

Dynamic games with incomplete information*

1 Perfect Bayesian Equilibrium (PBE)

We have now covered static and dynamic games of complete information and static games of incomplete information. The next step is to focus on dynamic games of incomplete information. When solving games of this type we will need to invoke Bayes' rule since players later in the game will have additional information, which is why we spent time discussing Bayes' rule. The solution concept that we will use for games of this type will be the perfect Bayesian equilibrium. On the one hand, perfect Bayesian equilibrium refines the Bayes-Nash equilibrium concept by ruling out noncredible threats. However, it also rules out some of the SPNE that rely on "noncredible threats" when there is imperfect information. So, perfect Bayesian equilibrium can be viewed as a stronger equilibrium concept than the previous ones.

This basic game captures many types of card games, such as Bridge, Spades, and Poker, in which one player does not know what cards the other player(s) is holding. When playing games of this type people often use both the knowledge of the entire game as well as the actions that have previously occurred in the game to update their beliefs about which node in the information set they are at.

1.1 Definition and structure of a PBE

With a PBE we will still require that all players choose strategies that are best responses to the other player's strategies. However, when there is a player who has multiple decision nodes within an information set we now require that this player specifies a *belief* about which node in the information set he is at. The belief is simply a probability. Note that these probabilities (or beliefs) must follow the laws of probability – no probabilities greater than 1 or less than zero, and the probabilities for all decision nodes within an information set must sum to 1. Thus the first new requirement is that beliefs for the uninformed players must be specified – exactly how these beliefs (or probabilities) are specified will be discussed shortly.


The second requirement that we make is that given the players' beliefs, the strategy choices must be sequentially rational. Thus, each player must be acting optimally at each information set given his beliefs and the other players' subsequent strategies (the strategies that follow the information set). So the second requirement basically says that strategies must now also be best responses to beliefs, in addition to best responses to other players' strategies.

The third and fourth requirements for a PBE specify how the beliefs must be updated. At information sets along the equilibrium path (along the equilibrium path means that the information set is reached when the equilibrium is played) beliefs are determined by Bayes' rule and the players' equilibrium strategies. These first 3 requirements constitute what is known as a *weak perfect Bayesian equilibrium* (WPBE). A fourth requirement is that off the equilibrium path beliefs are also determined by Bayes' rule and the players' equilibrium strategies where possible. The 4 requirements together define a *strong perfect Bayesian equilibrium* (SPBE).

1.2 WPBE and SPBE


Now we will differentiate between a WPBE and a SPBE. Look at the following game:

*Based on Chapter 4 of Gibbons (1992).



Player 2 cannot tell which node he is at if player 1 chooses L or M. There are 2 SPNE. One is that Player 1 chooses R and Player 2 chooses R', and the other is that Player 1 chooses L and Player 2 chooses L'. However, L' strictly dominates R', so Player 1 knows that if he chooses L he will get 2 (choosing M yields a payoff of 0 and R yields a payoff of 1). Player 2 knows this as well, and so his belief is that Player 1 chooses L with probability 1. Thus, Player 2 has updated his belief about which strategy Player 1 is using if Player 2 gets to make a decision. A weak perfect Bayesian equilibrium for this game is that Player 1 chooses L, Player 2 believes that Player 1 chooses L with probability 1, and Player 2 chooses L'. Note that this equilibrium also satisfies requirement 4 because there are no off-the-equilibrium path information sets.

Let's look at another game to illustrate the difference between the weak and strong Perfect Bayesian equilibrium concepts.



First begin by analyzing the subgame that begins at Player 2's decision node.


		Player 3	
		L'	R'
Player 2	L	2, 1	3, 3
	R	1, 2	1, 1

The Nash equilibrium to this subgame is Player 2 chooses L and Player 3 chooses R' . Player 1 knows this, and chooses D . So D, L, R' is a SPNE to the game, and if Player 3 has a belief that Player 2 chooses Left with probability 1 (which Player 3 should have since L is a strictly dominant strategy for Player 2), then requirements 1-3 are satisfied for this to be a weak perfect Bayesian equilibrium. Again, requirement 4 is satisfied because there are no off the equilibrium path information sets.

Now, consider the potential equilibrium where Player 1 chooses A , Player 2 chooses L , Player 3 believes that Player 2 chooses R with probability 1, and Player 3 chooses L' . Note that Player 1 is playing a best response to the strategies L and L' by Players 2 and 3 (Player 1 receives 1 if he plays D and 2 if he plays A). Player 3 is playing a best response given his beliefs about Player 2's actions (if he believes Player 2 is choosing R then Player 3 does better by choosing L'). Player 2 is choosing his strictly dominant strategy of L , and even if he switched his strategy to R he would still receive 0, so he is playing a best response to A, R' . Thus, this set of strategies and beliefs satisfies the first 3 requirements and is a weak perfect Bayesian equilibrium. However, the 4th requirement is not satisfied because Player 3's belief is inconsistent with the fact that Player 2 has a dominant strategy to play L . To implement this consistency requirement, Player 3 must believe Player 2 plays L with probability 1, but then L' is NOT an optimal response (R' is the optimal response) and we are now led back to D, L, R' with Player 3 believing that Player 2 chooses L with probability 1.

2 Separating and Pooling equilibria

There are many instances in which one player knows his own type and then takes an action and another player cannot observe the first player's type but only his action. These types of games generally fall under the category of signaling games, because the action taken by the first mover may (or may not) signal which type the first mover is. In a pooling equilibrium, all types choose the same action (or send the same signal). In a separating equilibrium, different types choose different actions. Thus, in a pooling equilibrium the player without the information on type is unable to update his belief about which type of player chose which action since all types are choosing the same action. We will consider the following game:



Note how the game plays out. Nature first determines the type of the first mover (the Sender). With probability of 0.5 the Sender is type t_1 and with probability 0.5 the Sender is type t_2 . The Sender knows which type he is. Each Sender type can choose either L or R . The Receiver observes only the choice of L or R and not the Sender's actual type. Based upon the observation of L or R the Receiver can then choose U or D . Payoffs then follow, with the Sender's (first mover's) payoff listed first and the Receiver's (second mover's) payoff listed second. There are two potential pooling equilibria and two potential separating equilibria. The two potential pooling equilibria involve either (1) both types t_1 and t_2 choosing L or (2) both types t_1 and t_2 choosing R . The two potential separating equilibria involve either (1) type t_1 choosing R and type t_2 choosing L or (2) type t_1 choosing L or type t_2 choosing R . We will discuss these potential equilibria in detail.

2.1 Potential Separating Equilibria

As mentioned there are two potential separating equilibria. Note that all of these equilibria will consist of (1) a strategy for the Sender (which is an action if the Sender is type t_1 and an action if the Sender is type t_2), (2) a set of beliefs for the Receiver about which decision node in the information set the Receiver is at, and (3) a strategy for the Receiver (which is an action if L is observed and an action if R is observed). We begin by analyzing the one where type t_1 chooses R and type t_2 chooses L .

2.1.1 Type t_1 chooses R and type t_2 chooses L

We begin by specifying the *potential* strategy choice of the Sender. Suppose the Sender uses the separating strategy: R if t_1 and L if t_2 . If this is the case, what does the Receiver believe? Note that when forming the Receiver's beliefs it is as if the Receiver knows precisely which equilibrium is being played. Thus, what is the probability that the Sender is type t_1 if the Receiver observes R ? We can abbreviate "the probability that the Sender is type t_1 if the Receiver observes R " as $\Pr(t_1|R)$. In this particular potential equilibrium, only the t_1 type chooses R . Thus, if the Receiver observes a choice of R it should believe with 100% probability that it was type t_1 who chose R . Thus, we have $\Pr(t_1|R) = 1$. Now, what is $\Pr(t_2|R)$? Since type t_2 is NEVER choosing R in this potential separating equilibrium, the Receiver should believe that the probability a type t_2 chose R is equal to 0, or $\Pr(t_2|R) = 0$. We are not yet done with beliefs – we still need to specify the Receiver's beliefs about which node he is at when a choice of L is observed. What is $\Pr(t_1|L)$? If the Receiver observes L it knows with certainty that, in this potential equilibrium, it was type t_2 who chose L . Thus, $\Pr(t_1|L) = 0$ and $\Pr(t_2|L) = 1$. We are now done specifying the Receiver's beliefs.

The Receiver must now specify a strategy, so an action if he observes L and an action if he observes R . If L is observed the Receiver knows it is type t_2 , and also knows that if he chooses U he will receive 4 and if he chooses D he will receive 1. Since $4 > 1$ the Receiver chooses U . If R is observed the Receiver knows it is type t_1 , and also knows that if he chooses U he will receive 1 and if he chooses D he will receive 0. Since $1 > 0$ the Receiver chooses U . Thus, a *potential* separating PBE of the game is:

$$\begin{array}{l}
 t_1 \text{ choose } R \\
 t_2 \text{ choose } L \\
 \Pr(t_1|L) = 0 \\
 \Pr(t_2|L) = 1 \\
 \Pr(t_1|R) = 1 \\
 \Pr(t_2|R) = 0 \\
 \text{choose } U \text{ if } L \text{ observed} \\
 \text{choose } U \text{ if } R \text{ observed}
 \end{array}
 \begin{array}{l}
 \{\text{Sender's strategy}\} \\
 \{\text{Receiver's beliefs}\} \\
 \{\text{Receiver's strategy}\}
 \end{array}$$

Again, as of now this is a *potential* separating PBE of the game. We need to make sure that (1) the Receiver is playing a best response to the Sender's strategy and his (the Receiver's) beliefs and (2) the Sender is playing a best response to the Receiver's strategy. We have already done part (1) in constructing the Receiver's strategy. However, we still need to check part (2). Under the proposed equilibrium, type t_1 receives 2; if type t_1 were to switch to L he would receive 1 (because the Receiver is choosing U if L is observed) and so type t_1 does not wish to deviate. Under the proposed equilibrium type t_2 receives 2; if type t_2 were to switch to R he would receive 1 (because the Receiver is choosing U if R is observed) and so type t_2 does not wish to deviate. Thus, since no player wishes to deviate, the proposed PBE is a separating PBE of the game.

2.1.2 Type t_1 chooses L and type t_2 chooses R

Again, begin with the *potential* separating PBE of the game. Type t_1 chooses L and type t_2 chooses R . The Receiver's beliefs if L is observed are that $\Pr(t_1|L) = 1$ and $\Pr(t_2|L) = 0$ because in this equilibrium only type t_1 chooses L . The Receiver's beliefs if R is observed are that $\Pr(t_1|R) = 0$ and $\Pr(t_2|R) = 1$ because in this equilibrium only type t_2 chooses R . If L is observed the Receiver gets 3 if U is chosen and 0 if D is chosen. Thus, the Receiver would choose U if L is observed. If R is observed the Receiver gets 0

if U is chosen and 2 if D is chosen and so chooses D . A *potential* separating PBE is:

$$\begin{array}{ll}
 t_1 \text{ choose } L & \\
 t_2 \text{ choose } R & \{\text{Sender's strategy}\} \\
 \\
 \Pr(t_1|L) = 1 & \\
 \Pr(t_2|L) = 0 & \\
 \Pr(t_1|R) = 0 & \{\text{Receiver's beliefs}\} \\
 \Pr(t_2|R) = 1 & \\
 \\
 \text{choose } U \text{ if } L \text{ observed} & \\
 \text{choose } D \text{ if } R \text{ observed} & \{\text{Receiver's strategy}\}
 \end{array}$$

Again, we need to check to see if either type t_1 or type t_2 would deviate. In the proposed equilibrium type t_1 receives 1 (because the Receiver chooses U if L is chosen); if type t_1 were to switch to R he would receive 0 (because the Receiver chooses D if R is observed). Thus, type t_1 would not wish to deviate. For type t_2 , in the proposed equilibrium he receives 1; if type t_2 were to switch to L he would receive 2 (because the Receiver chooses U if L is chosen). Thus, type t_2 WOULD deviate from the proposed equilibrium, so the proposed equilibrium is NOT a separating PBE to this game. Thus, there is no separating equilibrium where type t_1 chooses L and type t_2 chooses R .

2.2 Potential Pooling Equilibria

We now shift our focus to pooling equilibria. With pooling equilibria all types choose the same action so that the uninformed party (the Receiver in our game) cannot condition his belief upon the action chosen. There are two potential pooling equilibria in our game – one where both types t_1 and t_2 choose L and another where both types t_1 and t_2 choose R .

2.2.1 Both types choose L

Suppose that both Sender types choose L . If this is the case then what is the probability that the sender is type t_1 if the Receiver observes L ? It is just the starting (or initial or prior) probability of type t_1 being drawn by nature, which in this example is 0.5. Thus, $\Pr(t_1|L) = 0.5$. Now, what is the probability that the Sender is type t_2 if the Receiver observes L ? Again, it is just the initial probability of 0.5. Thus, $\Pr(t_2|L) = 0.5$. That is the easy part – since it is a pooling equilibrium there is no updating to be done on the action upon which the Senders pool. However, we still need to specify $\Pr(t_1|R)$ and $\Pr(t_2|R)$. But there really is no good reason for any particular probability at this point, so we just let $\Pr(t_1|R) = q$ and $\Pr(t_2|R) = (1 - q)$ for now.

Now, what is the Receiver's best response if L is observed? If the Receiver chooses U then he gets 3 half of the time and 4 the other half of the time, so his expected value is $3 * \frac{1}{2} + 4 * \frac{1}{2} = \frac{7}{2}$ if he chooses U . If he chooses D he gets 0 half of the time and 1 the other half of the time, so his expected value is $0 * \frac{1}{2} + 1 * \frac{1}{2} = \frac{1}{2}$. So if the Receiver observes L the Receiver will choose U (note that technically U is a strictly dominant strategy for the Receiver if L is observed).

What is the Receiver's best response if R is observed? If the Receiver chooses U then he gets 1 with probability q and he gets 0 with probability $1 - q$, so his expected value is $1 * q + 0 * (1 - q) = q$. If the Receiver chooses D then he gets 0 with probability q and 2 with probability $1 - q$, so his expected value from choosing D is $0 * q + 2 * (1 - q) = 2 - 2q$. When will his payoff from choosing U be greater than his payoff from choosing D ? When $q \geq 2 - 2q$, or when $q \geq \frac{2}{3}$. Thus, if the Receiver believes (for whatever reason – remember, this is off the equilibrium path) that $q \geq \frac{2}{3}$ then the Receiver will choose U , while if $q < \frac{2}{3}$ the

Receiver will choose D . So, our proposed pooling PBE is:

$$\begin{array}{l}
 t_1 \text{ choose } L \\
 t_2 \text{ choose } L \\
 \Pr(t_1|L) = \frac{1}{2} \\
 \Pr(t_2|L) = \frac{1}{2} \\
 \Pr(t_1|R) \geq \frac{1}{3} \\
 \Pr(t_2|R) < \frac{1}{3} \\
 \text{choose } U \text{ if } L \text{ observed} \\
 \text{choose } U \text{ if } R \text{ observed}
 \end{array}
 \begin{array}{l}
 \{\text{Sender's strategy}\} \\
 \\
 \{\text{Receiver's beliefs}\} \\
 \\
 \{\text{Receiver's strategy}\}
 \end{array}$$

or alternatively:

$$\begin{array}{l}
 t_1 \text{ choose } L \\
 t_2 \text{ choose } L \\
 \Pr(t_1|L) = \frac{1}{2} \\
 \Pr(t_2|L) = \frac{1}{2} \\
 \Pr(t_1|R) \leq \frac{1}{3} \\
 \Pr(t_2|R) > \frac{1}{3} \\
 \text{choose } U \text{ if } L \text{ observed} \\
 \text{choose } D \text{ if } R \text{ observed}
 \end{array}
 \begin{array}{l}
 \{\text{Sender's strategy}\} \\
 \\
 \{\text{Receiver's beliefs}\} \\
 \\
 \{\text{Receiver's strategy}\}
 \end{array}$$

We can check to see if either or neither or both of these are pooling PBE (note the difference in the two potential equilibria is in the inequality sign for the Receiver's beliefs). Checking the first one, would type t_1 deviate from L to R ? In the proposed equilibrium, where he chooses L , he receives 1. If he deviates to R , he receives 2 (because he plays R and the Receiver is playing U). Thus, we can rule out the first proposed pooling PBE already.

What about the second proposed pooling PBE where both play senders play L ? In the proposed equilibrium Sender type t_1 chooses L and receives 1. If he switches to R , he receives 0 (because the Receiver is choosing D). So type t_1 does not wish to deviate. For Sender type t_2 , he receives 2 when he chooses L in the proposed equilibrium. If he switches to R , he receives 1 (because the Receiver is choosing D – actually, in this case he receives 1 if he chooses R regardless of what the Receiver chooses).

So this second proposed equilibrium is a pooling equilibrium. The key is that the Receiver must believe that $\Pr(t_2|R) > \frac{2}{3}$.

2.2.2 Both types choose R

Suppose that both Sender types choose R . Again, there is no chance for the Receiver to update his beliefs about which node he is at if he observes R , so we have $\Pr(t_1|R) = 0.5$ and $\Pr(t_2|R) = 0.5$. Also, we do not know what his beliefs are if he observes L (since he should never observe L in this equilibrium), so for now specify $\Pr(t_1|L) = p$ and $\Pr(t_2|L) = (1 - p)$. If R is chosen and the Receiver chooses U he gets 1 half of the time and 0 half of the time, so his expected value is $0 * \frac{1}{2} + 1 * \frac{1}{2} = \frac{1}{2}$. If R is chosen and the Receiver chooses D he gets 0 half of the time and 2 the other half, so his expected value is $0 * \frac{1}{2} + 2 * \frac{1}{2} = 1$. Thus, the Receiver would choose D if R is observed. If L is observed and the Receiver chooses U he gets 3 with probability p and he gets 4 with probability $(1 - p)$, so his expected value is $3p + 4 * (1 - p) = 4 - p$. If L is observed and the Receiver chooses D he gets 0 with probability p and 1 with probability $(1 - p)$, so his expected value is $0 * p + 1 * (1 - p) = 1 - p$. His expected payoff from choosing U will be greater than his expected payoff from choosing D if $4 - p > 1 - p$, or $4 > 1$. All this means is that if L is observed the Receiver will choose U – we should have already known this because earlier we noted that U was a strictly dominant strategy if the Receiver observed L . Thus, it does not matter what the Receiver's beliefs are if L

is observed – the Receiver will always choose U . So a *potential* pooling PBE is:

$$\begin{array}{l}
 t_1 \text{ choose } R \\
 t_2 \text{ choose } R \\
 \Pr(t_1|L) \geq 0 \\
 \Pr(t_2|L) \leq 1 \\
 \Pr(t_1|R) = \frac{1}{2} \\
 \Pr(t_2|R) = \frac{1}{2} \\
 \text{choose } U \text{ if } L \text{ observed} \\
 \text{choose } D \text{ if } R \text{ observed}
 \end{array}
 \begin{array}{l}
 \{\text{Sender's strategy}\} \\
 \{\text{Receiver's beliefs}\} \\
 \{\text{Receiver's strategy}\}
 \end{array}$$

Note that the Receiver's beliefs state that no matter what the relative probabilities are between $\Pr(t_1|L)$ and $\Pr(t_2|L)$ the Receiver will always choose U (again, because it is strictly dominant).

Finally, is this potential pooling PBE actually an equilibrium? It is easy to see that it is not – we know that type t_2 receives 1 if he chooses R and the Receiver chooses D . However, since the Receiver is choosing U if L is observed, then type t_2 could switch to L and receive 2, so type t_2 would deviate from the proposed strategy. There is no need to check if type t_1 would deviate because we know at least one player type will deviate so the proposed equilibrium cannot be an equilibrium.

2.3 Summing up Sender-Receiver

We found that there was one separating PBE:

$$\begin{array}{l}
 t_1 \text{ choose } R \\
 t_2 \text{ choose } L \\
 \Pr(t_1|L) = 0 \\
 \Pr(t_2|L) = 1 \\
 \Pr(t_1|R) = 1 \\
 \Pr(t_2|R) = 0 \\
 \text{choose } U \text{ if } L \text{ observed} \\
 \text{choose } U \text{ if } R \text{ observed}
 \end{array}
 \begin{array}{l}
 \{\text{Sender's strategy}\} \\
 \{\text{Receiver's beliefs}\} \\
 \{\text{Receiver's strategy}\}
 \end{array}$$

and we found that there were a class of pooling PBE:


$$\begin{array}{l}
 t_1 \text{ choose } L \\
 t_2 \text{ choose } L \\
 \Pr(t_1|L) = \frac{1}{2} \\
 \Pr(t_2|L) = \frac{1}{2} \\
 \Pr(t_1|R) \leq \frac{2}{3} \\
 \Pr(t_2|R) > \frac{2}{3} \\
 \text{choose } U \text{ if } L \text{ observed} \\
 \text{choose } D \text{ if } R \text{ observed}
 \end{array}
 \begin{array}{l}
 \{\text{Sender's strategy}\} \\
 \{\text{Receiver's beliefs}\} \\
 \{\text{Receiver's strategy}\}
 \end{array}$$

The reason I write "class" is because there are many different sets of beliefs that will lead to this pooling equilibrium. If $q = \frac{1}{8}$ and $(1 - q) = \frac{7}{8}$ then the above equilibrium is a pooling equilibrium. If $q = \frac{1}{9}$ and $(1 - q) = \frac{8}{9}$ then this is also a pooling equilibrium. So there are a lot of pooling equilibria, as long as $q \leq \frac{2}{3}$.

3 Signaling games

The perfect Bayesian equilibrium concept is useful for analyzing games of asymmetric information. One type of game we will consider is a signaling game. In a signaling game, one player has information that is unobservable to the other player and can take actions to signal to the other player what the information is. Consider a used-car market, where the sellers are able to observe the quality of the car but the buyers

may not be able to observe the quality. In this market, the seller receives either a good car or a bad car (the probability is determined by chance or nature) and then determines whether or not to make an offer to the buyer. The buyer then observes that a car has been offered for sale at some price p , but is unable to determine whether it is a good car or a bad car upon inspection because the seller incurs a cost c of making the bad car look like a good car. The buyer has a value of V if the car is a good car and W if the car is a bad car. We assume that $V > p > W > 0$ for this example. The extensive form of this game looks like:



There are two types of equilibria that we will discuss in games of this type. They are pooling equilibria and separating equilibria. In a pooling equilibrium, the action taken by the informed agent does not allow the uninformed agent to discern the object's type. In a separating equilibrium, the action taken by the informed agent does allow the uninformed agent to discern the object's type.

3.1 Pooling equilibrium

Consider the used-car market game. One method of finding an equilibrium in games of this type (when the actual payoff amounts are unspecified and left as variables) is to state what equilibrium you want to find and then determine the conditions needed for your proposed equilibrium to actually be an equilibrium. Suppose we want that the seller offers both good and bad cars to the buyer and that the buyer purchases the car. Now we can start at the back and work forward. The buyer must choose to buy or not buy. Since the buyer does not know whether or not he is being offered a good car or a bad car (this is a pooling equilibrium) the buyer must compare the expected value of buying a car conditional on seeing one for sale. We let $\Pr(\text{good}|\text{offer})$ be the probability that the buyer is getting a good car conditional on seeing an offer and $\Pr(\text{bad}|\text{offer})$ be the probability that the buyer is getting a bad car conditional on seeing an offer. The buyer's expected value of buying a car is then:

$$E[\text{Buy}] = \Pr(\text{good}|\text{offer}) * (V - p) + \Pr(\text{bad}|\text{offer}) * (W - p)$$

The buyer compares this with the expected value of not buying a car, which is 0. Thus, the buyer will purchase a car if the expected value of buying a car is positive (this is why we assume that $p > W$ – if $p \leq W$ then there is not much of a decision for the buyer), or alternatively, if $\Pr(\text{good}|\text{offer}) * (V - p) \geq \Pr(\text{bad}|\text{offer}) * (W - p)$. Again, we are proposing that the seller offers both types of cars for sale. This will only occur if the buyer chooses to buy and $\Pr(\text{good}) * p \geq 0$ and $\Pr(\text{bad}) * (p - c) \geq 0$. If the buyer chooses Not Buy, then the seller, having offered both cars, will have a payoff of $\Pr(\text{Bad}) * (-c) < 0$. We know that $\Pr(\text{good}) * p \geq 0$, but $\Pr(\text{bad}) * (p - c) \geq 0$ only if $p \geq c$. Thus, if the cost of making the bad car look like a good car is higher than the price, then the seller would rather not incur the cost because he is lowering his profit even if the buyer buys. So we now know that we will have an equilibrium where the seller offers all cars for sale and the buyer buys a car if the buyer's expected value of buying is greater than his expected value of not buying and the seller's cost of making the bad car look good is lower than the price of the bad car. We have the strategies down, so now all we need is the buyer's beliefs. The buyer must specify his beliefs about the node at which he is, and in this game the buyer receives no new information based on the seller's actions (it is a pooling equilibrium), so $\Pr(\text{good}|\text{offer}) = \Pr(\text{good})$ and $\Pr(\text{bad}|\text{offer}) = \Pr(\text{bad})$. Thus, a pooling equilibrium in this game is that the seller offers a car if it is good, the seller offers a car if it is bad, the buyer buys and has beliefs $\Pr(\text{good}|\text{offer}) = \Pr(\text{good})$ and $\Pr(\text{bad}|\text{offer}) = \Pr(\text{bad})$. Again, this is only an equilibrium if the buyer's expected value condition holds and $c \leq p$.

3.2 Separating equilibrium

A separating equilibrium holds if the actions of the informed agent allows the uninformed agent to discern the object's type. Thus, in a separating equilibrium we might want the seller to offer good cars for sale, not offer bad cars for sale, and the buyer to buy the good car with the belief that $\Pr(\text{good}|\text{offer}) = 1$. So we would have:

Seller's strategy: Offer if good, do not offer if bad
 Buyer's beliefs: $\Pr(\text{good}|\text{offer}) = 1$ and $\Pr(\text{bad}|\text{offer}) = 0$
 Buyer's strategy: Buy

Note that this potential equilibrium differs on two points from the previous one. Obviously, the seller is not offering bad cars for sale. A slightly more subtle change is in the buyer's beliefs, as the buyer now has a belief that the seller will only offer good cars for sale.

When will this be an equilibrium? For the buyer, choosing buy is optimal if $V - p \geq 0$, which it is by assumption. For the seller, given that the buyer is choosing buy he must be better off if he offers a good car for sale (he is, since $p > 0$) and he must be better off NOT offering a bad car for sale (which he will be if $c > p$). Again, the proposed set of strategies and beliefs is a strong perfect Bayesian equilibrium (there are no off the equilibrium path information sets, so if it is a weak perfect Bayesian equilibrium it must be a strong perfect Bayesian equilibrium).

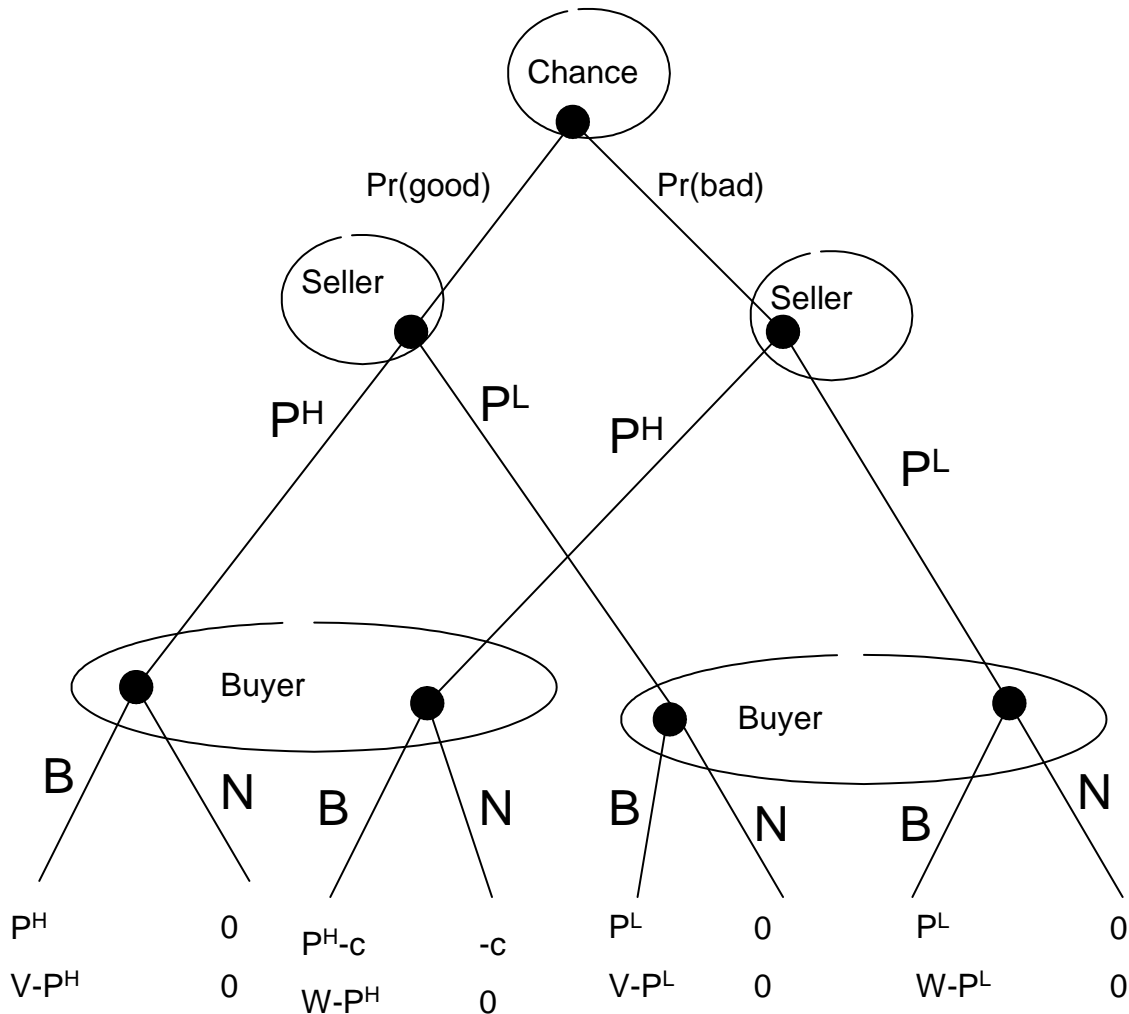
3.2.1 Can "anything" be an equilibrium if the right conditions hold?

Well, I suppose if certain conditions hold then anything can be an equilibrium. Given the structure we are currently using, it seems unlikely that there would be an equilibrium where the seller offered the bad car for sale but not the good car (actually, we would need to change our primary assumption, that $V > p > W > c$). However, consider the pooling equilibrium we found. If $E[\text{Buy}] < E[\text{Not Buy}]$, then the buyer would not purchase the car. Thus, there may not be a pooling equilibrium and only the good cars might be sold. This is contrary to Akerlof's lemons market,¹ but you must remember that sellers of good cars and sellers of bad cars had different values for keeping their cars in Akerlof's model. In the simple model I have drawn up, the value of a good car and the value of a bad car have the same value to the seller if he does not offer the car for sale. Also, in Akerlof's model there is no cost to making the bad car look good. If there is no cost then what we may see is that the buyers choose to Not Buy (if $\Pr(\text{good}|\text{offer}) * (V - p) + \Pr(\text{bad}|\text{offer}) * (W - p) < 0$) and the sellers may choose to offer all cars for sale at p .

¹Akerlof, G. (1970). The Market for Lemons: Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics*, 84:3, 488-500.

3.3 Choosing different prices in the used-car market

We can move one step further in this game. The seller can now offer two different prices in the used-car market, p^H and p^L , where $V > p^H > W > p^L > 0$. Now, instead of the buyer conditioning his belief on whether or not the seller is offering the car for sale the buyer conditions his belief on the price of the car. The game tree looks like the one below, where $B = Buy$ and $N = Not Buy$.



Can we find a separating perfect Bayesian equilibrium where the buyer always purchases and the seller offers the good car at P^H and the bad car at P^L ? The buyer will buy since $V - P^H > 0$ and $W - P^L > 0$. The seller prefers to charge P^H for the good car rather than P^L because $P^H > P^L$. The key then is the relationship between $P^H - c$ and P^L . If $P^H - c > P^L$ AND the buyer is choosing buy the seller would wish to charge P^H for the bad car. If $P^L > P^H - c$ then the seller will choose P^L for the bad car. Thus, if $P^L > P^H - c$ and $V > p^H > W > p^L > 0$ we have a perfect Bayesian equilibrium where:

Seller's strategy: Charge P^H if good, charge P^L if bad.

Buyer's beliefs: $\Pr(good|P^H) = 1$, $\Pr(bad|P^H) = 0$; $\Pr(good|P^L) = 0$, $\Pr(bad|P^L) = 1$

Buyer's strategy: Buy if observe P^H , buy if observe P^L

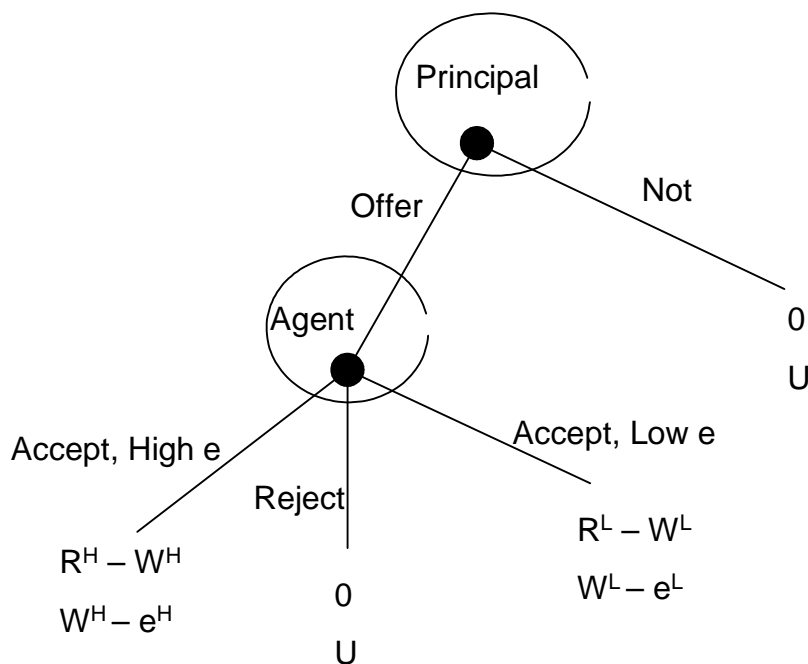
4 Principal-Agent problems

Another standard problem that we can analyze using this framework is that of the problem between the principal of a firm (its owner(s)) and the agent (manager). Typically, it is impossible or extremely costly for

principals to monitor the effort level of the agents. Thus, there is a fear that agents might provide less than their highest effort (this is why some contracts have incentive clauses in them – think about profit sharing between the agent and the firm, or sports contracts with performance goals – the idea is to tie the agent’s payment to the outcome with the hope that this will induce the agent to give high effort).

4.1 Observable Effort

We begin the discussion of the principal-agent problem by assuming that the agent’s effort is observable. The game begins with the principal deciding whether or not to make a contract offer to the agent. If the principal does not make a contract offer then the principal receives 0 and the agent receives some reservation utility U . If the principal makes a contract offer then it is the agent’s move. The contract specifies that the agent will receive W^H if he exerts high effort and W^L if he exerts low effort. The agent has 3 choices – he can choose to accept the contract and exert high effort, accept the contract and exert low effort, or reject the contract. If he rejects the contract then the principal receives 0 and the agent receives U . If he accepts the contract and exerts low effort then the agent receives $W^L - e^L$, where e^L is the effort cost of exerting low effort and the principal receives $R^L - W^L$, where R^L is the revenue the principal receives if the agent exerts low effort. If the agent accepts the contract and exerts high effort then he receives $W^H - e^H$, where e^H is the effort cost of exerting high effort and the principal receives $R^H - W^H$, where R^H is the revenue the principal receives if the agent exerts high effort. Again, note that the agent’s effort level is perfectly observable in this game. The game tree is:



What restrictions are necessary on the parameters to have a SPNE (there are no information sets that contain multiple nodes, so no need to specify beliefs for a WPBE or a SPBE) of the game be that the principal offers the contract and the agent accepts and exerts high effort? From the principal’s point of view, we need that $R^H - W^H > 0$ because if it is not then the principal could be better off choosing to not offer the contract. It would also be helpful if $R^H - W^H > R^L - W^L$ because then the principal would prefer the agent to exert high effort. If the principal prefers that the agent exert low effort then it is fairly easy to ensure this by simply setting $W^H = W^L$ so that the agent receives the same payment regardless of which effort level is chosen. From the agent’s point of view, two things need to happen. One is that $W^H - e^H > W^L - e^L$, so that the agent finds it more profitable to exert high effort rather than low effort. It also needs to be the case that $W^H - e^H > U$ so that the agent chooses to accept the contract rather than reject the contract. It is possible that $W^H - e^H > W^L - e^L$ but that exerting effort for this agent is too

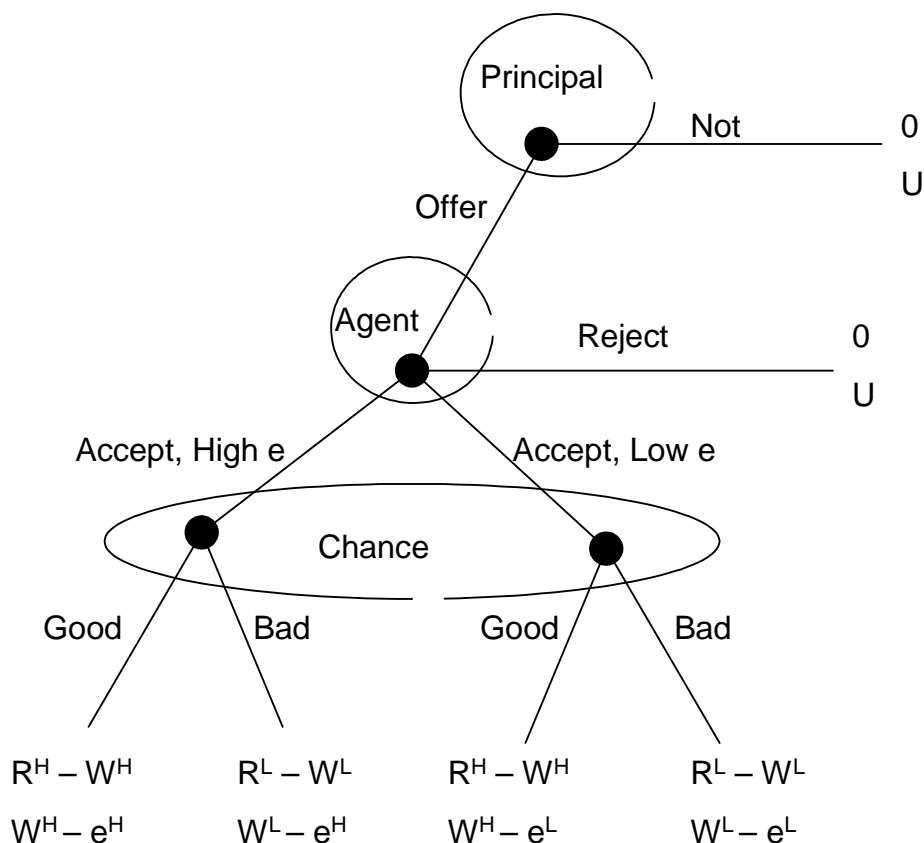
costly relative to his opportunity cost (U) so the agent would simply choose to reject the contract. These 2 conditions are relevant in many types of problems and are called the incentive compatibility constraint and the participation constraint.

Incentive compatibility constraint: The principal must structure the contract such that it gives the agent the incentive to act in the principal's best interest (in this example, choosing high effort over low effort).

Participation constraint: The principal must structure the contract such that participation by the agent is better than non-participation.

4.2 Unobservable effort

Suppose now that the principal cannot observe effort and can only observe outcome (either R^H or R^L). Moreover, there is a possibility that the principal receives R^L even if the agent exerts high effort e^H and a chance that the principal receives R^H even if the agent exerts low effort e^L . Since the principal can only observe outcome he bases the contract on the observed outcome – if he observes R^H then the agent receives W^H and if he observes R^L then the agent receives W^L . The game is as follows:



Again, we can work through the participation and incentive compatibility constraints to determine what the parameter restrictions need to be to ensure a particular equilibrium. Suppose we want the principal to offer a contract and the agent to accept the contract and put forth high effort. Assuming the agent is risk neutral, and that the agent's effort does not affect the probability of the good and bad states, the agent's incentive compatibility constraint is:

$$\Pr(\text{Good}) * (W^H - e^H) + \Pr(\text{Bad}) * (W^L - e^H) \geq \Pr(\text{Good}) * (W^H - e^L) + \Pr(\text{Bad}) * (W^L - e^L)$$

and the agent's participation constraint is:

$$\Pr(\text{Good}) * (W^H - e^H) + \Pr(\text{Bad}) * (W^L - e^H) \geq U.$$

Rewriting the IC constraint:

$$\begin{aligned}
\Pr(\text{Good}) * (W^H - e^H) + (1 - \Pr(\text{Good})) * (W^L - e^H) &\geq \Pr(\text{Good}) * (W^H - e^L) + (1 - \Pr(\text{Good})) * (W^L - e^L) \\
-e^H \Pr(\text{Good}) - e^H + e^H \Pr(\text{Good}) &\geq -e^L \Pr(\text{Good}) - e^L + e^L \Pr(\text{Good}) \\
-e^H &\geq -e^L \\
e^L &\geq e^H
\end{aligned}$$

Thus, because the agent's effort does not affect the state of the world the principal will be unable to offer the agent a contract that is incentive compatible.

Now supposing that the agent's effort affects the probability of the good and bad states, the agent's incentive compatibility constraint is:

$$\Pr(\text{Good}|e^H) * (W^H - e^H) + (1 - \Pr(\text{Good}|e^H)) * (W^L - e^H) \geq \Pr(\text{Good}|e^L) * (W^H - e^L) + (1 - \Pr(\text{Good}|e^L)) * (W^L - e^L)$$

where $\Pr(\text{Good}|e^H)$ is the probability of the good state when high effort is chosen and $\Pr(\text{Good}|e^L)$ is the probability of the good state when low effort is chosen. The agent's participation constraint for exerting high effort is now:

$$\Pr(\text{Good}|e^H) * (W^H - e^H) + \Pr(\text{Bad}|e^H) * (W^L - e^H) \geq U.$$

The principal also has a constraint that must be met:

$$\Pr(\text{Good}|e^H) * (R^H - W^H) + \Pr(\text{Bad}|e^H) * (R^L - W^L) \geq \Pr(\text{Good}|e^L) * (R^H - W^H) + \Pr(\text{Bad}|e^L) * (R^L - W^L).$$

Thus, if these conditions are satisfied then the principal will offer the agent a contract and the agent will accept the offer and exert high effort.

It is possible to make the game more realistic on many levels. One method of doing so would involve an agent who may be of high or low type where the type is unobservable to the principal. The principal might then offer two contracts (separating equilibrium) – one to get the high type to exert high effort and one to get the low type to exert low effort. It is also possible the principal would offer only one contract (pooling equilibrium).