Epistemic Game Theory: Incomplete Information

Part II: Correct Beliefs

Christian W. Bach

University of Liverpool & EPICENTER

Correct Beliefs Assumption and Equilibrium

- In games with complete information, it has been shown that common belief in rationality together with a correct beliefs assumption epistemically characterizes Nash Equilibirum.
- One possible way of fleshing out the correct beliefs assumption are simple belief hierarchies.
- In **Incomplete Information Part I** the notion of common belief in rationality has been generalized to incomplete information.
- In this part simple belief hierarchies are extended to games with incomplete information.
- It turns out that they are equivalent to a solution concept called Generalized Nash Equilibrium.

Simple Belief Hierarchies

- With complete information a simple belief hierarchy (SBH) is generated by a tuple of conjectures (or mixed choices) $(\sigma_i)_{i \in I}$, where $\sigma_i \in \Delta(C_i)$ for all $i \in I$.

- An important feature of a simple belief hierarchy is that i believes his opponents to be correct about all the beliefs i holds.

- With more than 2 players, two further conditions arise:
  - i believes that any opponent j's belief about a third player k is the same as i's belief about k. (PROJECTION)
  - i belief about his opponents' choices are independent. (INDEPENDENCE)

- For incomplete information, all of these conditions need to be tailored to i's extended basic space of uncertainty:

$$(\times_{j \neq i} C_j) \times (\times_{j \neq i} U_j)$$

Example: What is Barbara's favourite Colour?

Story:

- *Barbara* and *you* are going together to another party.
- *You* wonder what colour you should wear.
- *You* prefer *blue* (4) to *green*, *green* (3) to *red*, *red* (2) to *yellow* (1), and dislike most to wear the same colour (0) as *Barbara*.
- However, *you* drank so much at the last party, that you forgot *Barbara*'s colour preferences.
- *You* are still certain about *Barbara* also disliking most to wear the same colour (0) as you.
- Also, *you* remember that *Barbara* either prefers *red* (4) to *yellow*, *yellow* (3) to *blue*, *blue* (2) to *green* (1); **or** *blue* (4) to *yellow*, *yellow* (3) to *green*, *green* (2) to *red* (1).
- Question: Which colours can *you* **rationally** choose for tonight's party under **common belief in rationality**?

Example: What is Barbara's favourite Colour?

- The game in one-person perspective form:

$\Gamma^0_y(u_y)$

	blue	green	red	yellow
blue	0	4	4	4
green	3	0	3	3
red	2	2	0	2
yellow	1	1	1	0

$\Gamma^0_B(u^r_B)$

	blue	green	red	yellow
blue	0	2	2	2
green	1	0	1	1
red	4	4	0	4
yellow	3	3	3	0

$\Gamma^0_B(u^b_B)$

	blue	green	red	yellow
blue	0	4	4	4
green	1	0	1	1
red	2	2	0	2
yellow	3	3	3	0

- Suppose the following **epistemic model** of this game:
  - $T_{you} = \{t^1_y, t^2_y, t^3_y\}$ and $T_{Barbara} = \{t^1_B, t^2_B\}$,
  - $b_{you}[t^1_y] = (red, t^1_B, u^r_B)$,
  - $b_{you}[t^2_y] = (blue, t^1_B, u^b_B)$,
  - $b_{you}[t^3_y] = 0.6 \cdot (blue, t^1_B, u^b_B) + 0.4 \cdot (green, t^2_B, u^r_B)$,
  - $b_{Barbara}[t^1_B] = (green, t^2_y, u_y)$,
  - $b_{Barbara}[t^2_B] = (blue, t^1_y, u_y)$.

- *Your* type t^1_y believes that *Barbara* chooses *red* and has utility function u^r_B.

- Also, t^1_y believes that *Barbara* believes that *you* believe *Barbara* chooses blue and has utility function u^b_B.

- Thus, *you* believe *Barbara* to be incorrect about your (first-order) beliefs.

Example: What is Barbara's favourite Colour?

- The game in one-person perspective form:

$\Gamma_y^0(u_y)$

	blue	*green*	*red*	*yellow*
blue	0	4	4	4
green	3	0	3	3
red	2	2	0	2
yellow	1	1	1	0

$\Gamma_B^0(u_B^r)$

	blue	*green*	*red*	*yellow*
blue	0	2	2	2
green	1	0	1	1
red	4	4	0	4
yellow	3	3	3	0

$\Gamma_B^0(u_B^b)$

	blue	*green*	*red*	*yellow*
blue	0	4	4	4
green	1	0	1	1
red	2	2	0	2
yellow	3	3	3	0

- Suppose the following **epistemic model** of this game:
  - $T_{you} = \{t_y^1, t_y^2, t_y^3\}$ and $T_{Barbara} = \{t_B^1, t_B^2\}$,
  - $b_{you}[t_y^1] = (red, t_B^1, u_B^r)$,
  - $b_{you}[t_y^2] = (blue, t_B^1, u_B^b)$,
  - $b_{you}[t_y^3] = 0.6 \cdot (blue, t_B^1, u_B^b) + 0.4 \cdot (green, t_B^2, u_B^r)$,
  - $b_{Barbara}[t_B^1] = (green, t_y^2, u_y)$,
  - $b_{Barbara}[t_B^2] = (blue, t_y^1, u_y)$.
- *Your* type t_y^2 believes that *Barbara* chooses *blue* and has utility function u_B^b.
- Also, t_y^2 believes that *Barbara* believes that *you* believe *Barbara* chooses blue and has utility function u_B^b.
- In fact, t_y^2 even believes that *Barbara* believes that *your* type is t_y^2.
- Thus, *you* believe that *Barbara* is correct about beliefs of yours – even about your entire belief hierarchy.

Example: What is Barbara's favourite Colour?

- The game in one-person perspective form:

$\Gamma^0_y(u_y)$

	blue	green	red	yellow
blue	0	4	4	4
green	3	0	3	3
red	2	2	0	2
yellow	1	1	1	0

$\Gamma^0_B(u^r_B)$

	blue	green	red	yellow
blue	0	2	2	2
green	1	0	1	1
red	4	4	0	4
yellow	3	3	3	0

$\Gamma^0_B(u^b_B)$

	blue	green	red	yellow
blue	0	4	4	4
green	1	0	1	1
red	2	2	0	2
yellow	3	3	3	0

- Suppose the following **epistemic model** of this game:
  - $T_{you} = \{t^1_y, t^2_y, t^3_y\}$ and $T_{Barbara} = \{t^1_B, t^2_B\}$,
  - $b_{you}[t^1_y] = (red, t^1_B, u^r_B)$,
  - $b_{you}[t^2_y] = (blue, t^1_B, u^b_B)$,
  - $b_{you}[t^3_y] = 0.6 \cdot (blue, t^1_B, u^b_B) + 0.4 \cdot (green, t^2_B, u^r_B)$,
  - $b_{Barbara}[t^1_B] = (green, t^2_y, u_y)$,
  - $b_{Barbara}[t^2_B] = (blue, t^1_y, u_y)$.
- The belief hierarchy induced by t^2_y is completely generated by the two (marginal) conjectures: $\sigma_y = (green, u_y)$ and $\sigma_B = (blue, u^b_B)$.
- Accordingly: *your* belief about *Barbara*'s choice and utility function is σ_B; *you* believe that *Barbara*'s belief about *your* choice and utility function is σ_y; *you* believe that *Barbara* believes that *your* belief about *Barbara*'s choice and utility function is σ_B; etc.
- This belief hierarchy is simple and it is generated by the tuple of marginal conjectures (σ_y, σ_B).

Outline

- Simple Belief Hierarchy
- Generalized Nash Equilibrium
- Charaterization

Simple Belief Hierarchy

Conjectures and Marginal Conjectures

- Let $i \in I$ be some player.
- A conjecture of i is a belief about his opponents' choices and utility functions, denoted as $\mu_i \in \Delta\big(\times_{j \neq i} (C_j \times U_j)\big)$.
- A marginal conjecture about player i is a belief about i's choice and utility function, denoted as $\sigma_i \in \Delta(C_i \times U_i)$.
- A conjecture of i induces a marginal conjecture $\mathrm{marg}_{C_j \times U_j} \mu_i$ about every opponent $j \neq i$.
- Note that a first-order belief of a type $t_i \in T_i$ constitutes a conjecture of that player:

$$\mathrm{marg}_{C_j \times U_j} b_i[t_i] \in \Delta\big(\times_{j \neq i} (C_j \times U_j)\big)$$

Belief Hierarchies based on Marginal Conjectures

Definition 1

Let Γ be a game with incomplete information, $\mathcal{M}^{\Gamma}$ an epistemic model of Γ, $i \in I$ some player, $t_i \in T_i$ some type of player i, and $(\sigma_j)_{j \in I} \in \times_{j \in I}\big(\Delta(C_j \times U_j)\big)$ a tuple of marginal conjectures. The induced belief hierarchy of t_i is called *generated* by $(\sigma_j)_{j \in I}$, whenever:

- player i's 1^{st}-order belief: $\prod_{j \neq i} \sigma_j$,
- player i's 2^{nd}-order belief: i believes that every opponent $j \neq i$ holds 1^{st}-order belief $\prod_{k \neq i} \sigma_k$,
- player i's 3^{rd}-order belief: i believes that every opponent $j \neq i$ believes that every opponent $k \neq j$ holds 1^{st}-order belief $\prod_{l \neq k} \sigma_l$,
- etc.

Simple Belief Hierarchy

Definition 2

Let Γ be a game with incomplete information, $\mathcal{M}^{\Gamma}$ an epistemic model of Γ, $i \in I$ some player, and $t_i \in T_i$ some type of player i. Type t_i holds a *simple belief hierarchy*, if there exists a tuple of marginal conjectures $(\sigma_j)_{j \in I} \in \times_{j \in I} \big(\Delta(C_j \times U_j)\big)$ such that the induced belief hierarchy of t_i is generated by $(\sigma_j)_{j \in I}$.

Decision Rule with SBH

Definition 3

Let Γ be a game with incomplete information, $i \in I$ some player, $c_i \in C_i$ some choice of player i, and $u_i \in U_i$ some utility function of player i. The choice c_i is *rational under common belief in rationality and a simple belief hierarchy* given u_i, if there exists an epistemic model $\mathcal{M}^\Gamma$ of Γ with some type $t_i \in T_i$ of player i such that

- t_i expresses common belief in rationality,
- t_i holds a simple belief hierarchy,
- c_i is optimal for (t_i, u_i).

Example: What is Barbara's favourite Colour?

- The game in one-person perspective form:

$\Gamma_y^0(u_y)$

	blue	green	red	yellow
blue	0	4	4	4
green	3	0	3	3
red	2	2	0	2
yellow	1	1	1	0

$\Gamma_B^0(u_B^r)$

	blue	green	red	yellow
blue	0	2	2	2
green	1	0	1	1
red	4	4	0	4
yellow	3	3	3	0

$\Gamma_B^0(u_B^b)$

	blue	green	red	yellow
blue	0	4	4	4
green	1	0	1	1
red	2	2	0	2
yellow	3	3	3	0

- Suppose the following **epistemic model** of this game:
  - $T_{you} = \{t_y\}$ and $T_{Barbara} = \{t_B\}$,
  - $b_{you}[t_y] = 0.5 \cdot (red, t_B, u_B^r) + 0.5 \cdot (blue, t_B, u_B^b)$,
  - $b_{Barbara}[t_B] = (green, t_y, u_y)$.

- The belief hierarchy induced by t_y is completely generated by the two (marginal) conjectures $\sigma_y = (green, u_y)$ and $\sigma_B = 0.5 \cdot (red, u_B^r) + 0.5 \cdot (blue, u_B^b)$ and therefore simple.

- Note that this simple belief hierarchy expresses inherent payoff uncertainty.

- Indeed *you* assign probability 0.5 to *Barbara*'s utility function u_B^r and 0.5 to u_B^b: *you* are thus inherently uncertain about *Barbara*'s utility function.

- Moreover, *you* believe that this payoff uncertainty is transparent between *Barbara* and *you*.

- Besides, observe that t_y actually expresses common belief in rationality.

Generalized Nash Equilibrium

Equilibrium for Incomplete Information

Definition 4

Let Γ be a game with incomplete information. A tuple of marginal conjectures $(\sigma_j)_{j\in I} \in \times_{j\in I}\big(\Delta(C_j \times U_j)\big)$ constitutes a *Generalized Nash Equilibrium* of Γ, whenever for all $i \in I$ and for all $(c_i, u_i) \in C_i \times U_i$ such that $\sigma_i(c_i, u_i) > 0$ it is the case that:

$$\sum_{c_{-i}\in C_{-i}} \prod_{j\in I\setminus\{i\}} \text{marg}_{C_j}\sigma_j(c_j) \cdot u_i(c_i, c_{-i})$$
$$\geq \sum_{c_{-i}\in C_{-i}} \prod_{j\in I\setminus\{i\}} \text{marg}_{C_j}\sigma_j(c_j) \cdot u_i(c_i', c_{-i})$$

for all $c_i' \in C_i$.

Remark: for the special case of complete information, Generalized Nash Equilibrium coincides with Nash Equilibrium.

Decision Rule with GNE

Definition 5

Let Γ be a game with incomplete information, $i \in I$ some player, $c_i \in C_i$ some choice of player i, and $u_i \in U_i$ some utility function of player $i \in I$. The choice c_i is *rational under generalized Nash equilibrium* given u_i, if there exists a generalized Nash equilibrium $(\sigma_j)_{j \in I}$ of Γ such that

$$\sum_{c_{-i} \in C_{-i}} \prod_{j \in I \setminus \{i\}} \mathrm{marg}_{C_j} \sigma_j(c_j) \cdot u_i(c_i, c_{-i})$$

$$\geq \sum_{c_{-i} \in C_{-i}} \prod_{j \in I \setminus \{i\}} \mathrm{marg}_{C_j} \sigma_j(c_j) \cdot u_i(c_i', c_{-i})$$

for all $c_i' \in C_i$.

Example: What is Barbara's favourite Colour?

- The game in one-person perspective form:

$\Gamma^0_y(u_y)$

	blue	green	red	yellow
blue	0	4	4	4
green	3	0	3	3
red	2	2	0	2
yellow	1	1	1	0

$\Gamma^0_B(u^r_B)$

	blue	green	red	yellow
blue	0	2	2	2
green	1	0	1	1
red	4	4	0	4
yellow	3	3	3	0

$\Gamma^0_B(u^b_B)$

	blue	green	red	yellow
blue	0	4	4	4
green	1	0	1	1
red	2	2	0	2
yellow	3	3	3	0

- Consider the two (marginal) conjectures $\sigma_y = (green, u_y)$ and $\sigma_B = 0.5 \cdot (red, u^r_B) + 0.5 \cdot (blue, u^b_B)$.

- σ_y only assigns positive probability to *green*.
- Observe that *your* choice of *green* is optimal for u_y given the marginal belief $0.5 \cdot red + 0.5 \cdot blue$ on *Barbara*'s choices.

- σ_B assigns positive probability to *red* as well as to *blue*.
- Observe that *Barbara*'s choice of *red* is optimal for u^r_B given the marginal belief $green$ on *your* choices as well as that *Barbara*'s choice of *blue* is optimal for u^b_B given the marginal belief $green$ on *your* choices.

- Therefore, the tuple (σ_y, σ_B) forms a generalized Nash equilibrium.
- *Your* choice of *green* is rational under the generalized Nash equilibrium (σ_y, σ_B) given u_y.
- *Barbara*'s choice of *red* is rational under the generalized Nash equilibrium (σ_y, σ_B) given u^r_B and *Barbara*'s choice of *blue* is rational under the generalized Nash equilibrium (σ_y, σ_B) given u^b_B.

Existence

Theorem 6

For every finite game with incomplete information there exists a Generalized Nash Equilibrium.

Characterization

Fixing a SBH ensures that CBR iff GNE

Lemma 7

Let Γ be a game with incomplete information, $\mathcal{M}^\Gamma$ an epistemic model of Γ, $(\sigma_j)_{j\in I} \in \times_{j\in I}\big(\Delta(C_j \times U_j)\big)$ some tuple of marginal conjectures, $i \in I$ some player, and $t_i \in T_i$ some type of player i that holds a simple belief hierarchy generated by $(\sigma_j)_{j\in I}$. The type t_i expresses common belief in rationality, if and only if, $(\sigma_j)_{j\in I}$ forms a Generalized Nash Equilibrium of Γ.

Proof of the *Only If* Direction of Lemma 7

- Consider some opponent $k \neq i$ of player i.
- As t_i believes in k's rationality, t_i only assigns positive probability to triples (c_k, t_k, u_k) such that c_k is optimal for (t_k, u_k).
- Since t_i's belief hierarchy is generated by $(\sigma_j)_{j \in I}$, t_i's marginal conjecture on $C_k \times U_k$ is given by σ_k and t_i believes k's belief about C_{-k} to be $\prod_{j \neq k} \text{marg}_{C_j} \sigma_j$.
- It follows that, for all $(c_k, u_k) \in \text{supp}(\sigma_k)$ it is the case that

$$\sum_{c_{-k} \in C_{-k}} \prod_{j \in I \setminus \{k\}} \text{marg}_{C_j} \sigma_j(c_j) \cdot u_k(c_k, c_{-k}) \geq \sum_{c_{-k} \in C_{-k}} \prod_{j \in I \setminus \{k\}} \text{marg}_{C_j} \sigma_j(c_j) \cdot u_k(c'_k, c_{-k})$$

for all $c'_k \in C_k$.

Proof of the *Only If* Direction of Lemma 7

- Due to his simple belief hierarchy, t_i believes each of his opponents to hold marginal conjecture σ_i about i.

- As t_i believes each of his opponents to believe in i's rationality, t_i only assigns positive probability to opponents' types that in turn only assign positive probability to triples (c_i, t_i', u_i) such that c_i is optimal for (t_i', u_i).

- Since t_i's belief hierarchy is generated by $(\sigma_j)_{j\in I}$, t_i believes that any opponent's marginal conjecture on $C_i \times U_i$ is given by σ_i and t_i believes that any opponent's type believes that i holds $\prod_{j\neq i} \mathrm{marg}_{C_j}\sigma_j$ as belief about C_{-i} .

- It follows that, for all $(c_i, u_i) \in \mathrm{supp}(\sigma_i)$ it is the case that

$$\sum_{c_{-i}\in C_{-i}} \prod_{j\in I\setminus\{i\}} \mathrm{marg}_{C_j}\sigma_j(c_j) \cdot u_i(c_i, c_{-i}) \geq \sum_{c_{-i}\in C_{-i}} \prod_{j\in I\setminus\{i\}} \mathrm{marg}_{C_j}\sigma_j(c_j) \cdot u_i(c_i', c_{-i})$$

  for all $c_i' \in C_i$.

- Consequently, for every $j \in I$, it is the case that σ_j only assigns positive probability to pairs (c_j, u_j) such that

$$\sum_{c_{-j}\in C_{-j}} \prod_{l\in I\setminus\{j\}} \mathrm{marg}_{C_l}\sigma_l(c_l) \cdot u_j(c_j, c_{-j}) \geq \sum_{c_{-j}\in C_{-j}} \prod_{l\in I\setminus\{l\}} \mathrm{marg}_{C_l}\sigma_l(c_l) \cdot u_j(c_j', c_{-j})$$

  for all $c_j' \in C_j$.

- Therefore, the tuple $(\sigma_j)_{j\in I}$ of marginal conjectures constitutes a generalized Nash equilibrium.

Epistemic Characterization of Generalized Nash Equilibrium

Theorem 8

Let Γ be a game with incomplete information, $i \in I$ some player, $c_i \in C_i$ some choice of player i, and $u_i \in U_i$ some utility function of player i. The choice c_i is rational under common belief in rationality and a simple belief hierarchy given u_i, if and only if, c_i is rational under Generalized Nash Equilibrium given u_i.

Example: The Moonlight Serenade

- *You* had a fight with *Barbara* and contemplate about three ways of apologizing to her:
  - perform a moonlight serenade outside her house,
  - bring her a box of her chocolate,
  - send your common friend Chris to apologize for you.

- When the doorbell rings *Barbara* can open up or ignore the bell.

- *Your* preferences are captured by the following decision problem:

$\Gamma_y(u_y)$

	open	*ignore*
serenade	4	0
chocolate	0	4
Chris	3	3

Example: The Moonlight Serenade

- *You* are uncertain about *Barbara*'s preferences and whether she will be in an angry or a forgiving mood.
- *Her* preferences are captured by the following two decision problems:

$\Gamma_B(u_B^{an})$		*serenade*	*chocolate*	*Chris*
	open	0	0	0
	ignore	1	1	1

$\Gamma_B(u_B^{for})$		*serenade*	*chocolate*	*Chris*
	open	0	1	0
	ignore	1	0	1

Example: The Moonlight Serenade

- The game in one-person perspective form:

$\Gamma_y(u_y)$

	open	ignore
serenade	0	4
chocolate	4	0
Chris	3	3

$\Gamma_B(u_B^{an})$

	serenade	chocolate	Chris
open	0	0	0
ignore	1	1	1

$\Gamma_B(u_B^{for})$

	serenade	chocolate	Chris
open	0	1	0
ignore	1	0	1

- Application of GISD to the game:
  - In $\Gamma_y(u_y)$ each of *your* three choices is optimal for some belief about *Barbara*'s choices.
  - In $\Gamma_B(u_B^{an})$ open is strictly dominated by ignore: delete open from $\Gamma_B(u_B^{an})$.
  - In $\Gamma_B(u_B^{for})$ each of *Barbara*'s two choices are optimal for some belief about *your* choices.

- It follows that $GISD = GISD_{you} \times GISD_{Barbara}$

$$= \{(serenade, u_y), (chocolate, u_y), (Chris, u_y)\} \times \{(ignore, u_B^{an}), (open, u_B^{for}), (ignore, u_B^{for})\}$$

- Consequently, *you* can rationally pick each of *your* three choices under common belief in rationality given *your* (only) utility function.

- *Barbara* can rationally only pick ignore under common belief in rationality if *she* is angry, whereas *she* can pick both open and ignore under common belief in rationality if *she* is forgiving.

Example: The Moonlight Serenade

- The game in one-person perspective form:

$\Gamma_y(u_y)$

	open	*ignore*
serenade	4	0
chocolate	0	4
Chris	3	3

$\Gamma_B(u_B^{an})$

	serenade	*chocolate*	*Chris*
open	0	0	0
ignore	1	1	1

$\Gamma_B(u_B^{for})$

	serenade	*chocolate*	*Chris*
open	0	1	0
ignore	1	0	1

- Next, GNE is applied to the game.
- Consider the tuple (σ_y, σ_B) of marginal conjectures, where

$$\sigma_y = 0.75 \cdot (chocolate, u_y) + 0.25 \cdot (Chris, u_y)$$

and

$$\sigma_B = 0.25 \cdot (open, u_B^{for}) + 0.75 \cdot (ignore, u_B^{an})$$

- Observe that
  - Chocolate is optimal for *you* given u_y and σ_B as belief about *Barbara*'s choice.
  - Chris is also optimal for *you* given u_y and σ_B as belief about *Barbara*'s choice.
  - Open is optimal for *Barbara* given u_B^{for} and σ_y as belief about *your* choice.
  - Ignore is optimal for *Barbara* given u_B^{an} and σ_y as belief about *your* choice.
- Therefore, (σ_y, σ_B) constitutes a GNE.
- Consequently, chocolate and Chris are rational under GNE for u_y as well as open is rational under GNE for u_B^{for} and ignore is rational under GNE for u_B^{an}.

Example: The Moonlight Serenade

- The game in one-person perspective form:

$\Gamma_y(u_y)$

	open	*ignore*
serenade	4	0
chocolate	0	4
Chris	3	3

$\Gamma_B(u_B^{an})$

	serenade	*chocolate*	*Chris*
open	0	0	0
ignore	1	1	1

$\Gamma_B(u_B^{for})$

	serenade	*chocolate*	*Chris*
open	0	1	0
ignore	1	0	1

- Can *Barbara* also rationally ignore the doorbell under GNE if she is forgiving?

- Consider the tuple (σ_y, σ_B) of marginal conjectures, where

$$\sigma_y = 0.5 \cdot (chocolate, u_y) + 0.5 \cdot (Chris, u_y)$$

  and

$$\sigma_B = 0.25 \cdot (open, u_B^{for}) + 0.75 \cdot (ignore, u_B^{for})$$

- Observe that
  - Chocolate is optimal for *you* given u_y and σ_B as belief about *Barbara*'s choice.
  - Chris is also optimal for *you* given u_y and σ_B as belief about *Barbara*'s choice.
  - Open is optimal for *Barbara* given u_B^{for} and σ_y as belief about *your* choice.
  - Ignore is also optimal for *Barbara* given u_B^{for} and σ_y as belief about *your* choice.

- Therefore, (σ_y, σ_B) constitutes a GNE.

- Hence, *Barbara* can indeed also rationally ignore the doorbell under GNE if she is forgiving.

Example: The Moonlight Serenade

- The game in one-person perspective form:

$\Gamma_y(u_y)$

	open	*ignore*
serenade	4	0
chocolate	0	4
Chris	3	3

$\Gamma_B(u_B^{an})$

	serenade	*chocolate*	*Chris*
open	0	0	0
ignore	1	1	1

$\Gamma_B(u_B^{for})$

	serenade	*chocolate*	*Chris*
open	0	1	0
ignore	1	0	1

- Can *you* also rationally play the moonlight serenade under GNE given *your* (unique) utility function?
- Towards a contradiction, let (σ_y, σ_B) be a GNE such that serenade is optimal for $(\text{marg}_{C_{Barbara}} \sigma_B, u_y)$.
- Then, $\text{marg}_{C_{Barbara}} \sigma_B(open) > 0$, as serenade would otherwise be strictly worse than chocolate and Chris.
- As open can only possibly be optimal for *Barbara* if she is forgiving, $\sigma_B(open, u_B^{for}) > 0$ must hold.
- This implies that open must be optimal for $(\text{marg}_{C_{you}} \sigma_y, u_B^{for})$ and hence $\text{marg}_{C_{you}} \sigma_y(chocolate) > 0$.
- Consequently, chocolate must also be optimal for $(\text{marg}_{C_{Barbara}} \sigma_B, u_y)$.
- Serenade and chocolate can only both be optimal for $(\text{marg}_{C_{Barbara}} \sigma_B, u_y)$, if σ_B assigns probability 0.5 to open and 0.5 to ignore.
- Both serenade and chocolate would then yield an expected payoff of 2 which is strictly worse than the 3 that the choice of Chris induces contradicting the optimality of serenade and chocolate.
- Therefore, there does not exist a GNE in which *you* can rationally play the moonlight serenade given *your* (unique) utility function.

Example: The Moonlight Serenade

- The game in one-person perspective form:

$\Gamma_y(u_y)$

	open	ignore
serenade	4	0
chocolate	0	4
Chris	3	3

$\Gamma_B(u_B^{an})$

	serenade	chocolate	Chris
open	0	0	0
ignore	1	1	1

$\Gamma_B(u_B^{for})$

	serenade	chocolate	Chris
open	0	1	0
ignore	1	0	1

- GISD and GNE give the same solution for *Barbara*:
  - $GISD_{Barbara} = \{(ignore, u_B^{an}), (open, u_B^{for}), (ignore, u_B^{for})\}$
  - *Her* rational choices under GNE are ignore only given u_B^{an} and ignore as well as open given u_B^{for}.

- GISD is strictly refined by GNE for *you* though:
  - $GISD_{you} = \{(serenade, u_y), (chocolate, u_y), (Chris, u_y)\}$
  - However, *you* can only rationally choose chocolate as well as Chris – but not serenade – under GNE given *your* (only) utility function u_y.