a signal $\tilde{s}(\omega, i) \in \Sigma$, prompting her to adopt an *action plan* $\mu(i, \tilde{s}(\omega, i)) \equiv (\mu(a|i, \tilde{s}(\omega, i)))_{a\in A} \in \Theta_{|A|}$ under μ

A State-determined Distribution

- In aggregation, joint player-action distribution would be some $\Delta(\omega,\mu) \in \mathcal{D}$ such that for any $I' \in \mathscr{B}(I)$ and $a \in A$,

$$\begin{aligned}[\Delta(\omega,\mu)]\,(I' \times \{a\}) &\equiv \textstyle\int_{I'} \mu(a|i,\tilde{s}(\omega,i)) \cdot \tilde{\lambda}(di) \\ &= \textstyle\sum_{\sigma\in\Sigma} \int_{I'\cap\mathcal{I}(\sigma|\omega)} \mu(a|i,\sigma) \cdot \tilde{\lambda}(di)\end{aligned}$$

- Note $\Delta(\omega,\mu)$ could be understood as $\tilde{\lambda} \odot K(\omega,\mu)$ with each I-to-$\Theta_{|A|}$ mapping $K(\omega,\mu)$ satisfying

$$[K(\omega,\mu)](\sigma|i) \equiv \mu(a|i,\sigma) \qquad \text{for } i \in \mathcal{I}(\sigma|\omega)$$

- Though heuristically plausible, we do not claim $\Delta(\omega,\mu)$ to be almost sure *empirical* player-action distribution resulting from players in $I \equiv [0,1]$ "*independently*" carrying out μ under ω

Relevant Definitions

- When taking action $a \in A$ under state $\omega \in \Omega$ while everyone adopts strategy profile $\mu \in \mathcal{M}$, player $i \in I$ would receive

$$V(\omega, i, a, \mu) \equiv \tilde{u}(\omega, i, a, \Delta(\omega, \mu))$$

- However, player knows *not state* $\omega \in \Omega$ *but signal* $\tilde{s}(\omega, i) \in \Sigma$ sent to her alone——she is an (i, σ)-player when $\tilde{s}(\omega, i) = \sigma$

- Thus, define *unnormalized* average payoff to an (i, σ)-player as

$$U(i, \sigma, a, \mu) \equiv \int_{\mathcal{W}(\sigma|i)} V(\omega, i, a, \mu) \cdot \tilde{\gamma}(d\omega)$$

  For purpose of identifying actions $a \in A$ that maximize average payoff, whether or not to *normalize* would *not* matter

NG Equilibrium

- For any player $i \in I$, signal $\sigma \in \Sigma$, and strategy profile $\mu \in \mathcal{M}$, let

$$\mathbb{B}^0(i, \sigma, \mu) \equiv \left\{ \bar{\theta}^0_{|A|,j} : \ \bar{a}_j \in \text{argmax}_{a \in A} U(i, \sigma, a, \mu) \right\}$$

- A member θ of its *convex hall* $\mathbb{B}(i, \sigma, \mu)$ would be characterized by

$$\sum_{a \in A} \theta(a) \cdot U(i, \sigma, a, \mu) \geq \sum_{a \in A} \theta'(a) \cdot U(i, \sigma, a, \mu),$$

for any $\theta' \in \Theta_{|A|}$ i.e., (i, σ)-player's set of *optimal* action distributions in response to common μ

- A mixed strategy profile $\mu^* \in \mathcal{M}$ would be an *equilibrium* when

$$\mu^*(i, \sigma) \in \mathbb{B}\left(i, \sigma, \mu^*\right), \qquad \forall i \in I, \ \sigma \in \Sigma$$

Equilibrium Existence

- Define a correspondence $\mathbb{F} : \mathcal{M} \rightrightarrows \mathcal{M}$ so that for each $\mu \in \mathcal{M}$,

$$\mathbb{F}(\mu) \equiv \{\mu' \in \mathcal{M} : \ \mu'(i,\sigma) \in \mathbb{B}(i,\sigma,\mu), \quad \forall i \in I \ \sigma \in \Sigma\}$$

  It is set of all strategy profiles μ' whose every component $\mu'(i,\sigma)$ is (i,σ)-player's *best response* to given strategy profile μ

- A strategy profile $\mu^* \in \mathcal{M}$ would be an *equilibrium* if and only if it is a *fixed point* of $\mathbb{F}(\cdot)$ satisfying $\mu^* \in \mathbb{F}(\mu^*)$

- With various *compactness*, *convexity*, *continuity*, *nonemptiness*, and *upper hemi-continuity*, we can use Fan-Glicksberg fixed point theorem to show *existence* of NG equilibria

- In other words, $\mathcal{M}^*(\tilde{\lambda}) \neq \emptyset$

A Finite Game

- Consider $\Gamma_n(i_{[n]})$ with player profile $i_{[n]} \equiv (i_m)_{m=1,2,\ldots,n} \in I^n$
- It would inherit from $\Gamma(\tilde{\lambda})$ *same* state space Ω, signal space Σ, action space A, payoff function $\tilde{u}(\omega, i, a, \delta)$, state-player–to–signal mapping $\tilde{s}(\omega, i)$, and prior state distribution $\tilde{\gamma}$
- When players' actions form a profile $a_{[n]} \equiv (a_m)_{m=1,2,\ldots,n} \in A^n$, *empirical player-action distribution* faced by player i_m would be $\varepsilon(i_{[n],-m}, a_{[n],-m})$ and hence leading to her payoff

$$\tilde{u}\left(\omega, i_m, a_m, \varepsilon(i_{[n],-m}, a_{[n],-m})\right)$$

- At a particular n and a given player profile $i_{[n]}$ which is settled in *background*, our finite game could be very general

Mixed Strategy for Finite Game

- For m-th player with an i_m characteristic, her *mixed strategy* could be some $\mu_m \equiv (\mu_m(\sigma))_{\sigma\in\Sigma} \equiv (\mu_m(a|\sigma))_{\sigma\in\Sigma,a\in A} \in \mathcal{M}_0 \equiv (\Theta_{|A|})^{\Sigma}$, such that each $\mu_m(a|\sigma)$ would represent *chance* for action a to be taken when she receives a σ signal
- Let us use $\mu_{[n]} \equiv (\mu_m)_{m=1,2,\ldots,n} \in \mathcal{M}_0^n$ for strategy profile adopted by all players and $\mu_{[n],-m} \equiv (\mu_l)_{l\neq m} \in \mathcal{M}_0^{n-1}$ for profile of strategies adopted by all players except m-th one
- Average payoff $V_n(\omega, i_m, a_m, i_{[n],-m}, \mu_{[n],-m})$ to m-th player when she takes action a_m under state ω while her opponents form $i_{[n],-m}$ and adopt *strategy profile* $\mu_{[n],-m}$ would be

$$\sum_{a_{[n],-m}\in A^{n-1}} \left[\prod_{l\neq m} \mu_l\left(a_l|\tilde{s}(\omega, i_l)\right) \cdot \tilde{u}\left(\omega, i_m, a_m, \varepsilon(i_{[n],-m}, a_{[n],-m})\right)\right]$$

Equilibrium-related Concepts

- As m-th player does *not* see actual *state* ω *but signal* $\sigma_m \equiv \tilde{s}(\omega, i_m)$ sent her way, she should naturally care about average payoff $U_n\left(i_m, \sigma_m, a_m, i_{[n],-m}, \mu_{[n],-m}\right)$ defined as

$$\int_{\mathcal{W}(\sigma_m|i_m)} V_n\left(\omega, i_m, a_m, i_{[n],-m}, \mu_{[n],-m}\right) \cdot \tilde{\gamma}(d\omega),$$

  which is so far *unnormalized*

- For any $\epsilon \geq 0$, consider $\mu_{[n]} \in \mathcal{M}_0^n$ an ϵ-*equilibrium* for $\Gamma_n(i_{[n]})$ when for any $m = 1, 2, ..., n$, $\sigma_m \in \Sigma$, and $a' \in A$,

$$\begin{aligned}\textstyle\sum_{a_m \in A} \mu_m(a_m|\sigma_m) \cdot U_n(i_m, \sigma_m, a_m, i_{[n],-m}, \mu_{[n],-m}) \\ \geq U_n(i_m, \sigma_m, a', i_{[n],-m}, \mu_{[n],-m}) - \tilde{\gamma}(\mathcal{W}(\sigma_m|i_m)) \cdot \epsilon\end{aligned}$$

Equilibrium Interpretations

- Namely, m-th player's *unnormalized* average payoff

$$\sum_{a_m \in A} \mu_m(a_m|\sigma_m) \cdot U_n(i_m, \sigma_m, a_m, i_{[n],-m}, \mu_{[n],-m})$$

  of using $\mu_m(\cdot|\sigma_m)$ is better than $U_n(i_m, \sigma_m, a', i_{[n],-m}, \mu_{[n],-m})$ of taking any action a' except for some $\tilde{\gamma}(\mathcal{W}(\sigma_m|i_m)) \cdot \epsilon$ margin

- Sub-unitary weight $\tilde{\gamma}(\mathcal{W}(\sigma_m|i_m))$ in front of ϵ, which makes condition *more stringent* than when it were not there, matches *unnormalized* payoff definition

- Let $\mathcal{M}_n^*(i_{[n]}, \epsilon)$ be $\Gamma_n(i_{[n]})$'s set of ϵ-*equilibria*

- An ϵ-equilibrium here would be *same as* that in *traditional* sense

Randomly Generated Finite Games

- A version of *law of large numbers* says for any $\epsilon > 0$,

$$\lim_{n \longrightarrow +\infty} \tilde{\lambda}^n \left(\left\{ i_{[n]} \equiv (i_m)_{m=1,2,\ldots,n} \in I^n : \ \rho_I(\tilde{\lambda}, \varepsilon(i_{[n]})) < \epsilon \right\} \right) = 1,$$

  meaning that *empirical distribution* of randomly generated finite game's *player profile* i would converge to NG's *signature distribution* $\tilde{\lambda}$ in a *probabilistic* sense

- Note $[\varepsilon(i_{[n]})](\{i'\}) \equiv \sum_{m=1}^{n} \mathbf{1}(i_m = i')/n$ for any $i' \in I$ and ρ_I is *Prokhorov metric* for player distribution space $\mathscr{P}(I)$

- Though not directly used, this probabilistic closeness serves as a rationale for *mixed*-to-*mixed* and *mixed*-to-*pure* approximations

A Mixed-to-mixed Guarantee

- An NG strategy profile $\mu \in \mathcal{M}$ would *induce* a profile $\tilde{\nu}_{[n]}(\mu, i_{[n]})$ for n-player game $\Gamma_n(i_{[n]})$, so that any m-th player with an i_m characteristic would *behave as if* she were a player i_m in $\Gamma(\tilde{\lambda})$
- When n-player profile $i_{[n]}$ is *randomly sampled* from distribution $\tilde{\lambda}$ while μ^* is an *equilibrium* for NG $\Gamma(\tilde{\lambda})$, there would be a *big chance* for resulting $\tilde{\nu}_{[n]}(\mu^*, i_{[n]})$ *to be good* for $\Gamma_n(i_{[n]})$
- Precisely, suppose *state space* Ω is *finite* and utility function $\tilde{u}$ is sufficiently *continuous*

  Then, for any $\epsilon > 0$ and $\mu^* \in \mathcal{M}^*(\tilde{\lambda})$,

$$\lim_{n \longrightarrow +\infty} \tilde{\lambda}^n \left(\left\{ i_{[n]} \in I^n : \ \tilde{\nu}_{[n]}(\mu^*, i_{[n]}) \in \mathcal{M}^*_n(i_{[n]}, \epsilon) \right\} \right) = 1$$

A Key Convergence of External Environment

- A key is that a player in randomly generated $\Gamma_n(i_{[n]})$ would face an *external environment* that increasingly *resembles* one in NG $\Gamma(\tilde{\lambda})$
- Recall $\Delta(\omega, \mu)$ is joint player-action *distribution* corresponding to *state* $\omega \in \Omega$ and *strategy profile* $\mu \in \mathcal{M}$
- For any $\epsilon > 0$, we can show $\tilde{\lambda}^n(\mathcal{I}_n(\omega, \mu, \epsilon)) > 1 - \epsilon$ for large enough n where $\mathcal{I}_n(\omega, \mu, \epsilon)$ is

$$\{i_{[n]} \in I^n : \sum_{a_{[n],-m} \in A^{n-1}} \prod_{l \neq m} \mu(a_l | i_l, \tilde{s}(\omega, i_l)) \times \\ \times \mathbf{1}(a_{[n],-m} \in \mathcal{A}_{n-1}(\omega, \mu, \epsilon, i_{[n],-m})) > 1 - \epsilon, \text{ for any } m\},$$

while each $\mathcal{A}_{n-1}(\omega, \mu, \epsilon, i_{[n],-m})$ is

$$\left\{a_{[n],-m} \in A^{n-1} : \rho_{I \times A}\left(\Delta(\omega, \mu), \varepsilon(i_{[n],-m}, a_{[n],-m})\right) < \epsilon\right\}$$

A Source Enabler

- Convergence of *external environment* in turn stems from a result of *law of large numbers* sort——for any $\epsilon > 0$,

$$\lim_{n \longrightarrow +\infty} (\Delta(\omega, \mu))^n (\{(i_{[n]}, a_{[n]}) \in (I \times A)^n : \\ \max_{m=1}^{n} \ \rho_{I \times A} \left(\Delta(\omega, \mu), \varepsilon(i_{[n],-m}, a_{[n],-m})\right) < \epsilon\}) = 1$$

- By a *pure* strategy profile for NG $\Gamma(\tilde{\lambda})$, we mean some

$$\pi \equiv (\pi(i, \sigma))_{i \in I, \sigma \in \Sigma} \equiv (\pi(a|i, \sigma))_{a \in A, i \in I, \sigma \in \Sigma},$$

  which is an element of $\mathcal{P} \equiv \mathscr{M}\left(I \times \Sigma, \Theta^0_{|A|}\right)$, so that a with $\pi(a|i, \sigma)$=1 is *action* to be taken by an (i, σ)-player

- As $\tilde{\nu}_{[n]}(\pi, i_{[n]})$ is a pure n-player strategy profile for any $\pi \in \mathcal{P}$, our *mixed*-to-*mixed* result would bear *pure*-to-*pure* message as well

A Mixed-to-pure Guarantee

- Even a generally *mixed* equilibrium μ^* for NG $\Gamma(\tilde{\lambda})$ would, very likely in some μ^*-based *probabilistic* sense, help achieve a *pure* ϵ-equilibrium for an n-player game $\Gamma_n(i_{[n]})$ whose player profile $i_{[n]} \equiv (i_m)_{m=1,2,\ldots,n}$ is *randomly generated* from $\tilde{\lambda}$
- A precise description requires a few definitions
- Let $\mathcal{P}_n^*(i_{[n]}, \epsilon)$ be n-player game $\Gamma_n(i_{[n]})$'s set of *pure* ϵ-equilibria
- Given an n-player action-plan profile

$$\alpha_{[n]} \equiv (\alpha_m)_{m=1,2,\ldots,n} \equiv (\alpha_m(\sigma))_{m=1,2,\ldots,m,\sigma\in\Sigma} \in (A^{\Sigma})^n,$$

we can define a *pure* n-player *strategy profile* $\tilde{\pi}_{[n]}(\alpha_{[n]})$ so that

$$(\tilde{\pi}_m(\alpha))(a|\sigma) = 1 \qquad \text{if and only if} \qquad a = \alpha_m(\sigma)$$

A Joint Player–action-plan Distribution

- For any NG strategy profile $\mu \in \mathcal{M}$, define an $(I \times A^{\Sigma})$-distribution $\Psi(\mu) \in \mathscr{P}(I \times A^{\Sigma})$ so that for any player subset $I' \in \mathscr{B}(I)$ and signal-based action plan $\alpha \equiv (\alpha(\sigma))_{\sigma \in \Sigma} \in A^{\Sigma}$,

$$[\Psi(\mu)](I' \times \{\alpha\}) \equiv \int_{I'} \left[\prod_{\sigma \in \Sigma} \mu(\alpha(\sigma)|i, \sigma)\right] \cdot \tilde{\lambda}(di)$$

- Sampling from $\Psi(\mu)$ would amount to obtaining a *player* $i \in I$ following NG-defining distribution $\tilde{\lambda}$ and then for *every possible signal* $\sigma \in \Sigma$, obtaining an *action* $\alpha(\sigma) \in A$ following (i, σ, μ)-determined distribution $\mu(\cdot|i, \sigma)$

- For some NG equilibrium μ^*, our *mixed*-to-*pure* result would be based on random sampling from $\Psi(\mu^*)$ first of *players* and then of their corresponding *signal-based action plans*

The Precise Mixed-to-pure Statement

- Beyond demands on Ω and $\tilde{u}$, we also require $\mathcal{I}(\sigma|\omega)$'s for every state ω to be F-sigma sets or equivalently, at most *countable* unions of open, half-open-half-closed, and closed *intervals*
- A precise statement reads that for any $\epsilon > 0$,

$$\lim_{n \longrightarrow +\infty} (\Psi(\mu^*))^n \left(\left\{(i_{[n]}, \alpha_{[n]}) : \ \tilde{\pi}_{[n]}(\alpha_{[n]}) \in \mathcal{P}_n^*(i_{[n]}, \epsilon)\right\}\right) = 1,$$

for any NG equilibrium $\mu^* \in \mathcal{M}^*(\tilde{\lambda})$

- A key enabling property involving *external environment* is for $\epsilon > 0$,

$$\lim_{n \longrightarrow +\infty} (\Psi(\mu))^n(\{(i_{[n]}, \alpha_{[n]}) \in (I \times A^{\Sigma})^n : \max_{m=1}^n \rho_{I \times A^{\Sigma}}(\Psi(\mu), \varepsilon(i_{[n],-m}, \alpha_{[n],-m})) < \epsilon\}) = 1$$

More about Pure Equilibria

- Earlier *mixed*-to-*mixed* and current *mixed*-to-*pure* convergence would convey same *pure*-to-*pure* message when NG equilibrium μ^* is some *pure* π^* to start with

- When state space Ω is *finite* and game is *anonymous*, NG $\Gamma(\tilde{\lambda})$ would indeed have *pure* equilibria π^*

- Our NG is considered *anonymous* when other players influence a given player through *marginal action distribution* only, i.e., for some $\tilde{u}^{\mathsf{a}} : \Omega \times I \times A \times \mathscr{P}(A)$, payoff satisfies

$$\tilde{u}(\omega, i, a, \delta) \equiv \tilde{u}^{\mathsf{a}}(\omega, i, a, \delta|_A)$$

- Proof involves *purification* relying on $\tilde{\lambda}$'s *atomlessness*

Concluding Remarks

- We have formulated a Bayesian NG that
  (i) allows players to have *correlated* signals
  (ii) has n-player counterparts that are quite *general*
- We have relied on *finiteness* of signal space Σ, action space A, and many times that of state space Ω——relaxations should be welcome
- Our *mixed*-to-*pure* result, though having somehow relaxed $|I| \times |\Sigma| \gg |A|$ for pure-equilibrium existence at price of being *approximate*, is in *fixed*-NG–to–*random*-finite-game direction
- It might be beneficial to learn from *statistics* on how to deal with a *given finite* Bayesian game——1st to derive *convergence rates*?