Bayesian Games

Mihai Manea

MIT

Partly based on lecture notes by Muhamet Yildiz.

Bayesian Games

A Bayesian game is a list (N, A, Θ, T, u, p)

- N: set of players
- $A = (A_i)_{i \in N}$: set of action profiles
- Θ: set of payoff parameters
- T_i: set of types for player i; $T = \prod_{i \in N} T_i$
- $u_i : \Theta \times A \to \mathbb{R}$: payoff function of player i
- $p_i(\cdot|t_i) \in \Delta(\Theta \times T_{-i})$: belief of type t_i

Each player i knows his own type t_i but does not necessarily know θ or other players' types. . . belief $p_i(\cdot|t_i)$.

The game has a *common prior* if there exists $\pi \in \Delta(\Theta \times T)$ such that

$$p_i(\cdot|t_i) = \pi(\cdot|t_i), \forall t_i \in T_i, \forall i \in N.$$

Infinite Hierarchies of Beliefs

When a researcher models incomplete information, there is often no ex-ante stage or explicit information structure in which players observe signals and make inferences. At the modeling stage, each player *i* has an *infinite hierarchy of beliefs*

- a *first-order belief* $\tau_i^1 \in \Delta(\Theta)$ about payoffs (and other aspects of the world)
- a *second-order belief* $\tau_i^2 \in \Delta\left(\Theta \times \Delta(\Theta)^{N \setminus \{i\}}\right)$ about θ and other players' first-order beliefs τ_{-i}^1
- a *third-order belief* τ_i^3 about correlations in player *i*'s second-order uncertainty τ_i^2 and other players' second-order beliefs $\tau_{-i}^2 \ldots$

Formal Definition

For simplicity, consider two players.

Suppose that Θ is a Polish (complete separable metric) space.

Player i has beliefs about θ, about other's beliefs about θ,...

$$\begin{aligned}
X_0 &= \Theta \\
X_1 &= X_0 \times \Delta(X_0) \\
&\vdots \\
X_n &= X_{n-1} \times \Delta(X_{n-1}) \\
&\vdots
\end{aligned}$$

$\tau_i = (\tau_i^1, \tau_i^2, \ldots) \in \prod_{n=0}^{\infty} \Delta(X_n)$: *belief hierarchy* of player i

$H_i = \prod_{n=0}^{\infty} \Delta(X_n)$: set of i's hierarchies of beliefs

Every X_n is Polish. Endow X_n with the weak topology.

Interpretation of Type Space

Harsanyi's (1967) parsimonious formalization of incomplete information through a *type space* (Θ, T, p) naturally generates an infinite hierarchy of beliefs for each $t_i \in T_i$, which is consistent by construction:

$$\text{first-order belief:} \qquad h_i^1(\cdot|t_i) = \text{marg}_\Theta p(\cdot|t_i) = \sum_{t_{-i}} p(\theta, t_{-i}|t_i)$$

$$\text{second-order belief:} \qquad h_i^2\left(\theta, \hat{h}_{-i}^1|t_i\right) = \sum_{t_{-i}|h_{-i}^1(\cdot|t_{-i})=\hat{h}_{-i}^1} p(\theta, t_{-i}|t_i)\ldots$$

A type $t_i \in T_i$ in a space (Θ, T, p) models a belief hierarchy $(\tau_i^1, \tau_i^2, \ldots)$ if $h_i^n(\cdot|t_i) = \tau_i^n$ for each n.

Coherency

How expressive is Harsanyi's language?

Is there (T, p) s.t. $\{h_i(\cdot|t_i)|t_i \in T_i\} = H_i$?

Hierarchies should be *coherent*:

$$\text{marg}_{X_{n-2}} \tau_i^n = \tau_i^{n-1}.$$

Different levels of beliefs should not contradict one another.

H_i^0: set of i's coherent hierarchies.

Proposition 1 (Brandenburger and Dekel 1993)

There exists a homeomorphism $f_i : H_i^0 \to \Delta(\Theta \times H_{-i})$ *s.t.*

$$\text{marg}_{X_{n-1}} f_i(\cdot|\tau_i) = \tau_i^n, \forall n \geq 1.$$

Common Knowledge of Coherency

Is there (T, p) s.t. $\{h_i(\cdot|t_i)|t_i \in T_i\} = H_i^0$?

We need to restrict attention to hierarchies of beliefs under which there is *common knowledge of coherency*:

- $H_i^1 = \{\tau_i \in H_i^0 | f_i(H_{-i}^0|\tau_i) = 1\}$
- $H_i^2 = \{\tau_i \in H_i^1 | f_i(H_{-i}^1|\tau_i) = 1\}$
- ...
- $H_i^* = \bigcap_{k \geq 0} H_i^k$

Proposition 2 (Brandenburger and Dekel 1993)

There exists a homeomorphism $g_i : H_i^* \to \Delta(\Theta \times H_{-i}^*)$ s.t.

$$\text{marg}_{X_{n-1}} g_i(\cdot|\tau_i) = \tau_i^n, \forall n \geq 1.$$

H_i^*: universal type space

The Interim Game

For any Bayesian game $\mathcal{B} = (N, A, \Theta, T, u, p)$, define the *interim game* $IG(\mathcal{B}) = \left(\hat{N}, \hat{S}, U\right)$,

$$\begin{aligned}
\hat{N} &= \cup_{i \in N} T_i \\
\hat{S}_{t_i} &= A_i \\
\hat{U}_{t_i}(\hat{s}) &= E_{p_i(\cdot|t_i)}\left[u_i\left(\theta, \hat{s}\right)\right] \equiv \sum_{(\theta, t_{-i})} p_i(\theta, t_{-i}|t_i) u_i\left(\theta, \hat{s}_{t_i}, \hat{s}_{t_{-i}}\right), \forall t_i \in \hat{N},
\end{aligned}$$

where $\hat{s} = \left(\hat{s}_{t_i}\right)_{t_i \in \hat{N}}$.

Assume finite $\Theta \times T$ to avoid measurability issues.

Th Ex Ante Game

For a Bayesian game $\mathcal{B} = (N, A, \Theta, T, u, \pi)$ with a common prior π, the *ex-ante game* $G(\mathcal{B}) = (N, S, U)$ is given by

$$S_i = A_i^{T_i} \ni s_i : T_i \to A_i$$
$$U_i(s) = E_\pi[u_i(\theta, s(t))].$$

Bayesian Nash Equilibrium

Strategies of player i in $\mathcal{B}$ are mappings $s_i : T_i \to A_i$ (measurable when T_i is uncountable).

Definition 1

In a Bayesian game $\mathcal{B} = (N, A, \Theta, T, u, p)$, a strategy profile $s : T \to A$ is a *Bayesian Nash equilibrium* (BNE) if it corresponds to a Nash equilibrium of $IG(\mathcal{B})$, i.e., for every $i \in N, t_i \in T_i$

$$E_{p_i(\cdot|t_i)}\left[u_i\left(\theta, s_i\left(t_i\right), s_{-i}\left(t_{-i}\right)\right)\right] \geq E_{p_i(\cdot|t_i)}\left[u_i\left(\theta, a_i, s_{-i}\left(t_{-i}\right)\right)\right], \forall a_i \in A_i.$$

Interim rather than ex ante definition preferred since in models with a continuum of types the ex ante game has many spurious equilibria that differ on probability zero sets of types.

Connections to the Complete Information Games

When i plays a best-response type by type, he also optimizes ex-ante payoffs (for any probability distribution over T_i). Therefore, a BNE of $\mathcal{B}$ is also a Nash equilibrium of the ex-ante game $G(\mathcal{B})$.

$BNE(\mathcal{B})$: Bayesian Nash equilibria of $\mathcal{B}$

Proposition 3

For any Bayesian game $\mathcal{B}$ with a common prior π,

$$BNE(\mathcal{B}) \subseteq NE(G(\mathcal{B})).$$

If $\pi(t_i) > 0$ for all $t_i \in T_i$ and $i \in N$, then

$$BNE(\mathcal{B}) = NE(G(\mathcal{B})).$$

Existence of Bayesian Nash Equilibrium

Theorem 1

Let $\mathcal{B} = (N, A, \Theta, T, u, p)$ be a Bayesian game in which

- *A_i is a convex, compact subset of a Euclidean space*
- *$u_i : \Theta \times A \to \mathbb{R}$ is continuous and concave in a_i*
- *Θ is a compact metric space*
- *T is finite.*

Then $\mathcal{B}$ has a BNE in pure strategies.

Concavity is used instead of quasi-concavity to ensure (quasi-)concavity of payoffs in the interim game. . . integrals (sums) of quasi-concave functions are not always quasi-concave.

Upper-hemicontinuity of BNE with respect to parameters, including the beliefs p, can be established similarly to the complete information case.

Ex-ante Rationalizability

In a Bayesian game $\mathcal{B} = (N, A, \Theta, T, u, p)$, a strategy $s_i : T_i \to A_i$ is *ex-ante rationalizable* if s_i is rationalizable in the ex-ante game $G(\mathcal{B})$.

Appropriate solution concept if the ex-ante stage is real. Imposes restrictions on players' interim beliefs: all player *i*'s types have the same beliefs about other players' actions.

An Example

$$\Theta = \{\theta, \theta'\}$$
$$T_1 = \{t_1, t_1'\}, T_2 = \{t_2\}$$
$$p(\theta, t_1, t_2) = p(\theta', t_1', t_2) = 1/2$$

θ	L	R
U	$1, \varepsilon$	$-2, 0$
D	$0, 0$	$0, 1$

θ'	L	R
U	$-2, \varepsilon$	$1, 0$
D	$0, 0$	$0, 1$

$G(\mathcal{B})$:

	L	R
UU	$-1/2, \varepsilon$	$-1/2, 0$
UD	$1/2, \varepsilon/2$	$-1, 1/2$
DU	$-1, \varepsilon/2$	$1/2, 1/2$
DD	$0, 0$	$0, 1$

$$S^\infty(G(\mathcal{B})) = \{(DU, R)\}$$

Interim Independent Rationalizability

In a Bayesian game $\mathcal{B} = (N, A, \Theta, T, u, p)$, an action a_i is *interim independent rationalizable* for type t_i if a_i is rationalizable in the interim game $IG(\mathcal{B})$.

Implicit assumption on the interim game: common knowledge that the belief of a player i about (θ, t_{-i}) is independent of his belief about other players' actions. His belief about (θ, t_{-i}, a_{-i}) is derived from $p_i(\cdot|t_i) \times \mu_{t_i}$ for some $\mu_{t_i} \in \Delta\left(A_{-i}^{T_{-i}}\right)$. We take expectations with respect to $p_i(\cdot|t_i)$ in defining $IG(\mathcal{B})$ before considering t_i's beliefs about other players' actions.

t_i believes that θ and a_j are *independent* conditional on t_j.

Example

$$\Theta = \{\theta, \theta'\}$$
$$T_1 = \{t_1, t_1'\}; T_2 = \{t_2\}$$
$$p(\theta, t_1, t_2) = p(\theta', t_1', t_2) = 1/2$$

θ	L	R
U	$1, \varepsilon$	$-2, 0$
D	$0, 0$	$0, 1$

θ'	L	R
U	$-2, \varepsilon$	$1, 0$
D	$0, 0$	$0, 1$

$IG(\mathcal{B})$: $\hat{N} = \{t_1, t_1', t_2\}$; t_1 chooses rows, t_2 columns, and t_1' matrices

U :

	L	R
U	$1, \varepsilon, -2$	$-2, 0, 1$
D	$0, \varepsilon/2, -2$	$0, 1/2, 1$

D :

	L	R
U	$1, \varepsilon/2, 0$	$-2, 1/2, 0$
D	$0, 0, 0$	$0, 1, 0$

$$S_{t_1}^{\infty}(IG(\mathcal{B})) = S_{t_1'}^{\infty}(IG(\mathcal{B})) = \{U, D\}; S_{t_2}^{\infty}(IG(\mathcal{B})) = \{L, R\}.$$

Interim Correlated Rationalizability

Dekel, Fudenberg, and Morris (2007) allow t_i to have correlated beliefs regarding θ and a_{t_j}.

For each $i \in N, t_i \in T_i$, set $S_i^0[t_i] = A_i$ and define $S_i^k[t_i]$ for $k \geq 1$

$$a_i \in S_i^k[t_i] \iff a_i \in \arg\max_{a_i'} \int u_i(\theta, a_i', a_{-i}) d\pi(\theta, t_{-i}, a_{-i})$$

for some $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that

$$\text{marg}_{\Theta \times T_{-i}} \pi = p_i(\cdot|t_i) \text{ and } \pi\left(a_{-i} \in S_{-i}^{k-1}[t_{-i}]\right) = 1.$$

The set of *interim correlated rationalizable* (ICR) actions for type t_i is

$$S_i^\infty[t_i] = \bigcap_{k=0}^{\infty} S_i^k[t_i].$$

ICR and Higher Order Beliefs

Proposition 4

$S_i^k[t_i]$ depends only on t_i's k-order beliefs about θ. ICR is invariant with respect to belief hierarchies.

Sufficient to show that $S_i^1[t_i]$ depends only on t_i's marginal on θ...
If $a_i \in S_i^1[t_i]$, there exists $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ with $\text{marg}_{\Theta \times T_{-i}} \pi = p_i(\cdot|t_i)$ s.t.

$$\begin{aligned} a_i &\in \arg\max_{a_i'} \int u_i(\theta, a_i', a_{-i}) d\pi(\theta, t_{-i}, a_{-i}) \\ &= \arg\max_{a_i'} \int u_i(\theta, a_i', a_{-i}) d\pi^*(\theta, a_{-i}) \\ &= \arg\max_{a_i'} \int u_i(\theta, a_i', \sigma_{-i}(\theta)) d\pi^{**}(\theta) \end{aligned}$$

where π^* and π^{**} are π's marginals on $\Theta \times A_{-i}$ and Θ, resp. and $\sigma_{-i}(\theta)$ represents the marginal of π on A_{-i} conditional on θ. Conversely, if $a_i \in \arg\max_{a_i'} \int u_i(\theta, a_i', \sigma_{-i}(\theta)) d\pi^{**}(\theta)$ and $\pi^{**} = \text{marg}_\Theta p_i(\cdot|t_i)$...

BNE, IIR and Higher Order Beliefs

$\Theta = \{\theta, \theta'\}$, $N = \{1, 2\}$. Actions and payoffs:

θ	L	R
U	$1, 0$	$0, 0$
D	$.6, 0$	$.6, 0$

θ'	L	R
U	$0, 0$	$1, 0$
D	$.6, 0$	$.6, 0$

Type space $T = \{t_1, t_1'\} \times \{t_2, t_2'\}$ with common prior

θ	t_2	t_2'
t_1	$1/6$	$1/12$
t_1'	$1/12$	$1/6$

θ'	t_2	t_2'
t_1	$1/12$	$1/6$
t_1'	$1/6$	$1/12$

Every action can be played in a BNE, e.g.,

$$\begin{aligned} s_1^*(t_1) &= U, s_1^*\left(t_1'\right) = D \\ s_2^*(t_2) &= L, s_2^*\left(t_2'\right) = R. \end{aligned}$$

Second BNE with flipped actions. Each action is played by each type in a BNE, so all actions are interim independent rationalizable.

A Different Type Space

$\hat{T} = \{(\hat{t}_1, \hat{t}_2)\}$ with common prior $p(\theta, \hat{t}) = p(\theta', \hat{t}) = 1/2$; same payoffs:

θ	L	R
U	$1, 0$	$0, 0$
D	$.6, 0$	$.6, 0$

θ'	L	R
U	$0, 0$	$1, 0$
D	$.6, 0$	$.6, 0$

The only rationalizable action for player 1 in this game is D. In any BNE $\hat{t}_1$ must play D.

The two games represent the same hierarchy of beliefs! Each type $t_i \in T_i$ in the first game assigns probability 1/2 on θ.

The first type space induces correlation between θ and a_2 (via the correlation between t_2 and θ), second does not. ICR allows this sort of correlation by definition and does not eliminate any action in this example.

BNE and IIR are sensitive to the specification of the strategic environment, ICR depends only on belief hierarchies.

Coordination Game

- $N = \{1, 2\}$
- $\Theta = \{2/5, -2/5\}$

	α	β
α	θ, θ	$\theta - 1, 0$
β	$0, \theta - 1$	$0, 0$

- Complete information with $\theta = 2/5$
  - multiple equilibria
  - (α, α) is Pareto-dominant
- An incomplete information game
  - the two states equally likely, player 1 learns θ
  - if $\theta = 2/5$ player 1 sends an email, which is lost with probability 1/2
  - players send "confirmation of receipt" message until a message is lost; each message is lost with probability 1/2

Rubinstein's E-mail Game

	α	β
α	θ, θ	$\theta - 1, 0$
β	$0, \theta - 1$	$0, 0$

- $T_1 = \{-1, 1, 3, 5, \ldots\}$
- $T_2 = \{0, 2, 4, \ldots\}$
- t_i total number of messages sent or received by player i

Beliefs derived from common prior

$$\begin{aligned} p(\theta = -2/5, t_1 = -1, t_2 = 0) &= 1/2 \\ p(\theta = 2/5, t_1 = 2m - 1, t_2 = 2m - 2) &= 1/2^{2m} \\ p(\theta = 2/5, t_1 = 2m - 1, t_2 = 2m) &= 1/2^{2m+1} \end{aligned}$$

Rationalizable Actions

	α	β
α	θ, θ	$\theta - 1, 0$
β	$0, \theta - 1$	$0, 0$

β is uniquely rationalizable for every type, $S_i^\infty[t] = \{\beta\}$ for all t

- type $t_1 = -1$ knows that $\theta = -2/5$, so α is strictly dominated by β and $S_1^\infty[t_1 = -1] = \{\beta\}$
- if $\theta = 2/5$, β is best-response if probability the opponent plays β is at least $2/5$
- $p(\theta = -2/5, t_1 = -1 | t_2 = 0) = 2/3 > 2/5 \Rightarrow S_2^\infty[0] = \{\beta\}$.
- If $S_i^\infty[t] = \{\beta\}$, then $p(t|t+1) = 2/3 > 2/5$, $S_j^\infty[t+1] = \{\beta\}$.

Type $t \geq 0$ knows that $\theta = 2/5$, that the other player knows this, and so on, up to order t. For high t, beliefs about θ approach the common knowledge case with $\theta = 2/5$, where both actions are rationalizable.

Contagion: far away types lead to different behavior for similar types.

A Different Game

$\Theta = \{2/5, 6/5\}$

	α	β
α	θ, θ	$\theta - 1, 0$
β	$0, \theta - 1$	$0, 0$

Beliefs derived from common prior

$$
\begin{aligned}
p(\theta &= 6/5, t_1 = -1, t_2 = 0) = 1/2 \\
p(\theta &= 2/5, t_1 = 2m-1, t_2 = 2m-2) = 1/2^{2m} \\
p(\theta &= 2/5, t_1 = 2m-1, t_2 = 2m) = 1/2^{2m+1}
\end{aligned}
$$

α is uniquely rationalizable for every type, $S_i^\infty[t] = \{\alpha\}$ for all t

- $t_1 = -1$ knows that $\theta = -2/5$, so α strictly dominates β and $S_1^\infty[t_1 = -1] = \{\alpha\}$.
- if $\theta = 2/5$, α is best-response if opponent plays α with probability at least $3/5$
- type $t \geq 0$ places probability $2/3 > 3/5$ on $t - 1$...

A General Model

Fix

- N: finite set of players
- Θ^*: set of parameters (fundamentals)
- $A = A_1 \times ... \times A_n$: finite action space
- $u_i : \Theta^* \times A \to \mathbb{R}$ payoff function of $i \in N$, assumed continuous

Consider all Bayesian games $B = (N, \Theta, T, A, u, p)$ where $\Theta \subseteq \Theta^*$.

- hierarchy of beliefs

$$h_i(t_i|B) = \left(h_i^1(t_i|B), h_i^2(t_i|B), \ldots\right)$$

- through h, every type space can be embedded continuously in the universal type space T^*.
- topology on T^*

$$h_i\left(t_i^m|B^m\right) \to h_i(t_i|B) \iff \left[h_i^k\left(t_i^m|B^m\right) \to h_i^k(t_i|B), \forall k\right],$$

where the last convergence is in the weak topology.

Upper Hemicontinuity

Theorem 1 (Dekel, Fudenberg, and Morris 2006)

$S^\infty[t]$ is upper-hemicontinuous in t.

I.e., $\forall \left(t_i^m, B^m\right)$, (t_i, B) with $h_i\left(t_i^m|B^m\right) \to h_i(t_i|B)$,

$$\left[a_i \in S_i^\infty\left[t_i^m|B^m\right], \forall \text{large } m\right] \Rightarrow a_i \in S_i^\infty[t_i|B].$$

Corollary 1

If $S_i^\infty[t_i|B] = \{a_i\}$, then

$$h_i\left(t_i^m|B^m\right) \to h_i(t_i|B) \Rightarrow S_i^\infty\left[t_i^m|B^m\right] = \{a_i\} \text{ for large } m.$$

A researcher who gathers information about higher order beliefs can eventually confirm that the set of rationalizable actions with respect to his evidence is a subset of ICR for the actual type. Not true for other solution concepts.

Structure Theorem

Theorem 2 (Weinstein and Yildiz 2007)

Assume that Θ^* is rich, so that each action is strictly dominant at some θ. Then for any $a_i \in S_i^\infty[t_i]$, there exists $t_i^m \to t_i$ such that

$$S_i^\infty\left[t_i^m\right] = \{a_i\}, \forall m.$$

Proof relies on an extension of the contagion argument.

Comments

- In a game of complete information, every rationalizable strategy is sensitive to common-knowledge assumptions whenever there are multiple rationalizable strategies.
- ICR does not have a proper refinement that is upper-hemicontinuous with respect to belief hierarchies.
- ICR provides the only robust prediction with respect to higher order beliefs. It characterizes the predictions that may be verified by a researcher who can observe arbitrarily precise noisy signals about arbitrarily high but finite orders of beliefs.

An Epistemic Model

- Fix a finite Bayesian game $B = (N, A, \Theta, u, T, p)$.
- Information structure (Ω, I, π)
  - Ω: finite set of states
  - $I = (I_1, \ldots, I_n)$: profile of information partitions of Ω
  - $I_i(\omega)$: information set of i that contains ω
  - $\pi = (\pi_{i,\omega})_{i \in N, \omega \in \Omega}$: profile of beliefs $\pi_{i,\omega} \in \Delta(I_i(\omega))$.
  - $I_i(\omega) = I_i(\omega') \Rightarrow \pi_{i,\omega} = \pi_{i,\omega'}$
- Epistemic model for B: $(\Omega, I, \pi, \boldsymbol{\theta}, \boldsymbol{t}, \boldsymbol{a})$
  - $\boldsymbol{\theta} : \Omega \to \Theta$
  - $\boldsymbol{t} : \Omega \to T$
  - $\boldsymbol{a} : \Omega \to A$
  - $\boldsymbol{t}_i$ and $\boldsymbol{a}_i$ are constant over information sets of i
  - $\pi_{i,\omega} \circ (\boldsymbol{\theta}, \boldsymbol{t}_{-i})^{-1} = p_i(\cdot | \boldsymbol{t}_i(\omega))$ for all ω, i

Common Knowledge of Rationality = ICR

Definitions 2

Rationality is common knowledge in $(\Omega, I, \pi, \boldsymbol{\theta}, \boldsymbol{t}, \boldsymbol{a}) \iff$

$$\boldsymbol{a}_i(\omega) \in B_i\left(\pi_{i,\omega} \circ (\boldsymbol{\theta}, \boldsymbol{a}_{-i})^{-1}\right), \forall i \in N, \omega \in \Omega.$$

a_i is *consistent with common knowledge of rationality* for $t_i \iff$ there exists $(\Omega, I, \pi, \boldsymbol{\theta}, \boldsymbol{t}, \boldsymbol{a})$ and $\omega \in \Omega$ s.t.

- $\boldsymbol{t}_i(\omega) = t_i$; $\boldsymbol{a}_i(\omega) = a_i$;
- rationality is common knowledge in $(\Omega, I, \pi, \boldsymbol{\theta}, \boldsymbol{t}, \boldsymbol{a})$

Theorem 3 (Dekel, Fudenberg, and Morris 2007)

a_i^* is consistent with common knowledge of rationality for t_i^* iff $a_i^* \in S_i^\infty\left[t_i^*\right]$.