Notes on

Chapter 1. *Bayesian Games in Normal Form*

UNCERTAINTY AND CONTRACTS

3rd Year

Degree in Economics
Double Bachelor's Degree in Business and Economics
2022/2023

Iñaki Aguirre

Department of Economic Analysis

University of the Basque Country UPV/EHU

Chapter 1. **Bayesian Games in Normal Form**

Chapter 9, *Bayesian Games* (9.1, 9.2, 9.3 and 9.4) *M. Osborne, An Introduction to Game Theory*

Introduction

In the subject *Market Power and Strategy*, we assume complete information. That is, each player has to play in a game with perfect knowledge about her rival’s preferences and strategy spaces. In this chapter, we relax this assumption.

Underlying the notion of Nash equilibrium is that each player holds the correct belief about the other players’ actions. To do so, a player must know the game she is playing; in particular, she must know the other players’ preferences. In many contexts, the agents are not perfectly informed about their rivals’ characteristics: bargainers may not know each others’ valuations of the object of negotiation, firms may not know each others' cost functions, a monopolistic firm may not know consumers preferences, etcetera. In some situations, a participant may be well informed about her opponents' characteristics, but may not know how well these opponents are informed about her own characteristics. In this chapter, we describe the model of a "Bayesian game", which generalizes the notion of a strategic game to allow us to analyze any situation in which each player is imperfectly informed about an aspect of her environment that is relevant to her choice of an action.

1.1. **Motivational examples**

We start with a couple of examples that serve to illustrate the main ideas in a Bayesian game. We will define the notion of Nash equilibrium separately for each game. In the next section, we will define the general model of a Bayesian game and the notion of Nash equilibrium for such a game.

Example 1: (Battle of the Sexes) *Bach or Stravinsky?*

Two people wish to go out together. Two concerts are available: one of music by Bach, and one of music by Stravinsky. One person prefers Bach and the other prefers Stravinsky. If they go to different concerts, each of them is equally unhappy listening to the music of either composer. Player 1 (the one that prefers Bach) is the row player and player 2 (who prefers Stravinsky) is the column player. The game in normal form is:

	$Bach$	$Stravinsky$
$Bach$	($\underline{2}$, $\underline{1}$)	(0, 0)
$Stravinsky$	(0, 0)	($\underline{1}$, $\underline{2}$)

In this game, there are two Nash equilibria in pure strategies: (*Bach*, *Bach*) and (*Stravinsky*, *Stravinsky*).

Example 2 (273.1): *Bach or Stravinsky? Variant of BoS with imperfect information 1*

Consider a variant of *BoS* in which player 1 is unsure whether player 2 prefers to go out with her or prefers to avoid her, whereas player 2, as before, knows player 1's preferences. Assume that player 1 thinks that with probability ½ player 2 wants to go out with her and with probability ½ player 2 wants to avoid her. That is, player 1 thinks that with probability ½ she is playing the game on the left in the next figure and with probability ½ she is playing the game on the right.

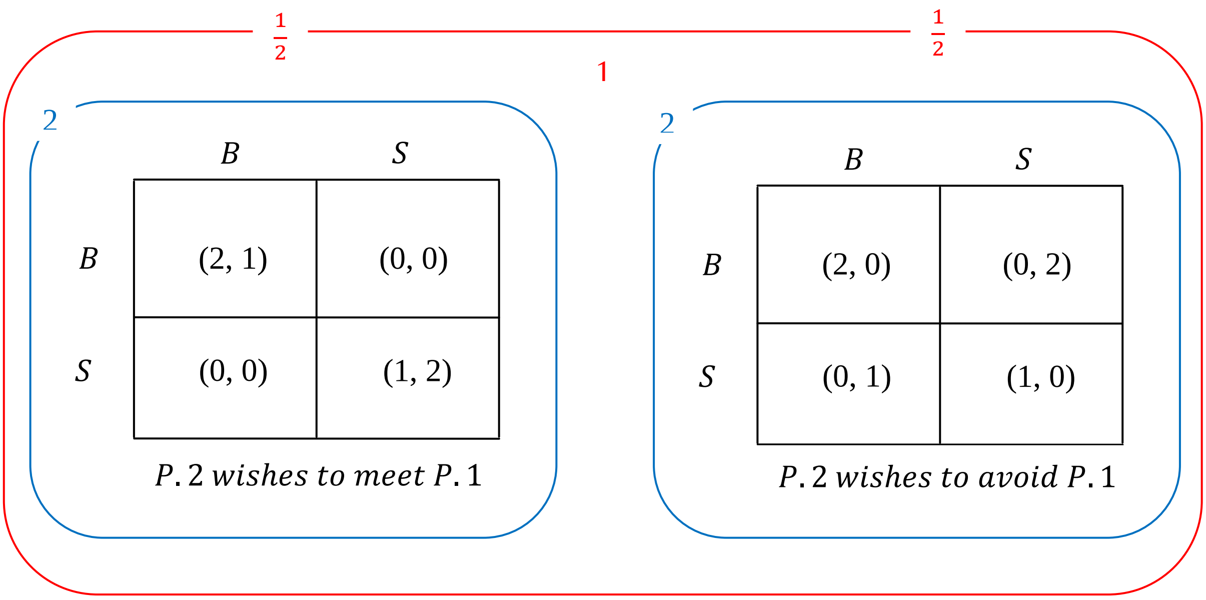

We can think of there being two *states* (of Nature), one in which payoffs are given in the left table and one in which payoffs are given in the right table. Player 2 knows the state (she knows whether she wishes to meet or to avoid player 1) whereas player 1 does not know; player 1 assigns probability ½ to each state.

The notion of Nash equilibrium for a strategic game models a steady state in which each player's beliefs about the other players' actions are correct, and each player acts optimally, given her beliefs. We want to generalize this notion to the current situation.

From player 1's point of view, player 2 has two possible *types*, one whose preferences are given in the left table and one whose preferences are given in the right table. Player 1 does not know player 2's type, so to choose an action rationally she needs to form a belief about the action of each type. Given these beliefs and her belief about the likelihood of each type, she can calculate her expected payoff of each of her actions. We next calculate the expected payoff of each one of player 1's actions corresponding to each combination of actions of the two types of player 2.

	(B, B)	(B, S)	(S, B)	(S, S)
B	2	1	1	0
S	0	$\frac{1}{2}$	$\frac{1}{2}$	1

For this situation, we define a pure strategy Nash equilibrium to be a triple of actions, one for player 1 and one for each type of player 2, with the property that

- the action of player 1 is optimal, given the actions of the two types of player 2 (and player 1's belief about the state).

- the action of each type of player 2 is optimal, given the action of player 1.

Now we obtain best responses of player 1 (against the possible actions of the two types of player 2) and best responses of each type of player 2 (against the actions of player 1).

	(*B*, *B*)	(*B*, *S*)	(*S*, *B*)	(*S*, *S*)
B	2	1	1	0
S	0	$\frac{1}{2}$	$\frac{1}{2}$	1

	B	*S*
B	(, 1)	(, 0)
S	(, 0)	(, 2)

$P.2$ *wishes to meet* $P.1$

	B	*S*
B	(, 0)	(, 2)
S	(, 1)	(, 0)

$P.2$ *wishes to avoid* $P.1$

It is easy to show that (*B*, (*B*, *S*)), where the first component is the action of player 1 and the other component is the pair of actions of the two types of player 2, is a Nash equilibrium.

Given that the actions of the two types of player 2 are (*B*, *S*), player 1's action *B* is optimal (that is, it maximizes her expected payoff); given that player 1 chooses *B*, *B* is optimal for the type who wishes to meet player 1 and S is optimal for the type who wishes to avoid player 1.

s_1	BR_2	s_2	BR_1
B	(*B*, *S*)	(*B*, *B*)	*B*
S	(*S*, *B*)	(*B*, *S*)	*B*
		(*S*, *B*)	*B*
		(*S*, *S*)	*S*

We can interpret the actions of the two types of player 2 to reflect player2's intentions in the hypothetical situation *before* she knows the state. We can tell the following story. Initially player 2 does not know the state: she is informed of the state by a *signal* that depends on the state. Before receiving this signal, she plans an action for each possible signal. After receiving the signal, she carries out her planned action for hat signal. We can tell a similar story for player 1. To be consistent with her not knowing the state when she takes an action, her signal must be uninformative; it must be the same in each state. Given her signal, she is unsure of the state; when choosing an action she takes into account her belief about the likelihood of each state given her signal.

Example 2 (273.1): *Bach or Stravinsky? Variant of BoS with imperfect information 2*

Consider a variant of *BoS* in which neither player knows whether the other wants to go out with her. Specifically, suppose that player 1 thinks that with probability ½ player 2 wants to go out with her, and with probability ½ player 2 wants to avoid her, and player 2 thinks that with probability 2/3 player 1 wants to go out with her and with probability 1/3 player 1 wants to avoid her. As before, assume that each player knows her own preferences.

We can model this situation by introducing 4 states, one for each possible configuration of preferences. We refer to these states as:

- *yy*: each player wants to go out with the other.
- *yn*: player 1 wants to go out with player 2 but player 2 wants to avoid player 1.
- *ny*: player 1 wants to avoid player 2 and player 2 wants to go out with player 1.
- *nn*: both players want to avoid the other.

The fact that player 1 does not know player 2's preferences means that she cannot distinguish between states *yy* and *yn*, or between states *ny* and *nn*. Similarly, player 2 cannot distinguish between states *yy* and *ny*, or between states *yn* and *nn*. We can model the players' information by assuming that each player receives a *signal* before choosing an action. Player 1 receives the same signal, say y_1, in states *yy* and *yn*, and a different signal, say n_1, in states *ny* and *nn*; player 2 receives the same signal, say y_2, in states *yy* and *ny*, and a different signal, say n_2, in states *yn* and *nn*. After player 1 receives the signal y_1, she is referred to as type y_1 of player 1 (who wishes to go out with player 2); after she receives the signal n_1, she is referred to as type n_1 of player 1 (who wishes to avoid player 2). In a similar way, player 2 has two types: y_2 and n_2.

Type y_1 of player 1 believes that the probability of each of the states *yy* and *yn* is ½; type n_1 of player 1 believes that the probability of each of the states *ny* and *nn* is ½. Type y_2 of player 2 believes that the probability of state *yy* is 2/3 and that of state *ny* is 1/3; type n_2 of player 2 believes that the probability of state *yn* is 2/3 and that of state *nn* is 1/3. We can represent the game as:

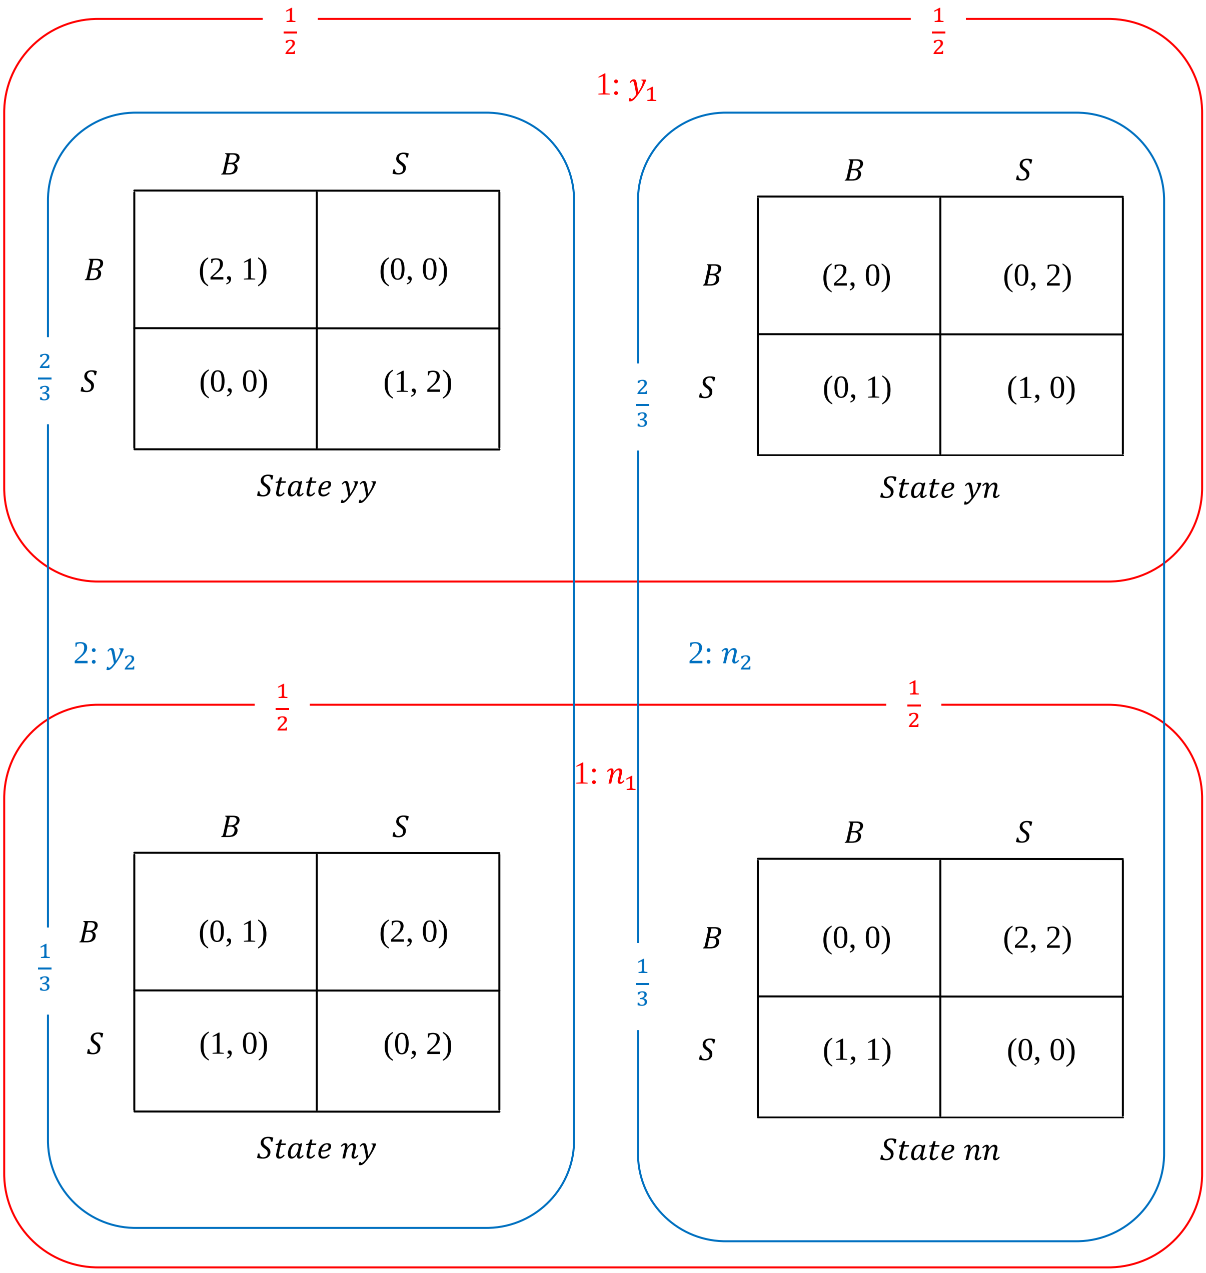

y_1:

	(B, B)	(B, S)	(S, B)	(S, S)
B	2	1	1	0
S	0	$\frac{1}{2}$	$\frac{1}{2}$	1

n_1:

	(B, B)	(B, S)	(S, B)	(S, S)
B	0	1	1	2
S	1	$\frac{1}{2}$	$\frac{1}{2}$	0

y_2:

	B	S
(B, B)	1	0
(B, S)	$\frac{2}{3}$	$\frac{2}{3}$
(S, B)	$\frac{1}{3}$	$\frac{4}{3}$
(S, S)	0	2

n_2:

	B	S
(B, B)	0	2
(B, S)	$\frac{1}{3}$	$\frac{4}{3}$
(S, B)	$\frac{2}{3}$	$\frac{2}{3}$
(S, S)	1	0

s_1	BR_2
(B, B)	(B, S)
(B, S)	(B, S)
	(S, S)
(S, B)	(S, B)
	(S, S)
(S, S)	(S, B)

s_2	BR_1
(B, B)	(B, S)
(B, S)	(B, B)
(S, B)	(B, B)
(S, S)	(S, B)

Nash equilibria: $((B, B), (B, S))$ and $((S, B), (S, S))$.

In each of these examples, a Nash equilibrium is a list of actions, one for each type of each player, such that the action of each type of each player is a best response to the actions of all the types of the other player, given the player's beliefs about the state after she observes her signal. We may define a Nash equilibrium in each example to be a Nash equilibrium of the strategic game in which the set of players is the set of all types of all players in the original situation.

In the next section, we define the general notion of a Bayesian game and the notion of Nash equilibrium in such game.

1.2. **General definitions**

1.2.1. *Bayesian games*

A strategic game with imperfect information is called a “Bayesian game”. As in a strategic game, the decision-makers are called *players*, and each player is endowed with a set of *actions*.

A key element in the specification of the imperfect information is the set of *states*. Each state is a complete description of one collection of the players’ relevant characteristics, including both their preferences and their information. For every collection of characteristics that some player believes to be possible, there must be a state. For instance, consider the first example of BoS and assume that player 2 wishes to meet player 1. In this case, the reason for including in the model the state in which player 2 wishes to avoid player 1, is that player 1 believes such a preference to be possible.

At the start of the game, a state is realized. The players do not observe this state. Rather, each player receives a *signal* that may give her some information about the state. Denote the signal player i receives in state w by $\tau_i(w)$. The function τ_i is called player i’s *signal function*. If, for example, $\tau_i(w)$ is different for each value of w, then player i knows, given her signal, the state that has occurred; after receiving her signal, she is perfectly informed about all the players’ relevant characteristics. At the other extreme, if $\tau_i(w)$ is the same for all states, then player i’s signal conveys no information about the state. If $\tau_i(w)$ is constant over some subsets of the set of states, but is not the same for all states, then player i’s signal conveys partial information. For example, if there are three states, w_1, w_2 and w_3, and $\tau_i(w_1) \neq$

$\tau_i(w_2) = \tau_i(w_3)$, then when the state is w_1 player *i* knows that it is w_1, whereas when it is either w_2 or w_3 she knows only that it is one of these two states.

We refer to player *i* in the event that she receives the signal t_i as **type $\boldsymbol{t_i}$** of player *i*. Each type of each player holds a *belief* about the likelihood of the states consistent with her signal. If, for example, $t_i = \tau_i(w_1) = \tau_i(w_2)$, then type t_i of player *i* assigns probabilities to w_1 and w_2.

Each player may care about the actions chosen by the other players, as in a strategic game with perfect information, and also about the state. The players may be uncertain about the state, so we need to specify their preferences regarding probability distributions over pairs (*a*, *w*) consisting of an action profile *a* and a state *w*. We assume that each player's preferences over such probability distributions are represented by the expected value of a *Bernoulli payoff function*. We specify each player *i*'s preferences by giving a Bernoulli payoff function u_i over pairs (*a*, *w*).

In summary, a Bayesian game is defined as follows.

Definition 1. *Bayesian game*

A Bayesian game consists of

- a set of **players**,

- a set of **states**

and for each player

- a set of **actions**

- a set of **signals** that she may receive and a signal function that associates a signal with each state

- for each signal she may receive, a **belief** about the states consistent with the signal (a probability distribution over the set of states with which the signal is associated)

- a **Bernoulli payoff function** over pairs (a, w), where a is an action profile and w is a state, the expected value of which represents the player's preferences among lotteries over the set of such pairs.

Example 2 (273.1): *Bach or Stravinsky? Variant of BoS with imperfect information 1*

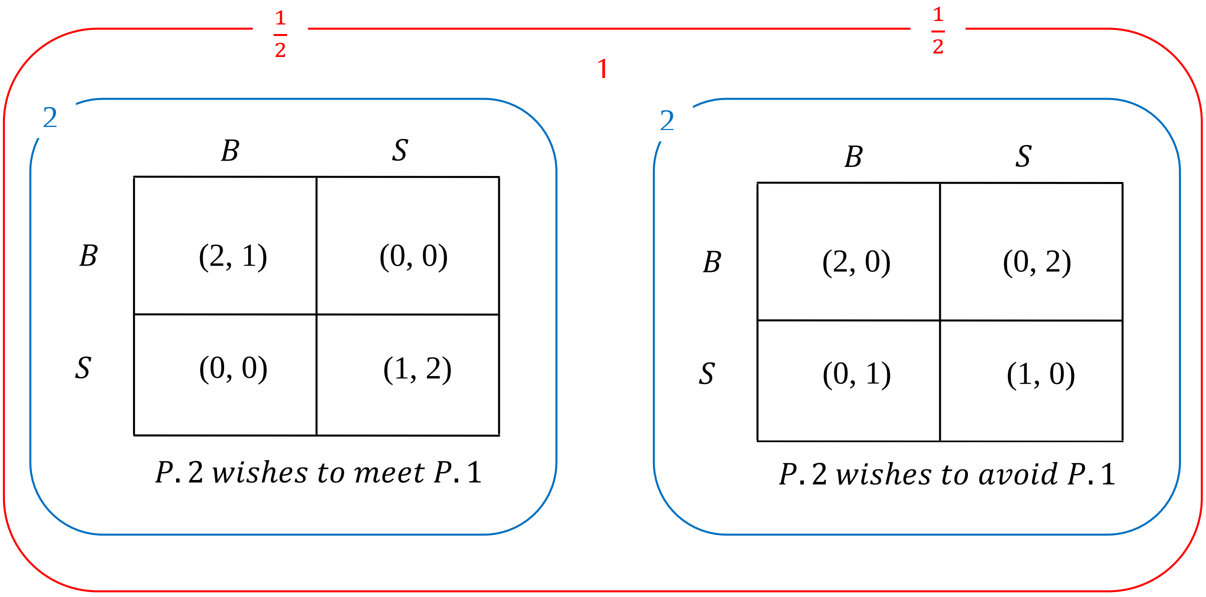

Players: 1 and 2

States: The set of states is {*meet*, *avoid*}

Actions: The set of actions of each player is {*B*, *S*}

Signals: Player 1 may receive a single signal, say z; her signal function τ_1 satisfies $\tau_1(meet) = \tau_1(avoid) = z$ (one type of player 1). Player 2 receives one of two signals, say

m and v; her signal function τ_2 satisfies $\tau_2(meet) = m$ and $\tau_2(avoid) = v$. (two types of player 2)

Beliefs: Player 1 assigns probability ½ to each state after receiving the signal *z*. Player 2 assigns probability 1 to the state meet after receiving the signal *m*, and probability 1 to the state *avoid* after receiving the signal *v*.

Payoffs: The payoffs $u_i(a, meet)$ of each player *i* for all possible action pairs are given in the left table and the payoffs $u_i(a, avoid)$ are given in the right table.

Example 2 (273.1): *Bach or Stravinsky? Variant of BoS with imperfect information 2*

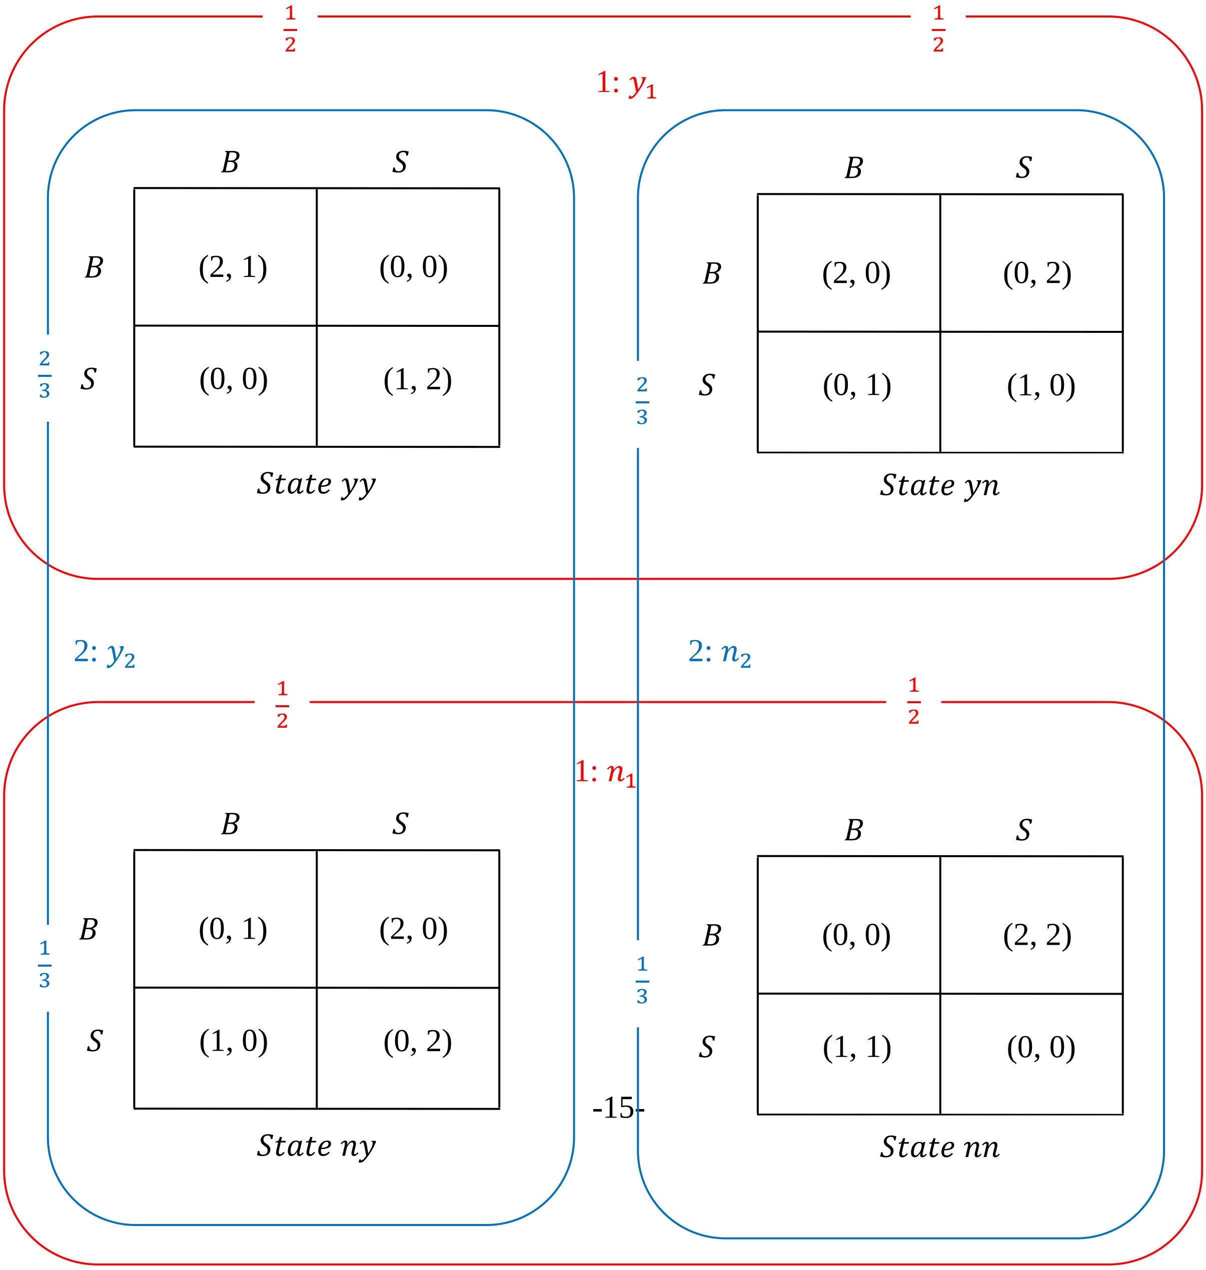

Players: 1 and 2

States: The set of states is {*yy*, *yn*, *ny*, *nn*}

Actions: The set of actions of each player is $\{B, S\}$

Signals: Player 1 receives one of two signals, y_1 and n_1; her signal function τ_1 satisfies $\tau_1(yy) = \tau_1(yn) = y_1$ and $\tau_1(ny) = \tau_1(nn) = n_1$ (two types of player 1). Player 2 receives one of two signals, say y_2 and n_2; her signal function τ_2 satisfies $\tau_2(yy) = \tau_2(ny) = y_2$ and $\tau_2(yn) = \tau_2(nn) = n_2$ (two types of player 2).

Beliefs: Player 1 assigns probability ½ to each of the states *yy* and *yn* after receiving the signal y_1 and probability ½ to each of the states *ny* and *nn* after receiving the signal n_1. Player 2 assigns probability 2/3 to the state *yy* and probability 1/3 to the state *ny* after receiving the signal y_2, and probability 2/3 to the state *yn* and probability 1/3 to the state *nn* after receiving the signal n_2.

Payoffs: The payoffs $u_i(a, w)$ of each player *i* for all possible action pairs and states are given in previous figure.

1.2.2. *Nash equilibrium*

In a Bayesian game each type of each player chooses an action. In a Nash equilibrium of such a game, the action chosen by each type of each player is optimal (that is, maximizes her expected payoff), given the actions chosen by every type of every other player.

Example 3: *Fighting an opponent of unknown strength*

Two people are involved in a dispute. Person 1 does not know whether person 2 is strong or weak; she assigns probability α to person 2's being strong. Person 2 is fully informed. Each person can either fight or yield. Each person's preferences are represented by the expected value of a Bernoulli payoff function that assigns the payoff of 0 if she yields (regardless of the other person's action) and a payoff of 1 if she fights and her opponent yields; if both people fight, then their payoff are (-1, 1) if person 2 is strong and (1, -1) if person 2 is weak. Formulate this situation as a Bayesian game and find its Nash equilibria if $\alpha < \frac{1}{2}$ and if $\alpha > \frac{1}{2}$.

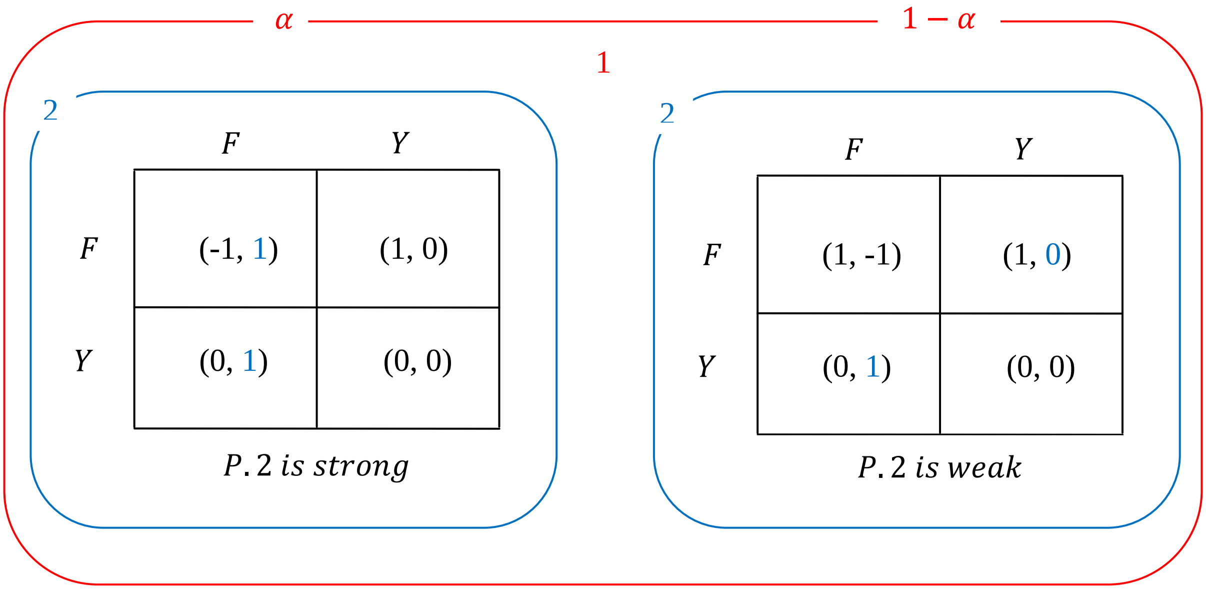

Players: Person 1 and person 2

States: The set of states is { $strong$, }

Actions: The set of actions of each player is $\{F, Y\}$

Signals: Person 1 receives one signal m that is not informative; her signal function τ_1 satisfies $\tau_1(strong) = \tau_1(weak) = m$ (one type of player 1). Person 2 receives one of two signals, say s and w; her signal function τ_2 satisfies $\tau_2(strong) = s$ and $\tau_2(weak) = w$ (two types of player 2)

Beliefs: Player 1 assigns probability α to state $strong$ and $1-\alpha$ to state $weak$ after receiving the signal m. Player 2 assigns probability 1 to state $strong$ after receiving the signal s. After receiving the signal w, she assigns probability 1 to state $weak$.

Payoffs: The profits of each firm for all possible action pairs and any possible state are as appear in payoffs matrices.

Strategies: $S_1 = \{F, Y\}$ and $S_2 = \{(F,F),(F,Y),(Y,F),(Y,Y)\}$.

	(*F*, *F*)	(*F*, *Y*)	(*Y*, *F*)	(*Y*, *Y*)
F	$1-2\alpha$	$1-2\alpha$	1	1
Y	0	0	0	0

$\alpha < \frac{1}{2}$

s_1	BR_2	s_2	BR_1
F	(*F*, *Y*)	(*F*, *F*)	*F*
Y	(*F*, *F*)	(*F*, *Y*)	*F*
		(*Y*, *F*)	*F*
		(*Y*, *Y*)	*F*

NE: (*F*, (*F*, *Y*))

$\alpha > \frac{1}{2}$

s_1	BR_2	s_2	BR_1
F	(F, Y)	(F, F)	Y
Y	(F, F)	(F, Y)	Y
		(Y, F)	F
		(Y, Y)	F

NE: $(Y, (F, F))$.

$\alpha = \frac{1}{2}$

s_1	BR_2	s_2	BR_1
F	(F, Y)	(F, F)	F, Y
Y	(F, F)	(F, Y)	F, Y
		(Y, F)	F
		(Y, Y)	F

NE: $(F, (F, Y))$ and $(Y, (F, F))$.

Example 3: *Fighting an opponent of unknown strength 2*

Assume now that person 2 does not either know whether person 1 is medium strength or super strong; she assigns probability β to person 1's being medium strength. Each person can either fight or yield. If player 1 is medium strength payoffs are (-1, 1) when both fight and person 2 is strong and (1, -1) when both fight and person 2 is weak. If player 1 is super strong, payoff are (1, -1) when both fight independently whether person 2 is strong or weak. In the rest of cases payoffs are as those in the previous game. That is, if one person fights, and other does not, then she obtains a payoff 1. If one person does not fight, independently of the rival's behavior, then she obtains 0.

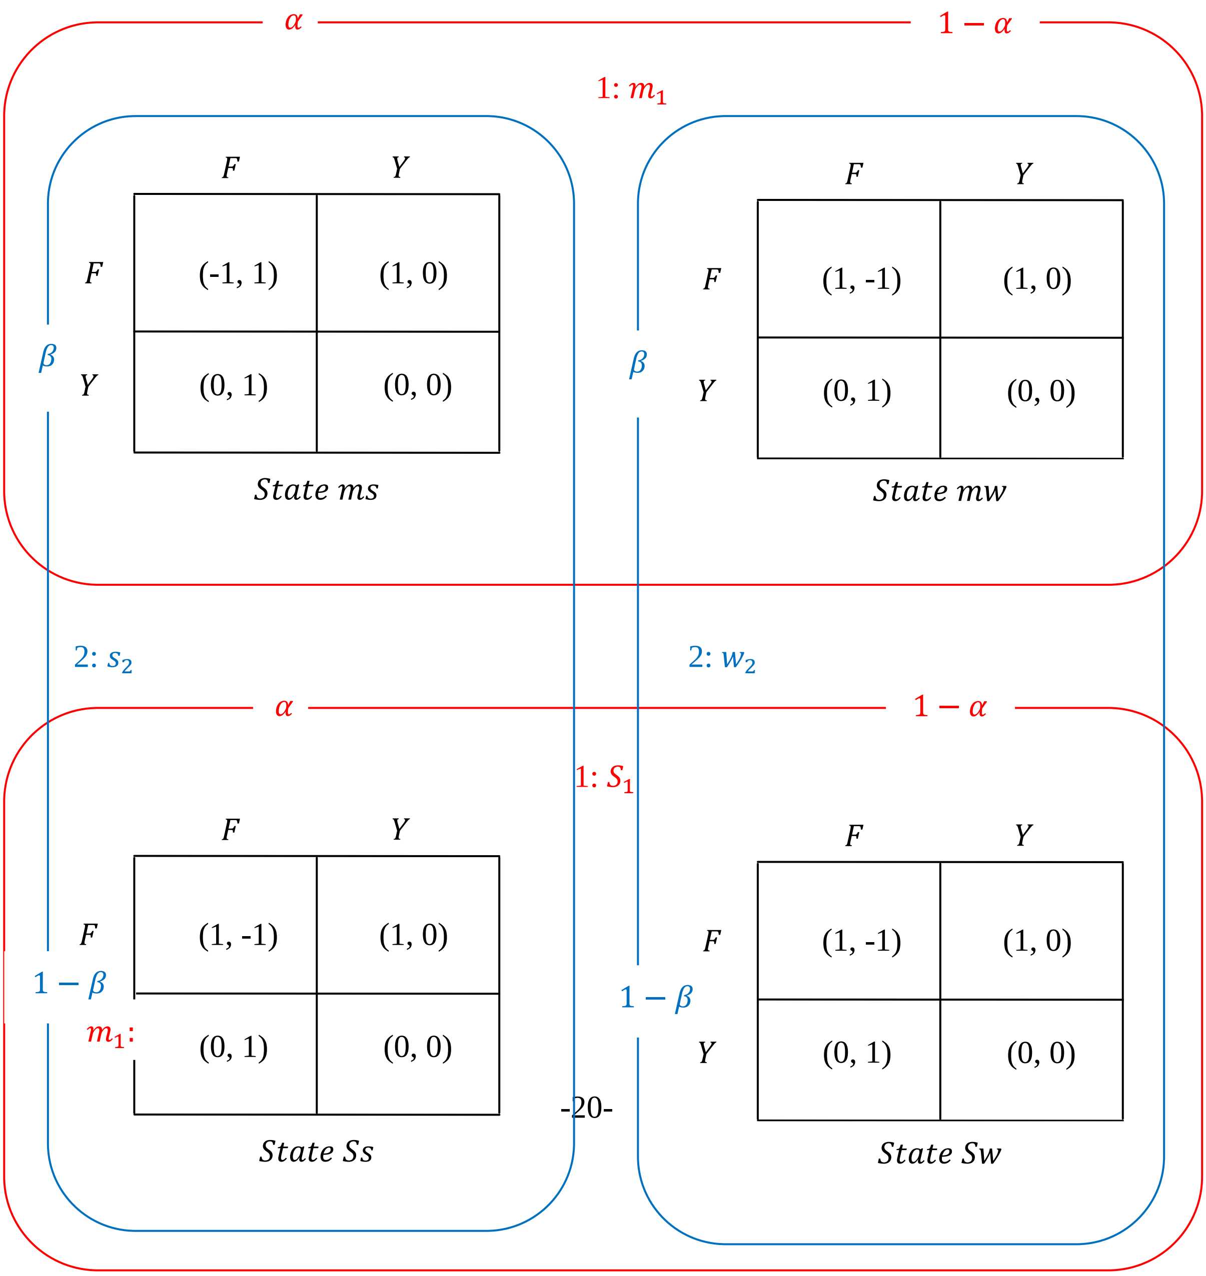

Represent the Bayesian game and identify the main elements of the Bayesian game. Obtain the Nash equilibria to any $\alpha, \beta \in [0,1]$. (Problem 5).

1.3. **Two examples concerning information**

The notion of a Bayesian game may be used to study how information patterns affect the outcome of strategic interaction. Here we consider two examples.

1.3.1. *More information may hurt*

A decision-maker in a single-person decision-problem cannot be worse off if she has more information: if she wishes, she can ignore the information. In a game the same is not true: if a player has more information and the other players know that she has more information, then she may be worse off.

Consider, for example, the following two-player Bayesian game where $0 < \epsilon < \frac{1}{2}$.

Player 1: $\frac{1}{2}$, $\frac{1}{2}$ — Player 2: $\frac{1}{2}$, $\frac{1}{2}$

	L	M	R
T	$(1, 2\epsilon)$	$(1, 0)$	$(1, 3\epsilon)$
B	$(2, 2)$	$(0, 0)$	$(0, 3)$

State w_1

	L	M	R
T	$(1, 2\epsilon)$	$(1, 3\epsilon)$	$(1, 0)$
B	$(2, 2)$	$(0, 3)$	$(0, 0)$

State w_2

In this game, there are two states, and neither player knows the state. Each player assigns probability of $\frac{1}{2}$ to each state. The expected payoffs are:

	L	M	R
T	$(1, 2\epsilon)$	$(1, 3/2\epsilon)$	$(1, 3/2\epsilon)$
B	$(2, 2)$	$(0, 3/2)$	$(0, 3/2)$

(B, L) is the only Nash equilibrium.

Consider now the game in which player 2 knows what the state is. So the new game is:

$\frac{1}{2}$ $\frac{1}{2}$

1

2

	L	M	R
T	$(1, 2\epsilon)$	$(1, 0)$	$(1, 3\epsilon)$
B	$(2, 2)$	$(0, 0)$	$(0, 3)$

State w_1

2

	L	M	R
T	$(1, 2\epsilon)$	$(1, 3\epsilon)$	$(1, 0)$
B	$(2, 2)$	$(0, 3)$	$(0, 0)$

State w_2

	LL	LM	LR	ML	MM	MR	RL	RM	RR
T	1	1	1	1	1	1	1	1	1
B	2	1	1	1	0	0	1	0	0

	L	M	R
T	$(1, 2\epsilon)$	$(1, 0)$	$(1, 3\epsilon)$
B	$(2, 2)$	$(0, 0)$	$(0, 3)$

	L	M	R
T	$(1, 2\epsilon)$	$(1, 3\epsilon)$	$(1, 0)$
B	$(2, 2)$	$(0, 3)$	$(0, 0)$

NE: $(T, (R, M))$

Therefore, player 2 is worse off when she knows the state than when she does not know the state.

1.3.2. *Infection*

The notion of a Bayesian game may be used to model not only situations in which players are uncertain about the others' preferences, but also situations in which they are uncertain about each others' *knowledge*. Consider the next game.

Example 4: *Infection*

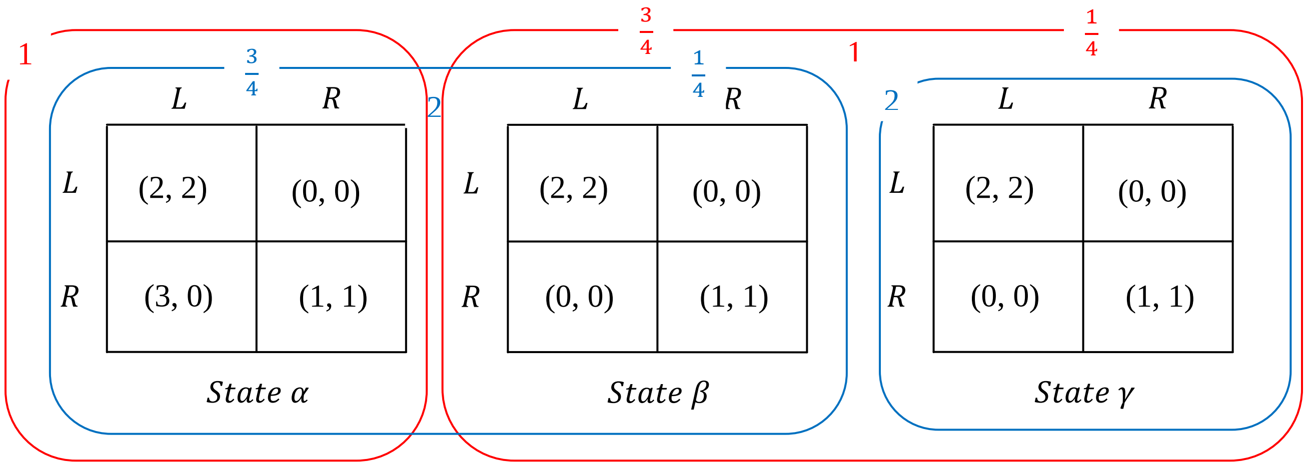

Note that player 2's preferences are the same in all three states, and player 1's preferences are the same in states β and γ. In particular, in state γ, each player knows the other player's preferences, and player 2 knows that player 1 knows her preferences. The defect in the players' information in state γ is that *player 1 does not know whether player 2 knows her preferences*: player 1 knows only that the state is either β or γ, and in state β player 2 does not know whether the state is α or β, and hence does not know player 1's preferences (because player 1's preferences in these two states differ).

There are two types of player 1 (the type who knows that the state is α and the type who knows that the state is β or γ) and two types of player 2 (the type who knows that the state is

α or β and the type who knows that the state is γ). We next obtain the best response of each type of each player.

$t_{1\alpha}$:

	(L, L)	(L, R)	(R, L)	(R, R)
L	2	2	0	0
R	3	3	1	1

$t_{1\beta or \gamma}$:

	(L, L)	(L, R)	(R, L)	(R, R)
L	2	$\frac{3}{2}$	$\frac{1}{2}$	0
R	0	$\frac{1}{4}$	$\frac{3}{4}$	1

$t_{2\alpha or \beta}$:

	L	R
(L, L)	2	0
(L, R)	$\frac{3}{2}$	$\frac{1}{4}$
(R, L)	$\frac{1}{2}$	$\frac{3}{4}$
(R, R)	0	1

$t_{2\gamma}$:

	L	R
(L, L)	2	0
(L, R)	0	1
(R, L)	2	0
(R, R)	0	1

s_1	BR_2	s_2	BR_1
(L, L)	(L, L)	(L, L)	(R, L)
(L, R)	(L, R)	(L, R)	(R, L)
(R, L)	(R, L)	(R, L)	(R, R)
(R, R)	(R, R)	(R, R)	(R, R)

NE: $((R, R)\ (R, R))$

Example 5: *Infection* 2

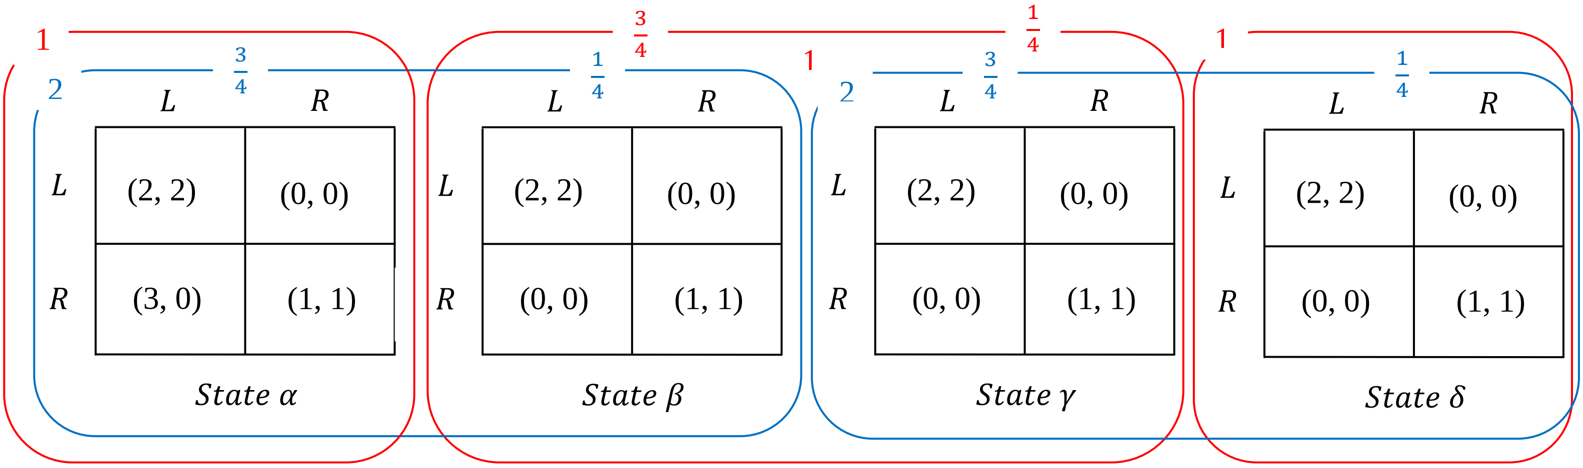

Consider state δ. In this state, player 2 knows player 1's preferences (because she knows that the state is either γ or δ, and in both states player 1's preferences are the same). What player 2 does not know is whether player 1 knows whether player 2 knows player 1's preferences. In this game, we have 3 types of player 1 (the type who knows that the state is α, the type who knows that the state is β or γ, and the type who knows that the state is δ) and two types of player 2 (the type who knows that the state is α or β, and the type who knows that the state is γ or δ).

$t_{1\alpha}$:

	(L, L)	(L, R)	(R, L)	(R, R)
L	2	2	0	0
R	3	3	1	1

$t_{1\beta or \gamma}$:

	(L, L)	(L, R)	(R, L)	(R, R)
L	2	$\frac{3}{2}$	$\frac{1}{2}$	0
R	0	$\frac{1}{4}$	$\frac{3}{4}$	1

$t_{1\delta}$:

	(L, L)	(L, R)	(R, L)	(R, R)
L	2	0	2	0
R	0	1	0	1

$t_{2\alpha or \beta}$:

	L	R
(L, L, L)	2	0
(L, L, R)	2	0
(L, R, L)	$\frac{3}{2}$	$\frac{1}{4}$
(L, R, R)	$\frac{3}{2}$	$\frac{1}{4}$
(R, L, L)	$\frac{1}{2}$	$\frac{3}{4}$
(R, L, R)	$\frac{1}{2}$	$\frac{3}{4}$
(R, R, L)	0	1
(R, R, R)	0	1

$t_{2\gamma or \delta}$:

	L	R
(L, L, L)	2	0
(L, L, R)	$\frac{3}{2}$	$\frac{1}{4}$
(L, R, L)	$\frac{1}{2}$	$\frac{3}{4}$
(L, R, R)	0	1
(R, L, L)	2	0
(R, L, R)	$\frac{3}{2}$	$\frac{1}{4}$
(R, R, L)	$\frac{1}{2}$	$\frac{3}{4}$
(R, R, R)	0	1

s_1	BR_2
(L, L, L)	(L, L)
(L, L, R)	(L, L)
(L, R, L)	(L, R)
(L, R, R)	(L, R)
(R, L, L)	(R, L)
(R, L, R)	(R, L)
(R, R, L)	(R, R)
(R, R, R)	(R, R)

s_2	BR_1
(L, L)	(R, L, L)
(L, R)	(R, L, R)
(R, L)	(R, R, L)
(R, R)	(R, R, R)

NE: $((R, R, R)\ (R, R))$

1.4. **Illustration: Cournot's duopoly game with imperfect information**

1.4.1. *Imperfect information about cost*

Two firms compete in selling a homogeneous product; one firm does not know the other firm's cost function. We next study how this lack of information will affect the firms' behavior.

Assume that both firms can produce the good at constant unit cost (that is, marginal cost is constant and there are no fixed costs). Assume also that they both know that firm 1's unit cost is c, but only firm 2 knows its own unit cost; firm 1 believes that firm 2's cost is c_L with probability θ and c_H with probability $1-\theta$, with $0<\theta<1$ and $c_L<c_H$. We can model this problem as a Bayesian game.

The information structure in this game is similar to that in Example 2.1 (*BoS* with imperfect information 1)

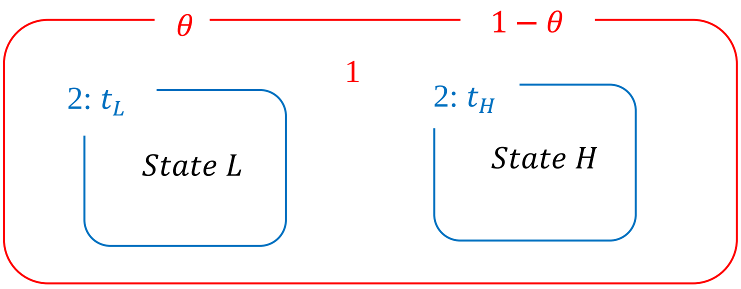

We next describe the Bayesian game.

Players: Firm 1 and firm 2.

States: $\{L, H\}$.

Actions: Each firm's set of actions is the set of nonnegative outputs.

Signals: Firm 1's signal function τ_1 satisfies $\tau_1(L) = \tau_1(H) = t_1$ (its signal is the same in both states; one type of firm 1); firm 2's signal function τ_1 satisfies $\tau_2(L) \neq \tau_2(H)$, with $\tau_2(L) = t_L$ and $\tau_2(H) = t_H$ (its signal is perfectly informative of the state; two types of firm 2).

Beliefs: After receiving the (non informative) signal t_1, the single type of firm 1 assigns probability θ to state L and probability $1-\theta$ to state H. Each type of firm 2 assigns probability 1 to the single state consistent with its signal. That is, after receiving the signal t_L, firm 2 assigns probability 1 to the state L, and after receiving the signal t_H, firm 2 assigns probability 1 to the state H.

Payoffs: The firms' Bernoulli payoffs are their profits; if the actions chosen are (q_1, q_2) and the state is I (either L or H), then firm 1's profit is $[p(q_1 + q_2) - c]q_1$ and firm 2's profit is $[p(q_1 + q_2) - c_I]q_2$, where $p(q_1 + q_2)$ is the market price.

A Nash equilibrium of this game is a triple $(q_1^*, (q_L^*, q_H^*))$, where q_1^* is the output of firm 1, q_L^* is the output of type t_L of firm 2 (that is, firm 2 when it receives the signal $\tau_2(L)$), and q_H^* is the output of type t_H of firm 2 (that is, firm 2 when it receives the signal $\tau_2(H)$), such that:

- q_1^* maximizes firm 1's (expected) profit given the output q_L^* of t_L of firm 2 and the output q_H^* of type t_H of firm 2,
- q_L^* maximizes the profit of type t_L of firm 2 given the output q_1^* of firm 1 and
- q_H^* maximizes the profit of type t_H of firm 2 given the output q_1^* of firm 1.

To find a Cournot-Nash equilibrium, we first obtain the firms' best response functions. Given firm 1's beliefs, its best response $b_1(q_L, q_H)$ to (q_L, q_H) solves:

$$\max_{q_1 \geq 0} \theta[p(q_1 + \ q_L) - c]q_1 + (1-\theta)[p(q_1 + \ q_H) - c]q_1.$$

Firm 2's best response $b_L(q_1)$ to q_1 when its cost is c_L solves:

$$\max_{q_L \geq 0}[p(q_1 + \ q_L) - c_L]q_L$$

and firm 2's best response $b_H(q_1)$ to q_1 when its cost is c_H solves:

$$\max_{q_H \geq 0}[p(q_1 + \ q_H) - c_H]q_H.$$

A Nash equilibrium is a combination of strategies $(q_1^*, (q_L^*, q_H^*))$ such that:

$q_1^* = b_1(q_L^*, q_H^*)$, $q_L^* = \ b_L(q_1^*)$ and $q_H^* = \ b_H(q_1^*)$.

Example 6. *Cournot imperfect information and linear demand*

Consider a Cournot duopoly game where the inverse demand function is $p(Q) = \alpha - Q$ for $Q \leq \alpha$ and $p(Q) = 0$ for $Q > \alpha$. Assuming that c_L and c_H are such that there is a Nash equilibrium in which all outputs are positive, obtain such equilibrium. Compare this equilibrium with the Nash equilibrium of the (perfect information) game in which firm 1 knows that firm 2's unit cost is c_L and with the Nash equilibrium of the game in which firm 1 knows that firm 2's unit cost is c_H.

To find a Cournot-Nash equilibrium, we first obtain the firms' best response functions. Given firm 1's beliefs, its best response $b_1(q_L, \ q_H)$ to $(q_L, \ q_H)$ solves:

$$\max_{q_1 \geq 0} \theta[p(q_1 + \ q_L) - c]q_1 + (1-\theta)[p(q_1 + \ q_H) - c]q_1.$$

$$\frac{\partial \pi_1}{\partial q_1} = 0$$

$$\rightarrow \theta[p(q_1 + \ q_L) + q_1 p'(q_1 + \ q_L) - c] + (1-\theta)[p(q_1 + \ q_H) + q_1 p'(q_1 + \ q_H) - c] = 0$$

$$\rightarrow \theta[\alpha - q_1 - q_L - q_1 - c] + (1-\theta)[\alpha - q_1 - q_H - q_1 - c] = 0$$

$$\rightarrow \alpha - 2q_1 - c - [\theta q_L + (1-\theta)q_H] = 0$$

$$(b_1(q_L, q_H) = \max\{\frac{\alpha - c - [\theta q_L + (1-\theta)q_H]}{2}, 0\})$$

$$b_1(q_L, q_H) = \frac{\alpha - c - [\theta q_L + (1-\theta)q_H]}{2}$$

Firm 2's best response $b_L(q_1)$ to q_1 when its cost is c_L solves:

$$\max_{q_L \geq 0}[p(q_1 + q_L) - c_L]q_L$$

$$\frac{\partial \pi_L}{\partial q_L} = 0 \rightarrow [p(q_1 + q_L) + q_L p'(q_1 + q_L) - c_L = 0 \rightarrow \alpha - q_1 - q_L - q_L - c_L = 0$$

$$(\rightarrow b_L(q_1) = \max\{\frac{\alpha - c_L - q_1}{2}, 0\})$$

$$\rightarrow b_L(q_1) = \frac{\alpha - c_L - q_1}{2}$$

Firm 2's best response $b_H(q_1)$ to q_1 when its cost is c_H solves:

$$\max_{q_H \geq 0}[p(q_1 + q_H) - c_H]q_H$$

$$\frac{\partial \pi_H}{\partial q_H} = 0 \rightarrow [p(q_1 + q_H) + q_H p'(q_1 + q_H) - c_H = 0 \rightarrow \alpha - q_1 - q_H - q_H - c_H = 0$$

$$(\rightarrow b_H(q_1) = \max\{\frac{\alpha - c_H - q_1}{2}, 0\})$$

$$\rightarrow b_H(q_1) = \frac{\alpha - c_H - q_1}{2}$$

A Nash equilibrium is a combination of strategies $(q_1^*, (q_L^*, q_H^*))$ such that:

$$q_1^* = b_1(q_L^*, q_H^*) = \frac{\alpha - c - [\theta q_L^* + (1-\theta)q_H^*]}{2}$$

$$q_L^* = b_L(q_1^*) = \frac{\alpha - c_L - q_1^*}{2}$$

$$q_H^* = b_H(q_1^*) = \frac{\alpha - c_H - q_1^*}{2}$$

$$q_1^* = b_1(q_L^*, q_H^*) = \frac{\alpha - c - [\theta \frac{\alpha - c_L - q_1^*}{2} + (1-\theta)\frac{\alpha - c_H - q_1^*}{2}]}{2}$$

$$\rightarrow q_1^* = \frac{\alpha - 2c + \theta c_L + (1-\theta)c_H + q_1^*}{4}$$

$$\rightarrow q_1^* = \frac{\alpha - 2c + \theta c_L + (1-\theta)c_H}{3}$$

$$q_L^* = b_L(q_1^*) = \frac{\alpha - c_L - \frac{\alpha - 2c + \theta c_L + (1-\theta)c_H}{3}}{2}$$

$$\rightarrow q_L^* = \frac{2\alpha - 3c_L + 2c - \theta c_L - (1-\theta)c_H}{6}$$

$$\rightarrow q_L^* = \frac{\alpha - 2c_L + c}{3} - \frac{(1-\theta)(c_H - c_L)}{6}$$

$$q_H^* = b_H(q_1^*) = \frac{\alpha - c_H - \frac{\alpha - 2c + \theta c_L + (1-\theta)c_H}{3}}{2}$$

$$\rightarrow q_H^* = \frac{2\alpha - 3c_H + 2c - \theta c_L - (1-\theta)c_H}{6}$$

$$\rightarrow q_H^* = \frac{\alpha - 2c_H + c}{3} + \frac{\theta(c_H - c_L)}{6}$$

Perfect Information

If firm 1 knows that firm 2’s unit cost is c_L, equilibrium outputs are:

$$q_1^* = \frac{\alpha - 2c + c_L}{3} \qquad q_2^* = \frac{\alpha - 2c_L + c}{3}$$

If firm 1 knows that firm 2's unit cost is c_H, equilibrium outputs are:

$$q_1^* = \frac{\alpha - 2c + c_H}{3} \qquad q_2^* = \frac{\alpha - 2c_H + c}{3}$$

Therefore, in comparison with the perfect information case, under imperfect information the low-cost firm 2 would produce a lower quantity and the high-cost firm 2 would produce a greater quantity.

1.4.2. *Imperfect information about both cost and information*

Now assume that firm 2 does not know whether firm 1 knows firm 2's cost. That is, suppose that one circumstance that firm 2 believes to be possible is that firm 1 knows its cost (although in fact it does not). Because firm 2 thinks this circumstance to be possible, we need four states to model this situation which we call $L0$, $H0$, $L1$, and $H1$, with the following interpretation.

$L0$: firm 2's cost is low and firm 1 does not know whether it is low or high.

$H0$: firm 2's cost is high and firm 1 does not know whether it is low or high.

$L1$: firm 2's cost is low and firm 1 knows it is low.

$H1$: firm 2's cost is high and firm 1 knows it is high.

Firm 1 receives one of three possible signals, 0, l, and h. The states $L0$ and $H0$ generate the signal 0 (firm 1 does not know firm 2's cost), the state $L1$ generates the signal l (firm 1 knows firm 2's cost is low), and the state $H1$ generates the signal h (firm 1 knows firm 2's cost is high). Firm 2 receives one of two possible signals, L, in states $L0$ and $L1$, and H, in states $H0$ and $H1$. Denote by θ (as before) the probability assigned by type 0 of firm 1 to firm 2's cost being c_L, and by μ the probability assigned by each type of firm 2 to firm 1's knowing firm 2's cost (the case $\mu = 0$ is equivalent to the one considered in subsection 1.4.1).

The information structure in this game is as follows

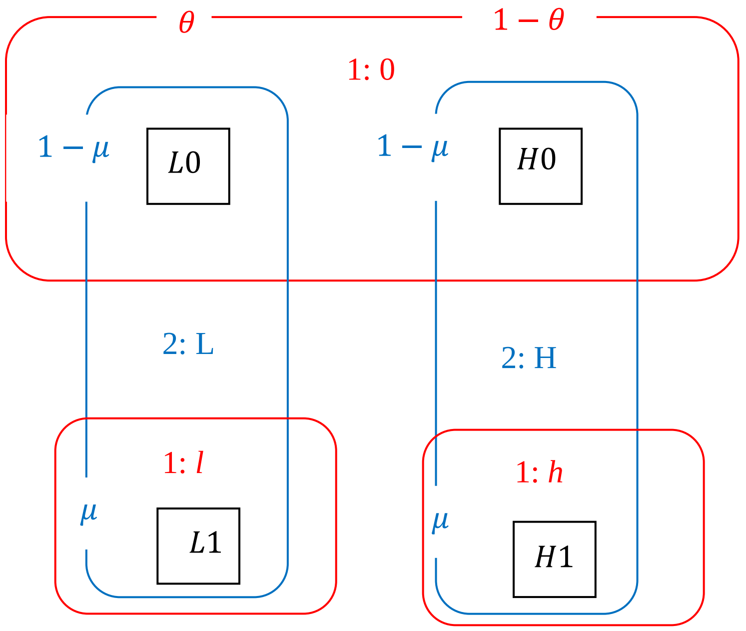

A Bayesian game that models the situation is defined as follows.

Players: Firm 1 and Firm 2.

States: $\{L0, L1, H0, H1\}$, where the first letter in the name of the state indicates firm 2's cost and the second letter indicates whether firm 1 knows (1) or does not know (0) firm 2's cost.

Actions: Each firm's set of actions is the set of its possible (nonnegative) outputs.

Signals: Firm 1 gets one of the signals 0, l, and h, and its signal function τ_1 satisfies $\tau_1(L0) = \tau_1(H0) = 0$, $\tau_1(L1) = l$, and $\tau_1(H1) = h$. Firm 2 gets the signal L or H and its signal function τ_2 satisfies $\tau_2(L0) = \tau_2(L1) = L$ and $\tau_2(H0) = \tau_2(H1) = H$.

Beliefs: After receiving the (non informative) signal 0 firm 1 assigns probability θ to state $L0$ and probability $1-\theta$ to state $H0$; after receiving the signal l firm 1 assigns probability 1 to state $L1$; after receiving the signal h firm 1 assigns probability 1 to state H. After receiving the signal L firm 2 assigns probability μ to state $L1$ and probability $1-\mu$ to state $L0$; after receiving the signal H firm 2 assigns probability μ to state $H1$ and probability $1-\mu$ to state $H0$.

Payoff functions: The firms' Bernoulli payoffs are their profits; if the actions chosen are (q_1, q_2), then firm 1's profit is $[p(q_1+\ q_2)-c]q_1$ and firm 2's profit is $[p(q_1+\ q_2)-c_L]q_2$ in states $L0$ and $L1$, and $[p(q_1+\ q_2)-c_H]q_2$ in states $H0$ and $H1$.

Example 7. *Cournot, imperfect information about cost and information and linear demand*

Write down the maximization problems that determine the best response function of each type of each player. Denote by q_0, q_l, and q_h the outputs of types 0, l, and h of firm 1, and by q_L and q_H the outputs of types L and H of firm 2. Suppose that the inverse demand function is $p(Q)=\alpha-Q$ for $Q\leq\alpha$ and $p(Q)=0$ for $Q>\alpha$. Assuming that c_L and c_H are such that there is a Nash equilibrium in which all outputs are positive, obtain such equilibrium. Check that when $\mu=0$ the equilibrium output of type 0 of firm 1 is equal to the equilibrium output of firm 1 corresponding to exercise 6, and that the equilibrium outputs of the two types of firm 2 are the same as the ones corresponding to that exercise. Check also that when $\mu=1$ the equilibrium outputs of type *l* of firm 1 and type *L* of firm 2 are the same as the equilibrium outputs when there is perfect information and the costs are *c* and c_L, and that the equilibrium outputs of type *h* of firm 1 and type *H* of firm 2 are the same as the equilibrium outputs when there is perfect information and the costs are *c* and c_H. Show that for $0<\mu<1$, the equilibrium outputs of type L and H of firm 2 lie between their values when $\mu = 0$ and when $\mu = 1$.

The best response $b_0(q_L, q_H)$ of type 0 of firm 1 is the solution of:

$$\max_{q_0 \geq 0} \theta[p(q_0 + \ q_L) - c]q_0 + (1-\theta)[p(q_0 + \ q_H) - c]q_0.$$

$\frac{\partial \pi_0}{\partial q_0} = 0$

$$\rightarrow \theta[p(q_0 + \ q_L) + q_0 p'(q_0 + \ q_L) - c] + (1-\theta)[p(q_0 + \ q_H) + q_0 p'(q_1 + \ q_H) - c] = 0$$

$$\rightarrow \theta[\alpha - q_0 - \ q_L - q_0 - c] + (1-\theta)[\alpha - q_0 - \ q_H - q_0 - c] = 0$$

$$\rightarrow \alpha - 2q_0 - c - [\theta q_L + (1-\theta) q_H] = 0$$

$$(b_0(q_L, q_H) = \max\{\frac{\alpha - c - [\theta q_L + (1-\theta) q_H]}{2}, 0\})$$

$$b_0(q_L, q_H) = \frac{\alpha - c - [\theta q_L + (1-\theta) q_H]}{2}$$

The best response $b_l(q_L, q_H)$ of type *l* of firm 1 is the solution of:

$$\max_{q_l \geq 0}[p(q_l + \ q_L) - c]q_l$$

$\frac{\partial \pi_l}{\partial q_l} = 0$

$$\rightarrow p(q_l + \ q_L) + q_l p'(q_l + \ q_L) - c = 0$$

$$\rightarrow \alpha - q_l - \ q_L - q_l - c = 0$$

$$(b_l(q_L, q_H) = \max\{\frac{\alpha - c - q_L}{2}, 0\})$$

$$b_l(q_L, q_H) = \frac{\alpha - c - q_L}{2}$$

The best response $b_h(q_L, q_H)$ of type h of firm 1 is the solution of:

$$\max_{q_h \geq 0} [p(q_h + \ q_H) - c]q_h$$

$\frac{\partial \pi_h}{\partial q_h} = 0$

$$\rightarrow p(q_h + \ q_L) + q_h p'(q_h + \ q_H) - c = 0$$

$$\rightarrow \alpha - q_h - \ q_H - q_h - c = 0$$

$$(b_h(q_L, q_H) = \max\{\frac{\alpha - c - q_H}{2}, 0\})$$

$$b_h(q_L, q_H) = \frac{\alpha - c - q_H}{2}$$

The best response $b_L(q_0, q_l, q_h)$ of type L of firm 2 is the solution of:

$$\max_{q_L \geq 0} (1 - \mu)[p(q_0 + \ q_L) - c_L]q_L + \mu[p(q_l + \ q_L) - c_L]q_L.$$

$\frac{\partial \pi_L}{\partial q_L} = 0$

$$\rightarrow (1 - \mu)[p(q_0 + \ q_L) + q_L p'(q_0 + \ q_L) - c_L] + \mu[p(q_l + \ q_L) + q_L p'(q_l + \ q_L) - c_L] = 0$$

$$\rightarrow (1 - \mu)[\alpha - q_0 - \ q_L - q_L - c_L] + \mu[\alpha - q_l - \ q_L - q_L - c_L] = 0$$

$$\rightarrow \alpha - 2q_L - c_L - [(1 - \mu)q_0 + \mu q_l] = 0$$

$$b_L(q_0, q_l, q_h) = \max\{\frac{\alpha - c_L - [(1 - \mu)q_0 + \mu q_l]}{2}, 0\})$$

$$b_L(q_0, q_l, q_h) = \frac{\alpha - c_L - [(1 - \mu)q_0 + \mu q_l]}{2}$$

The best response $b_H(q_0, q_l, q_h)$ of type H of firm 2 is the solution of:

$$\max_{q_H \geq 0} (1 - \mu)[p(q_0 + \ q_H) - c_H]q_H + \mu[p(q_h + \ q_H) - c_H]q_H.$$

$\frac{\partial \pi_H}{\partial q_H} = 0$

$$\rightarrow (1-\mu)[p(q_0 + \ q_H) + q_H p'(q_0 + \ q_H) - c_H] + \mu[p(q_h + \ q_H) + q_H p'(q_h + \ q_H) - c_H]$$

$$= 0$$

$$\rightarrow (1-\mu)[\alpha - q_0 - \ q_H - q_H - c_H] + \mu[\alpha - q_h - \ q_H - q_H - c_H] = 0$$

$$\rightarrow \alpha - 2q_H - c_H - [(1-\mu)q_0 + \mu q_h] = 0$$

$$b_H(q_0, q_l, q_h) = \max\{\frac{\alpha - c_H - [(1-\mu)q_0 + \mu q_h]}{2}, 0\})$$

$$b_H(q_0, q_l, q_h) = \frac{\alpha - c_H - [(1-\mu)q_0 + \mu q_h]}{2}$$

A Nash equilibrium is a strategy profile $((q_0^*, q_l^*, q_h^*), (q_L^*, q_H^*))$ such that:

$$q_0^* = b_0(q_L^*, q_H^*) = \frac{\alpha - c - [\theta q_L^* + (1-\theta)q_H^*]}{2}$$

$$q_l^* = b_l(q_L^*, q_H^*) = \frac{\alpha - c - q_L^*}{2}$$

$$q_h^* = b_h(q_L^*, q_H^*) = \frac{\alpha - c - q_H^*}{2}$$

$$q_L^* = b_L(q_0^*, q_l^*, q_h^*) = \frac{\alpha - c_L - [(1-\mu)q_0^* + \mu q_l^*]}{2}$$

$$q_H^* = b_H(q_0^*, q_l^*, q_h^*) = \frac{\alpha - c_H - [(1-\mu)q_0^* + \mu q_h^*]}{2}$$

$$q_L^* = \frac{\alpha - c_L - [(1-\mu)q_0^* + \mu \frac{\alpha - c - q_L^*}{2}]}{2} = \frac{2\alpha - 2c_L - [2(1-\mu)q_0^* + \mu(\alpha - c - q_L^*)]}{4}$$

$$q_L^* = \frac{2\alpha - 2c_L - [2(1-\mu)q_0^* + \mu(\alpha - c)]}{4-\mu}$$

$$q_H^* = \frac{\alpha - c_H - [(1-\mu)q_0^* + \mu \frac{\alpha - c - q_H^*}{2}]}{2} = \frac{2\alpha - 2c_H - [2(1-\mu)q_0^* + \mu(\alpha - c - q_H^*)]}{4}$$

$$q_H^* = \frac{2\alpha - 2c_H - [2(1-\mu)q_0^* + \mu(\alpha - c)]}{4-\mu}$$

$$q_0^* = \frac{\alpha - c}{2} - \frac{1}{2}\left\{\theta \frac{2\alpha - 2c_L - [2(1-\mu)q_0^* + \mu(\alpha - c)]}{4-\mu} \right.$$

$$\left. + (1-\theta)\frac{2\alpha - 2c_H - [2(1-\mu)q_0^* + \mu(\alpha - c)]}{4-\mu}\right\}$$

$$2(4-\mu)q_0^* = (\alpha - c)(4-\mu) - 2\alpha\theta + 2\theta c_L + 2\theta(1-\mu)q_0^* + \theta\mu(\alpha - c) - 2(1-\theta)\alpha$$

$$+ 2(1-\theta)c_H + 2(1-\theta)(1-\mu)q_0^* + (1-\theta)\mu(\alpha - c)$$

$$2(4-\mu)q_0^* - 2(1-\mu)q_0^*$$

$$= (\alpha - c)(4-\mu) - 2\alpha\theta + 2\theta c_L + \mu(\alpha - c) - 2(1-\theta)\alpha + 2(1-\theta)c_H$$

$$6q_0^* = 4(\alpha - c) - 2\alpha + 2\theta c_L + 2(1-\theta)c_H$$

$$6q_0^* = 2\alpha - 4c + 2\theta c_L + 2(1-\theta)c_H$$

$$q_0^* = \frac{\alpha - 2c + \theta c_L + (1-\theta)c_H}{3}$$

$$\rightarrow q_L^* = \frac{1}{3}[\alpha - 2c_L + c - \frac{2(1-\theta)(1-\mu)(c_H - c_L)}{4-\mu}]$$

$$\rightarrow q_H^* = \frac{1}{3}[\alpha - 2c_H + c + \frac{2\theta(1-\mu)(c_H - c_L)}{4-\mu}]$$

$$\rightarrow q_l^* = \frac{1}{3}[\alpha - 2c + c_L + \frac{(1-\theta)(1-\mu)(c_H - c_L)}{4-\mu}]$$

$$\rightarrow q_h^* = \frac{1}{3}[\alpha - 2c + c_H - \frac{\theta(1-\mu)(c_H - c_L)}{4-\mu}]$$

When $\mu = 0$

$$q_0^* = \frac{\alpha - 2c + \theta c_L + (1-\theta)c_H}{3}$$

$$q_L^* = \frac{1}{3}[\alpha - 2c_L + c - \frac{(1-\theta)(c_H - c_L)}{2}]$$

$$q_H^* = \frac{1}{3}[\alpha - 2c_H + c + \frac{\theta(c_H - c_L)}{2}]$$

$$q_l^* = \frac{1}{3}[\alpha - 2c + c_L + \frac{(1-\theta)(c_H - c_L)}{4-\mu}]$$

$$q_h^* = \frac{1}{3}[\alpha - 2c + c_H - \frac{\theta(c_H - c_L)}{4}]$$

So that q_0^* is equal to the equilibrium output of firm 1 in exercise 6, and q_L^* and q_H^* are the same as the equilibrium outputs of the two types of firm 2 in that exercise.

When $\mu = 1$, then

$$q_0^* = \frac{\alpha - 2c + \theta c_L + (1-\theta)c_H}{3}$$

$$q_L^* = \frac{1}{3}[\alpha - 2c_L + c]$$

$$q_H^* = \frac{1}{3}[\alpha - 2c_H + c]$$

$$q_l^* = \frac{1}{3}[\alpha - 2c + c_L]$$

$$q_h^* = \frac{1}{3}[\alpha - 2c + c_H]$$

So that q_l^* and q_L^* are the same as the equilibrium outputs when there is perfect information and the costs are c and c_L, and q_h^* and q_H^* are the same as the equilibrium outputs when there is perfect information and the costs are c and c_H.

For an arbitrary value of μ, we have:

$$q_L^* = \frac{1}{3}[\alpha - 2c_L + c - \frac{2(1-\theta)(1-\mu)(c_H - c_L)}{4-\mu}]$$

$$q_H^* = \frac{1}{3}[\alpha - 2c_H + c + \frac{2\theta(1-\mu)(c_H - c_L)}{4-\mu}]$$

To show that for $0 < \mu < 1$ the values of these variables are between their values when $\mu = 0$ and when $\mu = 1$, we need to show that

$$0 \leq \frac{2(1-\theta)(1-\mu)(c_H - c_L)}{4-\mu} \leq \frac{(1-\theta)(c_H - c_L)}{2}]$$

$$0 \leq \frac{2\theta(1-\mu)(c_H - c_L)}{4-\mu} \leq \frac{\theta(c_H - c_L)}{2}$$

Which holds since $c_H \geq c_L$, $\theta \geq 0$ and $0 \leq \mu \leq 1$.