20

Sequential Rationality with Incomplete Information

To understand the role of proper subgames and their use for subgame perfect equilibrium, it is illuminating to consider the following entry game of complete information as follows. Player 1 is a potential entrant to an industry that has a monopolist incumbent, player 2. If player 1 stays "out" (O), then the incumbent earns a profit of 2, while the potential entrant gets 0. The entrant's other option is to "enter" (E), which gives the incumbent a chance to respond. If the incumbent chooses to "accommodate entry" (A), then the entrant gets 1 while the incumbent gets 0. The incumbent's other option is to "fight entry" (F), in which case the payoff for each player is -1. The extensive form of this game is described in figure 5.1.

To find the Nash equilibria of this game it is useful to look at the matrix (normal form) game that describes the entry game:

	F	A
O	$0, 2$	$0, 2$
E	$-1, -1$	$1, 0$

A quick observation reveals that the game has two pure strategy Nash equilibria, which are (O, F) and (E, A). If, however, we consider the Subgame perfect

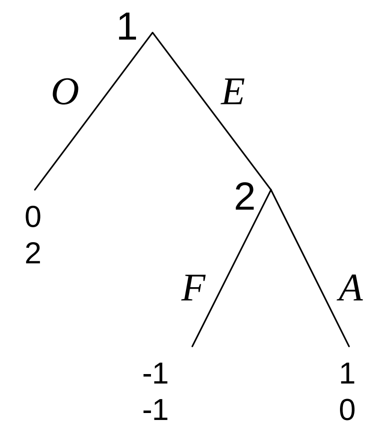

FIGURE 20.1.

equilibrium concept, backward induction clearly implies that upon entry player 2 will strictly prefer to accommodate entry, and therefore player 1 should enter anticipating a payoff of 1 rather than staying out and getting 0. Thus, Subgame perfection implies sequential rationality and picks only one of the two Nash equilibria as the unique Subgame Perfect equilibrium, (E, A), where form 1 enters and form 2 accommodates entry.

Now consider a straightforward variant of this game that includes incomplete information. In particular, imagine that the entrant may have a technology that is superior to that of the incumbent, in which case he would gain a lot by entering, and the incumbent would lose a lot from this entry. However, the entrant may have an inferior technology, in which case he would not gain by entering, and the incumbent would lose only a little if entry occurred. A particular case of this story can be captured by the following sequence of events:

1. Nature choose the entrant's type which can be "weak" or "strong" with equal probability. That is, $t_1 \in \{w, s\}$, and $\Pr\{t_1 = s\} = 0.5$. The entrant knows his type but the incumbent only knows that $\Pr\{t_1 = s\} = 0.5$.

2. The Entrant chooses between E and O as before, and the incumbent observes the entrant's choice.

3. After observing the action of the entrant, and if the entrant enters, the incumbent can choose between A and F as before.

The payoffs are different for each of the realization of player 1's type, and are given in the extensive form of the game that is depicted in Figure 5.2.

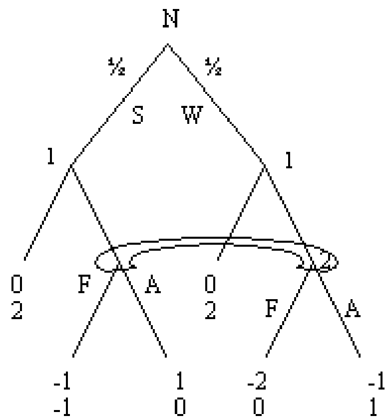

When considering the normal form of this game, player 1 has four pure strategies that follow from the fact that a strategy for player 1 is a type-dependent action, and there are two types and two actions. We define a strategy for player 1 as $s_1^s s_1^w$, where $s_1^t \in \{O, E\}$ is what a type t of player 1 chooses. Thus, the pure strategy set of player 1 is:

$$s_1^s s_1^w \in S_1 = \{OO, OE, EO, EE\} \ .$$

For example, $s_1^s s_1^w = OE$ means that player 1 chooses O when his type is s, and he chooses E when his type is w. Since player 2 has only one information set that follows entry, and two actions in that information set, then he has two pure strategies, $s_2 \in \{A, F\}$.

To convert this extensive form game to a Normal form matrix game, we need to compute the expected payoffs *from each pair of pure strategies*, where expectations are over the randomizations caused by nature. Since player 1 has four pure strategies and player 2 has two, there will be eight entries into the normal form matrix. For example, consider the pair of strategies $(s_1^s s_1^w, s_2) = (OE, A)$. The payoffs of the game will be determined by one of the following two outcomes:

1. Nature's choice of $t_1 = s$, in which case player 1 plays O and the payoffs are $(0, 2)$. This happens with probability $\Pr\{t_1 = s\} = 0.5$.

2. Nature's choice of $t_1 = w$, in which case player 1 plays E and player 2 plays A. The payoffs for this outcome are $(-1, -1)$, and this happens with probability $\Pr\{t_1 = w\} = 0.5$.

From these two possible outcomes we can compute the expected payoff for players 1 and 2 as follows:

$$Eu_1 = 0.5 \cdot 0 + 0.5 \cdot (-1) = -0.5$$

and

$$Eu_2 = 0.5 \cdot 2 + 0.5 \cdot (-1) = 0.5.$$

Similarly, if the strategies are $(s_1^s s_1^w, s_2) = (EE, F)$ then the expected payoff for players 1 and 2 are $Eu_1 = 0.5{\cdot}(-1){+}0.5{\cdot}(-2) = -1.5$, and $Eu_2 = 0.5{\cdot}(-1){+}0.5{\cdot}0 = -0.5$.

In this way we can complete the matrix and obtain the following representation for this Bayesian game of incomplete information as follows:

	F	A
OO	$0, 2$	$0, 2$
OE	$-1, 1$	$-\frac{1}{2}, \frac{3}{2}$
EO	$-\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, 1$
EE	$-\frac{3}{2}, -\frac{1}{2}$	$0, \frac{1}{2}$

Now it is easy to find the pure strategy Bayesian Nash equilibria: *every Nash equilibrium of the Bayesian game matrix we have just calculated is a Bayesian Nash equilibrium of the Bayesian game of incomplete information.* Therefore, both (OO, F) and (EO, A) are pure strategy Bayesian Nash equilibria of the Bayesian game.[1]

[1] Notice that for player 1 OE is strictly dominated by OO and EE is strictly dominated by EO. Notice also that there are a continuum of mixed strategy Bayesian Nash equilibria where the incumbent plays F with probability $p \geq \frac{1}{2}$ and the entrant plays OO. There cannot be a mixed strategy Bayesian Nash equilibrium where the entrant mixes between OO and EO because then the incumbent's best response is A, in which case the entrant would strictly prefer to play EO over OO.

Interestingly, these two Bayesian Nash equilibria are tightly related to the two Nash equilibria of the complete information game that we analyzed earlier. The equilibrium (OO, F) is one where the incumbent "threatens" to fight, which causes the entrant to stay out regardless of his type, similar to the (O, F) equilibrium in the game of complete information in figure 5.1. The equilibrium (EO, A) is one where the incumbent accommodates entry, which causes the strong entrant to enter (getting 1 instead of 0), and the weak entrant to stay out (getting 0 instead of -1), similar to the (E, A) equilibrium in the game of complete information.

Not only are the equilibria similar, but there is a similar problem of credibility with the equilibrium (OO, F): player 2 threatens to fight, but if he finds himself in the information set that follows entry, he has a strict best response which is to accommodate entry. Thus, the Bayesian Nash equilibrium (OO, F) involves non-credible behavior of player 2.

Now we can ask, which of these two equilibria survive as a Subgame Perfect equilibrium in the extensive form game? Recall that the definition of a SPE is that in every *proper subgame*, the restriction of the strategies to that subgame must be a Nash equilibrium in the subgame. This, as we saw in chapter 3, means that players are playing mutual best responses both *on and off the equilibrium path*. However, looking at the extensive form game in figure 5.2, it is easy to see that there is only one proper subgame which is the complete game. Therefore, both, (OO, F) and (EO, A) survive as subgame perfect equilibria!

This example shows us that the very appealing concept of Subgame perfect equilibrium has no bite for some games of incomplete information. At first, this may seem somewhat puzzling. However, the problem is that SPE restricts attention to best responses *within subgames*, but when there is incomplete information, the induced information sets over types of other players often cause the only proper subgame to be the complete game.

The reason is that in the modified entry game we analyzed, even though player 2 observes the actions of player 1, the fact that player 1 has several types implies that there is no proper subgame that starts with player 2 making a move. This means that there are no proper subgames except for the whole game. Indeed, this problem of action nodes being linked through information sets that represent the uncertainty players have over the types of their opponents will carry over to many games of incomplete information.

In order to extend the "logic" of SPE to dynamic games of incomplete information, we need to impose more rigorous structure on our solution concepts in order for sequential rationality to be well defined. Our goal is to identify a structure of analysis that will cause the elimination of equilibria that involve such non-credible threats as in the modified entry game, and this is done in the next section. L15

21

Perfect Bayesian Equilibrium

To deal with the problem demonstrated above, we need to be able and identify a way to express the fact that following entry, it cannot be sequentially rational for player 2 to play fight. How can this be done? We need to make statements about the behavior of player 2 *within his information set* even though it is not itself the node of a proper subgame. This will essentially allow us to say something like "in this information set player 2 is not playing a best response, and therefore his behavior is not sequentially rational."

This is precisely what was missing from the definition of subgame perfect equilibrium – we were not able to asolate player 2 in his information set. However, to talk about the best response of player within his information set, we will have to ask what player 2 is playing a best response to. The answer follows from the definition of what a best response is: we have to *add beliefs* into the picture, so that we can talk about the beliefs of player 2, and then naturally talk about his best response to these beliefs.

More generally, to introduce sequential rationality at information sets that are not singletons, we need to endow players with beliefs *in every part of the game.* These will include information sets that are reached through the proposed actions of the players (the profile of strategies that we are scrutinizing as equilibrium behavior). However, these will also include information sets that *are not reached* given the proposed strategy profile, and have the players play best responses to these beliefs.

At this stage you may recall the notion of being *on* or *off the equilibrium path*, which we used when we discussed subgame perfect equilibrium. Here, we introduce the idea more formally as follows:

Definition: Let $s^* = (s_1^*, ..., s_n^*)$ be a Bayesian Nash equilibrium profile of strategies in a game of incomplete information. We say that an information set is *on the equilibrium path* if given s^* and the distribution of types, it is reached with positive probability. We say that an information set is *off the equilibrium path* if given s^* and the distribution of types it is reached with zero probability.

Notice that an important part of being on or off the equilibrium path is having an equilibrium profile of strategies to start with. This is why the definition includes a proposed equilibrium profile, s^*, from which we can discuss the information sets that are on the equilibrium path *given* s^*, and those that are not on the equilibrium path *given* s^*. To illustrate this point consider the entry game above illustrated again in figure 5.3, but now assume the more general formalization in which player 1's type is "strong" with probability p (not necessarily $p = 0.5$).

Assume that player 1 chooses the strategy EO, implying that he enters if he is strong and stays out when he is weak. In this case, all information sets are reached with positive probability: both information sets of player 1 are reached (with probability p and $1 - p$ respectively), and the information set of player 2 is reached after nature chooses a strong player 1, meaning that it is reached with probability p. Now, assume that player 1 chooses the strategy OO, implying that he never enters. In this case the information set of player 2 is not reached with positive probability. This suggests a more general fact about games of incomplete information: the actions of the player with private information will have an affect on which information sets of the *uninformed* player are reached with positive probability.

As we alluded to earlier, to talk about the optimal behavior of players in information sets, both on and off the equilibrium path, we need to refer to the beliefs that players have in their information sets. For this we add the following definition:

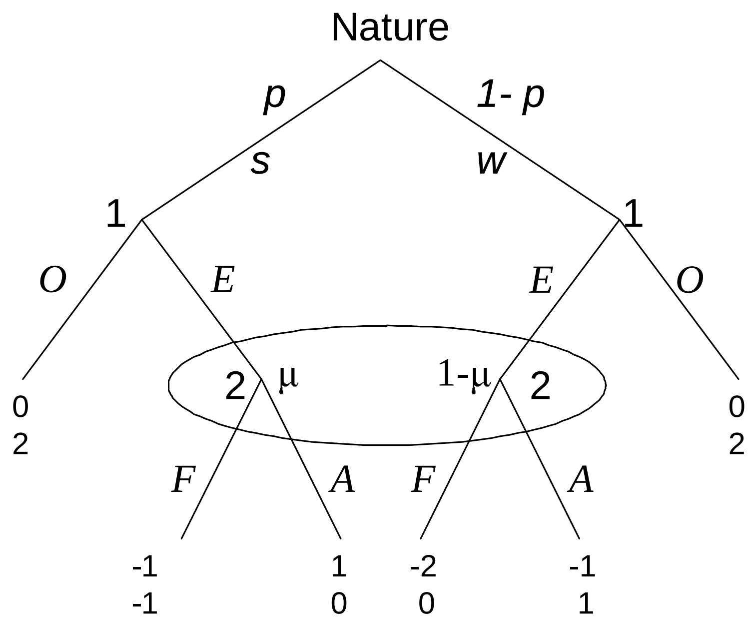

FIGURE 21.1.

Definition: A **system of beliefs** μ of an extensive form game is a probability distribution over decision nodes for every information set. That is, for every decision node x, $\mu(x) \in [0,1]$ satisfies:

$$\sum_{x' \in h(x)} \mu(x') = 1\,.$$

For example, in the game above a belief for player 2 is a probability distribution over nodes in his information set, where, say, μ is his belief that he is at the node corresponding to player 1 being Strong (s) and playing E, and $1-\mu$ is his belief that he is at the node corresponding to player 1 being Weak (W) and playing E. This allows us to lay out our first requirement of sequential rationality:

Requirement R1: In *each* information set the player moving in that information set will have a well defined belief over where he is. That is, the game will have a *system of belief.*

Now that we have established the need to have a system of beliefs, we need to ask the following question: how will the beliefs in a belief system be determined?

That is, can players have any beliefs they want to have, or should the beliefs be restricted by some elements that are not controlled by the player himself?

Just like we required the beliefs of players to be correct about the strategies of their opponents for games of complete information, here too we will add a very similar requirement. Opponents' behavior, however, is just one part of the constraints on beliefs, one that we can consider as an *endogenous constraint on beliefs*. It is endogenous in the sense that it is determined by the strategies of players, which are the "variables" that the players control. The second constraint on beliefs comes from the choices of nature through the distribution of types. This is an *exogenous constraint on beliefs*. It is exogenous in the sense that it is determined by nature, which is not something that the players control, but rather part of the environment.

To illustrate this consider the entry game above. By R1, we define player 2's belief in his information set as

$$\mu = \Pr\{\text{player 1 is "strong"} \mid \text{Enter}\}$$

and $1-\mu$ is defined accordingly (that is all we need for a system of beliefs in this game). But what probability should player 2 assign for μ? Let's say that player 2 believes that player 1 is using the strategy EO, so that he enters of he is strong and stays out of he is weak. What should player 2 believe if his information set is reached? The only consistent belief would be that payer 1 is strong, and therefore it must be that $\mu = 1$. This follows because with probability p Nature chooses a strong type for player 1, and with the strategy EO a strong player 1 always enters. The probability of reaching the node "weak followed by entry" inside the information set of player 2 is 0 because with probability $1-p$ Nature chooses a weak type for player 1, and with the strategy EO a weak player 1 never enters. Any other belief would not be consistent with the belief that player 1 is playing EO.

Following this argument, we are ready to state our second requirement for sequential rationality. Namely, given a *conjectured profile of strategies*, and given *nature's choices*, we want players' beliefs to be "correct" in information sets that are on the equilibrium path. Formally,

Requirement R2: Let $s^* = (s_1^*, ..., s_n^*)$ be a Bayesian Nash equilibrium profile of strategies. We require that in all information sets *on the equilibrium path* beliefs are consistent with *Bayes Rule.*

This is precisely how constraints on beliefs are formed, which include both the *exogenous constraints* that follow from Nature's probability distribution, and the *endogenous constraints* that follow from the beliefs about the other players' strategies from $s^* = (s_1^*, ..., s_n^*)$.

To see how this works more generally, consider the entry game in figure 5.3, and consider the strategy IO for player 1. As we discussed, the probability of reaching the node "strong followed by entry" inside the information set of player 2 is p. What should the belief μ be for any strategy of player 1? By definition, μ is the belief that player 2 assigns to "strong *conditional on* entry". Assume that player 1 is playing the following strategy: if $t_i = s$ then he plays E with probability σ_s and O with probability $1 - \sigma_s$. Similarly, if $t_i = w$ then he plays E with probability σ_w and O with probability $1 - \sigma_w$. Now, since Nature chooses $\Pr\{t_i = s\} = p$, by Bayes rule we must have,

$$\begin{aligned}
\mu &= \Pr\{\text{player 1 is “strong”} \mid \text{Enter}\} \\
&= \frac{\Pr\{\text{strong \& entry}\}}{\Pr\{\text{strong \& entry}\} + \Pr\{\text{weak \& entry}\}} \\
&= \frac{p \cdot \sigma_s}{p \cdot \sigma_s + (1-p) \cdot \sigma_w}
\end{aligned}$$

The pure strategy EO is just a special case of this mixed (behavioral) strategy that has $\sigma_s = 1$ and $\sigma_w = 0$.

Now consider the pure strategy OO (or $\sigma_s = \sigma_w = 0$). The probability of reaching player 2's information set is *not positive* since there is no way in which it is reached. That is, no type of player 1 plays E, and this implies that E is never chosen. If player 2 *believes* that player 1 chooses OO, and suddenly finds himself in his information set that follows entry, then Bayes formula above *does not apply* because given the suggested strategy, both the numerator and denominator are zero, and thus μ is not well defined. What then should determine μ? In other words, if we cannot apply Bayes rule because of the beliefs over strategies, what will we use to determine beliefs? For this we introduce the third requirement.

Requirement R3: At information sets that are *off the equilibrium path* any belief can be assigned where Bayes rule does not apply.

This means that when the moves of nature combined with the belief over the strategies of the other players do not impose constraints on beliefs, then indeed beliefs could be whatever we choose them to be. Looking back at the case where player 1 chooses the pure strategy OO in the entry game, then μ can be any number in the interval $[0, 1]$ since it is not constrained by Bayes rule.

Remark 12 *On a technical note it is worth pointing out that the more accurate requirement R3 is that beliefs can be anything off the equilibrium path unless they are somehow constrained by Bayes rule. How can it be that Bayes rule has bite when an information set is not reached? Consider the following example that does not have Nature, but by definition can have beliefs assigned to each information set, and thus can accommodate requirement R3:*

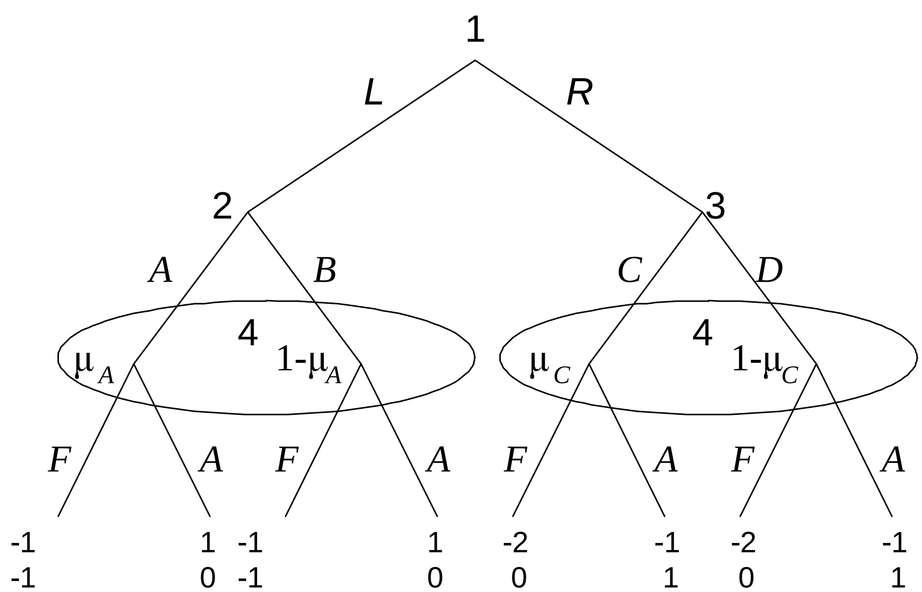

In this game, imagine that a profile of strategies has player 1 playing L, player 2 playing a mixed strategy $\sigma_A = \Pr\{A\}$*, and player 3 playing a mixed strategy* $\sigma_C = \Pr\{C\}$*. If these are the beliefs of player 4, then in this case the information set following player 2 is reached with probability 1, and the belief must be* $\mu_A = \sigma_A$*.*

However, the information set following player 3 is not reached with positive probability, and R3 as stated above allows for any beliefs in this information set. A more careful scrutiny would imply the following logic: if player 4 realizes that his information set following player 3 was reached, then he must assume that player 1 deviated from his strategy of playing L. But why should he believe that player 3 too deviated from his strategy σ_c*? Indeed, a natural belief would be that only player 1 deviated, and therefore we must have* $\mu_C = \sigma_C$*. This is the meaning of the more accurate statement of requirement R3 that sometimes information sets that are reached with zero probability can still have constraints on beliefs. We will not deal with games like in this example, but rather with games that are in the spirit of the entry game. Namely, if an information set is not reached then any beliefs can be formed since they will not be determined by some other players' strategies.* ■

All of the first three requirements were imposed to introduce well defined beliefs for every player at each of his information sets. Now we come to the final requirement of sequential rationality, which imposes constraints on behavior given the system of beliefs. Since we are interested in sequential rationality, then we want players to play a best response to their beliefs at every stage in the game. That is,

Requirement R4: Given their beliefs, players' strategies must be *sequentially rational*. That is, in every information set players will play a best response to there beliefs about where they are in the information set, and to their opponents' strategies.

We can now incorporate requirements R1-R4 to see that the non-credible equilibrium in the entry game with incomplete information is fragile. To see this, notice that once we specify a belief μ for player 2, then *for any* $\mu \in [0,1]$, it is a best response for player 2 to play A. This means that once we endow player 2 with a well defined belief (R1) then despite the fact that these beliefs are not restricted if player 1 chooses OO (by R3) the Bayesian Nash equilibrium (OO, F) has player 2 not playing a best response to any belief, which violates R4.

To write down R4 formally, consider player i with beliefs over information sets derived from the belief system μ, and with beliefs σ_{-i} about his opponents' strategies. Then, R4 requires that if h is an information set for player i, it must be true that he is playing a strategy s_i that satisfies,

$$E[u_i(s_i, \sigma_{-i}, t_i)|h, \mu] \geq E[u_i(s_i', \sigma_{-i}, t_i)|h, \mu] \text{ for all } s_i' \in S_i,$$

where expectations are given over the random beliefs of player i through μ.

Now that we have defined beliefs, and given beliefs the motion of sequential rationality, we need to combine all these together to define a coherent equilibrium concept:

Definition: A profile of strategies $s^* = (s_1^*, ..., s_n^*)$, together with a system of beliefs μ, constitute a *Perfect Bayesian Equilibrium* (PBE) for an n player game if they satisfy requirements 1 through 4 above.

This definition puts together our four requirements in a way that will guarantee sequential rationality. The next obvious question is, how do we find PBE for a game? The following fact is useful:

Fact: If $s^* = (s_1^*, ..., s_n^*)$ together with a system of beliefs μ, constitute a PBE for an n player game, then $s^* = (s_1^*, ..., s_n^*)$ is a BNE for that game.

The proof is quite straightforward. If this were not true, then it means that some player i has a better response against s_{-i}^*, which means that in some information set he can change his action and get a higher payoff than from following s_i^*. But this will contradict the fact that s^* is a PBE because in that information set i is not playing a best response (this goes against requirement 4).

What this fact implies is that if we first find all the profiles of strategies in the Bayesian game that are Bayesian Nash equilibria of the game, then we can systematically check for each Bayesian Nash equilibrium to see whether we can find a system of beliefs so that together they constitute a Perfect Bayesian equilibrium.

A first step in this process is the following observation. If a game has a BNE that has all the information sets being on the equilibrium path (all information sets are reached with positive probability), then we will get unique beliefs pinned down by Bayes rule due to requirement 2. This observation leads to the following result:

Proposition 20 *If a profile of (possibly mixed) strategies $\sigma^* = (\sigma_1^*, ..., \sigma_n^*)$ is a BNE of an n person game Γ, and if σ^* has* all the information sets reached with positive probability*, then σ^*, together with the belief system μ uniquely derived from σ^* by requirement 2, constitute a PBE for Γ.*

Proof: The logic of the proof is simple. Assume that the proposition was incorrect, that is, σ^* is a BNE but (σ^*, μ) is not a PBE. This implies that for some player i, playing according to σ_i^* is not a best response in (at least) one of his information sets, say h_i, against the beliefs derived from μ, and in that information set he has some better response $\sigma_i^{'}$ that is identical to σ_i^* except for the information set h_i. But since μ is derived from σ^*, then $\sigma_i^{'}$ will also be better than σ_i^* against σ_{-i}^*, contradicting the fact that σ^* is a BNE. ■

We now turn to some nice applications of the Perfect Bayesian equilibrium concept. L16

22
Signaling Games

"Signalling games" is the terminology used for a wide class of games whose origins come from the Nobel winning contribution developed by Michael Spence in his Ph.D. thesis, which investigated the role of education as an instrument that sends information to potential employers about a person's intrinsic abilities, and not necessarily what they have learned.

These games share the following structure:

- Nature chooses a type for player 1 that player 2 does not know, but care's about (common values).
- Player one has a rich action set in the sense that there are at least as many actions as there are types.
- Player one chooses an action first, and player 2 then responds *after* observing player 1's choice.
- Given player 2's belief about player 1's strategy, player 2 updates his prior after observing player 1's choice.

The reason these games are called signaling games is because of the potential role that player 1's actions can play. Namely, if in equilibrium each type of player

1 is playing a different choice then *in equilibrium*, the action of player 1 will reveal player 1's type to player 2 through the system of beliefs. Because of the signaling potential of player 1's strategies, these games have two important classes of equilibria:

1. **Pooling Equilibria:** These are equilibria in which *all the types of player 1 choose the same action*, thus revealing nothing to player 2. Player 2's beliefs must be derived from Bayes rule *only in the information set that is reached with probability 1*, and all other information sets are reached with probability 0. In those information sets, player 2 must hold beliefs that support his strategy, and the strategy of player 2 that is sequentially rational given his beliefs is what keeps player 1 from deviating from his "pooling" strategy.

2. **Separating Equilibria:** These are equilibria in which each type of player 1 *chooses a different action*, thus revealing himself in equilibrium to player 2. Player 2's beliefs are thus well defined by Bayes rule in all the information sets that are reached. If there are more actions than types for player 1, then player 2 has to have beliefs for the information sets that are not reached (the actions that no type of player 1 chooses), which in turn must support the strategy of player 2, and player 2's strategy supports the strategy of player 1.

The choice of words is not coincidental. In a **pooling equilibrium** all the types of player 1 *pool together* in the action set, and thus player 2 can learn nothing from the action of player 1. That is, his posterior belief after player 1 moves must be equal to is prior belief that comes immediately from the distribution of Nature's choices of types for player 1. In a **separating equilibrium** each types of player 1 *separates* from the others by choosing a unique action that no other type chooses. Thus, after observing what player 1 did, player 2 can infer *exactly what type player 1 is.*[1]

[1] There is a third class of equilibria called **hybrid equilibria**. In these, different types choose different *mixed strategies*, and the same information set (action) can be reached by different types. Thus, by Bayes rule, player 2 can learn something about player 1, but cannot infer exactly which type he is. These are more technically demanding in nature and are beyond the scope of this text. See Fudenberg and Tirole (1990), Chapter X for more on this.

The incomplete information entry-game that we analyzed in the previous section can be used to illustrate these two classes. In the Bayesian Nash equilibrium (OO, F), both types of player 1 chose "out", so player 2 learns nothing about player 1's type when it is his turn to play (in this case, he has no active action following player 1's decision to stay out). Thus, (OO, F) is a *pooling equilibrium*. In the Bayesian Nash equilibrium (IO, A), which is also a Perfect Bayesian equilibrium, player 1's action perfectly reveals his type: if player 2 sees entry, he believes with probability 1 that player 1 is strong, while of player 1 choose to stay out then player 2 believes with probability 1 that player 1 is weak. Thus, (IO, A) is a *separating equilibrium*.

22.1 The MBA Game

In this section we will analyze a very simple version of an education signaling game in the spirit of Spence's work, that sheds some light on the signaling value of education. To focus attention on the signalling value of education, we will ignore any productive value that education has. That is, we assume that a person learns nothing productive, but has to "suffer" the loss of time and the hard work of studying to get a diploma, in this case an MBA. The game proceeds as follows:

1. Nature chooses player 1's skill (a worker). It can be "high" or "low", and only player 1 knows his skill. Thus, his type set is $T_1 = \{H, L\}$. It is common knowledge that $p \equiv \Pr\{H\}$.

2. After player 1 learns his skill (type) he can choose whether to get a MBA diploma (D) or stay at his undergraduate level (U), so that his action set is $A_1 = \{D, U\}$. Getting a diploma requires some effort that is *type dependent*. Namely, a worker incurs a private cost c_{t_1}if he gets an MBA, and 0 of he does not. We assume that "smarter" people (i.e., high skilled workers) find it easier to study, which is captured by the values $c_H = 2$ and $c_L = 5$.

3. Player 2 is an employer, who can allocate the worker to one of two jobs in a project he is running. Specifically, player 2 can assign player 1 to be either a manager (M) or a blue-collar worker (B), so that her action set is

$A_2 = \{M, B\}$. The employer will retain the profit from the project, and must pay a wage to the worker depending on the job assignment. The market wage for a manager is 10, while the market wage for a blue-collar worker is 6.

4. Profit is determined by the combination of skill and job assignments. It is assumed that the MBA diploma adds nothing to productivity. A high skilled worker is relatively better at managing, while a low skilled worker is relatively better at blue-collar work. The employer's net profits of skill-assignment matches are given in the following matrix:

$$\begin{array}{cc|c|c|} & & \multicolumn{2}{c}{\overbrace{\quad\quad\quad}^{\text{Assignment}}} \\ & & M & B \\ \cline{3-4} \text{skill} \left\{ \right. & H & 10 & 5 \\ \cline{3-4} & L & 0 & 3 \\ \cline{3-4} \end{array}$$

Give the information about the game that is laid out in 1-4 above, the complete game tree is represented in figure 5.4. We depart from our previous methods of drawing information sets to a different one that is more accommodating for games that are more complex, say, than the entry game we have seen before. In this game we draw the information sets as dashed lines that connect the nodes that are in the same information set. Namely, the two nodes that follow the choice U are in one information set, and the two nodes that follow the choice D are in the second information set. In the analysis that follows, we refer to the first information set as I_U and to the second as I_D.

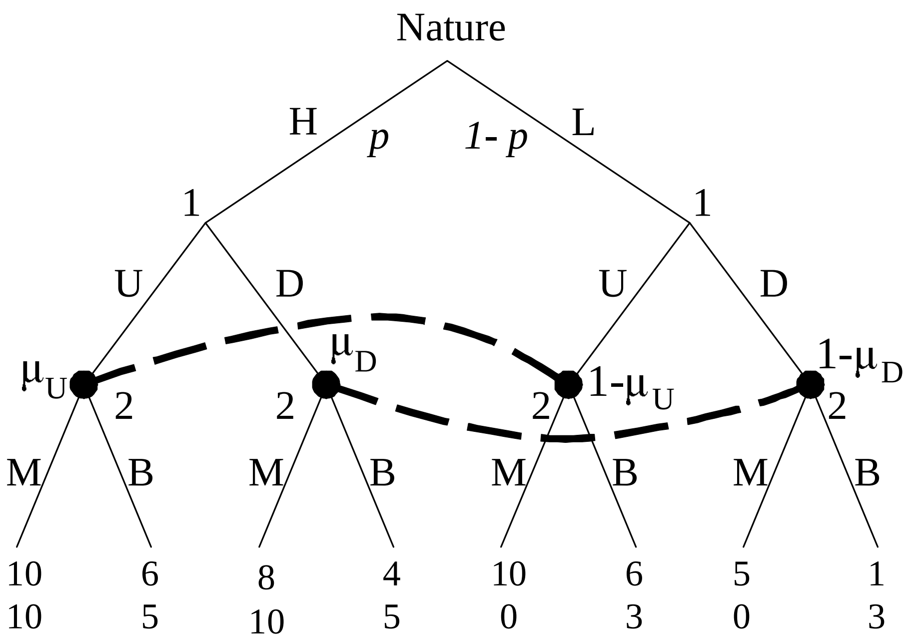

First, to define beliefs, let μ_U denote the belief of player 2 that player 1's type is H conditional on player 1 choosing U, and similarly let μ_D denote the belief of player 2 that player 1's type is H conditional on player 1 choosing D. These beliefs will be determined by the distribution of Natures choice, together with the beliefs that player 2 holds about the strategy that player 1 is playing. For equilibrium analysis, these beliefs will be determined according to requirements R2 and R3.

In general, if player 1 is using a mixed strategy where type H chooses U with probability σ^H and type L chooses U with probability σ^L, and if both σ^H and σ^L are strictly between 0 and 1(i.e., the two types are randomizing) then requirement 2 implies that by Bayes rule,

$$\mu_U = \frac{p\sigma^H}{p\sigma^H + (1-p)\sigma^L} ,$$

and

$$\mu_D = \frac{p(1-\sigma^H)}{p(1-\sigma^H) + (1-p)(1-\sigma^L)} .$$

If, however, both $\sigma^H = \sigma^L = 1$, then both types are choosing U, and beliefs are only well defined for So, for the first case we nee to determine what μ_U is.

We are now ready to see how to solve for PBE in the example of the MBA game above. Since each player has 2 information sets, and two actions in each of these sets, then each player has 4 pure strategies. Let player 1's strategy be denoted $a_1^H a_1^L$, where $a_1^t \in \{U, D\}$ denotes what player 1 does if he is type $t \in \{H, L\}$. Similarly, let $a_2^U a_2^D$ denote player 2's strategy where $a_2^k \in \{M, B\}$ denotes what player 2 does if he observes that player 1 chose $k \in \{U, D\}$.

Now, assume that Nature chooses according to $p = \frac{1}{4}$, so that we can derive the bi-matrix that is the normal-form representation of the MBA Bayesian game. As we have demonstrated earlier for the entry game, the payoffs in the matrix are calculated by taking each pair of pure strategies, observing which paths are played with the different probabilities due to Nature's choice, and then writing down the derived expected utilities from this pair of strategies. For example, if (UD, MB) are the pair of strategies, then with probability $\frac{1}{4}$ Nature chooses type H for player 1 who chooses U, and in response player 2 chooses M, yielding a payoff pair of (10,10). With probability $\frac{3}{4}$ Nature chooses type L for player 1 who chooses D, and in response player 2 chooses B, yielding a payoff pair of (1,3). The expected pair of payoffs for the players from the strategy (UD, MB) is therefore,

$$\frac{1}{4}(10, 10) + \frac{3}{4}(1, 3) = (3.25, 4.75)\,.$$

Similarly we calculate the expected payoffs for all the other 15 entries. Notice that when player 1 plays the same action for the different types (rows 1 and 4) then part of player 2's strategy is never used so there are repeat entries which reduces the number of calculations needed. The matrix representation is then,

Player 2

player 1	MM	MB	BM	BB
UU	$\underline{10, 2.5}$	$\underline{10, 2.5}$	$\overline{6, 3.5}$	$\underline{\overline{6, 3.5}}$
UD	$6.25, 2.5$	$\overline{3.25, 4.75}$	$5.25, 1.25$	$2.25, 3.5$
DU	$9.5, 2.5$	$8.5, 1.25$	$\underline{\overline{6.5, 4.75}}$	$4.5, 3.5$
DD	$5.75, 2.5$	$\overline{1.75, 3.5}$	$5.75, 2.5$	$\overline{1.75, 3.5}$

If we follow the method of underlining 1's best responses for each column and overlining player 2's best responses for each row, we immediately observe that there are two pure strategy Bayesian Nash equilibria: (UU, BB) and (DU, BM). To see whether these can be part of a PBE, we need to find a system of beliefs that support the proposed behavior, and that together with these strategies satisfy requirement R1 through R4.

From Proposition 16 above it follows that (DU, BM) can be part of a PBE. This follows because all of the information sets are reached with positive probability. In particular, the derived beliefs from (DU, BM) are $\mu_U = 0$ and $\mu_D = 1$.[2] It is easy to check that player 2 is playing a best response to these beliefs in each of his information sets, and that player 1 is playing a best response in each of his. So, (DU, BM) together with $\mu_U = 0$ and $\mu_D = 1$ constitute a PBE.

What about (UU, BB)? Unique beliefs are only derived for information set I_U since I_D is reached with zero probability. In particular, $\mu_U = \frac{1}{4}$ and μ_D is not well defined. Therefore, to see whether (UU, BB) can be part of a Perfect Bayesian equilibrium we need to see if there are beliefs μ_D that support B as a best response for player 2 in information set I_D.

For B to be a best response in the information set I_D, it must be the case that given the belief μ_D, the expected payoff from B is higher than the expected payoff from M. This can be written down as,

$$5 \cdot \mu_D + 3 \cdot (1 - \mu_D) \geq 10 \cdot \mu_D + 0 \cdot (1 - \mu_D),$$

which is true if and only if $\mu_D \leq \frac{3}{8}$. this implies that we *can* support (UU, BB) as part of a PBE. In particular, (UU, BB) together with belief $\mu_U = \frac{1}{4}$ and any belief satisfying $\mu_D \in [0, \frac{3}{8}]$ constitute a PBE.

There is something fundamentally different between the two pure strategy PBE that we have just derived. In the first, (DU, BM), different types of player 1 choose different actions, thus using there actions to *reveal* to player 2 their true type. In other words, this is a *separating Perfect Bayesian equilibrium*. In the second, (UU, BB), both types of player 1 do the same thing, and thus player 2 learns nothing from player 1's action, and this is a *pooling Perfect Bayesian equilibrium*. L17

[2] Using the equations we derived from Bayes rule above, this is the case where $\sigma^H = 0$ and $\sigma^L = 1$, and the resulting μ_U and μ_D follow.

22.2 Entry Deterrence

We turn to another application of the analysis of dynamic games of incomplete information to an important antitrust issue. Bain (1949) argued that incumbent firms can engage in *limit pricing* to deter entry, and thus try to gain control as a long run monopolist. This practice is defined as one in which a firm prices below marginal costs, which is obviously a losing strategy in the short run. But, if by doing this they run a competitor out of business and turn into a monopolist, then this might be a winning strategy in the long run.

The case for antitrust is clear: if some firm is faced with potential competition and is charging a price that is "too low" in the sense that the firm is losing money in the short run, and this behavior deters entry in the long run to capture future monopoly profits, then the form is engaging in limit pricing which is anti-competitive behavior.

This argument, however, should make one feel a bit uncomfortable: is it reasonable to punish firms for charging low prices? How can we convince ourselves that Bain's intuition is correct and that a firm may indeed choose to incur low run losses to generate long run profits? This was answered in a very nice paper by Paul Milgrom and John Roberts (1982). We present a simplified version of their analysis in the following game:

- Consider a market for some good that will last for 2 periods with demand $P = 5 - Q$ in each period.
- In period $t = 1$ only one firm, the incumbent (player 1) is in the market, and this incumbent will continue to produce in period $t = 2$.
- A potential entrant (player 2) observes what happened in the market in period $t = 1$ and then makes a choice of whether to enter at a fixed cost of $\frac{1}{2}$ and then compete against the incumbent in period $t = 2$, or to remain out of the market.
- Each firm has marginal costs c_i, where it is common knowledge that the entrant's marginal costs are $c_2 = 2$.

- There is asymmetric information with respect to the costs ("type") of the incumbent. In particular, the entrant only knows that $c_1 \in \{1, 2\}$ so that the incumbent can either be more efficient than the entrant (have low (L) costs of $c_L = 1$) or not (have high (H) costs, $c_H = 2$). It is common knowledge that $\Pr\{c_1 = 1\} = \frac{1}{2}$, but only the incumbent knows his true costs.

The timing of this game that captures this story is described as follows:

1. Nature choose the type of player 1, $c_1 \in \{c_L, c_H\} = \{1, 2\}$, each type with equal probability.

2. Player 1 (the incumbent) observes his costs and decides how much to produce in the first period, q_1^1, and the price in the first period is then $P = 5 - q_1^1$.

3. Player 2 (entrant) observes q_1^1 and decides whether to enter or not at a cost of $F = \frac{1}{2}$.

4. If player 2 stayed out, then in period $t = 2$ player 1 chooses q_1^2 and the price is $P = 5 - q_1^2$.

5. If player 2 enters, then in period $t = 2$ each player i simultaneously chooses q_i^2 and the price is $P = 5 - q_1^2 - q_2^2$ (Cournot competition).

Given the continuous action space of this game (quantities) it is not too useful to write down a game tree, but it can focus our attention on what we need to worry about. A way of considering such a game tree is depicted in figure 5.x.

Solving for the Perfect Bayesian equilibria of this game is not as simple as for the MBA game. The reason is that the strategy spaces are continuous, so we can't look for all the BNE in a bi-matrix and then check whether they can be supported as PBE. Thus, to solve for the PBE of this game, we go to the heart of what sequential rationality means, and solve the game using similar techniques to backwards induction. We proceed in two steps, first looking for separating PBE and then for pooling PBE.

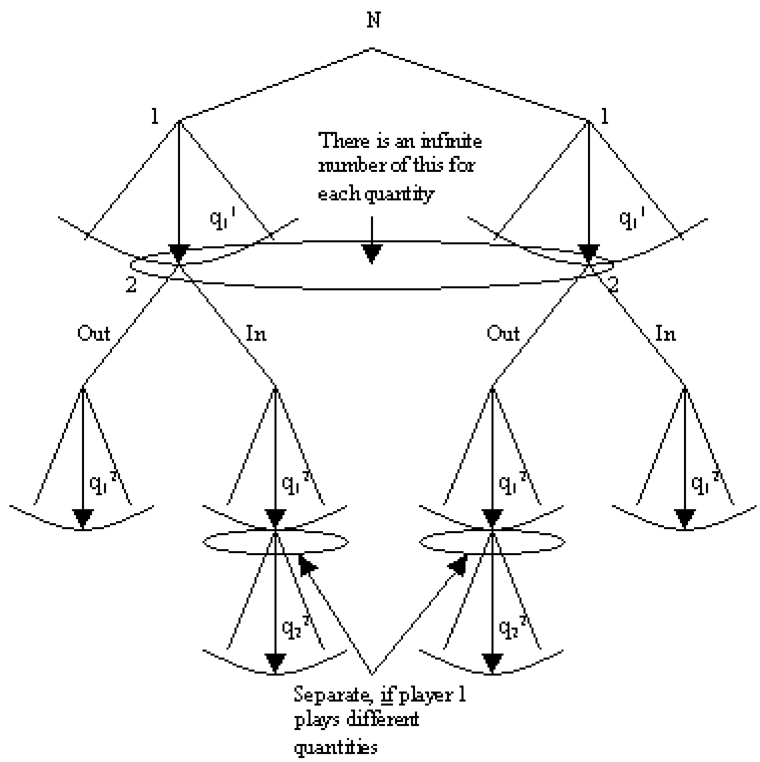
N
1
1
There is an infinite number of this for each quantity
q1'
q1'
2
2
Out
In
Out
In
q1²
q1²
q1²
q1²
q2²
q2²
Separate, if player 1 plays different quantities

FIGURE 22.1.

Separating Equilibria

We proceed to check for the existence and nature of separating PBE using a sequence of steps:

Step 1: What will q_1^2 be if firm 2 stays out?

The answer is quite simple: firm 1 is then a monopoly and maximizes:

$$\max_{q_1^2} (5 - q_1^2 - c_1)q_1^2$$

which yields the solution using the first order condition (which in this case is necessary and sufficient):

$$q_1^2 = \frac{5 - c_1}{2} = \begin{cases} 2 & \text{if } c_1 = 1 \\ 1.5 & \text{if } c_1 = 2 \end{cases}$$

Step 2: What will q_1^2 and q_2^2 be if firm 2 enters and firm 1's costs are known to all?

The answer is again simple: both firms play a Cournot game. Firm 1 solves

$$\max_{q_1^2} (5 - q_1^2 - q_2^2 - c_1)q_1^2$$

which yields the best response function using the first order condition (which is again necessary and sufficient):

$$BR_1 : q_1^2 = \frac{5 - q_2^2 - c_1}{2} \tag{(BR1)}$$

and similarly for firm 2 (but there is the fixed costs in the objective function that does not effect the first order condition):

$$\max_{q_2^2} (5 - q_1^2 - q_2^2 - c_1)q_2^2 - \frac{1}{2}$$

which yields the best response function using the first order condition (which is again necessary and sufficient):

$$BR_2 : q_2^2 = \frac{5 - q_1^2 - c_2}{2} = \frac{3 - q_1^2}{2} \tag{(BR2)}$$

There are two cases to be considered if c_1 is known to all:

Case 1: In this case $c_1 = 2$, and (BR1) and BR2) imply that $q_1^2 = q_2^2 = 1$, $P = 3$, and profits for the two firms in the second period are $\pi_1^2 = (3-2)1 = 1$, while $\pi_2^2 = (3-2)1 = \frac{1}{2} = \frac{1}{2}$.

Case 2: In this case $c_1 = 1$, and (BR1) and BR2) imply that $q_1^2 = \frac{5}{3}$, $q_2^2 = \frac{2}{3}$, $P = \frac{8}{3}$, and profits for the two firms in the second period are $\pi_1^2 = (\frac{8}{3} - 1)\frac{5}{3} = \frac{25}{9}$, while $\pi_2^2 = (\frac{8}{3} - 2)\frac{2}{3} - \frac{1}{2} = -\frac{1}{18}$.

Step 3: What quantities for each type will support firm 1's separating (signaling) behavior?

To answer this we consider logic of separating versus pooling equilibrium. In *any* separating equilibrium, it must be that different types of player 1 choose different actions, thus in any such equilibrium, player 2 learns the type of player 1. To continue our analysis, let us consider a separating PBE with the following notation: $q_1^{t\tau}$ denotes the choice of player 1 in period $t \in \{1, 2\}$ when he is of type $\tau \in \{H, L\}$. That is, $q_1^{t\tau}$ is the strategy of type τ in period t. Let $q_2^2(q_1^1)$ denote the choice of player 2 in period 2 in response to q_1^1 being played by firm 1 in period 1, which is the strategy of player 2. Finally, let $\mu(q_1^1) = \Pr\{c_1 = 2|q_1^1\}$ be player 2's belief in period 2 in response to q_1^1 being played by firm 1 in period 1. Using this notation we can write a PBE as[3]

$$\{\underbrace{q_1^{1L}, q_1^{1H}}_{\text{firm 1,t=1}}, \underbrace{q_1^{2L}, q_1^{2H}}_{\text{firm 1,t=2}}, \underbrace{q_2^2(q_1^1)}_{\text{firm 2,t=2}}, \underbrace{\mu(q_1^1)}_{\text{firm 2's beliefs}}\} \quad .$$

We can now establish a series of claims:

Claim 1: In any separating PBE, when firm 2 believes that $c_1 = 1$ then it will stay out and firm 1 will choose $q_1^{2L} = 2$, while if firm 2 believes that $c_1 = 2$ it will enter, and both firms produce $q_1^{2H} = q_2^2 = 1$.

[3] To be precise, $q_1^{2\tau}$ must be a function of whether or not form 2 decides to enter or not, and firm 2's strategy should be two elements; first whether to enter or not depending on q_1^1, and then upon entry what the choice $q_2^2(q_1^1)$ is. We are simplifying by collapsing both of these choices of firm 2 into the quantity choice, and we will let $q_2^2(q_1^1) = \mathbf{0}$ be the choice of no entry.

The proof just follows from the analysis above of the monopoly problem (Step 1) and the Cournot problem (Step 2), and the fact that we require firm 2's behavior to be sequentially rational. This is one of the nice features of separating PBE: even though player 2 *does not know* the type of player 1, in a separating equilibrium types are revealed, and player 2 must act *as if he knew* the type of player 1.

Claim 2: In any separating PBE, in the first period a high cost incumbent must produce $q_1^{1H} = 1.5$.

Proof: From claim 1, in any separating PBE it must be that following q_1^{1H} firm 2 will enter, and following q_1^{1L} firm 2 will stay out and let firm 1 be a monopolist (otherwise it would not be a separating PBE). If q_1^{1H} is indeed part of a PBE, then it must be that $q_1^{1H} = 1.5$. To see this, assume not. Then in period 1 firm 1 is making less than monopoly profits when its marginal costs are $c_1 = 2$, and in period 2 firm 1 is making Cournot profits. Now consider a deviation of firm 1 (when $c_1 = 2$) to the monopoly quantity of 1.5. In period 1 profits are higher, which means that for this deviation not to be profitable it must be that firm 1 gets *less than Cournot profits* in the second period. But this cannot happen because either: (i) firm 2's beliefs after the deviation remain $\Pr\{c_1 = 2\} = 1$ in which case they will play the same Cournot game, or (ii) firm 2 changes its beliefs to $\Pr\{c_1 = 2\} < 1$ in which case firm 1 will make higher than Cournot profits (either firm 2 will stay out or it will play Cournot against an unknown rival and produce less than 1 depending on its beliefs). This concludes that if $q_1^{1H} \neq 1.5$ then firm 1 has a profitable deviation to $q_1^{1H} = 1.5$.■

So, we have established from our analysis that if a separating PBE exists then it must satisfy $q_1^{2L} = 2$, $q_1^{2H} = 1$, $q_2^2(q_1^{2L}) = 0$, $q_2^2(q_1^{2H}) = 1$, $\mu(q_1^{1L}) = 0$ and $\mu(q_1^{1H}) = 1$. From Claim 2 we established that it must satisfy $q_1^{1H} = 1.5$. We are left to find two more things: first, we must define beliefs for all other quantities $q_1^1 \notin \{q_1^{1H}, q_1^{1L}\}$, and we have to find q_1^{1L}. If we find these values in a way that strategies and beliefs satisfy requirements R1-R4, then we have found a separating PBE.

Step 4: Setting off-equilibrium path beliefs.

Here we will use a "trick" that is common for games with continuous strategy spaces, and it is similar to what we did for the MBA game. Recall that we want the separating PBE to work in such a way that each type of player 1 $\tau \in \{L, H\}$ will stick to his strategy $q_1^{1\tau}$, rather than deviating to some other quantity q_1^1. To do this, we can make the continuation game following any such deviation to be *as worse as possible for player 1*. When does this happen? Precisely when player 2 enters and believes that player 1 has high costs. This follows because when this player 2 acts in this way, player 1 faces the most severe second period competition, and this is the most undesirable outcome for player 1. So, the easiest way to prevent deviations and keep player 1 *on the equilibrium path*, is by setting beliefs that make *off the equilibrium path* very unattractive, and this will follow from:

$$\mu(q_1^1) = \begin{cases} 1 & \text{if } q_1^1 \neq q_1^{1L} \\ 0 & \text{if } q_1^1 = q_1^{1L} \end{cases}$$

These beliefs cause a unique best response for player 2 which is not to enter if $q_1^1 = q_1^{1L}$ and to enter and produce $q_2^2 = 1$ if $q_1^1 \neq q_1^{1L}$. So, we have set beliefs, and in return we know what $q_2^2(q_1^1)$ must be to satisfy sequential rationality.

Step 5: What should q_1^{1L} be?

We know that once we have calculated all of the equilibrium components we found above, then for q_1^{1L} to satisfy the missing piece of the puzzle, it must satisfy the following two important conditions:

1. When firm 1 is a L type, it prefers to choose q_1^{1L} over any other quantity, in particular over q_1^{1H}
2. When firm 1 is a H type, it prefers to choose q_1^{1H} over q_1^{1L}

The first condition is that type L is playing a best response. We can think of this as first, a L type does not want to choose some other random quantity, and second, a L type does *not want to "imitate"* a H type. The second condition just says that a H type does not want to imitate a L type. What about other quantities? Using the belief system we defined above, claim 2 already implies that a H type prefers

$q_1^{1H} = 1.5$ to *any other non-monopoly profit that induces entry*. The reason we have to care about imitating a L type is because q_1^{1L} *induces exit, not entry*.

We call these two conditions, that each type prefers choosing his designated action rather than imitating some other type, *incentive compatibility conditions*. The meaning is precisely that each type has an incentive to choose what the equilibrium prescribes for him, and not choose what the equilibrium prescribes for the other types. this is something that was true for the PBE that we found in the MBA game, and is a general set of constraints for signalling games.

To solve for q_1^{1L} we can start with a **natural candidate**: *monopoly quantity for the* L *type*. From the analysis of step 1 above this would be $q_1^{1L} = 2$, and the first period profits would be 4. If this is a separating PBE, this choice by firm 1 induces form 2 to stay out, and thus firm 1 again gets a profit of 4. Now we must check for incentive compatibility for the low type: would the L type of firm 1 want to deviate and choose q_1^{1H}? By construction of beliefs the answer is clearly no since it loses both in the first period (by deviating from monopoly profits) and in the second since it induces entry. So we are fine with L playing a best response.

What about incentive compatibility of type H? If he follows his strategy, he produces 1.5 in the first period and gets monopoly profits of 2.25, and this induces entry that results in Cournot profits of 1 in the second period, so total profits are 3.25. From the analysis above, he would not deviate to any other *entry inducing* quantity. We are left to check whether he would want to deviate and choose $q_1^{1L} = 2$ instead of $q_1^{1H} = 1.5$. If he indeed deviates he gets total profits of:

$$\pi = \overbrace{(5-2-2)2}^{\text{1st period}} + \overbrace{(5-1.5-2)1.5}^{\text{2nd period}} = 4.25 > 3.25\,!$$

This happens because by imitating a L type, the H type is sacrificing monopoly profits in the first period to get them in the second period, and it turns out that the sacrifice is not as bad as playing Cournot in the second period. So, the suggested $q_1^{1L} = 2$ violates incentive compatibility of the H type.[4]

[4] Note that if we introduce discounting into this model then this conclusion might change. In particular, getting monopoly profits tomorrow is not as good as getting them today, so the sacrifice may be enough to keep the H type from deviating. This will happen if there is enough discounting of second period profits.

Finding q_1^{1L} : The intuition from the previous unsuccessful attempt is that for incentive compatibility to hold we need an H type to *suffer enough from deviating.* This can be done by making the deviation less attractive, which means we need to make q_1^{1L} larger. This follows because if we make q_1^{1L} smaller, it is closer to 1.5, which means higher profits for a H type. This would make the deviation of a H type even more attractive since it is closer to his monopoly profits.

Fortunately, we don't need to make any guesses about q_1^{1L}: In our analysis above we have already *pinned down all the other components* of *any separating PBE.* So, to see if such a PBE exists all we need is to find a q_1^{1L} that satisfies the two incentive compatibility constraints. First, we need to make sure that a H type does not want to deviate to q_1^{1L}. This will be satisfied if the first period profits from deviating, plus the second period monopoly profits he gets (because firm 2 stays out after q_1^{1L}), are no more than what he gets from producing monopoly first and Cournot later. This can be written as,

$$(5 - q_1^{1L} - \overbrace{2}^{c_1^H})q_1^{1L} + \overbrace{2.25}^{\substack{\text{2nd period} \\ \text{monopoly}}} \leq 3.25$$

which is reduced to,

$$q_1^{1L} \geq \frac{3}{2} + \sqrt{\frac{5}{4}} = 2.618 > 2\,!$$

That is, there is a *lower bound* on q_1^{1L} that captures the smallest sacrifice needed to deter the H type from deviating from q_1^{1H} to q_1^{1L}. Similarly, we have to make sure that the L type will not want to deviate from q_1^{1L} to some other quantity. This would put an upper bound on q_1^{1L}. To find this upper bound we need to figure out what would be the *best deviation* for a type L. Given the belief system we are imposing, any deviation from q_1^{1L} will induce entry, and firm 2 will produce $q_2^2 = 1$ from the analysis above. So, a type L who deviates should play a best response to $q_2^2 = 1$ in the second period. This is calculate by:

$$\max_{q_1^2} (5 - q_1^2 - \overbrace{1}^{q_2^2} - \overbrace{1}^{c_1^L})q_1^2$$

which is maximized by $q_1^2 = 1.5$, yielding a second period profit of 2.25. Since *any deviation* of a L type will be followed by this BR, then a L type gains the most from deviating to his monopoly quantity of $q_1^1 = 2$ with a first period profit of 4. This implies that the best deviation for firm 1 when it is a L type yields a profit of 6.25, and incentive compatibility of the L type is guaranteed if,

$$(5 - q_1^{1L} - 1)q_1^{1L} + 4 \geq 6.25$$

which reduces to,

$$q_1^L \leq 3.323.$$

Summing up the complete analysis we have done above allows us to state the following conclusion on separating PBE of this game:

Result: In the entry deterrence game there is a continuum of separating PBE. In particular, any profile of strategies $\{q_1^{1L}, q_1^{1H}, q_1^{2L}, q_1^{2H}, q_2^2(q_1^1)\}$ together with beliefs $\mu(q_1^1)$ that satisfies

$$\begin{aligned}
2.618 &\leq q_1^L \leq 3.323, \\
q_1^{1H} &= 1.5, \\
q_1^{2L} &= 2, \\
q_1^{2H} &= 1, \\
\mu(q_1^1) &= \begin{cases} 1 & \text{if } q_1^1 \neq q_1^{1L} \\ 0 & \text{if } q_1^1 = q_1^{1L} \end{cases}, \\
q_2^2(q_1^1) &= \begin{cases} 1 & \text{if } q_1^1 \neq q_1^{1L} \\ 0 & \text{if } q_1^1 = q_1^{1L} \end{cases},
\end{aligned}$$

is a separating PBE.

It is worth noting something special about all these PBE: for a low cost firm to deter entry, it must choose a quantity that is higher than it's short-run profit maximizing (monopoly) profits. Why? Because to *credibly* scare off opponents, it must choose something that a high cost firm would not find profitable to do. In other words, for a *signal to be credible*, it must be that no other type would want to use the signal. For a signal to be part of an equilibrium, it must be not too costly so that one type wants to use it, but costly enough so that other types do not. Note that a low cost firm is not charging below marginal costs here, but it is clearly charging below monopoly costs.

If we would extend the game to many more periods, and if by deterring entry in the first period the firm would be a monopoly for several periods thereafter, the lower bound we derived would have to be higher, as would the upper bound. This would eventually drive the first period separating quantity for the low type to be so high that price would be below marginal cost. This would be the extent of the sacrifice needed to make sure that the signal is credible. Thus, Bain's intuition survives the test of a formal model that is built to capture the essence of anti-competitive limit pricing.

Pooling Equilibria

Pooling equilibria are easier to find. For this game a pooling equilibrium will be $\{q_1^{1*}, q_1^{2L}, q_1^{2H}, q_2^2(q_1^1), \mu(q_1^1)\}$ because both types of player 1choose the same quantity q_1^{1*} in the first period. This immediately implies that beliefs must satisfy $\mu(q_1^{1*}) = \frac{1}{2}$ since it is common knowledge that nature chooses $\Pr\{c_1 = 1\} = \frac{1}{2}$. This implies that firm 2, if it would enter following q_1^{1*}, is in a static game of incomplete information, and we need to find the BNE (which is an incomplete information version of the Cournot equilibrium) for this second period. Thus, we find q_1^{2L}, q_1^{2H} and $q_2^2(q_1^{1*})$ by solving the following three maximization problems: For a type L firm:

$$\max_{q_1^{2L}} (5 - q_1^{2L} - q_2^2(q_1^{1*}) - 1)q_1^{2L}$$

which yields the FOC:

$$q_1^{2L} = \frac{4 - q_2^2(q_1^{1*})}{2}$$

Similarly, for a H type firm,

$$\max_{q_1^{2H}} (5 - q_1^{2H} - q_2^2(q_1^{1*}) - 2)q_1^{2H}$$

which yields the FOC:

$$q_1^{2H} = \frac{3 - q_2^2(q_1^{1*})}{2}$$

And finally, for firm 2 we have

$$\max_{q_2^2(q_1^{1*})} \frac{1}{2}\left[(5 - q_1^{2L} - q_2^2(q_1^{1*}) - 2)q_2^2(q_1^{1*})\right] + \frac{1}{2}\left[(5 - q_1^{2H} - q_2^2(q_1^{1*}) - 2)q_2^2(q_1^{1*})\right] - \frac{1}{2}$$

which yields the FOC:

$$q_2^2(q_1^{1*}) = \frac{3 - \frac{1}{2}(q_1^{2L} + q_1^{2H})}{2}$$

The BNE of this Cournot game is solving the three FOCs simultaneously, which yields,

$$q_2^2(q_1^{1*}) = \frac{5}{6},\ q_1^{2L} = \frac{19}{12},\ q_1^{2H} = \frac{13}{12}$$

which yields profits to firm 2 of $\frac{7}{18} > 0$. The profits to the H type are $(\frac{13}{12})^2$, and for the L type are $(\frac{19}{12})^2$.

This implies that in a pooling equilibrium, in which firm 2 does not learn the type of player 1 but rather sticks to the priors given by nature, then firm 2 will choose to enter. Now we have to set beliefs for off-equilibrium paths ($q_1^1 \neq q_1^{1*}$), which will determine the best response function of firm 2 $q_2^2(q_1^1)$ for information sets that are off the equilibrium path. We will then need to find the value of q_1^{1*} to complete the equilibrium.

As with the separating equilibrium, the best way to keep the different types of player 1 on the equilibrium path is by making off the equilibrium path very undesirable.. This is again done by setting the "worse" beliefs for player 1, which are $\mu(q_1^1) = 1$ for $q_1^1 \neq q_1^{1*}$, and on the equilibrium path Bayes rule implies that $\mu(q_1^{1*}) = \frac{1}{2}$. This immediately implies the best response function of firm 2,

$$q_2^2(q_1^1) = \begin{cases} 1 & \text{if } q_1^1 \neq q_1^{1*} \\ \frac{5}{6} & \text{if } q_1^1 = q_1^{1*} \end{cases}$$

We are left to find a q_1^{1*} from which neither type 1 would want to deviate. For this, we need to calculate the best deviation of each type, but given the belief system we already have this from the analysis above. The best deviation for both types is to choose monopoly quantities in the first period, followed by the best response to $q_2^2 = 1$ in the second period. These yielded a profit of 3.25 to the H type and 6.25 to the L type as we have calculated above. So, we need two inequalities satisfied for q_1^{1*} to be a best response for both types. For the H type,

$$(5-q_1^{1*}-2)q_1^{1*}+\left(\frac{13}{12}\right)^2 \geq 3.25$$

which reduces to

$$1.0833 \leq q_1^{1*} \leq 1.9167\,. \tag{ICH}$$

Similarly, for the L type,

$$(5-q_1^{1*}-1)q_1^{1*}+\left(\frac{19}{12}\right)^2 \geq 6.25$$

which reduces to

$$1.4931 \leq q_1^{1*} \leq 2.5069\,. \tag{ICL}$$

Thus, since q_1^{1*} must satisfy[5] both (ICH) and (ICL) we can summarize our analysis above as follows:

Result: In the entry deterrence game there is a continuum of pooling PBE. In particular, any profile of strategies $\{q_1^{1*}, q_1^{2L}, q_1^{2H}, q_2^2(q_1^1)\}$ together with beliefs $\mu(q_1^1)$ that satisfies

$$\begin{aligned}
1.4931 &\leq q_1^{1*} \leq 1.9167, \\
q_1^{2L} &= \frac{19}{12}, \\
q_1^{2H} &= \frac{13}{12}, \\
\mu(q_1^1) &= \begin{cases} 1 & \text{if } q_1^1 \neq q_1^{1*} \\ \frac{1}{2} & \text{if } q_1^1 = q_1^{1*} \end{cases}, \\
q_2^2(q_1^1) &= \begin{cases} 1 & \text{if } q_1^1 \neq q_1^{1*} \\ \frac{5}{6} & \text{if } q_1^1 = q_1^{1*} \end{cases},
\end{aligned}$$

is a pooling PBE.

[5] Note that the inequalities make sense: for each type the best deviation is to its monopoly quantity, which is 1.5 for the H type and 2 for the L type. Thus, for both types the range is around the monopoly quantity.

Unlike the separating equilibrium that had the feature of the incumbent's action revealing some important information to the entrant, this is not the case here. In fact, it is not that convincing that different types of the incumbent will choose the same actions, which makes the pooling equilibrium look somewhat artificial, and less appealing. As we will see in section 5.5, there are some convincing arguments that not only favor the separating equilibria as more reasonable, but actually often select one of the separating equilibria as the most reasonable.