Static Games with Incomplete Information: Payoff Uncertainty

P. Battigalli
Bocconi University
Game Theory: Analysis of Strategic Thinking

October 9, 2023

Abstract

Some justifications of solution concepts make sense under the assumption that the rules of the game and players' personal preferences are common knowledge. A situation of strategic interaction features **incomplete information** when this is not the case. We represent this with **games with payoff uncertainty**, whereby the payoff functions depend on a vector of parameters about which players have partial and asymmetric knowledge. The game features **private values** if there is common knowledge of the outcome function, and **interdependent values** otherwise. It is relatively straightforward to extend rationalizability and pure self-confirming equilibrium to allow for payoff uncertainty.

[These slides summarize and complement parts of Sections 8.1-3 and 8.7 of Ch. 8 of GT-AST]

Introduction

- Whether a solution or equilibrium concept is consistent with incomplete information is a matter of *interpretation*. We must look at the conceptual *motivations*:
- Standard *rationalizability* (iterated deletion of strictly dominated actions) is explicitly motivated as representing the behavioral implications of rationality and common belief in rationality *under complete information* (common knowledge of the payoff functions).
- *Nash equilibrium* can be motivated as an "*obvious way to play* the game": See deductive interpretation, and self-enforcing agreement interpretation. Also this makes sense *under the complete information assumption*.
- *Deductive interpretation of NE*: it makes sense when there is a unique rationalizable outcome, see above.

Incomplete Information and Self-Enforcing Agreements

- *Self-enforcing agreement interpretation of NE*: Again we need complete information (or maybe something "close" to it) in order to make sense of this interpretation.
- Consider the following game and the Pareto dominant agreement (t, ℓ). The *agreement is self-enforcing if there is common belief* (or "almost common belief") *that there is no incentive to deviate from* (t, ℓ).

	ℓ	r
t	100,100	0,99
b	99,0	99,99

- Would Rowena (row player) play t if she is not sure of the payoff function of Colin (column player)? What if she is not sure that Colin is sure of her payoff function? What if...?
- [But **note**: as long as each player knows her/his payoff function, there is *no need to assume complete information to make sense of Nash* equilibrium *as* description of rest points *of adaptive processes*.]

Environments with Incomplete Information

- Rules of the game $\Rightarrow$ *outcome/consequence function* $g : A \to Y$.
- Each player $i \in I$ ranks (lotteries over) outcomes according to (the expectation of) a vNM utility function $v_i : Y \to \mathbb{R}$.
- In environments with **incomplete information** there is *lack of common knowledge of* g (outcome function) and/or $(v_i)_{i \in I}$ (personal preferences).
- Such situation can be described with *parameterized payoff functions*

$$u_i : \Theta \times A \to \mathbb{R},$$

with

  - $\theta \in \Theta$ parameter affecting payoffs,

$$\theta = (\theta_0, (\theta_i)_{i \in I}) \in \Theta = \Theta_0 \times (\times_{i \in I} \Theta_i)$$

  - $i \in I$ *knows* only θ_i =private information of i about payoffs.

Interpretation, Distributed Knowledge

- *Intuition*: it is common knowledge that $\theta \in \Theta$, Θ_i represents what is commonly believed possible about i's traits known to him (e.g., tastes, abilities), the "larger" Θ_i the more uncertain are the other players about such traits.
- If Θ_i is a *singleton* ($i \in I$), that is, $\Theta_i = \left\{\bar{\theta}_i\right\}$, it means that *what i knows is common knowledge* (it is common knowledge that $\theta_i = \bar{\theta}_i$) and Θ_i can be neglected: indeed, $\Theta_0 \times \left(\times_{j \in I \setminus \{i\}} \Theta_j\right)$ and Θ have the same cardinality; hence, they are (intuitively) isomorphic.
- Θ_0 represents the *residual uncertainty* that would remain if the players could pool their private information.
- We often *focus* on the case where Θ_0 is a singleton: there is *no residual uncertainty* after pooling private information (in this case it is said that there is "**distributed knowledge**" of θ). Thus, we will often neglect Θ_0.

Private Values

We distinguish between the case of **private values**, where u_i depends only on θ_i, and **interdependent values**, where u_i may depend on the whole θ.

- **Private values:** *Common knowledge of outcome function* g, but lack of common knowledge of preferences $(v_i)_{i \in I}$:
  - (it is common knowledge that) each i knows his vNM utility function $v_i \Rightarrow$ parameterized representation $v_i : \Theta_i \times Y \to \mathbb{R}$.
  - Note: $\{w_i \in \mathbb{R}^Y : \exists \theta_i \in \Theta_i, w_i = v_{i,\theta_i}\}$ is the set of utility functions that each $j \neq i$ thinks i might have $\Rightarrow$ get

$$u_i(\theta_i, a) = v_i(\theta_i, g(a))$$

  - Note: *under private values* we may assume w.l.o.g. that there is *distributed knowledge of* θ (Θ_0 singleton).

Interdependent Values

- **Interdependent values:** *lack of common knowledge of outcome function* g, which may depend on θ_0 or on personal traits such as some players' "ability").
  - *common knowledge of preferences* $(v_i)_{i \in I}$ (simplest case) $\Rightarrow$ parameterized representation $g : \Theta \times A \longrightarrow Y$; note: $\{\gamma \in Y^A : \exists \theta \in \Theta, \gamma = g_\theta\}$ is the set of possible outcome functions $\Rightarrow$ get
$$u_i(\theta, a) = v_i(g(\theta, a)).$$
  - More generally, *if neither the outcome function nor preferences are common knowledge*, each v_i is parameterized by θ_i and
$$u_i(\theta, a) = v_i(\theta_i, g(\theta, a)).$$
  - *Interdependence*: The value for i depends on what j knows, e.g., a personal trait of j.

Cournot Oligopoly

Example

Cournot oligopoly *model (quantity setting): firm* $i = 1, ..., n$ *produces* $q_i \geq 0$ *units of homogeneous good*

- Inverse demand $P(Q) = [\bar{p} + \theta_0 - Q]_+$ (with $[x]_+ := \max\{0, x\}$, $Q = \sum_{i=1}^{n} q_i$)
- Cost function *of firm i:* $C_i(q_i, \theta_i) = \theta_i q_i$, $0 \leq q_i \leq \bar{q}$ ($\bar{q}$=common capacity)
- Common knowledge of risk neutrality *and of sets* $\Theta_0, \Theta_1, ..., \Theta_n$
- Payoff *of i:* $u_i(\theta_0, \theta_i, q_1, ..., q_n) = \left(\left[\bar{p} + \theta_0 - \sum_{j=1}^{n} q_j \right]_+ - \theta_i \right) q_i$
- *There are private values and distributed knowledge of* θ **if** *there is common knowledge of market demand (*Θ_0 *singleton)*

Team Production

Example

Team production: Team agents $i = 1, ..., n$, i *exerts effort* $e_i \geq 0$

- Cost of effort *(in units of output)* $C_i(e_i, k_i) = k_i e_i^2$, $k_i \in K_i \subseteq \mathbb{R}_+$
- Production function*:* $y = \prod_{i=1}^{n} e_i^{p_i}$, $p_i \in P_i \subseteq \mathbb{R}_+$
- $\theta_i = (k_i, p_i) \in K_i \times P_i = \Theta_i$
- Common knowledge of (output-)risk neutrality *and of sets* $\Theta_i = K_i \times P_i$
- Payoff function *of* i*:* $u_i(k_1, p_1, ..., k_n, p_n, e_1, ..., e_n) = \frac{1}{n} \prod_{j=1}^{n} e_j^{p_j} - k_i e_i^2$
- Private values iff *sets* $P_1, ..., P_n$ *are singletons (productivities are common knowledge), otherwise* interdependent *values*

Games with Payoff Uncertainty

- We can represent (simultaneous) strategic interaction under *incomplete information* with the mathematical structure

$$\hat{G} = \left\langle I, \Theta_0, \left(\Theta_i, A_i, u_i : \Theta \times A \to \mathbb{R}\right)_{i \in I} \right\rangle ;$$

  it is assumed that the interactive situation represented by $\hat{G}$ is common knowledge. This is called **game with payoff uncertainty**; θ_i is called the **information-type** of i.
- *Interpretation*: θ_0 affects the payoffs of somebody (if $\theta_0' \neq \theta_0''$, then $\exists i \in I, u_i(\theta_0', \cdot) \neq u_i(\theta_0'', \cdot)$). But part, or all, of i's private information θ_i *may be payoff irrelevant*. Yet even payoff-irrelevant information may be strategically relevant (e.g., θ_i may be the report to i by an art expert about the autenticity of a painting for sale).
- Take the obvious extension to payoff uncertainty of the definition of "**compact-continuous game**." To extend "**nice game**," add to the obvious properties the *convexity* (or connectedness) of each Θ_i .

Rationality and Common Belief in Rationality

- Games with payoff uncertainty are sufficient to describe certain aspects of strategic thinking, specifically, *rationality and common belief in rationality*.
- Write $\mathrm{B}_i(E)$ for "*i believes E*" (with prob. 1), and $\mathrm{B}(E) = \bigcap_{i \in I} \mathrm{B}_i(E)$ for "*everybody believes E*," R_i for "*i is rational*," $R = \bigcap R_i$ for "*everybody is rational*."
- What actions of i are consistent with R (rationality), $\mathrm{B}(R)$ (mutual belief in rationality), $\mathrm{B}(\mathrm{B}(R))$, $\mathrm{B}(\mathrm{B}(\mathrm{B}(R)))$ ... $R \cap \mathrm{CB}(R)$?

Example: Incomplete Information on One Side

Example

Possible payoff functions given by the following tables. Player 1 (Rowena) knows θ *while player 2* (Colin) does not *($\Theta \cong \Theta_1$)*

$\hat{G}^1$:

θ'	ℓ	r
t	4,0	2,1
b	3,1	1,0

θ''	ℓ	r
t	2,0	0,1
b	0,1	1,2

▶ $R_1 \Rightarrow [t \text{ if } \theta']$, *because* t dominates b given $\theta = \theta'$ *(recall, Row. knows* θ*)* $\Rightarrow (\theta', b)$ *is inconsistent with rationality (delete).*▶ $R_2 \cap B_2(R_1) \Rightarrow r$, *because* $u_2(\theta, x, \ell) < u_2(\theta, x, r)$ *for all* $(\theta, x) \neq (\theta', b)$ *(those consistent with* R_1*).*▶ $R_1 \cap B_1(R_2) \cap B_1(B_2(R_1)) \Rightarrow$ *Row. picks best reply to r given* θ
$\Rightarrow [b \text{ if } \theta = \theta'']$.

Envelope Game

Example

Players 1 and 2 receive an envelope. Envelope of i contains θ_i Euros, with $\theta_i = 1, ..., K$. Each player can offer to exchange (OE) by paying transaction cost $\varepsilon > 0$ (small). Exchange executed IFF both offer:

$\hat{G}^2$:

$a_i \backslash a_j$	OE	No
OE	$\theta_j - \varepsilon$	$\theta_i - \varepsilon$
No	θ_i	θ_i

Note: *A rational player i* offers to exchange only *if she assigns positive probability to event* $[\theta_j > \theta_i] \cap [a_j = OE]$.▶ $R_i \Rightarrow [a_i = No \text{ if } \theta_i = K]$ *because OE is dominated in this case.*▶ $R_i \cap \mathrm{B}_i(R_j) \Rightarrow [a_i = No \text{ if } \theta_i = K - 1]$ *because ...*▶ $R_i \cap \mathrm{B}_i(R_j) \cap \mathrm{B}_i(\mathrm{B}_j(R_i)) \Rightarrow [a_i = No \text{ if } \theta_i = K - 2]$ *because ...*▶ *It can be shown that:* $R \cap \mathrm{CB}(R) \Rightarrow (\forall \theta_i, a_i = No \text{ given } \theta_i)$ (*no-trade!*).

Rationalizability in Games with Payoff Uncertainty

- To ease notation, assume *distributed knowledge:* $\Theta \cong \times_{i\in I}\Theta_i$.
- Given conjecture $\mu^i \in \Delta(\Theta_{-i} \times A_{-i})$ and private information $\theta_i \in \Theta_i$, let

$$r_i(\mu^i, \theta_i) := \arg\max_{a_i\in A_i} \mathbb{E}_{\mu^i}\left(u_{i,\theta_i,a_i}\right)$$

where $u_{i,\theta_i,a_i} : \Theta_{-i} \times A_{-i} \to \mathbb{R}$ is the section of u_i at (θ_i, a_i); in the *finite support case*

$$\mathbb{E}_{\mu^i}\left(u_{i,\theta_i,a_i}\right) = \sum_{(\theta_{-i},a_{-i})\in \mathrm{supp}\mu^i} u\left(\theta_i, \theta_{-i}, a_i, a_{-i}\right)\mu^i\left(\theta_{-i}, a_{-i}\right)$$

- Let $C_i \subseteq \Theta_i \times A_i$ (with $\mathrm{proj}_{\Theta_i} C_i = \Theta_i$); interpretation: set of "surviving" pairs (see previous examples); $C_{-i} = \times_{j\neq i} C_j$, $\mathcal{C}$ collection of (closed) Cartesian products.
- Define the (monotone) **rationalization operator** $\rho : \mathcal{C} \to \mathcal{C}$.

$$\begin{aligned}\rho_i(C_{-i}) &= \left\{(\theta_i, a_i) \in \Theta_i \times A_i : \exists \mu^i \in \Delta(C_{-i}), a_i \in r_i(\mu^i, \theta_i)\right\}.\\ \rho(C) &= \times_{i\in I}\rho_i(C_{-i}).\end{aligned}$$

Behavioral Implications of RCBR

Assumptions about behavior and beliefs	Implications for $(\theta_i, a_i)_{i\in I}$
R	$\rho(\Theta\times A)$
$R\cap \mathrm{B}(R)$	$\rho^2(\Theta\times A)$
$R\cap \mathrm{B}(R)\cap \mathrm{B}^2(R)$	$\rho^3(\Theta\times A)$
...	...
$R\cap\left(\bigcap_{k=1}^m \mathrm{B}^k(R)\right)$	$\rho^{m+1}(\Theta\times A)$
...	...
$R\cap\left(\bigcap_{k=1}^\infty \mathrm{B}^k(R)\right) = R\cap \mathrm{CB}(R)$	$\rho^\infty(\Theta\times A)$

Theorem

If $\hat{G}$ is finite or compact-continuous, then

$$\rho^\infty(\Theta\times A) = \rho\left(\rho^\infty(\Theta\times A)\right) \quad \textit{and} \ \mathrm{proj}_\Theta \rho^\infty(\Theta\times A) = \Theta.$$

Furthermore, for each $C\in\mathcal{C}$, $C\subseteq\rho(C)$ implies $C\subseteq\rho^\infty(\Theta\times A)$.

Comments on RCBR and Rationalizability

- The previous theorem extends Theorems 2 and 3 of GT-AST from games with complete information to games with incomplete information (payoff uncertainty):
- $\rho^{\infty}(\Theta \times A) = \rho\left(\rho^{\infty}(\Theta \times A)\right)$ is the "fixed set property" of the rationalizable set: after countably many iterations there is no need to re-start the iterated deletion procedure.
- $\mathrm{proj}_{\Theta}\rho^{\infty}(\Theta \times A) = \Theta$ means that, for every $(\theta_i)_{i \in I} \in \Theta$, the *set of rationalizable actions for* information-type θ_i is *not empty*.
- $C \subseteq \rho(C) \Rightarrow C \subseteq \rho^{\infty}(\Theta \times A)$ means that every (Cartesian) subset of $\Theta \times A$ with the Best Reply Property is included in the rationalizable set.

- Fix Cartesian subset $C = \times_{i \in I} C_i$ with $C_i \subseteq \Theta_i \times A_i$. Let $C_{i,\theta_i} := \{a_i \in A_i : (\theta_i, a_i) \in C_i\}$ (section of set C_i at θ_i).

Definition

Mixed action α_i **dominates** a_i **given** θ_i **within** C, written $\alpha_i \gg_{(\theta_i, C)} a_i$, if $\mathrm{supp}\alpha_i \subseteq C_{i,\theta_i}$ and

$$\forall (\theta_{-i}, a_{-i}) \in C_{-i}, \;\; u_i(\theta_i, \theta_{-i}, \alpha_i, a_{-i}) > u_i(\theta_i, \theta_{-i}, a_i, a_{-i}).$$

Lemma

Fix a finite or compact-continuous $\hat{G}$; let $C = \times_{i \in I} C_i$ be non-empty and compact. For all $i \in I$ and $(\theta_i, a_i^*) \in C_i$ the following are equivalent:
(1) $\nexists \alpha_i$ s.t. $\alpha_i \gg_{(\theta_i, C)} a_i^*$ (a_i^* undominated given θ_i within C)
(2) $\exists \mu^i \in \Delta(C_{-i})$ s.t. $a_i^* \in \arg\max_{a_i \in C_{i,\theta_i}} \mathbb{E}_{\mu^i}(u_{i,\theta_i,a_i})$.

Iterated Dominance

- The previous result (with $C = \Theta \times A$) extends the Wald-Pearce Lemma on justifiability and dominance to simultaneous-moves games with incomplete information (payoff uncertainty).
- For each $C \in \mathcal{C}$ ($\forall i \in I$, $C_i \subseteq \Theta_i \times A_i$), define $\mathrm{ND}(C)$ as follows:

$$
\begin{aligned}
\mathrm{ND}_i(C) &= C_i \backslash \left\{ (\theta_i, a_i) \in C_i : \exists \alpha_i \in \Delta(C_{i,\theta_i}), \alpha_i \gg_{(\theta_i, C)} a_i \right\}, \\
\mathrm{ND}(C) &= \times_{i \in I} \mathrm{ND}_i(C).
\end{aligned}
$$

- Similarly, dominance by pure actions (very relevant for nice games) gives

$$
\begin{aligned}
\mathrm{ND}_{p,i}(C) &= C_i \backslash \left\{ (\theta_i, a_i) \in C_i : \exists a_i' \in C_{i,\theta_i}, a_i' \gg_{(\theta_i, C)} a_i \right\}, \\
\mathrm{ND}_p(C) &= \times_{i \in I} \mathrm{ND}_{p,i}(C).
\end{aligned}
$$

Rationalizability and Iterated Dominance

Extension of the equivalence *"rationalizable IFF iteratively undominated"* to games with payoff uncertainty (r is the point-rationalization operator obtained with deterministic conjectures):

Theorem

If $\hat{G}$ is finite or compact-continuous, for all $m = 1, 2, ..., \infty$,
$\rho^m (\Theta \times A) = \mathrm{ND}^m (\Theta \times A)$.

Theorem

If $\hat{G}$ is nice, point-rationalizability, rationalizability, and pure iterated dominance coincide, that is, for all $m = 1, 2, ..., \infty$,
$\mathrm{r}^m (\Theta \times A) = \rho^m (\Theta \times A) = \mathrm{ND}^m_p (\Theta \times A)$;
furthermore, the projections of these sets onto A ($proj_A \rho^m (\Theta \times A)$, $m \in \mathbb{N} \cup \{\infty\}$) are closed order-intervals (products of closed intervals).

Rationalizability in the Envelope Game

Example

Recall: *The game is symmetric, the envelope of i contains* €θ_i, *with* $\theta_i = 1, ..., K$. *Each player can offer to exchange (OE) by paying a* small *transaction cost* $\varepsilon > 0$. *Exchange executed IFF both offer:*

$\hat{G}^2$:

$a_i \backslash a_j$	*OE*	*No*
OE	$\theta_j - \varepsilon$	$\theta_i - \varepsilon$
No	θ_i	θ_i

Recall: *A rational player i* offers to exchange only *if she assigns positive probability to* $[\theta_j > \theta_i] \cap [a_j = OE]$. *With this, for each i*:▶ *OE is dominated for* $\theta_i = K$; *delete* $(\theta_i, a_i) = (K, OE)$.▶ *Given that* (K, OE) *is deleted for j, OE is dominated for* $\theta_i = K - 1$; *delete* $(K - 1, OE)$.▶ ...*Given that* (θ_j, OE) *is deleted for each* $\theta_j \in \{K - k + 1, ..., K\}$ $(1 \leq k < K)$, *OE is dominated for* $\theta_i = K - k$; *delete* $(K - k, OE)$...▶ $\rho^\infty (\Theta \times A) = \rho^K (\Theta \times A) = (\{1, ..., K\} \times \{No\})^2$ (*no type trades*).

Example

The Cournot model presented above (with $n = 2$) is a *nice game with payoff uncertainty*. Thus, look at best replies (B.R.) to *deterministic* conjectures.

Assume: $\theta_0 = 0$ commonly known, marg. cost $\theta_i \in [0, 1]$, $\bar{p} > 2$ (highest average cost much lower than $\bar{p}$), $\bar{q}$ large ($\bar{q} > \bar{p} - 1$).

The model has *private values* and is *symmetric*: $A_i = [0, \bar{q}]$, $\Theta_i = [0, 1]$, and each firm i's payoff depends on θ_i and q_{-i} in the same way.

Hence, common B.R. function

$$r(\theta_i, q_{-i}) = \left[\frac{\bar{p} - \theta_i}{2} - \frac{1}{2} q_{-i}\right]_+$$

where $\frac{\bar{p}-\theta_i}{2} = r(\theta_i, 0)$=monopolistic output for cost-type θ_i.

Rationalizability in a Duopoly with Incomplete Inform., II

Example

(Cont.) Let $\underline{r}(q_{-i}) = \left[\frac{\bar{p}-1-q_{-i}}{2}\right]_+$, $\bar{r}(q_{-i}) = \left[\frac{\bar{p}-q_{-i}}{2}\right]_+$ be the B.R. functions of, respectively, the *least efficient* ($\theta_i = 1$) and *most efficient* ($\theta_i = 0$) cost-type (see picture at the end).

Look at min and max output at each rationalizability step $k \in \mathbb{N}_0$:
$\underline{q}(k) = \underline{r}(\bar{q}(k-1))$ and $\bar{q}(k) = \bar{r}\left(\underline{q}(k-1)\right)$, with
$\underline{q}(0) = 0$, $\bar{q}(0) = \bar{q}$. Then, $\mathrm{proj}_A \rho^k(\Theta \times A) = \left[\underline{q}(k), \bar{q}(k)\right]^2$ with

$\underline{q}(1) = \underline{r}(\bar{q}) = \left[\frac{\bar{p}-1-\bar{q}}{2}\right]_+ = 0$, $\bar{q}(1) = \bar{r}(0) = \frac{\bar{p}}{2}$,

$\underline{q}(2) = \underline{r}\left(\frac{\bar{p}}{2}\right) = \frac{\bar{p}-2}{4}$, $\bar{q}(2) = \bar{r}(0) = \frac{\bar{p}}{2}$,

$\underline{q}(3) = \underline{r}\left(\frac{\bar{p}}{2}\right) = \frac{\bar{p}-2}{4}$, $\bar{q}(3) = \bar{r}\left(\frac{\bar{p}-2}{4}\right) = \frac{3\bar{p}+2}{8}$,

$\underline{q}(4) = \underline{r}\left(\frac{3\bar{p}+2}{8}\right) = \frac{5\bar{p}-10}{16}$, $\bar{q}(4) = \bar{r}\left(\frac{\bar{p}-2}{4}\right) = \frac{3\bar{p}+2}{8}$,

Show: $\mathrm{proj}_A \rho^\infty(\Theta \times A) = \left[\lim_{\ell\to\infty}(\underline{r}\circ\bar{r})^\ell(0), \bar{r}\left(\lim_{\ell\to\infty}(\underline{r}\circ\bar{r})^\ell(0)\right)\right]^2$
and compute $\underline{q}(\infty) = \underline{r}(\bar{q}(\infty))$, $\bar{q}(\infty) = \bar{r}\left(\underline{q}(\infty)\right)$.

Contextual Restrictions on Beliefs

- In many applications it is plausible to assume that the context makes some restrictions on players' beliefs about θ **transparent**, i.e., true and commonly believed (see examples in the book).
- It may also make sense to assume that restrictions on beliefs about both θ *and* behavior (i.e., on conjectures) are transparent.
- Such restrictions may depend on the information-type θ_i. We represent them with a restricted set of conjectures $\Delta_{i,\theta_i} \subseteq \Delta\left(\Theta_{-i} \times A_{-i}\right)$ (keep assuming distributed knowledge of θ) for each i and $\theta_i \Rightarrow$ profile of restricted sets $\Delta = \left(\Delta_{i,\theta_i}\right)_{i\in I,\theta_i\in\Theta_i}$.
- Modified rationalization operator (monotone!): for each $C \in \mathcal{C}$,

$$\rho_{i,\Delta}\left(C_{-i}\right) = \left\{(\theta_i, a_i) : \exists\mu^i \in \Delta_{i,\theta_i} \cap \Delta\left(C_{-i}\right), a_i \in r_i(\mu^i, \theta_i)\right\},$$

$$\rho_\Delta\left(C\right) = \times_{i\in I}\rho_{i,\Delta}\left(C_{-i}\right).$$

Directed Rationalizability

- The transparent restrictions (represented by) Δ "direct" the rationalizability procedure toward some results. Hence, the approach is called "**directed rationalizability**."
- Say that Δ represents **restrictions on exogenous beliefs** if, for every i and θ_i there is some $\bar{\Delta}_{i,\theta_i} \subseteq \Delta(\Theta_{-i})$ s.t. $\Delta_{i,\theta_i} = \left\{\mu^i \in \Delta(\Theta_{-i} \times A_{-i}) : \mathrm{marg}_{\Theta_{-i}}\mu^i \in \bar{\Delta}_{i,\theta_i}\right\}$: *only beliefs about exogenous θ_{-i} are restricted*. Then, for every information-type, the set of rationalizable actions is nonempty (*finiteness* for simplicity).

Theorem

Fix a finite game with payoff uncertainty $\hat{G}$ and and a profile Δ of restrictions about exogenous beliefs. Then

$$\rho_\Delta^\infty(\Theta \times A) = \rho_\Delta\left(\rho_\Delta^\infty(\Theta \times A)\right), \ \mathrm{proj}_\Theta \rho_\Delta^\infty(\Theta \times A) = \Theta.$$

Furthermore, for each $C \in \mathcal{C}$, $C \subseteq \rho_\Delta(C)$ implies $C \subseteq \rho_\Delta^\infty(\Theta \times A)$.

Self-Confirming Equilibrium with Incomplete Info., I

- Intuitively, the *SCE concept does not capture sophisticated strategic reasoning*. Players' *conjectures are only disciplined by long-run evidence*. Again, to ease notation, assume $\Theta \cong \times_{i \in I} \Theta_i$ (distributed knowledge).
- *As long as each player knows her payoff function (private values), we should get the same concept introduced in the previous lecture.* We will make this formal.
- Feedback is modeled by functions $f_i : \Theta \times A \longrightarrow M_i$ $(i \in I)$. Recall that $f_{i,\theta_i,a_i} : \Theta_{-i} \times A_{-i} \longrightarrow M_i$ is the **section** of f_i **at** (θ_i, a_i). If i observes m_i given (θ_i, a_i), she infers that the unknown profile (θ_{-i}, a_{-i}) must belong to the subset

$$f_{i,\theta_i,a_i}^{-1}(m_i) = \left\{ \left(\theta'_{-i}, a'_{-i}\right) : f_i\left(\theta_i.a_i, \theta'_{-i}, a'_{-i}\right) = m_i \right\}.$$

Self-Confirming Equilibrium with Incomplete Info., II

Definition

Fix a game with payoff uncertainty $\hat{G}$ and a profile of feedback functions $f = (f_i : \Theta \times A \to M_i)_{i \in I}$. A profile of actions and conjectures $\left(a_i^*, \mu^i\right)_{i \in I}$ is a (pure) **self-confirming equilibrium of** $\left(\hat{G}, f\right)$ **at** θ if, for each $i \in I$, (1, B.R.) $a_i^* \in r_i\left(\mu^i, \theta_i\right)$ and (2, CONF) $\mu^i\left(f_{i,\theta_i,a_i^*}^{-1}\left(f_i\left(a^*, \theta\right)\right)\right) = 1$.

- For any fixed $\theta \in \Theta$, let $\left(\hat{G}_\theta, f_\theta\right) := \left\langle I, (A_i, u_{i,\theta} : A \to \mathbb{R}, f_{i,\theta} : A \to M_i)_{i \in I}\right\rangle$ (a game with feedback as in the previous lecture).
- **Observation:** *Fix $a^* \in A$ and $\theta^* \in \Theta$ arbitrarily; if $\hat{G}$ has* private values, *then a^* is part of an SCE of $\left(\hat{G}, f\right)$ at θ^* IF AND ONLY IF a^* is part of an SCE of $\left(\hat{G}_{\theta^*}, f_{\theta^*}\right)$.*
- **Proof:** Let $\left(a_i^*, \mu^i\right)_{i \in I}$ be an SCE of $\left(\hat{G}, f\right)$ *at* θ^*. For each $i \in I$, let $\bar{\mu}^i = \mathrm{marg}_{A_{-i}} \mu^i \in \Delta\left(A_{-i}\right)$ (marginal of μ^i onto A_{-i}). Since each u_i is independent of θ_{-i}, $\left(a_i^*, \bar{\mu}^i\right)_{i \in I}$ must be an SCE of $\left(\hat{G}_{\theta^*}, f_{\theta^*}\right)$. ■

Incomplete Information and Properties of Feedback

- Recall, $f_{i,\theta_i,a_i}:\Theta_{-i}\times A_{-i}\longrightarrow M_i$ and $u_{i,\theta_i,a_i}:\Theta_{-i}\times A_{-i}\longrightarrow\mathbb{R}$ are the **sections** of f_i and u_i **at** (θ_i,a_i).
- Let $\mathcal{F}_{-i}(\theta_i,a_i)$ denote the "ex post information partition" of $\Theta_{-i}\times A_{-i}$ given (θ_i,a_i):

$$\mathcal{F}_{-i}(\theta_i,a_i)=\left\{C_{-i}\in 2^{\Theta_{-i}\times A_{-i}}:\exists m_i\in M_i,\, C_{-i}=f^{-1}_{i,\theta_i,a_i}(m_i)\right\}.$$

- $f_i:\Theta\times A\longrightarrow M_i$ satisfies
  - **own-action independence** of feedback about others (**OAI**) if $\mathcal{F}_{-i}(\theta_i,a_i')=\mathcal{F}_{-i}(\theta_i,a_i'')$ for all θ_i and all a_i',a_i'' *justifiable* (undominated) *for* θ_i;
  - **observable payoffs (OP)** if $f_{i,\theta_i,a_i}(\theta'_{-i},a'_{-i})=f_{i,\theta_i,a_i}(\theta''_{-i},a''_{-i})\Rightarrow u_{i,\theta_i,a_i}(\theta'_{-i},a'_{-i})=u_{i,\theta_i,a_i}(\theta''_{-i},a''_{-i})$ for all θ_i, a_i, θ'_{-i}, a'_{-i} and θ''_{-i}, a''_{-i} (u_{i,θ_i,a_i} is constant on each cell of $\mathcal{F}_{-i}(\theta_i,a_i)$, for all θ_i,a_i).

Properties of Feedback, SCE, and NE

- $\left(\hat{G}, f\right)$ satisfies OAI and OP if each f_i does ($i \in I$).
- **Examples:**
  - The Envelope Game and Cournot Game with $f_i = u_i$ for each $i \in I$ satisfy OP (obviously), but do *not* satisfy OAI.
  - The Cournot Game with known inverse demand function $P(\cdot)$ and $f_i(\theta, \cdot) = P(\cdot)$ for all i and θ satisfies OP and OAI.

Theorem

Suppose that $\left(\hat{G}, f\right)$ satisfies OAI and OP and fix $a^* \in A$ and $\theta^* \in \Theta$ arbitrarily; then a^* is part of an SCE at θ^* of $\left(\hat{G}, f\right)$ IF AND ONLY IF a^* is a Nash equilibrium of $\hat{G}_{\theta^*} = \left\langle I, \left(A_i, u_{i,\theta^*} : A \to \mathbb{R}\right)_{i \in I} \right\rangle$.

References

BATTIGALLI, P., E. CATONINI, AND N. DE VITO (2023): *Game Theory: Analysis of Strategic Thinking*. Typescript, Bocconi University.

BATTIGALLI, P. (2023): *Mathematical Language and Game Theory.* Typescript, Bocconi University.

SECTION OF $C_i \subseteq (\Theta_i \times A_i)$ AT $\bar{\theta}_i$
PROJECTION OF C_i ONTO Θ_i

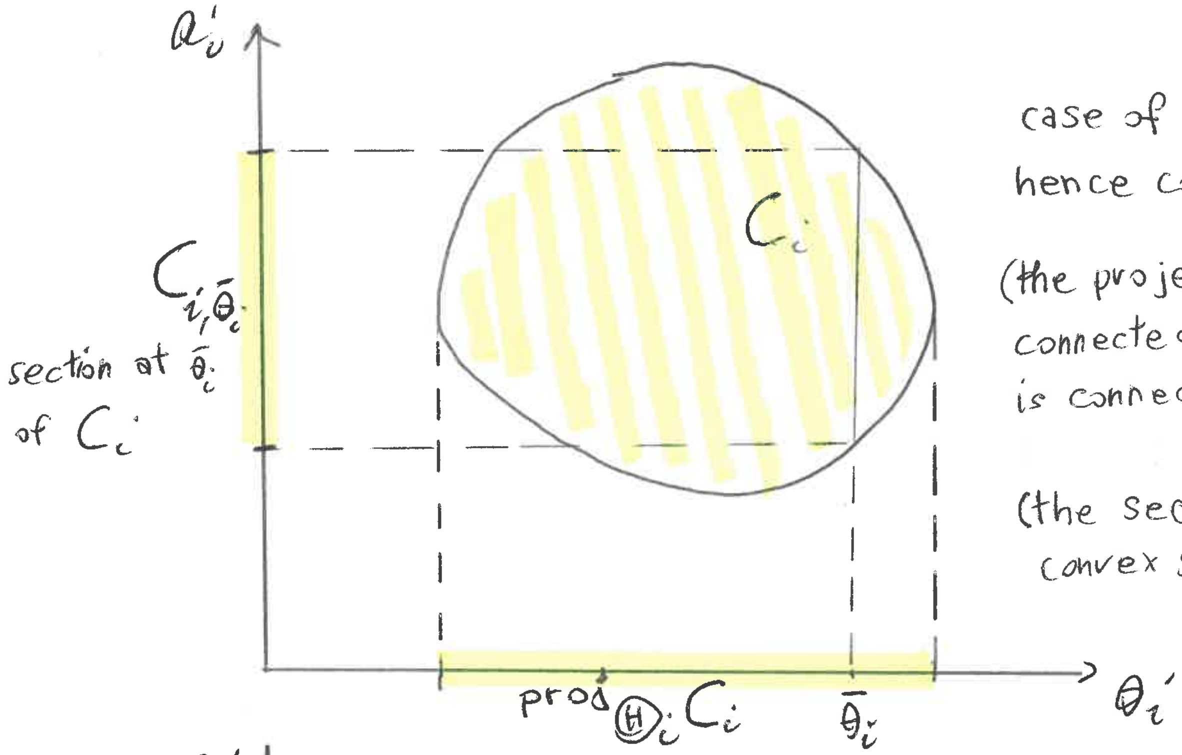

case of convex, hence connected C_i

(the projection of a connected/convex set is connected/convex)

(the section of a convex set is convex)

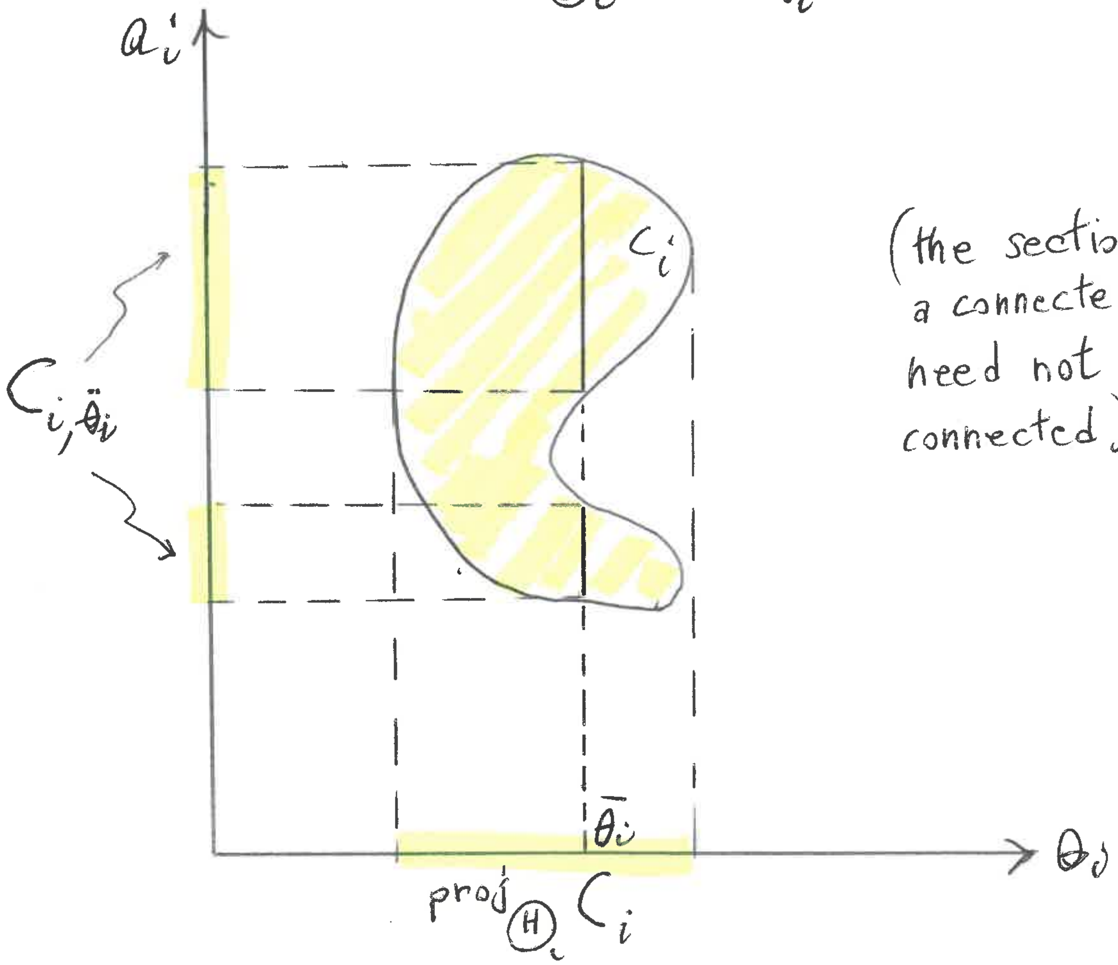

(the section of a connected set need not be connected)

Cournot duopoly with uncertainty about marginal cost of the competitor

$$\bar{q}(\infty) = \frac{1}{2}\left(\bar{P} - \underline{q}(\infty)\right)$$

$$\underline{q}(\infty) = \frac{1}{2}\left(\bar{P} - 1 - \bar{q}(\infty)\right)$$

Every $q_i \in [\underline{q}(\infty), \bar{q}(\infty)]$ is rationalizable for at least one $\theta_i \in [0,1]$

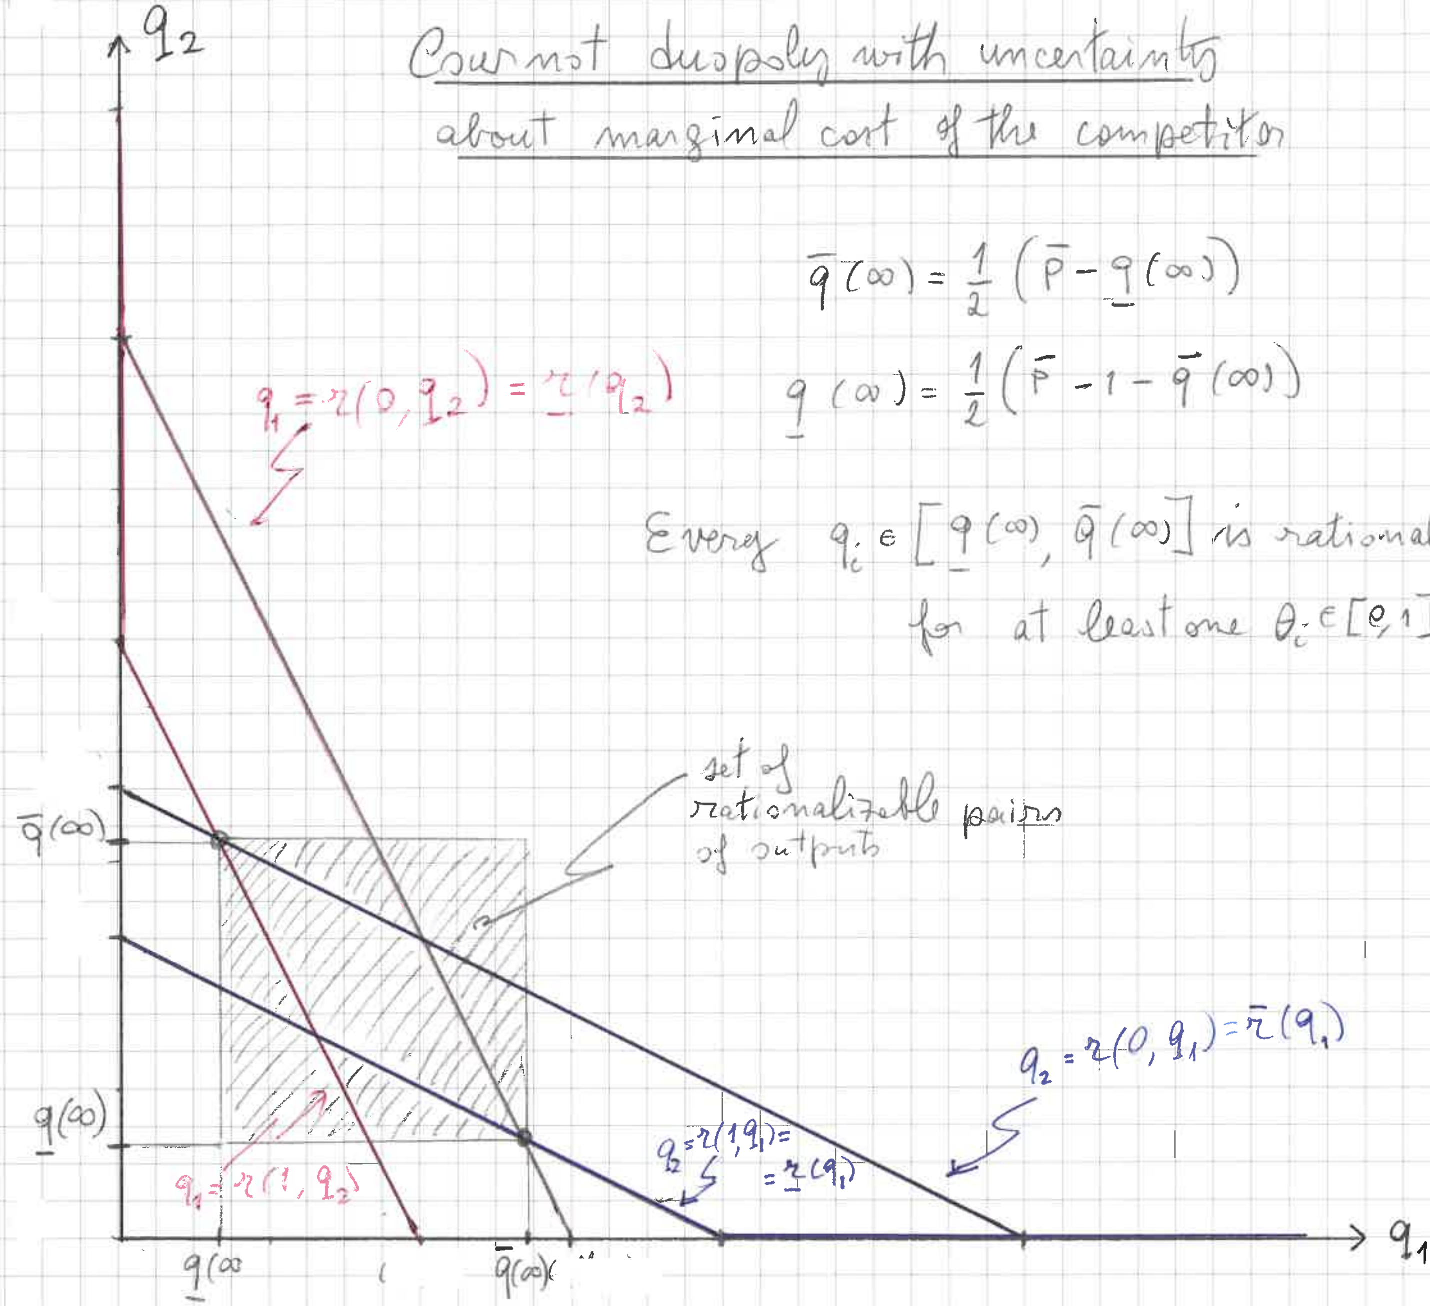

θ_i = marg. cost of firm i (private information)

$\theta_i \in [0, 1]$ $\theta_i = 0$ most efficient type $\theta_i = 1$ least efficient type

Inverse demand $P(q_1 + q_2) = \max\{0, \bar{P} - (q_1 + q_2)\}$

Ex ante symmetric nice game with payoff uncertainty