Epistemic Game Theory: Incomplete Information

Part II: Correct Beliefs

Christian W. Bach

University of Reading & EPICENTER Maastricht

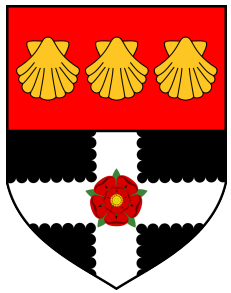

Correct Beliefs Assumption and Equilibrium

- 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 is the notion of simple belief hierarchy.
- In **Incomplete Information Part I** the notion of common belief in rationality has been generalized to incomplete information.
- In this part the notion of simple belief hierarchy is extended to games with incomplete information.
- It turns out that common belief in rationality and simple belief hierarchy are equivalent to a solution concept called Generalized Nash Equilibrium (which is a generalization of Nash Equilibrium).

Simple Belief Hierarchies

- With complete information a simple belief hierarchy (SBH) is generated by a tuple of probabiltiy distributions

$$(\sigma^i)_{i \in I} \in \times_{i \in I} \big(\Delta(C_i)\big).$$

- 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, 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:

$\Gamma_y^0(u_y)$

	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:
  - $T_{you} = \{t_y^1, t_y^2\}$ and $T_{Barbara} = \{t_B^1, t_B^2\}$,
  - $v_{you}[t_y^1] = v_{you}[t_y^2] = u_y$,
  - $v_{Barbara}[t_B^1] = u_B^r$ and $v_{Barbara}[t_B^2] = u_B^b$
  - $b_{you}[t_y^1] = (red, t_B^1)$ and $b_{you}[t_y^2] = (blue, t_B^2)$,
  - $b_{Barbara}[t_B^1] = (green, t_y^2)$ and $b_{Barbara}[t_B^2] = (blue, t_y^2)$.

- *Your* type t_y^1 believes that *Barbara* chooses *red* and has utility function u_B^r.
- Also, t_y^1 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:
  - $T_{you} = \{t_y^1, t_y^2\}$ and $T_{Barbara} = \{t_B^1, t_B^2\}$,
  - $v_{you}[t_y^1] = v_{you}[t_y^2] = u_y$,
  - $v_{Barbara}[t_B^1] = u_B^r$ and $v_{Barbara}[t_B^2] = u_B^b$
  - $b_{you}[t_y^1] = (red, t_B^1)$ and $b_{you}[t_y^2] = (blue, t_B^2)$,
  - $b_{Barbara}[t_B^1] = (green, t_y^2)$ and $b_{Barbara}[t_B^2] = (blue, t_y^2)$.

- *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 any belief of yours – i.e. about your entire belief hierarchy.

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:
  - $T_{you} = \{t_y^1, t_y^2\}$ and $T_{Barbara} = \{t_B^1, t_B^2\}$,
  - $v_{you}[t_y^1] = v_{you}[t_y^2] = u_y$,
  - $v_{Barbara}[t_B^1] = u_B^r$ and $v_{Barbara}[t_B^2] = u_B^b$
  - $b_{you}[t_y^1] = (red, t_B^1)$ and $b_{you}[t_y^2] = (blue, t_B^2)$,
  - $b_{Barbara}[t_B^1] = (green, t_y^2)$ and $b_{Barbara}[t_B^2] = (blue, t_y^2)$.

- The belief hierarchy induced by t_y^2 is completely generated by the two probability distributions: $\sigma^y = (blue, u_y)$ 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
- Fixed Beliefs on Utilities

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 $\mu_i \in \Delta\big(\times_{j \neq i} (C_j \times U_j)\big)$.
- Note that a first-order belief of a type $t_i \in T_i$ induces a conjecture of player i via his description map $b_i[t_i] \in \Delta\big((C_j \times T_j)_{j \in I \setminus \{i\}}\big)$
- A conjecture of i induces a marginal conjecture about every opponent $j \in I \setminus \{i\}$, that is $\text{marg}_{C_j \times U_j} \mu_i$.
- A marginal conjecture **about** player i is a belief about i's choice and utility function, denoted as $\sigma^i \in \Delta(C_i \times U_i)$.

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 1st-order belief: $\prod_{j \neq i} \sigma^j$,
- player i's 2nd-order belief: i believes that every opponent $j \neq i$ holds 1st-order belief $\prod_{k \neq i} \sigma^k$,
- player i's 3rd-order belief: i believes that every opponent $j \neq i$ believes that every opponent $k \neq j$ holds 1st-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,
- $v_i[t_i] = u_i$,
- c_i is optimal for t_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 $T_{Barbara} = \{t_B^1, t_B^2\}$,
  - $v_{you}[t_y] = u_y$,
  - $v_{Barbara}[t_B^1] = u_B^r$ and $v_{Barbara}[t_B^2] = u_B^b$
  - $b_{you}[t_y] = 0.5(red, t_B^1) + 0.5(blue, t_B^2)$,
  - $b_{Barbara}[t_B^1] = b_{Barbara}[t_B^2] = (green, t_y)$.
- 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: *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 and that green is optimal for t_y.

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:

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

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

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

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

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

for all $c_i' \in C_i$.

Example: What is Barbara's favourite Colour?

- The game in one-person perspective form:

$\Gamma_y^0(u_y)$

	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

- Consider the two (marginal) conjectures $\sigma^y = (green, u_y)$ and $\sigma^B = 0.5 \cdot (red, u_B^r) + 0.5 \cdot (blue, u_B^b)$.

- σ^y only assigns positive probability to *green*.
- Observe that *your* choice of *green* is optimal for u_y given the marginal belief $0.5 \cdot red + 0.5 \cdot blue$ about *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_B^r given the marginal belief $green$ about *your* choices as well as that *Barbara*'s choice of *blue* is optimal for u_B^b given the marginal belief $green$ about *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_B^r 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 pairs (c_k, t_k) such that c_k is optimal for $u_k[t_k]$.
- Since t_i's belief hierarchy is generated by $(\sigma^j)_{j \in I}$, the marginal conjecture of t_i 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

$$\sum_{c_{-k} \in C_{-k}} \prod_{j \in I \setminus \{k\}} \text{marg}_{C_j} \sigma^j(c_j) \cdot u_k[t_k](c_k, c_{-k})$$

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

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 pairs (c_i, t_i') such that c_i is optimal for $u_i[t_i']$.
- 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 $\prod_{j \neq i} \operatorname{marg}_{C^j} \sigma_j$ as belief about C_{-i} .
- It follows that, for all $(c_i, u_i) \in \operatorname{supp}(\sigma^i)$ it is the case that

$$\sum_{c_{-i} \in C_{-i}} \prod_{j \in I \setminus \{i\}} \operatorname{marg}_{C_j} \sigma^j(c_j) \cdot u_i[t_i'](c_i, c_{-i}) \geq \sum_{c_{-i} \in C_{-i}} \prod_{j \in I \setminus \{i\}} \operatorname{marg}_{C_j} \sigma^j(c_j) \cdot u_i[t_i'](c_i', c_{-i})$$

  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

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

  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 favourite 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

$\Gamma_B(u_B^{an})$

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

$\Gamma_B(u_B^{for})$

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

Fixed Beliefs on Utilities

Restricting the Reasoning about Payoff Uncertainty

- Possibly, some beliefs about the opponents' preferences seem more plausible than others.
- This can be captured by fixing the players' beliefs about their respective opponents' utility functions in their reasoning.
- Accordingly, a player believes that there are unique beliefs about the players' preferences that are transparent to everyone.
- What happens if a player expresses common belief in rationality and holds a simple belief hierarchy as well as reasons in line with fixed beliefs on utilities?

Fixed Beliefs on Utilities

Definition 9

Let Γ be a game with incomplete information, $\mathcal{M}^\Gamma$ an epistemic model of Γ, $p = (p^i)_{i \in I} \in \times_{i \in I} \big(\Delta(U_i)\big)$ a tuple of probability distributions, $i \in I$ some player, and $t_i \in T_i$ some type of player i.

- The type t_i expresses *1-fold belief in* p, if t_i assigns to every tuple $u_{-i} \in \times_{j \in I \setminus \{i\}} U_j$ the probability $\prod_{j \in I \setminus \{i\}} p^j(u_j)$.
- Let $k > 1$. The type t_i expresses *k-fold belief in* p, if t_i only assigns positive probability to opponents' types that express *(k-1)*-fold belief in p.
- The type t_i expresses *common belief in* p, if t_i expresses k-fold belief in p for all $k \geq 1$.

In other words, t_i's (marginal) belief hierarchy on utility functions is the simple belief hierarchy generated by some fixed profile p of probability distributions.

Rational Choice under CBR, SBH, and CBp

Definition 10

Let Γ be a game with incomplete information, $p = (p^i)_{i \in I} \in \times_{i \in I} \big(\Delta(U_i)\big)$ a tuple of probability distributions, $i \in I$ 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, a simple belief hierarchy, and common belief in p* 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,
- t_i expresses common belief in p,
- $v_i[t_i] = u_i$,
- c_i is optimal for t_i.

Example: What is Barbara's favourite Colour?

- Just before deciding on what colour to wear at tonight's party, you recall that Barbara drives a red car and also owns a red bike.
- Consequently, you deem it much more likely that Barbara's most preferred colour is red.
- More precisely, you are 80% confident that her utility function is u_B^r (*"red-peaked"-preferences*).
- You believe that Barbara believes this; you believe that Barbara believes that you believe this; etc.
- In other words, you express common belief in p, where

$$p = (1u_y, 0.8u_B^r + 0.2u_B^b)$$

Example: What is Barbara's favourite Colour?

- If you hold belief

$$0.8u_B^r + 0.2u_B^b$$

about Barbara's utilities and express CBR, then you must believe Barbara to choose red with probability 0.8 and blue or yellow with probabiltiy 0.2.

- Then, blue yields at least 3.2 and green precisely 3 as expected payoffs for you.

- Therefore, the only choice you can **rationally** make under **common belief in rationality** and **common belief in** p is blue.

- Your choice blue can also be supported by a belief hierarchy that on top of satisfying CBR and CBp qualifies as simple.

Example: What is Barbara's favourite Colour?

$\Gamma_y^0(u_y)$

	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 $T_{Barbara} = \{t_B, \hat{t}_B\}$,
  - $v_{you}[t_y] = u_y$,
  - $v_{Barbara}[t_B] = u_B^b$ and $v_{Barbara}[\hat{t}_B] = u_B^r$
  - $b_{you}[t_y] = 0.8(red, \hat{t}_B) + 0.2(yellow, t_B)$,
  - $b_{Barbara}[t_B] = b_{Barbara}[\hat{t}_B] = (blue, t_y)$.
- All types believe in rationality and believe in $p = (1u_y, 0.8u_B^r + 0.2u_B^b)$, and thus also express common belief in rationality and common belief in p.
- The belief hierarchy of t_y is generated by the pair of marginal conjectures

$$(\sigma^Y, \sigma^B) = \big((blue, u_y), 0.8(red, u_B^r) + 0.2(yellow, u_B^b)\big),$$

and therefore simple.