Chapter 6

Correct and Symmetric Beliefs with Incomplete Information

In Chapter 4 we have introduced the ideas of *correct* and *symmetric* beliefs for standard games, where the players are certain about the conditional preference relations of their opponents. Recall that we have formalized these ideas by the notions of *simple* and *symmetric belief hierarchies.* Together with common belief in rationality, these two restrictions on belief hierarchies led to the concepts of *Nash equilibrium* and *correlated equilibrium,* respectively.

In this chapter we will extend this analysis to games with incomplete information. We will start by defining simple belief hierarchies for the case of incomplete information, and show that this restriction, in combination with common belief in rationality, leads to a concept called *generalized Nash equilibrium.* Similarly, we define symmetric belief hierarchies for the case of incomplete information, and show that it will lead to the concept of *Bayesian equilibrium* when combined with common belief in rationality. Finally, we investigate the two concepts in the light of fixed beliefs on the players' utility functions. In Chapter 6 of the online appendix we study some economic applications.

6.1 Correct Beliefs

Like we did for standard games, we start by defining *simple belief hierarchies,* and explain why it describes situations where you believe the opponents to be *correct* about your beliefs. Subsequently, we merge this condition with *common belief in rationality,* and show that it leads to a concept called *generalized Nash equilibrium.* As the name suggests, it reduces to traditional Nash equilibrium if we apply it to games without incomplete information.

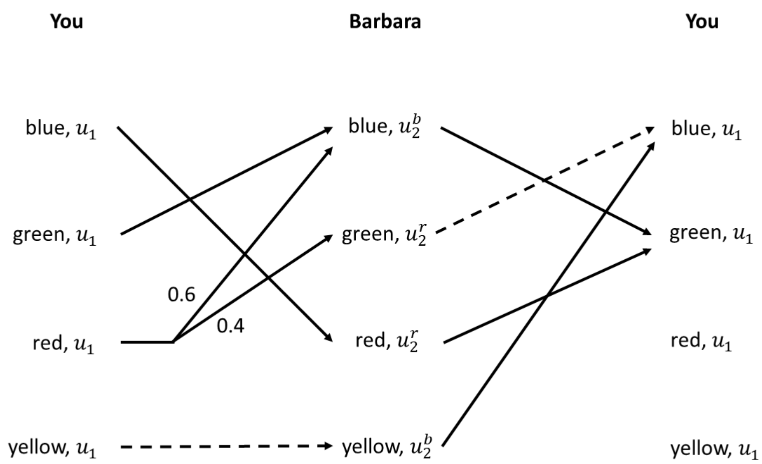

Figure 6.1.1 A beliefs diagram for "What is Barbara's favorite color?"

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

u_1

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

u_2^r

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

u_2^b

Table 6.1.1 Decision problems for "What is Barbara's favorite color?"

6.1.1 Simple Belief Hierarchies

Recall the definition of a simple belief hierarchy for standard games, in Section 4.1.1. It stated that your belief hierarchy is entirely induced by a single belief σ_1 about player 1's choice, a single belief σ_2 about player 2's choice, and so on. An important feature of such a belief hierarchy is that you believe that your opponents are *correct* about all the beliefs you hold. Moreover, if there are three or more players in the game, then you believe that opponent j's belief about a third player k will be the same as your own belief about player k, and that your belief about j's choice will be independent from your belief about k's choice.

If we turn to games with *incomplete information,* where you may be uncertain about the utility functions of some of the opponents, then a belief hierarchy does not only concern the players' choices but also the players' utility functions. To refresh your memory, consider again the beliefs diagram in Figure 5.2.1 for the example "What is Barbara's favorite color?". For easier reference, we have reproduced this beliefs diagram in Figure 6.1.1. Also, we reproduce the utilities for this example in Table 6.1.1, to make it easier for the reader.

In the beliefs diagram, consider the belief hierarchy that starts at your choice-utility pair (*blue,* u_1). It states that you believe that Barbara chooses *red* while having the utility function u_2^r. At the same time, you believe that Barbara believes that you believe that Barbara chooses *blue* while having the utility function u_2^b. As such, you believe that Barbara is incorrect about your first-order belief.

Compare this to the belief hierarchy that starts at your choice *green*: Not only do you believe that Barbara chooses *blue* while having the utility function u_2^b, you also believe that Barbara believes that you indeed believe this. Even more, you believe that Barbara believes that you indeed have the belief hierarchy that starts at your choice *green.* In other words, you believe that Barbara is correct about your *entire belief hierarchy.*

In fact, your belief hierarchy that starts at your choice *green* is completely generated by the single belief $\sigma_1 =$ (*green,* u_1) about you, and the single belief $\sigma_2 =$ (*blue,* u_2^b). Here, $\sigma_1 =$ (*green,* u_1) is the belief about your choice-utility pairs that assigns probability 1 to you wearing *green* while having the utility function u_1, whereas $\sigma_2 =$ (*blue,* u_2^b) is the belief about Barbara's choice-utility pairs that assigns probability 1 to Barbara wearing *blue* while having the utility function u_2^b. Indeed, in the belief hierarchy that starts at your choice *green* your belief about Barbara's choice and utility function is σ_2, you believe that Barbara's belief about your choice and utility function is σ_1, you believe that Barbara believes that your belief about Barbara's choice and utility function is σ_2, and so on. We say that your belief hierarchy is *simple,* and that it is *generated* by the combination of beliefs (σ_1, σ_2).

In general, simple belief hierarchies can be defined as follows.

Definition 6.1.1 (Simple belief hierarchy) *Let σ_1 be a probabilistic belief about player 1's choice and utility function, σ_2 a probabilistic belief about player 2's choice and utility function, and so on, until σ_n being a probabilistic belief about player n's choice and utility function. The belief hierarchy for player i generated by the beliefs $(\sigma_1, \sigma_2, ..., \sigma_n)$ is defined as follows:*

(1) in the first-order belief, player i assigns to every opponents' choice-utility combination $(c_j, u_j)_{j\neq i}$ the probability $\prod_{j\neq i} \sigma_j(c_j, u_j)$,

(2) in the second-order belief, player i believes with probability 1 that every opponent j assigns to every opponents' choice-utility combination $(c_k, u_k)_{k\neq j}$ the probability $\prod_{k\neq j} \sigma_k(c_k, u_k)$,

(3) in the third-order belief, player i believes with probability 1 that every opponent j believes with probability 1 that every opponent k assigns to every opponents' choice-utility combination $(c_l, u_l)_{l\neq k}$ the probability $\prod_{l\neq k} \sigma_l(c_l, u_l)$, and so on.

A belief hierarchy is called ***simple*** *if it is generated by a combinations of such beliefs $(\sigma_1, \sigma_2, ..., \sigma_n)$.*

The only difference with the case of standard games is thus that the beliefs now concern choices *and utility functions,* instead of only choices. The rest of the definition is exactly the same.

Question 6.1.1 *Consider the beliefs diagram in Figure 6.1.1. Which of Barbara's belief hierarchies is simple?*

In a sense, the simple belief hierarchy for you in Figure 6.1.1 was very special, because it reflects no uncertainty about Barbara's utility function. Indeed, in that simple belief hierarchy you are absolutely convinced that Barbara's utility function is u_2^b and no other, and you believe that this is transparent among Barbara and you.

However, there are also simple belief hierarchies that express inherent uncertainty about the opponents' utility functions. As an example, consider the beliefs diagram in Figure 6.1.2. Your belief hierarchy is simple, because it is generated by the combination of beliefs $\sigma_1 =$ (*green,* u_1) and

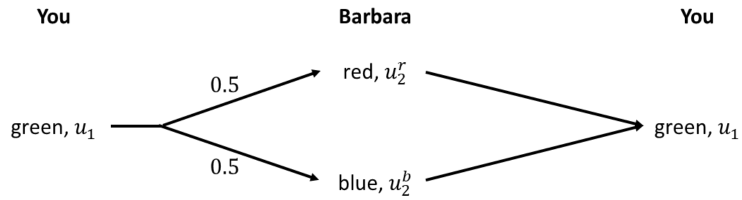

Figure 6.1.2 Simple belief hierarchy for "What is Barbara's favorite color?" with uncertainty about opponent's utility function

$\sigma_2 = (0.5) \cdot (red,\ u_2^r) + (0.5) \cdot (blue,\ u_2^b)$. In that belief hierarchy, you assign probability 0.5 to Barbara's utility functions u_2^r and u_2^b, which means that you are inherently uncertain about Barbara's conditional preference relation. Moreover, you believe this uncertainty to be transparent between Barbara and you. It may also be verified that this simple belief hierarchy expresses common belief in rationality.

Question 6.1.2 *Consider the example "What is Barbara's favorite color?". Construct a simple belief hierarchy for you that expresses common belief in rationality, and where you assign probability* 0.9 *to Barbara's utility function* u_2^r *and probability* 0.1 *to Barbara's utility function* u_2^b.

6.1.2 Relation with Generalized Nash Equilibrium

Suppose we combine the condition of a simple belief hierarchy with the conditions in common belief in rationality. What choices can you then rationally make? Recall that in the case of standard games, the resulting choices were those that are optimal in a *Nash equilibrium.* We will see that something similar will be true for games with incomplete information, if we replace Nash equilibrium by *generalized Nash equilibrium.*

Let us start from a simple belief hierarchy for player i generated by a combination of beliefs $(\sigma_1, ..., \sigma_n)$ where, for every player j, the belief σ_j is a probability distribution over j's choice-utility pairs. Assume, in addition, that this belief hierarchy expresses common belief in rationality. What properties should $(\sigma_1, ..., \sigma_n)$ have?

Fix an opponent j. As i believes in j's rationality, i's belief σ_j about opponent j's choice-utility pairs should only assign positive probability to pairs (c_j, u_j) where c_j is optimal for player j, given the utility function u_j, and given what i believes to be j's belief about the other players' choices. By construction, i's belief about j's belief about the choices of the other players is given by σ_{-j}, where $\sigma_{-j} = (\sigma_k)_{k \neq j}$ is the collection of beliefs about all players except j. By putting these two insights together, we conclude that σ_j should only assign positive probability to (c_j, u_j) if choice c_j is optimal for opponent j given the utility function u_j, and given the belief σ_{-j} about the other players.

Now, consider the belief σ_i about i's choice-utility pairs from i's viewpoint. By construction, σ_i is what i believes that an opponent believes about i's (that is, his own) choice-utility pairs. As i believes that every opponent believes in i's rationality, σ_i should only assign positive probability to (c_i, u_i) if c_i is optimal for player i given the utility function u_i, and given what i believes that the opponents believe about i's belief about the other players' choices. Again, by construction, i believes that the opponents believe that i's belief about the other players is given by σ_{-i}. By putting these two insights together, we conclude that σ_i should only assign positive probability to (c_i, u_i) if choice c_i is optimal for player i given the utility function u_i, and given the belief about the other players' choices induced by σ_{-i}.

Altogether, we see that if the simple belief hierarchy generated by $(\sigma_1, .., \sigma_n)$ expresses common belief in rationality, then for every player j, the belief σ_j only assigns positive probability to a choice-utility pair (c_j, u_j) if the choice c_j is optimal for player j given the utility function u_j, and given the belief about his opponents' choices induced by σ_{-j} . Belief combinations $(\sigma_1, ..., \sigma_n)$ with this property are called *generalized Nash equilibria.*

Definition 6.1.2 (Generalized Nash equilibrium) *Consider a combination $(\sigma_1, ..., \sigma_n)$ of beliefs where, for every player i, the belief σ_i is a probability distribution over i's choice-utility function pairs. The combination $(\sigma_1, ..., \sigma_n)$ is a* ***generalized Nash equilibrium*** *if for every player i, the belief σ_i only assigns positive probability to choice-utility pairs (c_i, u_i) where c_i is optimal for player i given the utility function u_i, and given the belief about the opponents' choices induced by σ_{-i} .*

Above we have thus seen that, if the simple belief hierarchy is generated by $(\sigma_1, ..., \sigma_n)$ and expresses common belief in rationality, then $(\sigma_1, ..., \sigma_n)$ is a generalized Nash equilibrium. In fact, the other direction is also true: If a simple belief hierarchy is generated by a generalized Nash equilibrium $(\sigma_1, ..., \sigma_n)$, then the belief hierarchy will express common belief in rationality. We thus arrive at the following general conclusion.

Theorem 6.1.1 (Relation with generalized Nash equilibrium) *Consider a simple belief hierarchy generated by a combination of beliefs $(\sigma_1, ..., \sigma_n)$ about choice-utility pairs. Then, the simple belief hierarchy expresses common belief in rationality, if and only if, $(\sigma_1, ..., \sigma_n)$ is a generalized Nash equilibrium.*

As an illustration, consider the beliefs diagram in Figure 6.1.2. As we have seen before, your belief hierarchy is simple and is generated by the combination of beliefs $\sigma_1 = (\textit{green}, u_1)$ and $\sigma_2 = (0.5)\cdot(\textit{red}, u_2^r) + (0.5)\cdot(\textit{blue}, u_2^b)$. In fact, it can be shown that this combination of beliefs is a generalized Nash equilibrium.

To see this, note that σ_1 assigns probability 1 to your choice-utility pair $(\textit{green}, u_1)$. The belief σ_2 assigns probability 0.5 to Barbara's choices *red* and *blue,* and under this belief it is optimal for you to wear *green* if your utility function is u_1. Thus, the associated optimality condition for player 1 in the definition of a generalized Nash equilibrium is satisfied.

On the other hand, σ_2 assigns positive probability to Barbara's choice-utility pairs $(\textit{red}, u_2^r)$ and $(\textit{blue}, u_2^b)$. If Barbara holds the belief σ_1, then she believes that you will wear *green* with probability 1. Under that belief, it would be optimal for Barbara to wear *red* if her utility function is u_2^r, whereas it would be optimal for her to wear *blue* if her utility function is u_2^b. Hence, the optimality condition for player 2 in generalized Nash equilibrium is again satisfied. Altogether, we conclude that that (σ_1, σ_2) is a generalized Nash equilibrium. By Theorem 6.1.1 it then follows that your simple belief hierarchy generated by (σ_1, σ_2) expresses common belief in rationality – something we already concluded above based on the beliefs diagram.

With Theorem 6.1.1 at hand, it is now also clear what are the choices that player i can rationally make if he holds a *simple* belief hierarchy that expresses common belief in rationality, and a utility function u_i. These are precisely the choices that are optimal for player i in a generalized Nash equilibrium $(\sigma_1, ..., \sigma_n)$ if he holds the utility function u_i.

To see this, suppose that player i holds a simple belief hierarchy that expresses common belief in rationality, and that a choice c_i is optimal for him under the utility function u_i. By Theorem 6.1.1 we know that the simple belief hierarchy must be induced by a generalized Nash equilibrium $(\sigma_1, ..., \sigma_n)$. In particular, player i holds the belief σ_{-i} about the opponents' choices. As such, the choice c_i must be optimal, given the belief σ_{-i} which is part of the generalized Nash equilibrium, and given the utility function u_i.

On the other hand, suppose that the choice c_i is optimal in a generalized Nash equilibrium $(\sigma_1, ..., \sigma_n)$ if player i holds the utility function u_i. We can then use Theorem 6.1.1 to conclude that the simple belief hierarchy induced by $(\sigma_1, ..., \sigma_n)$ expresses common belief in rationality. Thus, the choice c_i is optimal for player i if he holds a simple belief hierarchy that expresses common belief in rationality, and holds the utility function u_i. We therefore arrive at the following conclusion.

Theorem 6.1.2 (Choices optimal in a generalized Nash equilibrium) *For player i, a choice c_i is optimal for a utility function u_i and a simple belief hierarchy that expresses common belief in rationality, if and only if, c_i is optimal in a generalized Nash equilibrium $(\sigma_1, ..., \sigma_n)$ for the utility function u_i.*

Hence, if we want to find all choices that are possible if a player holds a *simple* belief hierarchy that expresses common belief in rationality, then we must concentrate on the generalized Nash equilibria in the game.

From Chapter 4 we know that a Nash equilibrium always exists for every standard game (without incomplete information) that contains finitely many choices. With this insight at hand, it then easily follows that a generalized Nash equilibrium will also always exist for every game with incomplete information that contains finitely many choices. The reason is simple: Fix a utility function u_i in U_i for every player i, and consider the standard game $\tilde{\Gamma}$ where every player i is believed to have this particular utility function u_i. From Chapter 4 we know that there is a Nash equilibrium $(\tilde{\sigma}_1, ..., \tilde{\sigma}_n)$ for $\tilde{\Gamma}$, where $\tilde{\sigma}_i$ is a probability distribution over i's choices. We can then construct a combination of beliefs $(\sigma_1, ..., \sigma_n)$ for the game with incomplete information, where for every player i and every choice c_i, the belief σ_i assigns probability $\tilde{\sigma}_i(c_i)$ to the choice-utility pair (c_i, u_i), with the distinguished utility function u_i.

Question 6.1.3 *Show that the combination of beliefs $(\sigma_1, ..., \sigma_n)$ is a generalized Nash equilibrium.*

In this way, we can always construct a generalized Nash equilibrium in which you are certain about the opponents' utility functions, and where you believe that the opponents are certain about your own utility function. We thus obtain the following general existence result.

Theorem 6.1.3 (Generalized Nash equilibrium always exists) *For every game with incomplete information and finitely many choices, there is always at least one generalized Nash equilibrium.*

Together with Theorem 6.1.1 we thus conclude that it is always possible to hold a simple belief hierarchy that expresses common belief in rationality.

6.1.3 Examples

We will now determine, for the two examples we have investigated in Chapter 5, which choices you can rationally make with a *simple* belief hierarchy that expresses common belief in rationality.

Example 6.1: What is Barbara's favorite color.

Recall the story from Chapter 5, and the decision problems in Table 6.1.1. In Section 5.4.3 of that chapter, we have seen that under common belief in rationality you can rationally wear *blue* or *green,*

You	20	40	60	80	100
20	5	0	0	0	0
40	−10	−5	0	0	0
60	−30	−30	−15	0	0
80	−50	−50	−50	−25	0
100	−70	−70	−70	−70	−35

u_1^{30}

You	20	40	60	80	100
20	15	0	0	0	0
40	10	5	0	0	0
60	−10	−10	−5	0	0
80	−30	−30	−30	−15	0
100	−50	−50	−50	−50	−25

u_1^{50}

You	20	40	60	80	100
20	25	0	0	0	0
40	30	15	0	0	0
60	10	10	5	0	0
80	−10	−10	−10	−5	0
100	−30	−30	−30	−30	−15

u_1^{70}

You	20	40	60	80	100
20	35	0	0	0	0
40	50	25	0	0	0
60	30	30	15	0	0
80	10	10	10	5	0
100	−10	−10	−10	−10	−5

u_1^{90}

Table 6.1.2 Decision problems for "Chris' drawings"

Barbara can only rationally wear *red* if her utility function is u_2^r, and she can rationally wear *blue* or *yellow* if her utility function is u_2^b. But which colors can you and Barbara rationally choose, for each of the possible utility functions, if you hold a *simple* belief hierarchy that expresses common belief in rationality. To answer this question we will rely on Theorem 6.1.2.

Consider the generalized Nash equilibrium $\sigma_1 = (green,\ u_1)$ and $\sigma_2 = (0.5)\cdot(red,\ u_2^r) + (0.5)\cdot(blue,\ u_2^b)$ we have investigated above. In that generalized Nash equilibrium your optimal choice is *green,* the optimal choice for Barbara is *red* if her utility function is u_2^r, whereas her optimal choice is *blue* if her utility function is u_2^b. On the basis of Theorem 6.1.2 we can thus conclude that under common belief in rationality with a *simple* belief hierarchy, you can rationally wear *green,* Barbara can rationally wear *red* if her utility function is u_2^r, whereas she can rationally wear *blue* if her utility function is u_2^b.

Consider next the belief combination (σ_1, σ_2) where $\sigma_1 = (blue,\ u_1)$ and $\sigma_2 = (yellow,\ u_2^b)$.

Question 6.1.4 *Explain why* (σ_1, σ_2) *is a generalized Nash equilibrium.*

Note that for you, *blue* is optimal for the belief σ_2 and the utility function u_1, and for Barbara, *yellow* is optimal for the belief σ_1 and the utility function u_2^b. In the light of Theorem 6.1.2 we thus see that under common belief in rationality with a *simple* belief hierarchy, you can rationally wear *blue,* whereas Barbara can rationally wear *yellow* if her utility function is u_2^b.

Overall, we conclude that under common belief in rationality with a *simple* belief hierarchy, you can rationally wear *blue* and *green,* Barbara can rationally wear *red* if her utility function is u_2^r, whereas she can rationally wear *blue* and *yellow* if her utility function is u_2^b. As these are precisely the choices that were possible under common belief in rationality, we see that the additional condition of a simple belief hierarchy does not further restrict the choices that you and Barbara can rationally make.

Example 6.2: Chris' drawings.

Recall the story from Chapter 5. For convenience, we have reproduced the decision problems in Table 6.1.2. We have seen in Section 5.4.4 of that chapter that under common belief in rationality, you can rationally bid 20, 40 or 60 if your valuation is 30, 50 or 70, whereas you can rationally bid 40, 60 or 80 if your valuation is 90, and similarly for Barbara.

Suppose we additionally insist on a *simple* belief hierarchy. What bids could you then rationally make for every possible valuation? Again, we rely on Theorem 6.1.2.

Consider first the combination of beliefs (σ_1, σ_2) where $\sigma_1 = (80,\, u_1^{90})$ and $\sigma_2 = (80,\, u_2^{90})$.

Question 6.1.5 *Explain why* (σ_1, σ_2) *is a generalized Nash equilibrium.*

Under the belief σ_2, you can rationally bid 20, 40 or 60 if your valuation is 30, 50 or 70, and you can rationally bid 80 if your valuation is 90. Thus, by Theorem 6.1.2, we see that under common belief in rationality with a simple belief hierarchy, you can rationally bid 20, 40 or 60 if your valuation is 30, 50 or 70, and you can rationally bid 80 if your valuation is 90.

Question 6.1.6 *Find a generalized Nash equilibrium where it is optimal for you to bid* 40 *if your valuation is* 90, *and another generalized Nash equilibrium where it is optimal for you to bid* 60 *if your valuation is* 90.

In the light of Question 6.1.6 and Theorem 6.1.2, we can thus conclude that under common belief in rationality with a simple belief hierarchy, you can rationally bid 40 or 60 if your valuation is 90. Altogether, we see that under common belief in rationality with a simple belief hierarchy, you can rationally bid 20, 40 or 60 if your valuation is 30, 50 or 70, and you can rationally bid 40, 60 or 80 if your valuation is 90.

These were exactly the bids you could rationally make under common belief in rationality. Therefore, also in this example the additional restriction of a simple belief hierarchy does not affect the choices you can rationaly make under common belief in rationality.

We will end this section by investigating a new example in which the additional condition of a simple belief hierarchy *does* alter the possible choices that you can rationally make under common belief in rationality.

Example 6.3: The moonlight serenade.

Recall from the example "Movie for two" in Chapter 4 that you had a fight with Barbara, some days ago. The plan to see each other at the cinema did not work out, and hence you have to think about a new strategy to make up with her. There is a full moon this evening – the perfect time to apologize to Barbara, and to let this fight behind you. Three possible plans come to your mind: You can either give a moonlight *serenade* in front of her door, or bring her a box of her favorite *chocolates,* or send Chris to apologize for you. The question is: When you ring the door bell, will Barbara *open* the door, or will she *ignore* the door bell?

Your conditional preference relation is as follows: If you believe that Barbara will open the door, then you would definitely prefer to give a serenade with your Spanish guitar. At the same time, you would be very disappointed to only arrive with a box of chocolates in this case, as you could have impressed her with a lovely Spanish song. As such, you would rather send Chris than offering a box of chocolates in this case.

On the other hand, if you believe that Barbara will ignore the door bell, you would be very disappointed if you would stand there with your Spanish guitar. For that reason, sending Chris would be better than intending to give a serenade. However, the best option in this case is to arrive with a box of chocolates, since you could still put it in her briefcase with a letter attached to it.

Finally, if you are inherently uncertain about Barbara opening the door or not, then your favorite plan would be to send Chris, as to avoid any big disappointment.

You	*open*	*ignore*
serenade	4	0
chocolates	0	4
Chris	3	3

u_1

Barbara	*serenade*	*chocolates*	*Chris*
open	0	0	0
ignore	1	1	1

u_2^a

Barbara	*serenade*	*chocolates*	*Chris*
open	0	1	0
ignore	1	0	1

u_2^f

Table 6.1.3 Decision problems for "The moonlight serenade"

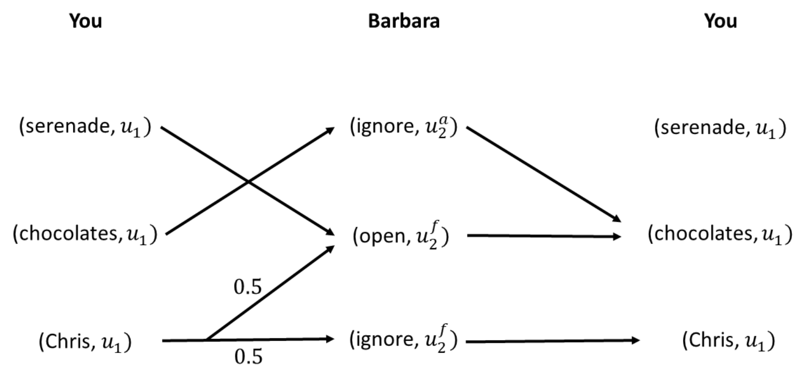

Figure 6.1.3 Beliefs diagram for "The moonlight serenade"

This evening, you are also not certain whether Barbara will be in an *angry* or a *forgiving* mood, and this may well affect her decision to open the door or not. More precisely, if Barbara is in an angry mood then she will simply prefer to ignore the door bell in any case. If she is in a forgiving mood, she would only consider opening the door if she deems it sufficiently likely that you bring her a box of chocolates. The reason is that you are a terrible singer, and therefore Barbara would rather not open the door if she thinks you are about to give a serenade. On the other hand, if she believes you sent Chris, then she will interpret this as an act of cowardice, not worthy of opening the door.

This story can be modelled by the decision problems in Table 6.1.3. Here, u_2^a and u_2^f denote the utility function for Barbara in case she is *angry* and *forgiving,* respectively.

Let us first investigate which plans you can rationally implement under common belief in rationality. If we apply the *generalized iterated strict dominance* procedure from Section 5.4, then we can only eliminate the choice *open* for Barbara at u_2^a, after which the procedure stops. Thus, under common belief in rationality, you can rationally implement any of the three plans, whereas Barbara can only rationally *ignore* the door bell if she is *angry*, and can rationally *open* the door or *ignore* the door bell in case she is *forgiving.* This conclusion is also confirmed by the beliefs diagram in Figure 6.1.3, where each of these choices is supported by a belief hierarchy that expresses common belief in rationality.

Suppose that we now additionally insist on a simple belief hierarchy. Which plans can you rationally implement then? From the beliefs diagram in Figure 6.1.3 it can be seen that your belief hierarchy starting at (*chocolates,* u_1) is simple and expresses common belief in rationality. Indeed, it is generated by the beliefs $\sigma_1 =$ (*chocolates,* u_1) and $\sigma_2 =$ (*ignore,* u_2^a). As your choice *chocolates* is optimal under this simple belief hierarchy, you can rationally plan to bring *chocolates* under common belief in rationality with a simple belief hierarchy.

What about your other two choices? Consider the combination of beliefs $\sigma_1 = (0.75) \cdot$ (*chocolates,* u_1) $+ (0.25) \cdot$ (*Chris,* u_1) and $\sigma_2 = (0.25) \cdot$ (*open,* u_2^f) $+ (0.75) \cdot$ (*ignore,* u_2^a).

Question 6.1.7 *Explain why* (σ_1, σ_2) *is a generalized Nash equilibrium.*

Since your choice *Chris,* together with your choice *chocolates,* is optimal in the generalized Nash equilibrium (σ_1, σ_2), it follows by Theorem 6.1.2 that you can rationally plan to send *Chris* under common belief in rationality with a simple belief hierarchy.

Question 6.1.8 *Explain why there is no generalized Nash equilibrium* (σ_1, σ_2) *where only your choice Chris is optimal.*

The question above thus explains why, in order to support the choice *Chris,* we had to construct a generalized Nash equilibrium where *Chris* was optimal together with some other choice.

This leaves the question: Can you rationally plan to give a *serenade* under common belief in rationality with a *simple* belief hierarchy? We will show that the answer is "no". To see this, we prove that there is no generalized Nash equilibrium (σ_1, σ_2) where the choice *serenade* is optimal for you.

Assume, contrary to what we want to show, that *serenade* would be optimal for you in a generalized Nash equilibrium (σ_1, σ_2). Then, σ_2 must assign positive probability to Barbara's choice *open.* As *open* can only be optimal for Barbara if her utility function is u_2^f, we conclude that σ_2 must assign positive probability to (*open,* u_2^f). This implies, in turn, that *open* must be optimal for Barbara under the belief σ_1 if her utility function is u_2^f. To make this possible, σ_1 must assign positive probability to *chocolates.* Hence, *chocolates* must be optimal for you under the belief σ_2. As such, we conclude that both *serenade* and *chocolates* must be optimal for you under the belief σ_2, which is only possible if σ_2 assigns probability 0.5 to both *open* and *ignore.* But then, both *serenade* and *chocolates* would yield an expected utility of 2, which is less than what *Chris* gives. As such, it cannot be that both *serenade* and *chocolates* are optimal under σ_2. We thus obtain a contradiction. Therefore, we conclude that there is no generalized Nash equilibrium where the choice *serenade* is optimal for you. By Theorem 6.1.2 it then follows that you cannot rationally plan to give a *serenade* under common belief in rationality with a simple belief hierarchy.

Here is the intuitive reason: Planning a *serenade* can only be optimal if you assign a high probability to Barbara opening the door. If you believe that Barbara is correct about your belief, as is the case in a simple belief hierarchy, you believe that Barbara believes that you indeed assign a high probability to Barbara opening the door. Hence, you must believe that Barbara believes that you will plan a *serenade,* or possibly plan to send *Chris.* In either case, you believe Barbara to *ignore* the door bell. But then, it cannot be optimal anymore to plan a *serenade.*

What about Barbara's choices? What choices can she rationally make, if she is *angry* or *forgiving,* under common belief in rationality with a simple belief hierarchy? Clearly, if she is angry she can only rationally choose *ignore.* For the case where she is forgiving, consider the generalized Nash equilibrium $\sigma_1 = (0.75) \cdot$ (*chocolates,* u_1) $+ (0.25) \cdot$ (*Chris,* u_1) and $\sigma_2 = (0.25) \cdot$ (*open,* u_2^f) $+ (0.75) \cdot$ (*ignore,* u_2^a) we have seen above. Since in this generalized Nash equilibrium it is optimal for Barbara to choose *open* if she is forgiving, we conclude on the basis of Theorem 6.1.2 that Barbara can rationally choose to *open* the door if she is *forgiving* under common belief in rationality with a simple belief hierarchy.

But can she rationally *ignore* the door bell if she is *forgiving* in this case? The answer is "yes". To see why, consider the combination of beliefs $\sigma_1 = (0.5) \cdot (\textit{chocolates, } u_1) + (0.5) \cdot (\textit{Chris, } u_1)$ and $\sigma_2 = (0.25) \cdot (\textit{open, } u_2^f) + (0.75) \cdot (\textit{ignore, } u_2^f)$.

Question 6.1.9 *Explain why* (σ_1, σ_2) *is a generalized Nash equilibrium.*

In this generalized Nash equilibrium it is optimal for Barbara to choose *ignore* if her utility function is u_2^f. Therefore, by Theorem 6.1.2, Barbara can rationally choose to *ignore* the door bell if she is *forgiving* under common belief in rationality with a simple belief hierarchy.

Summarizing, we see that under common belief in rationality with a *simple* belief hierarchy, you can rationally plan to bring *chocolates* or to send *Chris,* but not to give a *serenade,* whereas Barbara can rationally choose to *ignore* the door bell if she is *angry,* and she can rationally choose to *open* the door or to *ignore* the door bell if she is *forgiving.* In particular, if we add the condition of a *simple* belief hierarchy to common belief in rationality, then you can no longer rationally plan a moonlight *serenade* for Barbara.

6.2 Symmetric Beliefs

In this section we extend the idea of a *symmetric* belief hierarchy, as introduced in Section 4.2, to games with incomplete information. Similarly to standard games it reveals a symmetry between what you believe about the opponent's choice and utility function on the one hand, and what you believe that the opponent believes about your own choice and utility function on the other hand. As we did in Section 4.2, we show that symmetric belief hierarchies can be characterized by *common priors* on choice-type combinations. Combining the condition of common belief in rationality with that of a symmetric belief hierarchy leads to the concept of *Bayesian equilibrium.* More precisely, the choices that can rationally be made, for a given utility function, under common belief in rationality and a symmetric belief hierarchy are precisely those that are optimal in a Bayesian equilibrium for that utility function. In that sense, *Bayesian equilibrium* can be viewed as the counterpart to *correlated equilibrium* when we move from standard games to games with incomplete information. Indeed, if we apply Bayesian equilibrium to a standard game without incomplete information, then we obtain exactly the concept of correlated equilibrium. At the end, we add the condition of *one theory per choice* and show that it leads to the concept of *canonical* Bayesian equilibrium, similarly to how it transformed correlated equilibrium into *canonical* correlated equilibrium in standard games.

6.2.1 Symmetric Belief Hierarchies

As already mentioned above, a *symmetric* belief hierarchy reflects a certain degree of symmetry between your belief about the opponent's choice and utility function, and what you believe that the opponent believes about your own choice and utility function. Like we did in Section 4.2 for standard games, this idea can be formalized by stating that the belief hierarchy can be derived from a *symmetric weighted* beliefs diagram.

To see what we mean by this in a game with incomplete information, consider the beliefs diagram, and an associated weighted beliefs diagram, for "The moonlight serenade" in Figure 6.2.1. In the

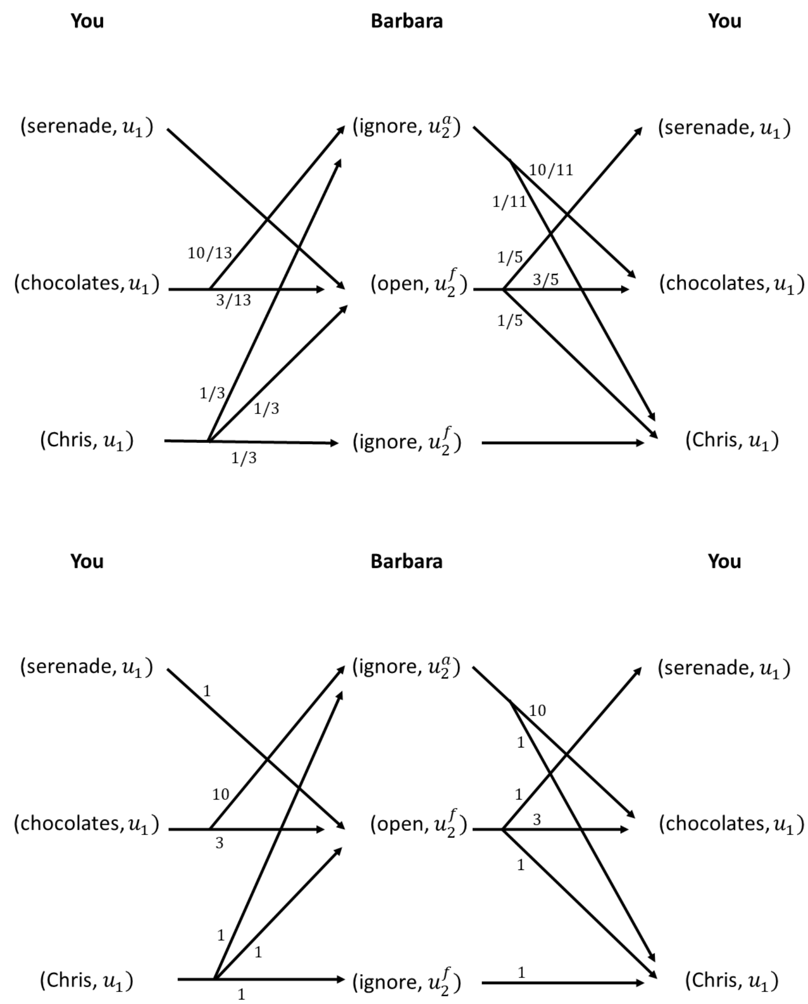

Figure 6.2.1 A beliefs diagram, and an associated weighted beliefs diagram, for "The moonlight serenade"

same way as in Section 4.2, it can be seen that the beliefs diagram on top of the figure is induced by the weighted beliefs diagram at the bottom. Consider, for instance, the two arrows that leave your choice *chocolates,* which carry the weights 10 and 3, respectively, in the weighted beliefs diagram. The relative weights are thus 10/13 and 3/13, which are exactly the associated probabilities in the beliefs diagram on top. In the same way, it can be verified that all the relative weights of the outgoing arrows in the weighted beliefs diagram correspond precisely to the probabilities of these outgoing arrows in the beliefs diagram on top.

Moreover, it turns out that the weighted beliefs diagram at the bottom is *symmetric*. Indeed, for every arrow from a choice-utility pair (c_1, u_1) of yours to a choice-utility pair (c_2, u_2) of Barbara, the symmetric counterpart, which is the arrow from (c_2, u_2) to (c_1, u_1), is also present. And *vice versa.* Moreover, every arrow always carries the same weight as its symmetric counterpart. Consider, for instance, the arrow from your choice-utility pair (*chocolate,* u_1) to Barbara's choice-utility pair (*ignore,* u_2^a), which carries the weight 10. Its symmetric counterpart, which is the arrow from (*ignore,* u_2^a) to (*chocolate,* u_1), carries the same weight 10. The same holds for all the arrows in the weighted beliefs diagram. For that reason, the weighted beliefs diagram is symmetric.

As the beliefs diagram on top is induced by the symmetric weighted beliefs diagram at the bottom, we say that all the belief hierarchies in it are *symmetric*. We thus see that the definition of a symmetric belief hierarchy is almost identical to the one in Section 4.2 for standard games. The only difference is that belief hierarchies now involve beliefs about choice-utility pairs, rather than about choices alone. We thus arrive at the following definition.

Definition 6.2.1 (Symmetric belief hierarchy) *(a) A* ***weighted beliefs diagram*** *starts from a beliefs diagram, removes the probabilities at the forked arrows (if there are any), and assigns to every arrow a from a choice-utility pair (c_i, u_i) to an opponents' choice-utility combination $(c_j, u_j)_{j\neq i}$ some positive weight, which we call $w(a)$.*

(b) Consider an arrow a from a choice-utility pair (c_i, u_i) to an opponents' choice-utility combination $(c_j, u_j)_{j\neq i}$. For every opponent j, the ***symmetric counterpart*** *to a is the arrow from the choice-utility pair (c_j, u_j) to the opponents' choice-utility combination $(c_k, u_k)_{k\neq j}$, using the same choice-utility pairs as a.*

(c) A weighted beliefs diagram is ***symmetric*** *if for every arrow a, each of its symmetric counterparts (one for every opponent) is also part of the diagram, and carries the same weight as a.*

(d) The weighted beliefs diagram induces a (normal) beliefs diagram in which the probability of an arrow a leaving a choice-utility pair (c_i, u_i) is equal to

$$p(a) = \frac{w(a)}{\sum_{arrows\ a'\ leaving\ (c_i,u_i)} w(a')}.$$

(e) A belief hierarchy is ***symmetric*** *if it is part of a beliefs diagram that is induced by a symmetric weighted beliefs diagram.*

Let us return to the beliefs diagram on top of Figure 6.2.1. We have seen that all the belief hierarchies in this beliefs diagram are symmetric, because the beliefs diagram is induced by a symmetric weighted beliefs diagram. Moreover, as all arrows are solid, all belief hierarchies involved express common belief in rationality. As each of your choices can be supported by one such belief hierarchy, we conclude that under common belief in rationality with a symmetric belief hierarchy, you can

	$((ser,\ t_1^{ser}), (open,\ t_2^{f,o}))$	$((cho,\ t_1^{cho}), (ignore,\ t_2^{a,i}))$
weights	1	10
probabilities	1/17	10/17

	$((cho,\ t_1^{cho}), (open,\ t_2^{f,o}))$	$((Chris,\ t_1^{Chris}), (ignore,\ t_2^{a,i}))$
weights	3	1
probabilities	3/17	1/17

	$((Chris,\ t_1^{Chris}), (open,\ t_2^{f,o}))$	$((Chris,\ t_1^{Chris}), (ignore,\ t_2^{f,i}))$
weights	1	1
probabilities	1/17	1/17

Table 6.2.1 Common prior for "The moonlight serenade"

rationally plan a *serenade,* to offer a box of *chocolates,* or to send *Chris.* Similarly, under common belief in rationality with a symmetric belief hierarchy, Barbara can rationally *ignore* the door bell if she is a*ngry* or *forgiving,* whereas she can rationally *open* the door if she is *forgiving.*

Recall that these are precisely the choices that were possible under common belief in rationality. As such, the additional condition of a symmetric belief hierarchy does not affect the choices that you and Barbara can rationally make under common belief in rationality. As we have seen, this was different for *simple* belief hierarchies: By additionally imposing a *simple* belief hierarchy, you could no longer rationally plan a *serenade* for Barbara.

6.2.2 Relation with Common Prior

Recall from Section 4.2.2 that in standard games, the symmetric belief hierarchies were exactly those that are induced by a *common prior* on choice-type combinations. We will see that the same is true for games with incomplete information, for identical reasons. Consider, as an example, the symmetric belief hierarchies in Figure 6.2.1. The reason why these belief hierarchies are symmetric is because they are all generated by the symmetric weighted beliefs diagram at the bottom.

The symmetric weights in this weighted beliefs diagram can be used to construct a *common prior* on choice-type combinations, as follows. Consider the types t_1^{ser}, t_1^{cho} and t_1^{Chris} for you, and the types $t_2^{a,i}$, $t_2^{f,o}$ and $t_2^{f,i}$ for Barbara, where all of your types have the utility function u_1, Barbara's type $t_2^{a,i}$ has the utility function u_2^a, and the other two types for Barbara have the utility function u_2^f. Then, we can identify the choice-utility pair (*serenade,* u_1) with the choice-type pair (*serenade,* t_1^{ser}), and similarly for the other choice-utility pairs in the beliefs diagram. The new beliefs diagram, where choice-utility pairs have been replaced by choice-type pairs, is called a *beliefs diagram in choice-type combinations.* In fact, the original beliefs diagram and the new one represent exactly the same belief hierarchies, only the labeling has changed. We can thus use both beliefs diagrams interchangeably.

In the beliefs diagram in choice-type combinations, take the arrow from (*chocolates,* t_1^{cho}) to (*ignore,* $t_2^{a,i}$), and its symmetric counterpart from (*ignore,* $t_2^{a,i}$) to (*chocolates,* t_1^{cho}). By symmetry, both carry the same weight, which in this case is 10. We can thus assign the weight $w((chocolates,\ t_1^{cho}), (ignore,\ t_2^{a,i})) = 10$ to the corresponding choice-type combination $((chocolates,\ t_1^{cho}), (ignore,\ t_2^{a,i}))$. Of course, we can do the same for the other arrows in the beliefs diagram, which results in the weights on the choice-type combinations in Table 6.2.1.

Note that the sum of all the weights on the different choice-type combinations is 17. If we divide all weights by 17, then we obtain the *probabilities* on the choice-type combinations in Table 6.2.1. This probability distribution π on the different choice-type combinations is called a *common prior.*

In the same way as in Chapter 4 for standard games, it can now be shown that this common prior π *induces* the beliefs diagram on top of Figure 6.2.1. Consider, for instance, the arrow a_1 from your choice-utility pair (*chocolates,* u_1) to Barbara's choice-utility pair (*ignore,* u_2^a). The corresponding choice-type combination is $((cho,\ t_1^{cho}), (ignore,\ t_2^{a,i}))$ to which the common prior π assings the probability 10/17. Similarly, the arrow a_2 from (*chocolates,* u_1) to (*open,* u_2^f) corresponds to the choice-type combination $((cho,\ t_1^{cho}), (open,\ t_2^{f,o}))$, to which the common prior π assigns the probability 3/17. As a_1 and a_2 are the only two arrows that leave (*chocolates,* u_1), the relative probabilities assigned by the common prior π to the arrows leaving (*chocolates,* u_1) are

$$p(a_1) = \frac{10/17}{10/17 + 3/17} = 10/13 \text{ and } p(a_2) = \frac{3/17}{10/17 + 3/17} = 3/13.$$

Note that these are precisely the probabilities assigned to these two arrows in the beliefs diagram of Figure 6.2.1. In that sense, the common prior π *induces* the probabilities assigned to the arrows leaving (*chocolates,* u_1) in the beliefs diagram.

It can be verified that the same holds for all other arrows in the beliefs diagram. Therefore, we conclude that the common prior π indeed induces the beliefs diagram in Figure 6.2.1.

The following definition should not come as a surprise, given what we have already seen in Section 4.2.2.

Definition 6.2.2 (Common prior on choice-type combinations) *Take a beliefs diagram, and identify every choice-utility pair (c_i, u_i) with a unique choice-type pair (c_i, t_i) where the type t_i has the utility function u_i. What results is a **beliefs diagram in choice-type combinations**.*

*(a) A **common prior on choice-type combinations** is a probability distribution π that assigns to every choice-type combination $(c, t) = (c_i, t_i)_{i \in I}$ so obtained a probability $\pi(c, t)$.*

*(b) The beliefs diagram in choice-type combinations is **induced by a common prior** π on choice-type combinations if for every choice-type combination $((c_i, t_i), (c_{-i}, t_{-i}))$ and every player i, the corresponding arrow a from (c_i, t_i) to (c_{-i}, t_{-i}) is present exactly when $\pi((c_i, t_i), (c_{-i}, t_{-i})) > 0$, and the probability of this arrow a is equal to*

$$p(a) = \frac{\pi((c_i, t_i), (c_{-i}, t_{-i}))}{\pi(c_i, t_i)}.$$

*(c) A belief hierarchy is **induced by a common prior** π on choice-type combinations if it is part of a beliefs diagram in choice-type combinations that is induced by π.*

Note that this definition is the same as the one we saw in Section 4.2.2. Recall that a beliefs diagram for standard games could contain the same choice c_i twice. In the beliefs diagram we could then call these choices c_i and c_i', and in the associated beliefs diagram in choice-type combinations, these two copies could be identified with different choice-type pairs (c_i, t_i) and (c_i, t_i').

Similarly, for games with incomplete information a beliefs diagram could contain different copies of the same choice-utility pair (c_i, u_i). In the associated beliefs diagram in choice-type combinations, these two copies would be identified with different choice-type pairs (c_i, t_i) and (c_i, t_i'), where both t_i and t_i' have the utility function u_i.

For the beliefs diagram of Figure 6.2.1 we have seen that every symmetric belief hierarchy that is present in this beliefs diagram is induced by a common prior on choice-type combinations. In fact, the method we have used to design a common prior that induces these symmetric belief hierarchies can be applied in general, for every symmetric belief hierarchy. As such, we can conclude that every symmetric belief hierarchy is induced by a common prior on choice-type combinations.

The other direction is also true: If a belief hierarchy is induced by a common prior, then it will always be symmetric. Altogether, we thus see that the symmetric belief hierarchies are precisely those that are induced by a common prior on choice-type combinations.

Theorem 6.2.1 (Relation with common prior) *A belief hierarchy is symmetric, if and only if, it is induced by a common prior on choice-type combinations.*

This result will simplify the analysis of symmetric belief hierarchies, as it shows that they can always be summarized by a common prior on choice-type combinations – an object that is compact and easy to work with.

6.2.3 Relation with Bayesian Equilibrium

We will now combine the conditions of common belief in rationality with the requirement of a symmetric belief hierarchy, and show that it leads to a concept called *Bayesian equilibrium.* Consider a symmetric belief hierarchy β_i for player i that expresses common belief in rationality. By symmetry, it follows from Theorem 6.2.1 that β_i is induced by a common prior π on choice-type combinations. That is, β_i is part of a beliefs diagram where all belief hierarchies present are induced by the common prior π. Moreover, since β_i expresses common belief in rationality, we can always find such a beliefs diagram where all arrows are solid. That is, for every choice-utility pair (c_j, u_j) in the beliefs diagram, the choice c_j is optimal for the utility function u_j and the belief about the opponents' choices given by the outgoing arrows and their associated probabilities.

Now, identify every choice-utility pair (c_j, u_j) in the beliefs diagram with a choice-type pair (c_j, t_j) where t_j has utility function u_j. Moreover, we may assume that different choice-utility pairs correspond to different types. Since the beliefs diagram is induced by the common prior π, the probability that every such type t_j assigns to the opponents' choice-type combination $(c_k, t_k)_{k \neq j}$ is given by

$$b_j(t_j)((c_k, t_k)_{k \neq j}) = \frac{\pi((c_j, t_j), ((c_k, t_k)_{k \neq j})}{\pi(c_j, t_j)}.$$

As in Section 4.2.3, let $\pi(\cdot \mid c_j, t_j)$ be the belief that type t_j has about the opponents' choice-type combinations, given by the formula above. We refer to $\pi(\cdot \mid c_j, t_j)$ as the belief that player j has, conditional on his choice-type pair (c_j, t_j).

Recall from above that all arrows in the beliefs diagram are solid. As such, for every choice-type pair (c_j, t_j) in the beliefs diagram, the choice c_j must be optimal for the utility function that the type t_j has, and the belief $\pi(\cdot \mid c_j, t_j)$ that the type t_j has. Moreover, as the beliefs diagram is induced by the common prior π, the choice-type pairs (c_j, t_j) that are present in the beliefs diagram are exactly those that receive positive probability by π, that is, where $\pi(c_j, t_j) > 0$.

By putting these two insights together, we arrive at the following conclusion: For every choice-type pair (c_j, t_j) with $\pi(c_j, t_j) > 0$, the choice c_j must be optimal given the utility function that t_j has, and given the belief $\pi(\cdot \mid c_j, t_j)$ that player j has, conditional on his choice-type pair (c_j, t_j). Common priors π with this property are called *Bayesian equilibria.*

Definition 6.2.3 (Bayesian equilibrium) *A common prior π on choice-type combinations is a* ***Bayesian equilibrium*** *if for every player i, and every choice-type pair (c_i, t_i) with $\pi(c_i, t_i) > 0$, the choice c_i is optimal for the utility function that t_i has, and the belief $\pi(\cdot \mid c_i, t_i)$ that player i has conditional on his choice-type pair (c_i, t_i).*

Note that this definition is almost the same as the one for *correlated equilibrium* in Section 4.2.3. The only difference is that in the present setup, different types may specify different utility functions. If we apply the definition of a Bayesian equilibrium to standard games without incomplete information, where only one utility function is possible for every player, then we get exactly the definition of correlated equilibrium. In that sense, Bayesian equilibrium may be viewed as the counterpart to correlated equilibrium in the context of incomplete information.

By our arguments above we have thus shown that every symmetric belief hierarchy β_i that expresses common belief in rationality must be induced by a common prior π on choice-type combinations that is a *Bayesian equilibrium.* In fact, the other direction is also true: If a belief hierarchy is induced by a Bayesian equilibrium π, then the belief hierarchy will be symmetric and it will express common belief in rationality. We thus arrive at the following conclusion.

Theorem 6.2.2 (Relation with Bayesian equilibrium) *A belief hierarchy is symmetric and expresses common belief in rationality, if and only if, the belief hierarchy is induced by a Bayesian equilibrium.*

As an illustration, consider the beliefs diagram in Figure 6.2.1. We have already seen that all belief hierarchies present are symmetric. Moreover, since only solid arrows appear in the beliefs diagram, all these belief hierarchies express common belief in rationality. By Theorem 6.2.2 we then conclude that every belief hierarchy in this beliefs diagram must be induced by a Bayesian equilibrium. In fact, it turns out that *all* belief hierarchies present are induced by the *same* Bayesian equilibrium, which is the common prior π from Table 6.2.1.

Question 6.2.1 *Show that the common prior π from Table 6.2.1 is a Bayesian equilibrium.*

Theorem 6.2.2 is useful for several reasons. First, it can be used to easily construct symmetric belief hierarchies that express common belief in rationality. To do so, we start by constructing a Bayesian equilibrium π, which is just a common prior on choice-type combinations within an epistemic model with certain optimality properties, and then derive all possible belief hierarchies from π. All these belief hierarchies are guaranteed to be symmetric, and to express common belief in rationality.

Moreover, it can be used to verify whether a given choice can rationally be made under common belief in rationality with a symmetric belief hierarchy. To see how, suppose that the choice c_i^* can rationally be made, for some utility function u_i^*, under common belief in rationality with a symmetric belief hierarchy. Then, by Theorem 6.2.2, there must be a Bayesian equilibrium π such that the choice c_i^* is optimal, given the utility function u_i^*, for the belief hierarchy induced by π.

Assume that the belief hierarchy starts at the choice-type pair (c_i, t_i). Then, the induced belief about the opponents' choice-type combinations is $\pi(\cdot \mid (c_i, t_i))$. Since the choice c_i^* is optimal, given the utility function u_i^*, for this belief hierarchy, the choice c_i^* must be optimal for the belief $\pi(\cdot \mid (c_i, t_i))$ and the utility function u_i^*. In this case, we say that c_i^* is optimal in Bayesian equilibrium for the utility function u_i^*.

Definition 6.2.4 (Choice optimal in a Bayesian equilibrium) *A choice c_i^* is* ***optimal in a Bayesian equilibrium*** *π for the utility function u_i^* if there is some choice-type pair (c_i, t_i) with $\pi(c_i, t_i) > 0$ such that c_i^* is optimal for the induced belief $\pi(\cdot \mid (c_i, t_i))$ and the utility function u_i^*.*

You	*open*	*ignore*
serenade	4	0
chocolates	0	4
Chris	3	3

u_1

Barbara	*serenade*	*chocolates*	*Chris*
open	0	0	0
ignore	1	1	1

u_2^a

Barbara	*serenade*	*chocolates*	*Chris*
open	0	0	0
ignore	1	0	1

u_2^f

Table 6.2.2 Decision problems for "The moonlight serenade with a twist"

Above we have argued that every choice c_i^* that can rationally be made for a given utility function u_i^* under common belief in rationality with a symmetric belief hierarchy must be optimal in a Bayesian equilibrium for this utility function u_i^*. In fact, the other direction is also true: If the choice c_i^* is optimal in a Bayesian equilibrium π for the utility function u_i^*, then c_i^* can rationally be made for the utility function u_i^* under common belief in rationality with a symmetric belief hierarchy.

The reason is simple: Suppose that c_i^* is optimal in a Bayesian equilibrium π for the utility function u_i^*. Then, by definition, there is some choice-type pair (c_i, t_i) with $\pi(c_i, t_i) > 0$ such that c_i^* is optimal for the induced belief $\pi(\cdot \mid (c_i, t_i))$ and the utility function u_i^*. By Theorem 6.2.2 we know that the belief hierarchy β_i that can be derived from π and that starts at (c_i, t_i) is symmetric and expresses common belief in rationality. Thus, the choice c_i^* is optimal for the utility function u_i^* and the belief hierarchy β_i that is symmetric and expresses common belief in rationality. We therefore arrive at the following conclusion.

Theorem 6.2.3 (Choices optimal in a Bayesian equilibrium) *A choice c_i can rationally be made for the utility function u_i under common belief in rationality with a symmetric belief hierarchy, if and only if, the choice c_i is optimal for the utility function u_i in a Bayesian equilibrium.*

To conclude this section, we will apply the theorem above to a new example which is only a slight variation of "The moonlight serenade".

Example 6.4: The moonlight serenade with a twist.

Recall the story and the decision problems from "The moonlight serenade". Suppose now that Barbara, if she is in a *forgiving* mood and believes that you bring *chocolates,* would be indifferent between *opening* the door and *ignoring* the door bell. Besides this little twist, everything else is the same as before. The new decision problems are thus given by Table 6.2.2. Note that the only difference lies in Barbara's utility from (*open, chocolates*) at u_2^f, which is now 0 instead of 1. As a consequence, *opening* the door is *only* optimal for Barbara if she in a *forgiving* mood and assigns probability 1 to you bringing *chocolates.*

It can still be shown that under common belief in rationality, you can rationally make *any* plan, whereas Barbara can rationally choose to *ignore* the door bell if she is *angry,* and she can rationally choose to *open* the door or *ignore* the door bell if she is *forgiving.* Indeed, it suffices to consider the beliefs diagram from Figure 6.1.3 for "The moonlight serenade". It may be verified that this beliefs diagram is still valid for the new example. As all the choices mentioned above are supported by a belief hierarchy in this beliefs diagram that expresses common belief in rationality, our statement above about the possible choices under common belief in rationality follows.

However, if we additionally insist on a *symmetric* belief hierarchy, then it will no longer be rational for you to plan a *serenade.* To show this we will rely on Theorem 6.2.3. Suppose, contrary to what we want to show, that *serenade* is optimal for you in a symmetric belief hierarchy that expresses common belief in rationality. Then, by Theorem 6.2.3, *serenade* must be optimal for you in a Bayesian equilibrium π, which is a probability distribution on choice-type combinations.

We first show that π must assign probability 0 to your choice *serenade.* Suppose not. Then, $\pi(\textit{serenade, } t_1) > 0$ for some type t_1. As π is a Bayesian equilibrium, *serenade* must be optimal for the conditional belief $\pi(\cdot \mid (\textit{serenade, } t_1))$. This is only possible if $\pi(\cdot \mid (\textit{serenade, } t_1))$ assigns a positive probability to Barbara's choice *open.* Hence, there must be some type t_2 for Barbara such that $\pi((\textit{open, } t_2) \mid (\textit{serenade, } t_1)) > 0$, which implies that $\pi((\textit{serenade, } t_1), (\textit{open, } t_2)) > 0$. In particular, $\pi(\textit{open, } t_2) > 0$. As π is a Bayesian equilibrium, it follows that *open* must be optimal for Barbara under the belief $\pi(\cdot \mid (\textit{open, } t_2))$. However, since $\pi((\textit{serenade, } t_1), (\textit{open, } t_2)) > 0$ it follows that $\pi((\textit{serenade, } t_1) \mid (\textit{open, } t_2)) > 0$ as well, which means that *open* cannot be optimal for Barbara under $\pi(\cdot \mid (\textit{open, } t_2))$. This yields a contradiction. Hence, we conclude that $\pi(\textit{serenade}) = 0$.

Next, we show that π must also assign probability 0 to your choice *Chris.* Suppose not. Then, $\pi(\textit{Chris, } t_1) > 0$ for some type t_1, and hence *Chris* must be optimal under the belief $\pi(\cdot \mid (\textit{Chris, } t_1))$. This is only possible if $\pi(\cdot \mid (\textit{Chris, } t_1))$ assigns a positive probability to Barbara's choice *open.* Hence, there must be some type t_2 for Barbara such that $\pi((\textit{open, } t_2) \mid \pi(\textit{Chris, } t_1)) > 0$, which implies that $\pi((\textit{Chris, } t_1), (\textit{open, } t_2)) > 0$. In particular, $\pi(\textit{open, } t_2) > 0$, which means that *open* must be optimal for Barbara under the belief $\pi(\cdot \mid (\textit{open, } t_2))$. However, since $\pi((\textit{Chris, } t_1), (\textit{open, } t_2)) > 0$ it follows that $\pi((\textit{Chris, } t_1) \mid (\textit{open, } t_2)) > 0$ as well, which means that *open* cannot be optimal for Barbara under $\pi(\cdot \mid (\textit{open, } t_2))$. This yields a contradiction. Hence, we conclude that $\pi(\textit{Chris}) = 0$.

From the two insights above we conclude that π must assign probability 1 to your choice *chocolates.* Recall that we are assuming that *serenade* is optimal in the Bayesian equilibrium π. Hence, there must be some type t_1 with $\pi(\textit{chocolates, } t_1) > 0$ such that *serenade* is optimal for you under the belief $\pi(\cdot \mid (\textit{chocolates, } t_1))$. But then, both *serenade* and *chocolates* must be optimal under the belief $\pi(\cdot \mid (\textit{chocolates, } t_1))$, which means, in particular, that they must yield the same expected utility under $\pi(\cdot \mid (\textit{chocolates, } t_1))$. This is only possible if $\pi(\cdot \mid (\textit{chocolates, } t_1))$ assigns probability 0.5 to Barbara's choices *open* and *ignore.* But then, sending *Chris* would be better than both *serenade* and *chocolates.* We thus obtain a contradiction.

Therefore, our assumption above that *serenade* is optimal in a Bayesian equilibrium π cannot be true. But then, it follows from Theorem 6.2.3 that under common belief in rationality with a *symmetric* belief hierarchy, you cannot rationally plan a *serenade.*

We can also use Theorem 6.2.3 to show that under common belief in rationality with a symmetric belief hierarchy, you can rationally choose *chocolates* and *Chris,* that Barbara can rationally choose *ignore* if she is *angry,* and she can rationally choose *open* and *ignore* if she is *forgiving.* Consider the common prior π on choice-type combinations given by

$$\pi((\textit{chocolates, } t_1), (\textit{ignore, } t_2^a)) = 0.75 \text{ and } \pi((\textit{chocolates, } t_1), (\textit{open, } t_2^f)) = 0.25,$$

where t_2^a has the utility function u_2^a, and t_2^f has the utility function u_2^f.

Question 6.2.2 *Show that π is a Bayesian equilibrium.*

Note that both *chocolates* and *Chris* are optimal for you under the belief $\pi(\cdot \mid (\textit{chocolates, } t_1))$, that *ignore* is optimal for Barbara under the belief $\pi(\cdot \mid (\textit{ignore, } t_2^a))$ and the utility function u_2^a, and that both *open* and *ignore* are optimal for Barbara under the belief $\pi(\cdot \mid (open,\ t_2^f))$ and the utility function u_2^f. Hence, by Theorem 6.2.3 we conclude that under common belief in rationality with a symmetric belief hierarchy, you can rationally choose *chocolates* and *Chris,* that Barbara can rationally choose *ignore* if she is *angry,* and she can rationally choose *open* and *ignore* if she is *forgiving.*

Overall, we see that if we add the condition of a *symmetric* belief hierarchy to common belief in rationality, then you can no longer rationally plan to give a *serenade* for Barbara, but it does not affect the other choices that were possible under common belief in rationality.

6.2.4 Relation with Generalized Nash Equilibrium

Recall from Theorem 6.2.2 that *Bayesian* equilibria correspond to *symmetric* belief hierarchies that express common belief in rationality. On the other hand, we know from Theorem 6.1.1 that *generalized Nash* equilibria relate to *simple* belief hierarchies that express common belief in rationality. But how do the two concepts relate to each other?

In Section 4.3.4 we have shown, for standard games, that every simple belief hierarchy is symmetric. In a similar fashion, it can be shown that the same is true for games with *incomplete information.*

Theorem 6.2.4 (Relation with simple belief hierarchies) *Every simple belief hierarchy is symmetric.*

If we combine the Theorems 6.1.1, 6.2.2 and 6.2.4 we conclude that every belief hierarchy that is induced by a generalized Nash equilibrium is also induced by a Bayesian equilibrium. As such, every choice that can rationally be made in a generalized Nash equilibrium can also rationally be made in a Bayesian equilibrium.

Theorem 6.2.5 (Generalized Nash equilibrium implies Bayesian equilibrium) *For every player i and every utility function u_i, all choices that can rationally be made for the utility function u_i in a generalized Nash equilibrium can also rationally be made for u_i in a Bayesian equilibrium.*

As an illustration, consider the example "The moonlight serenade". We have seen that under common belief in rationality with a simple belief hierarchy, you can rationally choose to bring a box of *chocolates* or to send *Chris.* Hence, by Theorem 6.1.1, the choices *chocolates* and *Chris* are optimal for you in a generalized Nash equilibrium. Indeed, we have shown in Section 6.1.3 that *chocolates* and *Chris* are both optimal for you in the generalized Nash equilibrium $\sigma_1 = (0.75) \cdot (chocolates,\ u_1) + (0.25) \cdot (Chris,\ u_1)$ and $\sigma_2 = (0.25) \cdot (open,\ u_2^f) + (0.75) \cdot (ignore,\ u_2^a)$.

By Theorem 6.2.5 we may thus conclude that your choices *chocolates* and *Chris* are optimal in a Bayesian equilibrium. This has been confirmed in Sections 6.2.1, 6.2.2 and 6.2.3, where we have shown that all of your choices are optimal in the Bayesian equilibrium π from Table 6.2.1.

The converse of Theorem 6.2.5 may not be true in general. To see why, consider again the example "The moonlight serenade". We have seen that planning a *serenade* is optimal in the Bayesian equilibrium π from Table 6.2.1, but in Section 6.1.3 we have shown this choice is not optimal in any generalized Nash equilibrium.

6.2.5 One Theory per Choice-Utility Pair

In Section 4.3 we investigated the *one theory per choice* condition for standard games. Intuitively, it states that in the belief hierarchy every choice is supported by a unique belief. This could be formalized by requiring that the belief hierarchy is induced by a beliefs diagram where every choice appears only once. That is, no copies of the same choice are allowed in the beliefs diagram. In Section 4.3.3 we saw that adding the *one theory per choice* condition to symmetric belief hierarchies that express common belief in rationality may have consequences for the choices that can rationally be made. On the other hand, adding the *one theory per choice* condition only to common belief in rationality does not alter the choices that can rationally be made.

We will see that similar definitions and conclusions apply to games with *incomplete information,* where the one theory per choice condition is replaced by *one theory per choice-utility pair.* Intuitively, the *one theory per choice-utility pair* condition states that in the belief hierarchy, there is for every *choice-utility pair* a unique belief that supports this choice for *this particular utility function.* In particular, for a given choice we may use two different beliefs to support this choice for two different utility functions. This condition may be formalized as follows.

Definition 6.2.5 (One theory per choice-utility pair) *A belief hierarchy uses* ***one theory per choice-utility pair*** *if it can be generated by a beliefs diagram where every choice-utility pair only appears once.*

This condition will "typically" be satisfied, and it is therefore no coincidence that all the beliefs diagrams we have seen so far in Chapters 5 and 6 meet this condition.

The question we now wish to address is: What choices can rationally be made if you hold a symmetric belief hierarchy that uses *one theory per choice-utility pair* and expresses common belief in rationality? We will see that this leads to the concept of *canonical Bayesian equilibrium,* similarly to how the same conditions led to *canonical correlated equilibrium* in standard games.

Consider a symmetric belief hierarchy β that uses one theory per choice-utility pair and expresses common belief in rationality. By Theorem 6.2.2 we know that β is induced by a Bayesian equilibrium π, which is a common prior on choice-type combinations $(c_i, t_i)_{i\in I}$.

By the *one theory per choice-utility pair* condition, there is for every choice-utility pair (c_i, u_i) that receives positive probability a *unique* belief hierarchy – and therefore a *unique* type $t_i^{c_i,u_i}$ – that supports the choice c_i for the utility function u_i. But then, we can replace every choice-type pair $(c_i, t_i^{c_i,u_i})$ that receives positive probability by the associated choice-utility pair (c_i, u_i). Moreover, the common prior π on choice-type combinations then induces a common prior $\hat{\pi}$ on *choice-utility* combinations given by

$$\hat{\pi}((c_i, u_i)_{i\in I}) := \pi((c_i, t_i^{c_i,u_i})_{i\in I})$$

for every choice-utility combination $(c_i, u_i)_{i\in I}$ that receives positive probability under π.

Since π is Bayesian equilibrium, we know that for every choice-type pair $(c_i, t_i^{c_i,u_i})$ with $\pi(c_i, t_i^{c_i,u_i}) > 0$, the associated choice c_i is optimal for the utility function u_i under the belief $\pi(\cdot \mid (c_i, t_i^{c_i,u_i}))$ conditional on $(c_i, t_i^{c_i,u_i})$. But then, by construction, the common prior $\hat{\pi}$ on choice-utility combinations induced by π satisfies the following optimality condition: For every choice-utility pair (c_i, u_i) with $\hat{\pi}(c_i, u_i) > 0$, the associated choice c_i is optimal for the utility function u_i under the belief $\hat{\pi}(\cdot \mid (c_i, u_i))$ conditional on (c_i, u_i). Common priors on choice-utility combinations with this property are called *canonical Bayesian equilibria.*

Definition 6.2.6 (Canonical Bayesian equilibrium) *A common prior $\hat{\pi}$ on choice-utility combinations is a* ***canonical Bayesian equilibrium*** *if for every player i, and every choice-utility pair (c_i, u_i) with $\hat{\pi}(c_i, u_i) > 0$, the associated choice c_i is optimal for the utility function u_i under the belief $\hat{\pi}(\cdot \mid (c_i, u_i))$ conditional on (c_i, u_i).*

By the arguments above it follows that every symmetric belief hierarchy that uses *one theory per choice-utility pair* and expresses common belief in rationality is induced by a *canonical* Bayesian equilibrium. It turns out that the other direction is also true: If the belief hierarchy is induced by a canonical Bayesian equilibrium, then the belief hierarchy is symmetric, uses one theory per choice-utility pair, and expresses common belief in rationality. We thus arrive at the following conclusion.

Theorem 6.2.6 (Relation with canonical Bayesian equilibrium) *A belief hierarchy is symmetric, uses one theory per choice-utility pair and expresses common belief in rationality, if and only if, it is induced by a canonical Bayesian equilibrium.*

We say that a choice c_i^* is *optimal in a canonical Bayesian equilibrium* $\hat{\pi}$ for the utility function u_i^* if there is a choice-utility pair (c_i, u_i) with $\hat{\pi}(c_i, u_i) > 0$ such that c_i^* is optimal for the utility function u_i^* under the belief $\hat{\pi}(\cdot \mid (c_i, u_i))$. By Theorem 6.2.6 it then follows that also in terms of optimal choices, the conditions in Theorem 6.2.6 lead to canonical Bayesian equilibrium.

Theorem 6.2.7 (Choices optimal in a canonical Bayesian equilibrium) *Player i can rationally make the choice c_i for the utility function u_i with a symmetric belief hierarchy that uses one theory per choice-utility pair and expresses common belief in rationality, if and only if, the choice c_i is optimal for the utility function u_i in a canonical Bayesian equilibrium.*

As an illustration, consider the example "The moonlight serenade". We have seen that under common belief in rationality with a symmetric belief hierarchy, you can rationally plan a *serenade,* to bring a box of *chocolates* or to send *Chris.* Indeed, each of these three choices is optimal for your utility function u_1 in the Bayesian equilibrium from Table 6.2.1. By Theorem 6.2.3 it thus follows that you can rationally make each of these choices under common belief in rationality with a symmetric belief hierarchy.

Note that the belief hierarchies that can be derived from this Bayesian equilibrium are present in the beliefs diagram from Figure 6.2.1. As every choice-utility pair only appears once in this beliefs diagram, we conclude that all the symmetric belief hierarchies that can be derived from the Bayesian equilibrium use *one theory per choice-utility pair.* As such, you can rationally make each of your choices under common belief in rationality with a symmetric belief hierarchy that uses *one theory per choice-utility pair.*

This conclusion can be confirmed by using Theorem 6.2.7. To see this, consider the Bayesian equilibrium π from Table 6.2.1, which is a common prior on choice-type combinations. The Bayesian equilibrium π can be transformed into a common prior $\hat{\pi}$ on choice-utility pairs, where

$$\hat{\pi}((\textit{serenade},\ u_1), (\textit{open},\ u_2^f)) = 1/17,\ \hat{\pi}((\textit{chocolates},\ u_1), (\textit{ignore},\ u_2^a)) = 10/17,$$
$$\hat{\pi}((\textit{chocolates},\ u_1), (\textit{open},\ u_2^f)) = 3/17,\ \hat{\pi}((\textit{Chris},\ u_1), (\textit{ignore},\ u_2^a)) = 1/17,$$
$$\hat{\pi}((\textit{Chris},\ u_1), (\textit{open},\ u_2^f)) = 1/17 \text{ and } \hat{\pi}((\textit{Chris},\ u_1), (\textit{ignore},\ u_2^f)) = 1/17.$$

Question 6.2.3 *Explain why $\hat{\pi}$ is a canonical Bayesian equilibrium.*

Moreover, it may be verified that *serenade, chocolates* and *Chris* are optimal for you under the utility function u_1 for the associated conditional beliefs $\hat{\pi}(\cdot \mid (\textit{serenade},\ u_1))$, $\hat{\pi}(\cdot \mid (\textit{chocolates},\ u_1))$ and $\hat{\pi}(\cdot \mid (\textit{Chris},\ u_1))$. Thus, each of your choices is optimal in a canonical Bayesian equilibrium. By Theorem 6.2.7 it then follows that you can rationally make each of your choices under common belief in rationality with a symmetric belief hierarchy that uses *one theory per choice-utility pair.*

In particular, adding the one theory per choice-utility pair condition does not alter the choices you can rationally make under common belief in rationality with a symmetric belief hierarchy. The same holds for the other examples we have investigated in Chapters 5 and 6. The reason is that in each of these examples, every symmetric belief hierarchy expressing common belief in rationality that has been used to support these choices used one theory per choice-utility pair.

But there are games with incomplete information where additionally imposing one theory per choice-utility pair can make a difference for the choices you can rationally make. It suffices to consider the standard game "Rock, paper, scissors" from Chapter 4. We have seen that under common belief in rationality with a symmetric belief hierarchy, you can rationally choose *rock, paper, scissors* and *bomb,* but that *bomb* can no longer rationally be chosen if we insist on *one theory per choice.*

Clearly, this game is a special case of a game with incomplete information, where there is only *one* possible utility function for both players. As *one theory per choice-utility pair* reduces to *one theory per choice* for such games, we see that additionally imposing *one theory per choice-utility pair* makes a difference for the choices you can rationally make under common belief in rationality with a symmetric belief hierarchy.

In Chapter 4, when investigating standard games, we have seen that for the choices you can rationally make under *common belief in rationality* it does not matter whether we additionally impose *one theory per choice* or not. See Theorem 4.3.1. Moreover, we saw in Theorem 4.3.5 that every *simple* belief hierarchy automatically satisfies the *one theory per choice* condition.

The same is true for games with incomplete information, for essentially the same reasons. We will therefore not repeat these reasons here. In particular, we see that for the choices you can rationally make under common belief in rationality, or for the choices you can rationally make under common belief in rationality with a simple belief hierarchy, it does not matter whether we additionally impose *one theory per choice-utility pair* or not. But for the choices that can rationally be made under common belief in rationality with a *symmetric* belief hierarchy this may matter.

6.3 *Fixed Beliefs on Utilities

In Chapter 5 we have investigated the idea of *fixed beliefs on utilities,* which means that you believe that there are unique beliefs on the players' utility functions which are transparent to everyone. In Section 5.5 we have added this condition to *common belief in rationality,* and provided a recursive procedure which yields precisely those choices that are possible under the new set of conditions. In this section we explore what happens if we add *fixed beliefs on utilities* to the conditions of common belief in rationality with a *simple* belief hierarchy and common belief in rationality with a *symmetric* belief hierarchy.

6.3.1 Definition

In a game with incomplete information, consider for every player j a probability distribution $p_j \in \Delta(U_j)$ on player j's possible utility functions. In Section 5.5 we have defined what it means, for a type t_i within an epistemic model, to express *common belief* in the collection $p = (p_j)_{j \in I}$ of probability distributions. Intuitively, it means that t_i's belief about the opponents' utilities is given by p, that t_i believes that every opponent's belief about the utility functions of the others is given by p, and so on.

If we add this condition to common belief in rationality with a simple, or symmetric, belief hierarchy, we arrive at the following definition.

Definition 6.3.1 (Rational choice with fixed beliefs on utilities) *Let $p = (p_i)_{i \in I}$ be a profile of beliefs on utility functions, and $u_i \in U_i$ a utility function. Then, player i can* ***rationally make the choice*** c_i ***with utility function*** u_i ***under common belief in rationality with a simple (symmetric) belief hierarchy and common belief in*** p *if there is an epistemic model $(T_i, v_i, b_i)_{i \in I}$ and a type $t_i \in T_i$ such that (a) t_i expresses common belief in rationality, (b) t_i's belief hierarchy is simple (symmetric), (c) t_i expresses common belief in p, (d) t_i has utility function u_i, and (e) c_i is optimal for t_i.*

The belief hierarchies we have used until now in Chapter 6 to support choices that can rationally be made under common belief in rationality with a simple (symmetric) belief hierarchy typically do *not* express common belief in p, for any profile p of beliefs on utility functions. Consider, for instance, the beliefs diagram in Figure 6.2.1. The belief hierarchy that starts at your choice *serenade* does not express common belief in p for any profile p of beliefs on utility functions. Indeed, in that belief hierarchy you assign probability 1 to Barbara being *forgiving,* but at the same time you believe that Barbara assigns probability 3/5 to the event that you assign probability 10/13 to Barbara being *angry.*

Question 6.3.1 *Consider the beliefs diagram in Figure 6.2.1. Is there any belief hierarchy in this beliefs diagram that expresses common belief in p, for some profile p of beliefs on utility functions?*

A natural question that arises is whether we can always find, for a given profile p of beliefs on utility functions, a simple, or symmetric, belief hierarchy that expresses common belief in rationality and common belief in p. The answer to both questions is "yes".

Theorem 6.3.1 (Existence) *Consider a game with incomplete information and a profile $p = (p_i)_{i \in I}$ of beliefs on utility functions. Then, for every player i there is a simple, and hence symmetric, belief hierarchy that expresses common belief in rationality and common belief in p.*

That is, for every game, and every profile p of beliefs on utility functions, it is always possible to simultaneously reason in accordance with common belief in rationality, in accordance with the principles of a simple belief hierarchy, and to reason within the bounds of these fixed beliefs p on the utility functions.

6.3.2 Examples

To conclude this section we will review the four examples we have investigated so far in this chapter. For every example we will fix a profile p of beliefs on utility functions, and see what choices you can rationally make with a simple or symmetric belief hierarchy that expresses common belief in rationality and common belief in p.

Example 6.5: What is Barbara's favorite color?

As in Section 5.5.4, let $p = (p_1, p_2)$ be the profile of beliefs on utility functions where p_1 assigns probability 1 to your unique utility function u_1, and p_2 assigns probabilities 0.8 and 0.2 to Barbara's utility functions u_2^r and u_2^b, respectively. We have seen in Section 5.5.4 that under common belief in rationality and common belief in p, you can only rationally wear *blue,* Barbara can only rationally wear *red* if her utility function is u_2^r, and she can only rationally wear *yellow* if her utility function is u_2^b.

Recall, by Theorem 6.3.1, that there is a simple belief hierarchy for you and one for Barbara that expresses common belief in rationality and common belief in p. Hence, with a *simple* belief hierarchy that expresses common belief in rationality and common belief in p, you can only rationally wear *blue,* Barbara can only rationally wear *red* if her utility function is u_2^r, and she can only rationally wear *yellow* if her utility function is u_2^b. And the same will be true if we replace *simple* belief hierarchy by *symmetric* belief hierarchy.

We thus see that in this example, adding the condition of a simple or symmetric belief hierarchy to common belief in rationality and common belief in p does not alter the colors that you and Barbara can rationally wear.

Example 6.6: Chris' drawings.

As in Section 5.4.4, let $p = (p_1, p_2)$ be the profile of beliefs on utility functions where p_1 assigns probability 0.25 to your valuation being $30, 50, 70$ or 90, respectively, and similarly for p_2. Recall from Section 5.4.4 that under common belief in rationality and common belief in p, you can only rationally bid 20 if your valuation is 30, you can rationally bid 20 or 40 if your valuation is 50, you can only rationally bid 40 if your valuation is 70, and you can rationally bid 40 or 60 if your valuation is 90.

It turns out that these bids can also rationally be made if we add the condition of a simple belief hierarchy. To see why, consider the belief σ_1 on your choice-utility pairs where

$$\sigma_1 = (0.25) \cdot (20, u_1^{30}) + (0.25) \cdot (20, u_1^{50}) + (0.25) \cdot (40, u_1^{70}) + (0.25) \cdot (40, u_1^{90})$$

and similarly for σ_2.

Question 6.3.2 *Explain why* (σ_1, σ_2) *is a generalized Nash equilibrium.*

Consider the belief hierarchy β_1 for you that is induced by the generalized Nash equilibrium (σ_1, σ_2). Then, by Theorem 6.1.1, β_1 is simple and expresses common belief in rationality. Moreover, by construction of σ_1 and σ_2, the belief hierarchy β_1 expresses common belief in p. If you hold the belief hierarchy β_1, then it is optimal for you to bid 20 if your valuation is 30 or 50, and to bid 40 if your valuation is 70 or 90. Therefore, we conclude that with a simple belief hierarchy that expresses common belief in rationality and common belief in p, you can rationally bid 20 if your valuation is 30 or 50, and rationally bid 40 if your valuation is 70 or 90.

Question 6.3.3 *Explain, in a similar way as above, that with a simple belief hierarchy which expresses common belief in rationality and common belief in* p, *you can also rationally bid* 40 *if your valuation is* 50, *and rationally bid* 60 *if your valuation is* 90.

Together with the insight above, we thus conclude that with a simple belief hierarchy that expresses common belief in rationality and common belief in p, you can only rationally bid 20 if your valuation is 30, you can rationally bid 20 or 40 if your valuation is 50, you can only rationally bid 40 if your valuation is 70, and you can rationally bid 40 or 60 if your valuation is 90. That is, if we add the condition of a simple belief hierarchy to the conditions of common belief in rationality and common belief in p, then the bids that can rationally be made do not change. Since every simple belief hierarchy is symmetric, the same holds if we replace *simple* belief hierarchy by *symmetric* belief hierarchy.

Example 6.7: The moonlight serenade.

In Section 6.1.3 we saw that under common belief in rationality, you can rationally plan to give a *serenade,* to bring a box of *chocolates,* or to send *Chris.* Moreover, Barbara can only rationally *ignore* the door bell if she is *angry,* whereas she can rationally *ignore* the door bell or *open* the door if she is *forgiving.*

Moreover, we have seen that if we add the condition of a *simple* belief hierarchy, then you can no longer rationally plan to give a *serenade,* whereas the other choices for you and Barbara will still be possible. In turn, all choices you and Barbara could rationally make under common belief in rationality will still be possible if we add the requirement that the belief hierarchy must be symmetric. This can be seen from the belief hierarchies in Table 6.2.1, which are all symmetric and express common belief in rationality.

Now suppose that you deem it quite likely that Barbara is *angry,* and you believe this to be transparent between Barbara and you. More precisely, let $p = (p_1, p_2)$ be the profile of beliefs on utility functions where p_1 assigns probability 1 to your unique utility function u_1, and p_2 assigns probabilities 0.8 and 0.2 to Barbara being *angry* and *forgiving,* respectively. Since Barbara can only rationally choose *ignore* if she is *angry,* you must assign probability at least 0.8 to Barbara *ignoring* the door bell. But then, your only optimal choice is to bring a box of *chocolates.* Barbara, anticipating on this, will certainly *open* the door if she is *forgiving.* Hence, under common belief in rationality and common belief in p, you can only rationally bring a box of *chocolates,* Barbara can only rationally choose to *ignore* the door bell if she is *angry,* and she can only rationally choose to *open* the door if she is *forgiving.*

Recall, by Theorem 6.3.1, that there must be a simple belief hierarchy for you and one for Barbara that expresses common belief in rationality and common belief in p. But then, it follows from the insight above that with a simple belief hierarchy which expresses common belief in rationality and common belief in p, you can only rationally bring a box of *chocolates,* Barbara can only rationally choose to *ignore* the door bell if she is *angry,* and she can only rationally choose to *open* the door if she is *forgiving.* Since every simple belief hierarchy is symmetric, the same holds if we replace *simple* belief hierarchy by *symmetric* belief hierarchy.

Example 6.8: The moonlight serenade with a twist.

In Section 6.2.3 we have seen that under common belief in rationality, you can rationally plan to give a *serenade,* to bring a box of *chocolates,* or to send *Chris.* Moreover, Barbara can only rationally *ignore* the door bell if she is *angry,* whereas she can rationally *ignore* the door bell or *open* the door if she is *forgiving.* However, if we add the requirement of a *symmetric* belief hierarchy, then you can no longer rationally plan a *serenade* for Barbara, but the other choices for you and Barbara will still be possible.

In fact, the same will be true if we replace *symmetric* belief hierarchy by *simple* belief hierarchy. To see this, note first that you cannot rationally plan a *serenade* with a simple belief hierarchy that expresses common belief in rationality, since every simple belief hierarchy is symmetric. To see that the other choices mentioned above can rationally be made under common belief in rationality with a simple belief hierarchy, consider the combination (σ_1, σ_2) of beliefs on choice-utility pairs where

$$\sigma_1 = (\textit{chocolates},\, u_1) \text{ and } \sigma_2 = (0.75) \cdot (\textit{ignore},\, u_2^a) + (0.25) \cdot (\textit{open},\, u_2^f).$$

Question 6.3.4 *Show that (σ_1, σ_2) is a generalized Nash equilibrium.*

Note that for you, it is optimal to bring *chocolates* or to send *Chris* under the belief σ_2. Moreover, if Barbara holds the belief σ_1, then it is optimal for her to *ignore* the door bell if she is *angry,* whereas it is optimal to *ignore* the door bell or to *open* the door if she is *forgiving.* Hence, these choices are all optimal in a generalized Nash equilibrium. By Theorem 6.1.2 we thus conclude that all these choices can rationally be made under common belief in rationality with a simple belief hierarchy.

As in the previous example, consider the profile $p = (p_1, p_2)$ of beliefs on utility functions where p_1 assigns probability 1 to your unique utility function u_1, and p_2 assigns probabilities 0.8 and 0.2 to Barbara being *angry* and *forgiving,* respectively.

Question 6.3.5 *Show that under common belief in rationality and common belief in* p*, you can only rationally bring chocolates, Barbara can only rationally ignore the door bell if she is angry, whereas she can rationally ignore the door bell or open the door if she is forgiving.*

In fact, if we add the requirement of a symmetric or simple belief hierarchy, the choices that can rationally be made will not change. To see this, consider the combination (σ_1, σ_2) of beliefs on choice-utility pairs where

$$\sigma_1 = (\textit{chocolates, } u_1) \text{ and } \sigma_2 = (0.8) \cdot (\textit{ignore, } u_2^a) + (0.2) \cdot (\textit{open, } u_2^f).$$

Then, it may be verified that (σ_1, σ_2) is a generalized Nash equilibrium. Let β_1 and β_2 be the belief hierarchies for you and Barbara, respectively, that are induced by this generalized Nash equilibrium (σ_1, σ_2). By Theorem 6.1.1 we know that these two belief hierarchies are simple and express common belief in rationality. Moreover, by construction of σ_1 and σ_2, the two belief hierarchies express common belief in p.

Note that under the belief σ_2 it is optimal for you to choose *chocolates.* Moreover, under the belief σ_1 it is optimal for Barbara to choose *open* or *ignore* if she is forgiving. Hence, with a simple belief hierarchy that expresses common belief in rationality and common belief in p, you can only rationally choose to bring *chocolates,* Barbara can only rationally *ignore* the door bell if she is *angry*, whereas she can rationally *ignore* the door bell or *open* the door if she is *forgiving.*

Since every simple belief hierarchy is symmetric, the same would hold if *simple* belief hierarchy is replaced by *symmetric* belief hierarchy.

6.4 Comparison of the Concepts

In Chapter 5 we have investigated the central concept of *common belief in rationality* for games with incomplete information. In Chapter 6 we have refined this concept by additionally imposing a *simple* belief hierarchy, a *symmetric* belief hierarchy, and a symmetric belief hierarchy using *one theory per choice-utility pair.* In Table 6.4.1 we summarize how the optimal choices under these various concepts can be characterized.

In this table we have ordered the concepts from least to most restrictive. When we talk about *optimal choices* in this table we always mean optimal with respect to a particular *utility function.* For instance, the choices that can rationally be made under common belief in rationality for a utility function u_i are those that survive the generalized iterated strict dominance procedure for this particular utility function u_i. Similarly for the other concepts in this table. In the table we did not incorporate

Common belief in rationality with ...	**Optimal choices are those that ...**
...	survive generalized iterated strict dominance
symmetric belief hierarchy	are optimal in a Bayesian equilibrium
symmetric belief hierarchy using one theory per choice-utility pair	are optimal in a canonical Bayesian equilibrium
simple belief hierarchy	are optimal in a generalized Nash equilibrium

Table 6.4.1 Comparison of the concepts in Chapters 5 and 6

Example	**Choices you can rationally make under common belief in rationality with ...** ...	a symmetric belief hierarchy	a symmetric belief hierarchy using one theory per choice-utility pair	a simple belief hierarchy
What is Barbara's favorite color? (Section 5.1.1)	blue and green (Section 5.4.3)	same	same	same (Section 6.1.3)
Chris' drawings (Section 5.4.4)	20,40,60 at $u_1^{30}, u_1^{50}, u_1^{70}$ 40,60,80 at u_1^{90} (Section 5.4.4)	same	same	same (Section 6.1.3)
Moonlight serenade (Section 6.1.3)	serenade, chocolates, Chris (Section 6.1.3)	same (Section 6.2.1)	same (Section 6.2.5)	chocolates, Chris (Section 6.1.3)
Moonlight serenade with a twist (Section 6.2.3)	serenade, chocolates, Chris (Section 6.2.3)	chocolates, Chris (Section 6.2.3)	same (Section 6.2.5)	same (Section 6.3.2)

Table 6.4.2 The four concepts in the various examples

the condition of fixed beliefs on utilities. But recall that each of the concepts in the table can be refined by imposing fixed beliefs on the players' utility functions, as we did in Sections 5.5 and 6.3.

If we compare this table with Table 4.4.1 in Chapter 4, we clearly see that the generalized iterated strict dominance procedure, Bayesian equilibrium, canonical Bayesian equilibrium and generalized Nash equilibrium are the incomplete information counterparts to the iterated elimination of strictly dominated choices, correlated equilibrium, canonical correlated equilibrium and Nash equilibrium, respectively. Indeed, the associated concepts are characterized by exactly the same conditions on the belief hierarchies. The only difference is that for games with incomplete information, such belief hierarchies also involve beliefs about the players' utility functions.

Finally, we provide in Table 6.4.2 an overview of the choices that each of these concepts selects for you in the various examples we have analyzed in Chapters 5 and 6. In the first column we list the four examples, and for each example we specify the section where it has been introduced. In the other four columns we list, for each of the four concepts, which choices it selects for you in the specific example.

We also state the section where this has been shown.

Note that in the examples "What is Barbara's favorite color?" and "Chris' drawings" we have not specified a section in the third and fourth column. The reason is that the results in these two columns follow from the results in the second and fifth column. Indeed, for both examples the choices you can rationally make under common belief in rationality are the same as those you can rationally make under common belief in rationality with a simple belief hierarchy. As such, these choices will also be the same as those you can rationally make under common belief in rationality with a symmetric belief hierarchy, with or without the one theory per choice-utility pair condition.

6.5 Proofs

6.5.1 Proofs of Section 6.1

Proof of Theorem 6.1.1. Consider the simple belief hierarchy for player i induced by the combination of beliefs $(\sigma_1, ..., \sigma_n)$. Suppose first that the belief hierarchy expresses common belief in rationality. Then, it follows from the arguments in Section 6.1.2 that $(\sigma_1, ..., \sigma_n)$ is a generalized Nash equilibrium.

Suppose next that $(\sigma_1, ..., \sigma_n)$ is a generalized Nash equilibrium. We will show, for every player, that the simple belief hierarchy generated by $(\sigma_1, ..., \sigma_n)$ expresses common belief in rationality. Construct an epistemic model $M = (T_i, v_i, b_i)_{i \in I}$ where $T_i = \{t_i^{u_i} \mid u_i \in U_i\}$, where $v_i(t_i^{u_i}) = u_i$ for every type $t_i^{u_i} \in T_i$, and where

$$b_i(t_i^{u_i})((c_j, t_j^{u_j})_{j \neq i}) := \prod\nolimits_{j \neq i} \sigma_j(c_j, u_j)$$

for every type $t_i^{u_i}$ and every opponents' choice-type combination $(c_j, t_j^{u_j})_{j \neq i}$. In particular, $t_i^{u_i}$ holds the belief σ_{-i} about the opponents' choices and utility functions, independent of the utility function that $t_i^{u_i}$ has.

Since for every player i, it is the case that all types hold the same belief σ_{-i} about the opponents' choices and utility functions, the belief hierarchy of every type $t_i^{u_i}$ will always be such that (i) it holds the belief σ_{-i} about the opponents' choices and utility functions, (ii) it believes, with probability 1, that every opponent j holds the belief σ_{-j} about his opponents' choices and utility functions, and so on. That is, the belief hierarchy of every type $t_i^{u_i}$ is simple, and is generated by the combination of beliefs $(\sigma_1, ..., \sigma_n)$.

We next show that every type in this epistemic model believes in the opponents' rationality. Consider a type $t_i^{u_i}$ and assume that $b_i(t_i^{u_i})$ assigns positive probability to some opponent's choice-type pair $(c_j, t_j^{u_j})$. Then, by construction, σ_j assigns positive probability to (c_j, u_j). As $(\sigma_1, ..., \sigma_n)$ is a generalized Nash equilibrium, the choice c_j is optimal for player j under the belief σ_{-j} if he holds the utility function u_j. Since the type $t_j^{u_j}$ indeed holds the belief σ_{-j} and the utility function u_j, we conclude that c_j is optimal for $t_j^{u_j}$. We thus conclude that $t_i^{u_i}$ believes in the opponents' rationality. Since this holds for *every* type $t_i^{u_i}$ in the epistemic model, we conclude that all types in the model express *common* belief in rationality.

Recall from above that the belief hierarchies held by the types in the model are exactly the simple belief hierarchies generated by $(\sigma_1, ..., \sigma_n)$. As such, all these simple belief hierarchies express common belief in rationality. This completes the proof. ■

Proof of Theorem 6.1.2. Follows from the arguments in Section 6.1.2. ■

Proof of Theorem 6.1.3. Follows from the arguments in Section 6.1.2. ■

6.5.2 Proofs of Section 6.2

Proof of Theorem 6.2.1. The proof is identical to the proof of Theorem 4.2.1, and is therefore omitted. ■

Proof of Theorem 6.2.2. (a) Suppose first that the belief hierarchy β_i is symmetric and expresses common belief in rationality. Then, we know by Theorem 6.2.1 that the belief hierarchy β_i is induced by a common prior π^* on choice-type combinations. Suppose that, within a beliefs diagram in choice-type representation, β_i starts at the choice-type pair (c_i^*, t_i^*). We say that a choice-type pair (c_j, t_j)

can be reached within one step from (c_i^*, t_i^*) if $\pi((c_j, t_j) \mid (c_i^*, t_i^*)) > 0$. Here, $\pi((c_j, t_j) \mid (c_i^*, t_i^*))$ denotes the probability that $\pi(\cdot \mid (c_i^*, t_i^*))$ assigns to all the opponents' choice-type combinations that contain (c_j, t_j). Say that a choice-type pair (c_j, t_j) can be reached within two steps from (c_i^*, t_i^*) if there is a pair (c_m, t_m) that can be reached within one step from (c_i^*, t_i^*) such that $\pi((c_j, t_j) \mid (c_m, t_m)) > 0$. For $k \geq 3$, we inductively define reachability with k steps as follows: Say that a choice-type pair (c_j, t_j) can be reached within k steps from (c_i^*, t_i^*) if there is a pair (c_m, t_m) that can be reached within $k-1$ steps from (c_i^*, t_i^*) such that $\pi((c_j, t_j) \mid (c_m, t_m)) > 0$.

For every player j (including i) let $(C_j \times T_j)^*$ be the sets of choice-type pairs that can be reached within finitely many steps from (c_i^*, t_i^*). Moreover, let π be the restriction of π^* to choice-type pairs in $(C_j \times T_j)^*$ given by

$$\pi((c_j, t_j)_{j \in I}) := \frac{\pi^*((c_j, t_j)_{j \in I})}{\sum\limits_{(c_j', t_j')_{j \in I} \in \times_{j \in I}(C_j \times T_j)^*} \pi^*((c_j', t_j')_{j \in I})}$$

for every $(c_j, t_j)_{j \in I} \in \times_{j \in I}(C_j \times T_j)^*$, and let $\pi((c_j, t_j)_{j \in I}) := 0$ otherwise.

Then, it may be verified that the belief hierarchy β_i is induced by the common prior π. We show that π is a Bayesian equilibrium.

For every player j, let T_j^* be the set of types that enter in $(C_j \times T_j)^*$. Assume, without loss of generality, that for every two choice-type pairs $(c_j, t_j), (c_j', t_j') \in (C_j \times T_j)^*$ with $c_j \neq c_j'$ we have that $t_j \neq t_j'$. Then, for every type $t_j \in T_j^*$ there is a unique choice $c_j[t_j] \in C_j^*$ such that $(c_j[t_j], t_j) \in (C_j \times T_j)^*$.

We create an epistemic model with sets of types T_j^* for every player j, where the beliefs of the types are given by

$$b_j(t_j)((c_m, t_m)_{m \neq j}) := \pi((c_m, t_m)_{m \neq j} \mid (c_j[t_j], t_j)) \tag{6.5.1}$$

for every $t_j \in T_j^*$, and every $(c_m, t_m)_{m \neq j} \in \times_{m \neq j}(C_m \times T_m)^*$. Moreover, the utility functions of the types are given by the utility functions they have in the beliefs diagram in choice-type representation.

Recall that the belief hierarchy β_i starts at the choice-type pair $(c_i^*, t_i^*) = (c_i[t_i^*], t_i^*)$. Then, by construction, the belief hierarchy β_i is the belief hierarchy induced by the type t_i^* within this epistemic model. We can always select the choice c_i^* such that c_i^* is optimal for t_i^*, as this does not affect the belief hierarchy β_i. Let us therefore assume, without loss of generality, that c_i^* is optimal for t_i^*. In other words, $c_i[t_i^*]$ is optimal for t_i^*.

We will now show that for every player j and every $t_j \in T_j{}^*$, the choice $c_j[t_j]$ is optimal for t_j. If $j = i$ and $t_i = t_i^*$, then we know this from our assumption above. Assume now that $t_j \neq t_i^*$. Then, $(c_j[t_j], t_j) \in (C_j \times T_j)^*$. Hence, in view of (6.5.1), there is a choice-type pair (c_m, t_m) reachable from (c_i^*, t_i^*) such that $b_m(t_m)(c_j[t_j], t_j) > 0$. As the belief hierarchy β_i expresses common belief in rationality, and β_i is the belief hierarchy held by the type t_i^*, we conclude that t_i^* expresses common belief in rationality. Since (c_m, t_m) is reachable from (c_i^*, t_i^*), it follows that t_m believes in j's rationality. As $b_m(t_m)(c_j[t_j], t_j) > 0$, it must thus be that $c_j[t_j]$ is optimal for t_j.

Now, take some player j and some $(c_j, t_j) \in C_j \times T_j^*$ with $\pi(c_j, t_j) > 0$. Then, $c_j = c_j[t_j]$. By our insights above, we thus know that $c_j[t_j]$ is optimal for t_j. By (6.5.1), the first-order belief of type t_j is $\pi(\cdot \mid (c_j, t_j))$. As c_j is optimal for t_j, it follows that c_j is optimal for utility function that t_j has and the induced first-order belief $\pi(\cdot \mid (c_j, t_j))$. We thus conclude that π is a Bayesian equilibrium. Hence, the belief hierarchy β_i is induced by a Bayesian equilibrium.

(b) Assume next that the belief hierarchy β_i is induced by a Bayesian equilibrium π. As π is a common prior on choice-type combinations, it follows by Theorem 6.2.1 that β_i is symmetric. It remains to show that β_i expresses common belief in rationality.

Suppose that β_i is generated within a beliefs diagram in choice-type representation, and that β_i starts at the choice-type pair (c_i^*, t_i^*). For every player j, let $(C_j \times T_j)^*$ be the set of choice-type pairs that enter in this beliefs diagram. Moreover, let T_j^* be the set of types that enter in the beliefs diagram. Similarly to part (a), we assume that for every $t_j \in T_j^*$ there is a unique choice $c_j[t_j]$ such that $(c_j[t_j], t_j) \in (C_j \times T_j)^*$.

We construct an epistemic model with sets of types T_j^* for every player j, and where the beliefs of the types are given by

$$b_j(t_j)((c_m, t_m)_{m \neq j}) := \pi((c_m, t_m)_{m \neq j} \mid (c_j[t_j], t_j)) \tag{6.5.2}$$

for every $t_j \in T_j^*$, and every $(c_m, t_m)_{m \neq j} \in \times_{j \neq m}(C_m \times T_m)^*$.

Recall that the belief hierarchy is induced by the Bayesian equilibrium π and starts at the choice-type pair (c_i^*, t_i^*). In view of (6.5.2), the belief hierarchy β_i is precisely the belief hierarchy held by the type t_i^*. We will now show that t_i^* expresses common belief in rationality. For this, it is sufficient to show that every type in the epistemic model above believes in the opponents' rationality.

For some player j, take a type $t_j \in T_j^*$ and an opponents' choice-type combination $(c_m, t_m)_{m \neq j} \in \times_{m \neq j}(C_m \times T_m^*)$ with $b_j(t_j)((c_m, t_m)_{m \neq j}) > 0$. Then, we know by (6.5.2) that $\pi((c_m, t_m)_{m \neq j} \mid (c_j[t_j], t_j)) > 0$. This implies that $\pi(c_m, t_m) > 0$ for every player $m \neq j$. Fix a player $m \neq j$. As π is a correlated equilibrium and $\pi(c_m, t_m) > 0$, we know that c_m is optimal for the utility function of t_m and the induced first-order belief $\pi(\cdot \mid (c_m, t_m))$. By (6.5.2) we know that t_m's first-order belief is $\pi(\cdot \mid (c_m, t_m))$. Therefore, c_m is optimal for the type t_m. We thus conclude that t_j believes in the opponents' rationality.

As such, every type in the epistemic model believes in the opponents' rationality. This, in turn, implies that every type expresses common belief in rationality. In particular, type t_i^* expresses common belief in rationality, which means that belief hierarchy β_i expresses common belief in rationality. This completes the proof. ■

Proof of Theorem 6.2.3. Follows from the arguments in Section 6.2.3. ■

Proof of Theorem 6.2.4. The proof is essentially identical to the proof of Theorem 4.3.5, and is therefore omitted. ■

Proof of Theorem 6.2.5. Follows from the arguments in Section 6.2.4. ■

Proof of Theorem 6.2.6. The proof is very similar to the proof of Theorem 4.3.3, and is therefore omitted. ■

Proof of Theorem 6.2.7. The proof is very similar to the proof of Theorem 4.3.4 and is therefore omitted. ■

Proof of Theorem 6.3.1. Like in the proof of Theorem 4.1.3, where we showed the existence of Nash equilibria, we rely on Kakutani's fixed point theorem. See the proofs section of Chapter 4 for the statement of Kakutani's fixed point theorem.

Consider the profile $p = (p_i)_{i \in I}$ of beliefs on utility functions, where $p_i \in \Delta(U_i)$ for every player i. For every player i, let $\Delta^{p_i}(C_i \times U_i)$ denote the set of beliefs on $C_i \times U_i$ where the induced belief on U_i is p_i. More precisely,

$$\Delta^{p_i}(C_i \times U_i) := \{\sigma_i \in \Delta(C_i \times U_i) \mid \sum_{c_i \in C_i} \sigma_i(c_i, u_i) = p_i(u_i) \text{ for all } u_i \in U_i\}.$$

By

$$A := \Delta^{p_1}(C_1 \times U_1) \times ... \times \Delta^{p_n}(C_n \times U_n)$$

we denote the set of all such belief combinations. Hence, A is a subset of some linear space $\mathbf{R}^X$. Moreover, it may be verified that the set A is nonempty, compact and convex.

We will now show that the set A contains a generalized Nash equilibrium $(\sigma_1, ..., \sigma_n)$. For every $(\sigma_1, ..., \sigma_n) \in A$, every player i and every utility function u_i, let $C_i^{opt}(u_i, \sigma_1, ..., \sigma_n)$ be the set of choices $c_i \in C_i$ that are optimal under the belief σ_{-i} for the utility function u_i. Let

$$\Delta^{p_i}(C_i^{opt}(\sigma_1, ..., \sigma_n)) := \{\sigma_i \in \Delta^{p_i}(C_i \times U_i) \mid \sigma_i(c_i, u_i) > 0 \text{ only if } c_i \in C_i^{opt}(u_i, \sigma_1, ..., \sigma_n)\}$$

be the set of beliefs on i's choice-utility pairs that induces the belief p_i on i's utilities, and that for every utility function u_i only assigns positive probability to choices that are optimal under σ_{-i} for the utility function u_i.

Define now the correspondence C^{opt} from A to A, which assigns to every belief combination $(\sigma_1, ..., \sigma_n) \in A$ the set of belief combinations

$$C^{opt}(\sigma_1, ..., \sigma_n) := \Delta^{p_1}(C_1^{opt}(\sigma_1, ..., \sigma_n)) \times ... \times \Delta^{p_n}(C_n^{opt}(\sigma_1, ..., \sigma_n)),$$

which is a subset of $\Delta^{p_1}(C_1 \times U_1) \times ... \times \Delta^{p_n}(C_n \times U_n)$, and hence is a subset of A.

It may be verified that the set $C^{opt}(\sigma_1, ..., \sigma_n)$ is nonempty and convex for every $(\sigma_1, ..., \sigma_n)$. It thus follows that the correspondence C^{opt} is convex-valued. It can also be shown that the correspondence C^{opt} is upper-semicontinuous, similarly to how we have shown it in the proof of Theorem 4.1.3. Altogether, we see that the set $A = \Delta^{p_1}(C_1 \times U_1) \times ... \times \Delta^{p_n}(C_n \times U_n)$ is nonempty, compact and convex, and that the correspondence C^{opt} from A to A is upper-semicontinuous and convex-valued. By Kakutani's fixed point theorem, it then follows that C^{opt} has at least one fixed point $(\sigma_1^*, ..., \sigma_n^*) \in A$. That is, there is some $(\sigma_1^*, ..., \sigma_n^*) \in A$ with

$$(\sigma_1^*, ..., \sigma_n^*) \in C^{opt}(\sigma_1^*, ..., \sigma_n^*).$$

By definition of C^{opt} this means that for every player i, we have that $\sigma_i^* \in \Delta^{p_i}(C_i^{opt}(\sigma_1^*, ..., \sigma_n^*))$. So, for every player i, the probability distribution σ_i^* induces the belief p_i on U_i, and only assigns positive probability to choice-type pairs (c_i, u_i) where c_i is optimal under σ_{-i}^* for the utility function u_i. This means, however, that $(\sigma_1^*, ..., \sigma_n^*)$ is a generalized Nash equilibrium.

For a given player i, consider the belief hierarchy β_i that is induced by the generalized Nash equilibrium $(\sigma_1^*, ..., \sigma_n^*)$. Then, we know by Theorem 6.1.1 that the belief hierarchy β_i is simple and expresses common belief in rationality. Moreover, by construction of $(\sigma_1^*, ..., \sigma_n^*)$, the belief hierarchy β_i expresses common belief in p. This completes the proof. ■

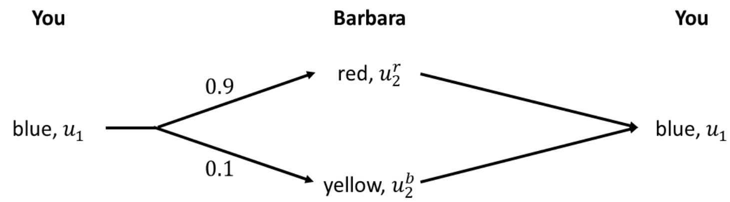

Figure 6.5.1 Beliefs diagram for Question 6.1.2

Solutions to In-Chapter Questions

Question 6.1.1. Note that Barbara's belief hierarchy that starts at her choice *blue* is the same as the one starting at her choice *red.* This is the only belief hierarchy for Barbara that is simple. It is generated by the beliefs $\sigma_1 = (green,\ u_1)$ and $\sigma_2 = (blue,\ u_2^b)$.

Question 6.1.2. Consider, for instance, your belief hierarchy as depicted in Figure 6.5.1.

Question 6.1.3. Suppose that σ_i assigns positive probability to a choice-utility pair (c_i, u_i). Then, $\tilde{\sigma}_i$ assigns positive probability to c_i. As $(\tilde{\sigma}_1, ..., \tilde{\sigma}_n)$ is a Nash equilibrium in the standard game $\tilde{\Gamma}$, where player i has utility function u_i, it follows that c_i is optimal for player i under the belief $\tilde{\sigma}_{-i}$ with the utility function u_i. Since σ_{-i} and $\tilde{\sigma}_{-i}$ induce the same belief about the opponents' choices, we conclude that c_i is optimal for player i under the belief σ_{-i} with the utility function u_i. Thus, the optimality condition in the definition of generalized Nash equilibrium is satisfied.

Question 6.1.4. Note that for you, *blue* is optimal for the belief σ_2 and the utility function u_1. Moreover, for Barbara, *yellow* is optimal for the belief σ_1 and the utility function u_2^b. Therefore, (σ_1, σ_2) is a generalized Nash equilibrium.

Question 6.1.5. For you, bidding 80 is optimal under the belief σ_2 if your utility function is u_1^{80}. Similarly, bidding 80 is optimal for Barbara under the belief σ_1 if her utility function is u_2^{80}. Hence, (σ_1, σ_2) is a generalized Nash equilibrium.

Question 6.1.6. It may first be verified that $\sigma_1 = (20,\ u_1^{30})$ and $\sigma_2 = (20,\ u_2^{30})$ is a generalized Nash equilibrium. If your valuation is 90, then your optimal bid under the belief σ_2 is 40.

Next, it may be verified that $\sigma_1 = (40,\ u_1^{50})$ and $\sigma_2 = (40,\ u_2^{50})$ is a generalized Nash equilibrium. If your valuation is 90, then your optimal bid under the belief σ_2 is 60.

Question 6.1.7. Note that under the belief σ_2, your choices *chocolates* and *Chris* are both optimal, as they both yield an expected utility of 3. Moreover, under the belief σ_1 the choice *open* is optimal for Barbara if her utility function is u_2^f and her choice *ignore* is optimal if her utility function is u_2^a. As such, (σ_1, σ_2) is a generalized Nash equilibrium.

Question 6.1.8. Suppose that the generalized Nash equilibrium (σ_1, σ_2) is such that only *Chris* is optimal for you. Then, necessarily, $\sigma_1 = (Chris,\ u_1)$. But then, the only choice for Barbara that is optimal under σ_1, no matter whether she is angry or forgiven, is *ignore.* Thus, σ_2 assigns probability 1 to Barbara's choice *ignore.* Hence, the only optimal choice for you under σ_2 is *chocolates,* and not *Chris.* This is a contradiction. We thus conclude that there is no generalized Nash equilibrium (σ_1, σ_2) where only *Chris* is optimal for you.

Question 6.1.9. Note that under the belief σ_2 your choices *chocolates* and *Chris* are both optimal, as they both yield an expected utility of 3. Moreover, under the belief σ_1 the choices *open* and *ignore* are both optimal for Barbara if her utility function is u_2^f, since they both yield an expected utility of 0.5. As such, (σ_1, σ_2) is a generalized Nash equilibrium.

Question 6.2.1. We must verify the optimality condition in the definition of a Bayesian equilibrium for the common prior π. Take some choice-type pair (c_i, t_i) that receives positive probability under π, for instance $(\textit{chocolates, } t_1^{cho})$. Then, the associated belief on the opponents' choice-type combinations conditional on $(\textit{chocolates, } t_1^{cho})$ is

$$\pi(\cdot \mid (\textit{chocolates, } t_1^{cho})) = (10/13) \cdot (\textit{ignore, } t_2^{a,i}) + (3/13) \cdot (\textit{open, } t_2^{f,o}),$$

which assigns probabilities 10/13 and 3/13 to Barbara choosing *ignore* and *open,* respectively. For this belief, it is indeed optimal for you to choose *chocolates* if your utility function is u_1. Thus, the optimality condition for $(\textit{chocolates, } t_1^{cho})$ is satisfied.

In the same way, it can be verified that the optimality condition for every other choice-type pair with positive probability is satisfied as well. Therefore, the common prior π is a Bayesian equilibrum.

Question 6.2.2. Note that $\pi(\cdot \mid (\textit{chocolates, } t_1))$ assigns probability 0.75 to *ignore* and probability 0.25 to *open.* As such, both *chocolates* and *Chris* are optimal for you under $\pi(\cdot \mid (\textit{chocolates, } t_1))$. In particular, *chocolates* is optimal for you under $\pi(\cdot \mid (\textit{chocolates, } t_1))$.

Note also that $\pi(\cdot \mid (\textit{ignore, } t_2^a))$ and $\pi(\cdot \mid (\textit{open, } t_2^f))$ both assign probability 1 to *chocolates.* Thus, *ignore* is optimal for Barbara under the belief $\pi(\cdot \mid (\textit{ignore, } t_2^a))$ and the utility function u_2^a that t_2^a has, and *open* is optimal for Barbara under the belief $\pi(\cdot \mid (\textit{open, } t_2^f))$ and the utility function u_2^f that t_2^f has. As such, we conclude that π is a Bayesian equilibrium.

Question 6.2.3. We must verify the optimality condition in the definition of a canonical Bayesian equilibrium for the common prior $\hat{\pi}$. Take some choice-utility pair (c_i, u_i) that receives positive probability under $\hat{\pi}$, for instance $(\textit{chocolates, } u_1)$. Then, the associated belief on the opponents' choice-utility combinations conditional on $(\textit{chocolates, } u_1)$ is

$$\pi(\cdot \mid (\textit{chocolates, } u_1)) = (10/13) \cdot (\textit{ignore, } u_2^a) + (3/13) \cdot (\textit{open, } u_2^f),$$

which assigns probabilities 10/13 and 3/13 to Barbara choosing *ignore* and *open,* respectively. For this belief, it is indeed optimal for you to choose *chocolates* if your utility function is u_1. Thus, the optimality condition for $(\textit{chocolates, } u_1)$ is satisfied.

In the same way, it can be verified that the optimality condition for every other choice-utility pair with positive probability is satisfied as well. Therefore, the common prior $\hat{\pi}$ is a canonical Bayesian equilibrum.

Question 6.3.1. It may be verified that no belief hierarchy in this beliefs diagram expresses common belief in p, for any profile p of beliefs on utility functions.

Question 6.3.2. We verify that (σ_1, σ_2) satisfies the optimality conditions in the definition of generalized Nash equilibrium. Take a choice-utility pair (c_i, u_i) to which σ_i assigns positive probability. For instance, $(40, u_1^{90})$. It may be verified that bidding 40 is optimal for you under the belief σ_2 if your valuation is 90. Similarly, the other optimality conditions can be verified. Therefore, (σ_1, σ_2) is a generalized Nash equilibrium.

Question 6.3.3. Consider the beliefs (σ_1, σ_2) on choice-utility pairs where

$$\sigma_1 = (0.25) \cdot (20, u_1^{30}) + (0.25) \cdot (40, u_1^{50}) + (0.25) \cdot (40, u_1^{70}) + (0.25) \cdot (60, u_1^{90})$$

and similarly for σ_2. It may be verified that (σ_1, σ_2) is a generalized Nash equilibrium.

Consider the belief hierarchy β_1 for you that is induced by the generalized Nash equilibrium (σ_1, σ_2). Then, by Theorem 6.1.1, β_1 is simple and expresses common belief in rationality. Moreover, by construction of σ_1 and σ_2, the belief hierarchy β_1 expresses common belief in p. If you hold the belief hierarchy β_1, then it is optimal for you to bid 40 if your valuation is 50, and to bid 60 if your valuation is 90. Therefore, we conclude that with a simple belief hierarchy that expresses common belief in rationality and common belief in p, you can rationally bid 40 if your valuation is 50, and rationally bid 60 if your valuation is 90.

Question 6.3.4. We verify the optimality conditions for generalized Nash equilibrium. Note first that for you, *chocolates* is optimal under the belief σ_2 and the utility function. Moreover, under the belief σ_1 Barbara's choice *ignore* is optimal for the utility function u_2^a, whereas her choice *open* is optimal for the utility function u_2^f. Hence, (σ_1, σ_2) is a generalized Nash equilibrium.

Question 6.3.5. You must assign probability 0.8 to Barbara being *angry.* Since Barbara can only rationally choose *ignore* if she is *angry,* you must assign probability at least 0.8 to Barbara *ignoring* the door bell. But then, your only optimal choice is to bring *chocolates.* Barbara, anticipating on this, will be indifferent between *open* and *ignore* if she is forgiving.

Problems

Problem 6.1: Chris' football stickers.

Remember from the example "Chris' drawings" that Chris decided to auction off the drawings he made when he was a teenager. The auction was a big success for Chris, and therefore he is going to auction the football stickers he gathered during those teenage years today. Again, you and Barbara are the only participants. Suppose that you and Barbara can bid either $20, 40, 60$ or 80 euros, and that the stickers collection goes to the person with the highest bid. In case of a tie, Chris will toss a coin to decide who gets the collection.

Different from before, the person who looses the auction must also pay his bid in euros to Chris. Such an auction is called an *all pay auction,* since everybody must pay his or her bid. If you bid b and win the auction, then your utility will be $w - b$, where w is your valuation for the stickers collection. If you loose by bidding b, then your utility will be $-b$. Similarly for Barbara.

Suppose you do not know the valuation that Barbara has, and that Barbara does not know the valuation that you have. However, it is known that your valuation and Barbara's valuation is either 30 or 50.

(a) Model this story as a game with incomplete information, by specifying the various decision problems for you. By symmetry, Barbara's decision problems will look the same.

(b) Which bids can you rationally make under common belief in rationality for each of your possible utility functions? Which procedure do you use?

(c) Show that each of the bids found in (b) can also rationally be made, for that particular utility function, under common belief in rationality with a *simple* belief hierarchy. Which theorem do you use?

(d) Without making any new computations, find the choices you can rationally make under common belief in rationality with a *symmetric* belief hierarchy, and the choices you can rationally make under common belief in rationality with a symmetric belief hierarchy using *one theory per choice-utility pair*, for each of your possible utility functions.

Consider the pair $p = (p_1, p_2)$ of beliefs on utilities, where p_1 assigns probability 0.5 to you having valuation 30 and 50, respectively, and similarly for p_2.

*__(e)__ Which bids can you rationally make under common belief in rationality and common belief in p, for each of your possible utility functions?

*__(f)__ Which bids can you rationally make under common belief in rationality and common belief in p with a *simple* belief hierarchy, for each of your possible utility functions? And with a *symmetric* belief hierarchy? And with a symmetric belief hierarchy using *one theory per choice-utility pair*? Explain your answer.

*__(g)__ Make a beliefs diagram that contains, for each of your possible utility functions u_1, and each of the choices c_1 you can rationally make for u_1 under common belief in rationality and common belief in p with a *simple* belief hierarchy, a simple belief hierarchy that expresses common belief in rationality and common belief in p, and such that c_1 is optimal for that belief hierarchy under u_1.

*__(h)__ Translate every simple belief hierarchy in this beliefs diagram into a generalized Nash equilibrium that respects p.

Problem 6.2: The ideal temperature.

Since a few weeks, Barbara and you share an office. There is, however, a problem: You both have very different preferences with respect to the ideal temperature in the office, and therefore it can be difficult to decide at which temperature to put the heating. It is well-known that your ideal temperature is 20 degrees Celcius, whereas Barbara's most preferred temperature is known to be extreme, but may differ from day to day. More precisely, Barbara's ideal temperature is either 16 degrees or 24 degrees, but you never know which of the two it will be when you go to the office.

To decide on the temperature you have invented the following: You both have a separate interface where you can enter one of the temperatures from $14, 16, 18, 20, 22$ or 24 degrees, or you can choose the option *home*, which means that you will go home and work there. If both you and Barbara have selected a temperature, the machine will calculate the average of the two, and that will be the temperature for the rest of the day. If only one person enters a temperature, the machine will adopt that temperature. If both of you choose *home*, the heating will not even turn on, as the office will remain empty.

You like working at the office, but only as long as the heating produces your ideal temperature. This is revealed in your utilities as follows: Your utility is 3 if you work at the office enjoying your ideal temperature, whereas your utility is 0 if you work at the office but the temperature if not ideal for you. If you work at home, your utility will be 2.

Barbara's conditional preferences are similar, but there is one major difference: Barbara does not like to work in the office alone. Indeed, if Barbara works at the office enjoying your presence and enjoying her ideal temperature, then her utility is 3. In all other cases her utility will be 0 if she works at the office. If she works at home, her utility will be 2.

(a) Model this story as a game with incomplete information, by specifying the decision problems for you and Barbara for each of the possible utility functions.

(b) What choices can you and Barbara rationally make under common belief in rationality for each of the possible utility functions? What procedure do you use?

(c) Make a beliefs diagram in which you support each of the choices found in (b) by a belief hierarchy that expresses common belief in rationality. In the beliefs diagram, only include the choices you found in (b), and only use solid arrows. Which of your belief hierarchies in the diagram are simple, and which are symmetric?

(d) Show that under common belief in rationality with a *simple* belief hierarchy you can only rationally make one choice. Which one? For this choice, construct a generalized Nash equilibrium for which that choice is optimal. For your other choices, explain why they cannot be rationally made under common belief in rationality with a simple belief hierarchy.

(e) Show that under common belief in rationality with a *symmetric* belief hierarchy you can only rationally make one choice. Which one? For this choice, construct a Bayesian equilibrium for which that choice is optimal. Can this choice also be supported by a canonical Bayesian equilibrium? For your other choices, explain why they cannot be rationally made under common belief in rationality with a symmetric belief hierarchy.

We will now consider the scenario where your belief about Barbara's ideal temperature is transparent between you and Barbara. More precisely, let $p = (p_1, p_2)$ be the pair of beliefs on utilities, where p_1 assigns probability 1 to your ideal temperature 20, and p_2 assigns probability 0.5 to Barbara having the ideal temperature 16 and 24, respectively.

*(**f**) Show that under common belief in rationality and common belief in p, there is only one choice you can rationally make, and there is only one choice that Barbara can rationally make for each of her utility functions. What do you expect to happen in this scenario?

*(**g**) In fact, there is only one belief hierarchy for you that expresses common belief in rationality and common belief in p. Make a beliefs diagram that represents this belief hierarchy. Is this belief hierarchy simple? Is it symmetric? Does it use one theory per choice-utility pair?

Problem 6.3: The tennis match.

Next week there will be a tennis match between Barbara and you, and Chris will be the referee. The question is: How many hours will you practice this week before the match? To keep things simple, suppose that Barbara and you can choose between practicing for 10 hours, for 20 hours, and not practicing at all. Clearly, you both would like to win, but training is costly because you could have been doing other things instead, which Barbara and you would deem even nicer.

More concretely, starting from a baseline utility of 0, winning the match would increase your utility by 30, whereas every hour you practice would decrease your utility by 1. Similarly for Barbara. However, you both disagree about your chances of winning. You are not so confident about your tennis skills, and think that you will only be able to win if you practice at least 10 hours more than Barbara. In all other cases you believe to lose. Barbara, on the other hand, is much more confident. If she is in a *confident* mood, she believes to win if she practices the same, or more, than you do. However, there are also days when she is *arrogant.* On those days she believes that she will always win, no matter how much you and Barbara practice.

Finally, you strongly dislike practicing more than necessary. Indeed, if you practice 20 hours more than Barbara, then you believe you could also have won by practicing less. In that case, your utility will decrease by 70. Barbara, on the other hand, does not mind practicing more than necessary.

When you decide on your practice schedule, you do not know whether Barbara is in a *confident* or *arrogant* mood.

(**a**) Model this story as a game with incomplete information, by writing down the decision problems for you and Barbara, for each of the possible utility functions.

(**b**) Explain why under common belief in rationality you can rationally choose each of your practice schedules. Make a beliefs diagram with solid arrows only where each of your choices is supported by a belief hierarchy that expresses common belief in rationality. Which of your belief hierarchies are simple? Which are symmetric?

(**c**) Consider the beliefs diagram in Figure 6.5.2. Explain why all belief hierarchies in this beliefs diagram are *symmetric,* by constructing a symmetric *weighted* beliefs diagram that induces it.

(**d**) Translate the symmetric weighted beliefs diagram from (c) into a common prior $\hat{\pi}$ on choice-utility combinations, and explain why $\hat{\pi}$ is a *canonical Bayesian equilibrium.*

(**e**) Explain why you can rationally choose each of your practice schedules under common belief in rationality with a *symmetric* belief hierarchy that uses *one theory per choice-utility pair.*

(**f**) Explain why you cannot rationally choose to practice for 20 hours under common belief in rationality with a *simple* belief hierarchy.

(**g**) Explain why under common belief in rationality with a *simple* belief hierarchy, you can rationally choose to practice for 0 or 10 hours.

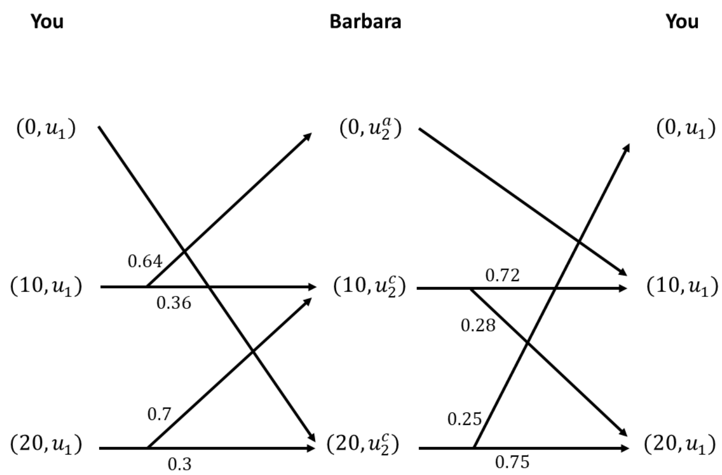

Figure 6.5.2 Beliefs diagram for Problem 6.3 (c)

Suppose now that you assign equal probability to Barbara being *confident* and *arrogant,* and that this belief is transparent between you and Barbara. That is, we consider the pair of beliefs on utilities $p = (p_1, p_2)$ where p_1 assigns probability 1 to your unique utility function, and p_2 assigns probability 0.5 to Barbara being *confident* and *arrogant,* respectively.

***(h)** Show that under common belief in rationality and common belief in p, you can only rationally plan one of the possible practice schedules. Which one? What is the probability by which you believe to win? What is the probability by which Barbara believes to win, depending on her mood?

Literature

Generalized Nash equilibrium. The notion of *generalized Nash equilibrium* is due to Bach and Perea (2020b, 2023b). In the latter paper it is shown that the incomplete information counterpart to Nash equilibrium is *not* Bayesian equilibrium, as it is often assumed in the literature. Indeed, it is shown that Bayesian equilibrium is characterized by the conditions of common belief in rationality and a common prior, just like *correlated* equilibrium in standard games. In that light, Bayesian equilibrium is the counterpart to correlated equilibrium, and not to Nash equilibrium. In reponse, Bach and Perea (2020b, 2023b) define the new concept of generalized Nash equilibrium, and show that it can be characterized by common belief in rationality and a simple belief hierarchy. More precisely, it is shown that a choice is optimal for a utility function in a generalized Nash equilibrium precisely when it can rationally be made for this utility function under common belief in rationality with a simple belief hierarchy. This resembles our Theorems 6.1.1 and 6.1.2. As these are precisely the conditions that characterize Nash equilibrium for standard games, the concept of generalized Nash equilibrium is the incomplete information counterpart to Nash equilibrium.

In Bach and Perea (2020b) it is shown that the concept of generalized Nash equilibrium can also be characterized by conditions on belief hierarchies that do *not* imply common belief in rationality. That is, not all layers of common belief in rationality are needed to arrive at generalized Nash equilibrium.

Bayesian equilibrium. The concept of *Bayesian equilibrium* has been introduced by Harsanyi (1967-68), and it has opened the door towards a systematic analysis of games with incomplete information. Harsanyi's original definition is somewhat different from ours, as it has been defined in a different framework. Instead of defining a Bayesian equilibrium as a probability distribution over choice-type combinations, as we do, he starts from a model with *Harsanyi types.* As explained in the literature sections of Chapters 3 and 4, a Harsanyi type h_i for player i specifies (i) a utility function for player i, (ii) a randomized choice for player i, and (iii) a probabilistic belief over the opponents' Harsanyi types. Similarly to how it works for epistemic models with types in this book, we can derive for every Harsanyi type a complete belief hierarchy over the players' choices and utility functions.

Harsanyi then defines a Bayesian equilibrium as a probability distribution over combinations of Harsanyi types such that, for every Harsanyi type h_i for player i selected with positive probability, every choice prescribed with positive probability must be optimal, given the prescribed utility function, and the induced belief about the opponents' choice combinations. If we replace Harsanyi types by types in our setting, and model the choices seperately from the types, then we essentially obtain our definition of a Bayesian equilibrium. The reason we define a Bayesian equilibrium in this way, and thus different from Harsanyi's definition, is that we want to define all of our concepts in a unified framework, to make the comparisons between the different concepts more direct and more transparent.

However, Bach and Perea (2023b) show that our definition is behaviorally equivalent to Harsanyi's definition. Indeed, in that paper it is shown that the choices that are optimal, for a given utility function, in a Bayesian equilibrium (as defined by Harsanyi) are precisely the choices that can rationally be made for that utility function under common belief in rationality with a common prior. Since we have seen that a common prior is equivalent to assuming a symmetric belief hierarchy, it follows from our Theorem 6.2.3 that these are exactly the choices that are optimal in a Bayesian equilibrium as we define it. Thus, if we are interested in the choices that players can rationally make for a given utility function, it does not matter whether we use Harsanyi's original definition or our definition of a Bayesian equilibrium.

Canonical Bayesian equilibrium. We have seen that the choices that can rationally be made under common belief in rationality with a symmetric belief hierarchy that uses *one theory per choice-utility pair* are precisely the choices that are optimal in a *canonical* Bayesian equilibrium. The definition of a canonical Bayesian equilibrium is based on the notion of incomplete information correlated equilibrium in Bergemann and Morris (2007), which in turn is based on one of the versions of correlated equilibrium for incomplete information in Forges (1993), and on the equivalent notion of *simplified* Bayesian equilibrium in Bach and Perea (2017). However, in that paper the concept is used for a different purpose: It is shown that the probability distributions on choice-utility combinations that are induced by Bayesian equilibria are exactly the simplified Bayesian equilibria.

Bayesian equilibria as Nash equilibria of a different game. Interestingly, the definition of a Bayesian equilibrium that typically appears in textbooks and papers is fundamentally different from both Harsanyi's definition and ours. The textbook definition fixes a common prior on the players' Harsanyi types, and lets every player choose a mapping which assigns to each of his possible Harsanyi types a (randomized) choice. The utility that a player gets is the expected utility induced by (i) the common prior on the players' Harsanyi types, and (ii) the players' choice mappings. A Bayesian equilibrium is then defined as a combination of choice mappings that constitutes a Nash equilibrium in this modified game with complete information.

Harsanyi (1967-68) has shown that the Nash equilibria in this modified game correspond one-to-one to the Bayesian equilibria of the original game with incomplete information. Harsanyi then uses this result to show the existence of Bayesian equilibria, relying on Nash's (1950, 1951) equilibrium existence theorem for games with complete information.

However, from a conceptual point of view this modified game seems very problematic. First, it seems unnatural for a player to select a choice for *every* Harsanyi type he can possible have, because the player *knows* which belief hierarchy and utility function he has, and therefore the player *knows* his Harsanyi type. Second, the modified game with complete information is much more complicated, and much less transparent, than the original game with incomplete information. It therefore seems more reasonable to use Harsanyi's original definition of a Bayesian equilibrium, or a definition that perfectly resembles Harsanyi's original definition, as we do in this book.

Due to Harsanyi's result mentioned above, which shows that Bayesian equilibria are equivalent to the Nash equilibria of some artificial transformed game, it is often argued that Bayesian equilibrium is the incomplete information counterpart to Nash equilibrium. However, as this chapter shows, this statement is false: Bayesian equilibrium is the incomplete information counterpart to *correlated* equilibrium.