Dynamic Games Of Incomplete Information

Introductory Comments:

To deal with static games of incomplete information, we introduced a new player called Nature, player types, and probabilities of player types. This allowed us to transform the game into a dynamic game of imperfect information. This device also works for dynamic games of incomplete information allowing us to transform them into dynamic games of imperfect information.

In static games of complete information, no new information is revealed to players during the course of the game, so players do not have to worry about contingency plans based on others' actions. In dynamic games of complete information, new information is revealed to at least one player by the actions of another, so contingency plans once again become important. For these games, the Nash equilibrium does not restrict the set the contingency plans a player can specify with its strategy, which tended to result in a multiplicity of equilibrium. In order to make more precise predictions, we used the subgame perfect equilibrium to rule out strategies that specify contingency plans that are not in a player's best interest to follow through with. Specifically, we required the equilibrium strategy to be Nash in all subgames, which effectively eliminated "incredible threats."

In static games of incomplete information, new information is revealed to players during the course of the game by Nature's choice of each player's type. We solved these games using the Bayesian Nash equilibrium, which is just a Nash equilibrium that recognizes a player's expected payoff is separable in types. What is important to realize about these static games of incomplete information is that our transformation actually produces a dynamic game of imperfect information. Therefore, given the multiplicity of equilibrium seen in several of our examples, it is natural to ask if we can get more precise predictions using subgame perfect reasoning. The answer is no because Nature doesn't care about a player's contingency plans. It just randomly assigns types, which is very different from when a player chooses a strategy to avoid an unfavorable outcome associated with the contingency plan specified by an opponent's strategy.

In static games of incomplete information, new information is revealed during the course of the game by Nature's choice of a player's type. New information is also revealed to at least one player by the choice of another player. So strategies for a dynamic game of incomplete information need to specify contingencies based on a player's type and may also have to specify contingencies based on the actions of another player. We can use the Bayesian Nash equilibrium to solve these games. More often than not there will be lots of strategy profiles that satisfy the Bayesian Nash definition. So we return to the question: can we get more precise predictions by using subgame perfect reasoning? The answer this time is yes, which is good news. The bad news is that while the appealing intuition of subgame perfection is apparent in dynamic games of incomplete information, the definition of a subgame is pretty much useless for solving these games. While we can define a static game of complete information, so that the Bayesian Nash and Nash equilibrium are essentially the same thing, we cannot specify dynamic games of incomplete information in order to use the subgame perfect equilibrium right out of the box. We need to do more.

Objective: Understand why the subgame perfect reasoning still makes sense, but subgame perfection cannot be used to solve dynamic games of incomplete information.

Figure 1 illustrates a two person dynamic game of incomplete information. There are two players: *A* and *B*. *A* has two types: 1 and 2 with probabilities *P* and 1-*P*. There is only one type for *B*. Nature starts the game by choosing *A*'s type. *A* then chooses whether to play *Left* or *Right*. If *A* chooses *Right*, *B* then gets to choose whether to play *Up* or *Down*, without seeing Nature's choice. The first payoff in parentheses is *A*'s, the second is *B*'s.

Question: What are the pure strategy Bayesian Nash equilibria for this game?

A has four strategies: S_A = {(*Left*, *Left*), (*Right*, *Left*), (*Left*, *Right*), (*Right*, *Right*)} where the first action in parentheses is chosen if it is assigned to type 1 by Nature and the second action in parentheses is chosen if it is assigned to type 2 by Nature. *B* has two strategies: S_B = {*Up*, *Down*}.

The normal form of the game when Nature chooses type 1 is:

		Player *B*	
		Up	Down
Player *A*: Type 1	(*Left*, *Left*)	50 20	50 20
	(*Right*, *Left*)	20 40	0 10
	(*Left*, *Right*)	50 20	50 20
	(*Right*, *Right*)	20 40	0 10

For type 2, the normal form is:

		Player *B*	
		Up	Down
Player *A*: Type 2	(*Left*, *Left*)	50 40	50 40
	(*Right*, *Left*)	50 40	50 40
	(*Left*, *Right*)	40 50	20 20
	(*Right*, *Right*)	40 50	20 20

Now to look for Bayesian Nash equilibria we must combine these two tables so it shows the expected payoff for each type of *A* and *B*. Before doing this it is important to recognize that *B* receives new information from *A*’s choice. Therefore, the probability *B* assigns to each type need not equal *P*, so lets call the probabilities at this point m and $1 - m$ instead.

Lets start with the strategy profile {(*Left*, *Left*), *Up*)}. The type 1 *A*’s expected payoff is simply 20, while type 2’s is 40 because both get to see Natures choice and by choosing *Left*, *B*’s strategy choice has no influence on the outcome of the game. Player *B*’s expected payoff will be $50m + (1 – m)50 = 50$.

Now lets consider the strategy profile {(*Right*, *Left*), *Up*)}. The type 1 *A*’s expected payoff is 40, while type 2’s is still 40. Player *B*’s expected payoff is $20m + (1 – m)50$. For strategy profile {(*Left*, *Right*), *Up*)}. The type 1 *A*’s expected payoff is 20, while type 2’s is 50. Player B’s expected payoff is $50m + (1 – m)40$.

Doing this for each possible strategy combination yields the expected payoffs summarized in Table 1.

Table 1

		Player B	
		Up	*Down*
Player A (Type 1, Type 2)	(*Left*, *Left*)	$50m + (1 – m)50$* (20,40)	$50m + (1 – m)50$* (20*,40*)
	(*Right*, *Left*)	$20m + (1 – m)50$* (40*,40)	$0m + (1 – m)50$ (10,40*)
	(*Left*, *Right*)	$50m + (1 – m)40$* (20,50*)	$50m + (1 – m)20$ (20*,20)
	(*Right*, *Right*)	$20m + (1 – m)40$* (40*,50*)	$0m + (1 – m)20$ (10,20)

Now lets put an * by each player and player type’s best response and look for the pure strategy Bayesian Nash equilibria. As you can see, there are two pure strategy Bayesian Nash equilibrium strategy profiles: {(*Left*, *Left*), *Down*} and {(*Right*, *Right*), *Up*}. These equilibrium profiles do not depend on m for this game, but in others they could.

Question: Is one of these strategy profiles more compelling than the other?

Consider the equilibrium {(*Left*, *Left*), *Down*}. Suppose *A* accidentally chooses (*Left*, *Right*) instead of its equilibrium strategy. Will *B* really choose *Down*? If it does, it earns 50 if *A* is type 1 and 20 if *A* is type 2. If it chooses *Up* instead, it earns 50 if *A* is a type 1 and 40 if *A* is a type 2. Therefore, it would be at least as well off and possibly better off by choosing *Up* instead of *Down*. Similar arguments can be made for (*Right*, *Left*) and (*Right*, *Right*). Therefore, it is hard to believe *B* would ever choose *Down* given the opportunity. If this is the case, {(*Left*, *Left*), *Down*} doesn’t really make much sense.

But this is exactly the type of reasoning we used to justify the subgame perfect equilibrium, so why don’t we just use it here?

Question: What are the subgames in this game?

Answer: There is only one subgame, which is the whole game. Since *B* does not know *A*'s type, we cannot be certain that it knows where it is at in the game tree. Since the only subgame is the whole game, the subgame perfect equilibria are identical to the Bayesian Nash equilibria, which are Nash equilibria.

What we really need here is a new equilibrium concept that captures the spirit of the subgame perfect equilibrium, but does not require multiple subgames to be useful.

Objective: Understand what a perfect Bayesian equilibrium is and how it generalizes the subgame perfect equilibrium to dynamic games of incomplete information.

A definitive statement for the perfect Bayesian equilibrium is rather elusive. What we need is something that captures the spirit of subgame perfection for dynamic games of imperfect information where a player's payoff is type separable. For general dynamic games of imperfect information, "trembling hand perfection" was developed by Selten and the sequential equilibrium was developed by Kreps and Wilson. However, for many classes of games we are interested in, these concepts are more than what we really need and often difficult to apply. Therefore, the perfect Bayesian equilibrium has emerged as a more user-friendly approach to solving different classes of games with imperfect information in the spirit of subgame perfection.

The specific conditions required by the perfect Bayesian equilibrium tend to depend on the class of games being analyzed. In all cases, the equilibrium is defined over a strategy and belief profile. Given this strategy and belief profile, two requirements that must be met are: (i) players' strategies must be sequentially rational given their beliefs and (ii) players' beliefs must be updated according to Bayes rule given the strategy profile whenever applicable. The first condition is analogous with requiring that a strategy be Nash for each and every subgame, but does not require subgames. The second condition is something new, but essentially says that players should be good statisticians when they form their beliefs about uncertain events.

Lets be more specific by considering Fudenberg and Tirole's definition for a perfect Bayesian equilibrium in multi-stage games with observed actions and incomplete information.

The game can be defined by

1. Players: A set N with elements denoted by $i = 1, 2, \ldots, n$ and Nature.
2. Player types: $q_i \in \Theta_i$ where $q = (q_1, q_2, \ldots, q_n)$ and $\Theta = \times_{i \in N} \Theta_i$.
3. Probabilities: $p(q) = \Pi_{i \in N}\, p_i(q_i)$ where $p_i(q_i)$ is the probability Nature chooses type q_i for player i. Note that this implies player types are independent.
4. Periods: $t = 0, 1, \ldots, T$ in which players simultaneously choose an action.
5. Actions: $a_i^t(h^t) \in A_i(h^t)$ such that $a^t = (a_1^t, a_2^t, \ldots, a_n^t)$ where $h^t = (a^0, a^1, \ldots, a^{t-1}) \in H^t$ is the history of play before period t.

6. Strategies: A mixed strategy for a player, call it s_i, can be defined as $s_i(a_i|h^t, q_i) \geq 0$ for all $a_i \in A_i(h^t)$ such that $\sum_{a_i \in A_i(h_t)} s_i(a_i \mid h^t, q_i)$, which is a probability distribution over a player's actions given the history of play and player's type. A mixed strategy profile is then just a collection of mixed strategies, one for each player: $s = (s_1, s_2, \ldots, s_n)$.
7. Beliefs: $m_{ij} = m_i(q_j| h^t, q_i)$ is player i's belief about the probability player j is type q_j given the history of play and player i's type. $m_i = m_i(q_{\sim i}|h^t, q_i)$ is player i's belief about the probability that all other players' types are $q_{\sim i} = (q_1, \ldots, q_{i-1}, q_{i+1}, \ldots, q_n)$ given the history of play and player i's type. A belief profile is then just a collection of beliefs, one for each player: $m = (m_1, m_2, \ldots, m_n)$.
8. Payoffs: $u_i(h^{T+1}, q)$ is player i's realized payoff given players' types and the history of the game. For any strategy and belief profile, we can calculate a player's expected payoff based on which period we are in, the player's type, and every other players' type: $u_i(s, m| h^t, q_i) = E(u_i(h^{T+1}, q)| s, m, h^t, q_i)$.

We are now ready to be a little more specific about the restrictions imposed by the perfect Bayesian equilibrium.

Definition: A perfect Bayesian equilibrium is a strategy and belief profile (s, m) such that

(i) $u_i(s, m| h^t, q_i) \geq u_i(s_i', s_{\sim i}, m| h^t, q_i)$ for all s_i', $h^t \in H^t$, $t \in (0,1,\ldots,T)$, $q_i \in \Theta_i$ and $i \in N$, and

(ii) $m_i(q_j \mid (h_t, a^t), q_i) = \dfrac{m_i(q_j \mid h_t, q_i)\, s_j(a_j^t \mid h_t, q_j)}{\sum_{\tilde{q}_j \in \Theta_j} m_i(\tilde{q}_j \mid h_t, q_i)\, s_j(a_j^t \mid h_t, \tilde{q}_j)}$ if $\sum_{\tilde{q}_j \in \Theta_j} m_i(\tilde{q}_j \mid h_t, q_i)\, s_j(a_j^t \mid h_t, \tilde{q}_j)$ > 0 for all $i \neq j$, $i \in N$, $j \in N$, $a_j^t(h^t) \in A_j(h^t)$, $h^t \in H^t$ and $t \in (0,1,\ldots,T)$.

Condition (i) is what we mean by sequential rationality. Condition (ii) is what we mean by saying beliefs should be updated according to Bayes rule whenever applicable.

The description of the game above and definition capture the four requirements outlined by Gibbons. But Fudenberg and Tirole go further in stating that the belief profiles should also satisfy:

(iii) $m_i(q_{\sim i} \mid h_t, q_i) = \prod_{j \in N, j \neq i} m_i(q_j \mid h_t, q_i)$ for all $i \in N$, $h^t \in H^t$ and $t \in (0,1,\ldots,T)$,

(iv) $m_i(q_j \mid (h_t, a^t), q_i) = m_i(q_j \mid (h_t, \hat{a}^t), q_i)$ if $a_j^t = \hat{a}_j^t$ for all $i \neq j$, $i \in N$, $j \in N$, $a_j^t \in A_j(h^t)$, $h^t \in H^t$ and $t \in (0,1,\ldots,T)$.

(v) $m_i(q_k \mid h_t, q_i) = m_j(q_k \mid h_t, q_j) = m(q_k \mid h_t)$ for all $i \neq j \neq k$, $i \in N$, $j \in N$, $k \in N$, , $h^t \in H^t$ and $t \in (0,1,\ldots,T)$.

Condition (iii) says that for any history of the game a player's beliefs about opponents' types are independent. That is, if player types are independent to start with, beliefs about those types remain independent regardless of the history of play.

Condition (iv) says that a player's beliefs about an opponent should not be influenced by the actions of some other opponent. If one opponent is just as ignorant about the type of a third as you are, that opponent's actions should not be able to convey any information about the third player's type.

Condition (v) implies that when types are independent two players should hold the same belief regarding the type of a third: two players cannot interpret the meaning of a third's actions differently.

In the games I have dealt with, I have never had to look beyond conditions (i) and (ii) to find the perfect Bayesian equilibrium. However, I do see how conditions (iii)-(v) could be useful when specifying beliefs where Bayes rule is not applicable.

Objective: Understand how to use the perfect Bayesian equilibrium to solve a game.

Let us start by considering the simple game in Figure 2, while focusing on the intuition of what conditions (i) and (ii) do for us.

The game is a dynamic game of incomplete information with two players denoted by *A* and *B*. *A* can choose *L* to end the game immediately or *R* to allow *B* to choose how the game ends. If *A* chooses *R*, *B* can end the game by choosing *U* or *D*. The catch here is that *B* does not know whether it is better off choosing *U* to end the game or *D* because *A* has some private information. To capture this notion of private information, we assume *A* is one of two types denoted by 1 and 2. We will also assume that the probability *A* is type 1 is $P = 1/3$, so the probability *A* is type 2 is $1 – P = 2/3$.

The Perfect Bayesian equilibrium requires us to look for a strategy and belief profile that (i) is sequentially rational and (ii) satisfies Bayes rule where applicable.

As with subgame perfection, the sequential rationality condition suggests working backwards. Let m be *B*'s belief that *A* is type 1 given *A* chooses *R*. If *B* chooses *U*, its expected payoff is $u_B(U|R) = m50 + (1\text{-}m)\,0 = m50$. Choosing *D* yields $u_B(D|R) = m0 + (1\text{-}m)\,50 = (1\text{-}m)50$. Therefore, we can describe *B*'s best response to *A* choosing *R* as

$$s_B(U|R) = \begin{cases} 1, & m > 0.5 \\ [0,1], & m = 0.5 \\ 0, & m < 0.5 \end{cases}$$

where $s_B(U|R)$ is the probability *B* chooses *U* given *A* chooses *R* and belief m.

Now we can ask what is *A*'s best response given the probability *B* chooses *U*: $s_B(U|R)$. If *A* is type 1, its expected payoff from choosing *R* is $u_A(R|1) = s_B(U|R)50 + (1\text{-}s_B(U|R))0$. It earns 40 from choosing *L*: $u_A(L|1) = 40$. Therefore, a type 1 *A*'s best response can be summarized as

$$s_A(R|1) = \begin{cases} 1, & s_B(U \mid R) > 0.8 \\ [0,1], & s_B(U \mid R) = 0.8 \\ 0, & s_B(U \mid R) < 0.8 \end{cases}$$

where $s_A(R|1)$ is the probability type 1 *A* chooses *R* given the probability *B* chooses *U*.

If *A* is type 2, its expected payoff from choosing *R* is $u_A(R|2) = s_B(U|R)50 + (1-s_B(U|R))0$. It earns 30 from choosing *L*: $u_A(L|2) = 30$. Therefore, a type 2 *A*'s best response can be summarized as

$$s_A(R|2) = \begin{cases} 1, & s_B(U \mid R) > 0.6 \\ [0,1], & s_B(U \mid R) = 0.6 \\ 0, & s_B(U \mid R) < 0.6 \end{cases}$$

where $s_A(R|2)$ is the probability type 2 *A* chooses *R* given the probability *B* chooses *U*.

Figure 3 illustrates the extent to which sequential rationality rules out various strategy and belief profiles. The solid line traversing its way from point (0, 0) to (1, 1) inclusive shows all potential equilibrium combinations of strategies and beliefs for B. For combinations on this line from [(0, 0), (0.5, 0.6)), A's only potential equilibrium strategy is $s_A(R|1) = s_A(R|2) = 0$. For point (0.5,0.6), A's potential equilibrium strategies are $s_A(R|1) = 0$ and $s_A(R|2) = [0,1]$. From ((0.5,0.6), (0.5,0.8)), A's potential equilibrium strategy is $s_A(R|1) = 0$ and $s_A(R|2) = 1$. For (0.5,0.8), A's potential equilibrium strategies are $s_A(R|1) = [0,1]$ and $s_A(R|2) = 1$. From ((0.5,0.8), (1,1)], A's potential equilibrium strategy is $s_A(R|1) = s_A(R|2) = 1$. So, we have certainly ruled out a lot of different possibilities, but there is still an infinite number remaining.

Now let us see how using Bayes rule can help us. Bayes rule implies

$$m = \begin{cases} [0,1], & \textit{for } s_A(R \mid 1) = s_A(R \mid 2) = 0 \\ \dfrac{s_A(R \mid 1)P}{s_A(R \mid 1)P + s_A(R \mid 2)(1-P)}, & \textit{otherwise} \end{cases}.$$

Now let us retrace our previous steps while asking whether or not each piece of the path is also consistent with Bayes rule.

From ((0.5,0.8), (1,1)], *A*'s potential equilibrium strategy is $s_A(R|1) = s_A(R|2) = 1$. If we plug this in above, we get $m = P$, which we have assumed is 1/3. But this is a contradiction because to be on this piece of the path $m \geq 0.5$. Therefore, Bayes rule allows us to rule out

$s_A(R|1) = s_A(R|2) = 1$, $s_B(U|R) = 1$, and $m > 0.5$

and

$s_A(R|1) = s_A(R|2) = 1$, $s_B(U|R) = (0.8, 1]$, and $m = 0.5$.

For (0.5,0.8), A's potential equilibrium strategy is $s_A(R|1) = [0,1]$ and $s_A(R|2) = 1$. If we plug this in above, we get $m = \frac{s_A(R|1)P}{s_A(R|1)P + (1-P)}$, which must equal 0.5 for $s_B(U|R) = 0.8$ to be a sequentially best response. This implies $s_A(R|1) = (1-P)/P$, which is (2/3)/(1/3) = 2 by assumption. Once again, we see a contradiction, so Bayes rule allows us to rule out

$s_A(R|1) = [0,1]$, $s_A(R|2) = 1$, $s_B(U|R) = 0.8$, and $m = 0.5$.

From ((0.5,0.6), (0.5,0.8)), A's potential equilibrium strategy is $s_A(R|1) = 0$ and $s_A(R|2) = 1$. If we plug this in above, we get the contradiction $m = 0$, so Bayes rule allows us to rule out

$s_A(R|1) = 0$, $s_A(R|2) = 1$, $s_B(U|R) = (0.6, 0.8)$, and $m = 0.5$.

For (0.5,0.6), A's potential equilibrium strategy is $s_A(R|1) = 0$ and $s_A(R|2) = [0,1]$. Plugging in above also yields the contradiction $m = 0$ (as long as $s_A(R|2) > 0$), so out goes

$s_A(R|1) = 0$, $s_A(R|2) = (0,1]$, $s_B(U|R) = 0.6$, and $m = 0.5$.

Finally, from [(0,0), (0.5,0.6)), A's potential equilibrium strategy is $s_A(R|1) = s_A(R|2) = 0$. If we plug that in above, we get $m = [0, 1]$ because Bayes rule doesn't apply. Alas, a strategy profile we cannot rule out. In terms of beliefs, Bayes rule rules nothing out, but sequential rationality does. What remains is

$s_A(R|1) = s_A(R|2) = 0$, $s_B(U|R) = 0$, and $m < 0.5$

and

$s_A(R|1) = s_A(R|2) = 0$, $s_B(U|R) = [0, 0.6]$, and $m = 0.5$.

Strictly speaking, we do not have a unique perfect Bayesian equilibrium. But from a practical standpoint, we do have a unique behavioral prediction that we can run with.

Objective: Understand the difference between a pooling and separating equilibrium.

Now let us consider the slightly more complicated Figure 4. Again we have two players A and B. A has private information, so B does not know its payoffs. Nature starts the game by choosing A's type: 1 with probability P and 2 with probability $1 - P$. Once Nature reveals A's type, A can choose either R or L. B then gets to choose either U or D after witnessing A's choice, but not Nature's choice.

The incentives of the game are such that *B* would like to know *A*'s type. *A* would also like *B* to know its type if it chooses *R*. If it chooses *L*, a type 2 would like *B* to know its type, but a type 1 *A* would not.

To solve this game for the perfect Bayesian equilibria, we should work backward starting with the question: what is *B*'s best response given its beliefs? To keep things easy, let us focus only on pure strategies.

Let m_R be *B*'s belief that *A* is type 1 given *R*, while m_L is *B*'s belief that *A* is type 1 given *L*. *B*'s pure strategy best response is:

$$s_B(U|R) = \begin{cases} 1, & m_R \geq 0.5 \\ 0, & m_R \leq 0.5 \end{cases}$$

and

$$s_B(U|L) = \begin{cases} 1, & m_L \geq 0.5 \\ 0, & m_L \leq 0.5 \end{cases}.$$

Now lets consider *A*'s pure strategy best response given $s_B(U|R)$ and $s_B(U|L)$. Type 1's is:

$$s_A(R|1) = \begin{cases} 1, & s_B(U \mid R) \geq \frac{3}{5} - \frac{2}{5} s_B(U \mid L) \\ 0, & s_B(U \mid R) \leq \frac{3}{5} - \frac{2}{5} s_B(U \mid L) \end{cases}.$$

Type 2's is:

$$s_A(R|2) = \begin{cases} 1, & s_B(U \mid R) \leq \frac{5}{2} s_B(U \mid L) - 1 \\ 0, & s_B(U \mid R) \geq \frac{5}{2} s_B(U \mid L) - 1 \end{cases}.$$

Figure 5 pulls all this information together where $s_B(U|R)$ is plotted on the vertical axis and $s_B(U|L)$ is plotted on the horizontal axis. Graphing $s_B(U \mid R) = \frac{3}{5} - \frac{2}{5} s_B(U \mid L)$ for type 1 and $s_B(U \mid R) = \frac{5}{2} s_B(U \mid L) - 1$ for type 2 allows us to show the different potential pure equilibrium strategies for *A* given $s_B(U|R)$ and $s_B(U|L)$. The potential pure equilibrium strategies for *B* can also be shown on the figure. Looking only at player's best response functions suggests four potential pure strategy equilibria that are sequentially rational:

1. $\{s_A(R|1)=0,\ s_A(R|2)=0,\ s_B(U|R)=0,\ s_B(U|L)=0,\ m_R < 0.5,\ m_L < 0.5\}$,

2. $\{s_A(R|1)=1,\ s_A(R|2)=0,\ s_B(U|R)=1,\ s_B(U|L)=0,\ m_R > 0.5,\ m_L < 0.5\}$,
3. $\{s_A(R|1)=1,\ s_A(R|2)=1,\ s_B(U|R)=1,\ s_B(U|L)=1,\ m_R > 0.5,\ m_L > 0.5\}$, and
4. $\{s_A(R|1)=0,\ s_A(R|2)=1,\ s_B(U|R)=0,\ s_B(U|L)=1,\ m_R < 0.5,\ m_L > 0.5\}$.

The question now is which of these if any can be eliminated using Bayes rule.

For candidate 1, Bayes rule is not applicable for m_R, so $m_R = [0,1]$. For m_L, Bayes rule implies $m_L = P$. Therefore, if $P < 0.5$, we cannot rule out this equilibrium. We can however rule it out if $P > 0.5$. Furthermore, we know that $m_R < 0.5$, for beliefs to be consistent with *B*'s best response. So,

$\{s_A(R|1)=0,\ s_A(R|2)=0,\ s_B(U|R)=0,\ s_B(U|L)=0,\ m_R < 0.5,\ m_L = P < 0.5\}$

is a set of perfect Bayesian equilibrium strategy and belief profiles.

For candidate 2, Bayes rule implies $m_R = 1$ and $m_L = 0$, which is consistent with *B*'s best response. Therefore,

$\{s_A(R|1)=1,\ s_A(R|2)=0,\ s_B(U|R)=1,\ s_B(U|L)=0,\ m_R = 1,\ m_L = 0\}$

is a perfect Bayesian equilibrium strategy and belief profile.

For candidate 3, Bayes rule implies $m_R = P$, but does not restrict $m_L = [0,1]$. For these beliefs to be consistent with *B*'s best response function $P > 0.5$ and $m_L > 0.5$. So,

$\{s_A(R|1)=1,\ s_A(R|2)=0,\ s_B(U|R)=1,\ s_B(U|L)=1,\ m_R = P > 0.5,\ m_L > 0.5\}$

is a set of perfect Bayesian equilibrium strategy and belief profile.

For candidate 4, Bayes rule implies $m_R = 0$ and $m_L = 1$, which is consistent with *B*'s best response. Therefore,

$\{s_A(R|1)=0,\ s_A(R|2)=1,\ s_B(U|R)=0,\ s_B(U|L)=1,\ m_R = 0,\ m_L = 1\}$

is a perfect Bayesian equilibrium strategy and belief profile.

Notice that given *P*, candidates 1 and 3 cannot both yield a perfect Bayesian equilibrium.

The other thing to note is that with candidates 1 and 3 no new information is revealed to *B* when *A* makes its choice because both types always choose the same actions. This is what is called a pooling equilibrium.

Alternatively, in the equilibria described by 2 and 4, player *A*'s actions fully reveal its type to player *B* which is referred to as a separating equilibrium. In a separating equilibrium, one player is able to learn everything it needs to know about the private information of another through that player's choice of actions. In a pooling equilibrium, *B* learns nothing about player *A* from its choices. There is also what is called a partially pooling/separating equilibrium where some, but not all, information about an opponent is revealed through the play of the game. For the game in Figure 4, this might occur if we considered mixed strategy equilibria in addition to pure strategy equilibria.

Objective: Understand the reasoning behind alternative equilibrium refinements.

In many of the dynamic games of incomplete information that we have looked at, there have been lots of potential perfect Bayesian equilibria. Again, we face the question: is one more compelling than the others?

The Bayesian Nash equilibrium really does nothing to restrict players' beliefs. The perfect Bayesian equilibrium requires players to be sequentially rational and restricts their beliefs using Bayes rule where applicable. Where not applicable, the perfect Bayesian equilibrium places few restrictions on players' beliefs (beliefs must support equilibrium best responses), which is one reason we can get lots of different equilibria. So if we are willing to place a few more restrictions on players' beliefs when Bayes rule is not applicable, we may be able to get rid of some of these equilibria. What restrictions are reasonable is the topic of a variety of papers on refinements to the Bayesian Nash equilibrium.

The basic question much of this literature asks is what types of players are most likely to benefit from deviating from an equilibrium strategy. If one type of player benefits from deviating in all circumstances that another type does and then some, we might expect it to be more likely to deviate. So, when we update our beliefs, we should take this into account.

To get a basic understanding of this line of reasoning lets work through some examples starting with Figure 6. In this game, there are two types for player *A*, which are chosen with equal probability. Player *A* can choose *L* to end the game immediately or *R* to allow *B* to end the game by playing *U*, *M*, or *D*. Player *B* must make this choice knowing only that player *A* chose *R*. The incentives of the game are such that both of *A*'s types want *B* to do the same thing when they choose *R*. *B* however wants to do different things depending on *A*'s type. It is also to *A*'s advantage for *B* to know its type if it is type 1, but not if it is type 2.

Now let us find the perfect Bayesian equilibria for some examples and ask which of these equilibria are supported by beliefs that do not seem reasonable.

In general, *B*'s best response function can be written as

$$s_B = \{ s_B(U),\ s_B(M),\ s_B(D)\} = \begin{cases} \{1,0,0\}, & \text{\textit{for} } m < 0.25 \\ \{s \ = [0,1], 1 - s \ ,0\}, & \text{\textit{for} } m = 0.25 \\ \{0,1,0\}, & \text{\textit{for} } 0.75 > m > 0.25 \\ \{0, s \ = [0,1], 1 - s \ \}, & \text{\textit{for} } m = 0.75 \\ \{1,0,0\}, & \text{\textit{for} } m > 0.75 \end{cases}$$

where $s_B(U)$, $s_B(M)$, and $s_B(D)$ are the probabilities *B* chooses *U*, *M*, and *D* and *m* is the probability *A* is type 1 given *R* is chosen. We can graph this best response on the simplex in Figure 7.

In general, *A*'s Best Response Function is

$$s_A^t = \{ s_A(R|t),\ s_A(L|t)\} = \begin{cases} \{1,0\}, & \text{for } p^t < 50s_B(M) + 100s_B(D) \\ \{f = [0,1], 1-f\}, & \text{for } p^t = 50s_B(M) + 100s_B(D) \\ \{0,1\}, & \text{for } p^t > 50s_B(M) + 100s_B(D) \end{cases}$$

where $s_A(R|t)$ and $s_A(L|t)$ are the probability *A* chooses *R* and *L* given its type is *t* and p^t is *A*'s payoff given its type is *t*.

Example 1: $p^1 = 60$ *and* $p^2 = 105$

Plugging these values into *A*'s best response function above allows us to derive an equation for the circumstances under which each type is indifferent:

$p^1 = 60 = 50s_B(M) + 100s_B(D) \Rightarrow s_B(M) = 1.2 - 2s_B(D)$ and
$p^2 = 105 = 50s_B(M) + 100s_B(D) \Rightarrow s_B(M) = 2.1 - 2s_B(D)$.

We can graph these equations in Figure 7 and identify *A*'s best response in each region of the simplex. Since $s_B(M) = 2.1 - 2s_B(D)$ falls completely outside the simplex, there are only two different regions identified by the indifference curve $s_B(M) = 1.2 - 2s_B(D)$. To the left of this line, type 1 *A* prefers *L*. To the right, type 1 *A* prefers *R*. Type 2 *A* will always prefer *L*.

So, what are the perfect Bayesian equilibria for this game?

If $s_A^1 = \{1,0\}$ and $s_A^2 = \{0,1\}$, Bayes rule implies $m = 1$ and $s_B = \{0,0,1\}$, which is consistent with $s_A^1 = \{1,0\}$ and $s_A^2 = \{0,1\}$.

If $s_A^1 = \{\phi = [0,1], 1 - \phi\}$ and $s_A^2 = \{0,1\}$, Bayes rule implies $m = 1$ and $s_B = \{0,0,1\}$, which is consistent with $s_A^1 = \{\phi = 1, 1 - \phi = 0\}$ and $s_A^2 = \{0,1\}$ only if $\phi = 1$. Nothing new here.

If $s_A^1 = \{0,1\}$ and $s_A^2 = \{0,1\}$, Bayes rule implies $m = [0,1]$, which only works if $m = 0.75$ and $s_B = \{0,\ s = [0.8,1],\ 1 - s\}$; $0.75 > m > 0.25$ and $s_B = \{0, 1, 0\}$; $m = 0.25$ and $s_B = \{s = [0,1],\ 1 - s,\ 0\}$; or $m < 0.25$ and $s_B = \{1, 0, 0\}$.

Therefore, the perfect Bayesian equilibria for the game are:
(i) $[s_A^1 = \{1,0\},\ s_A^2 = \{0,1\},\ s_B = \{0,0,1\},\ m = 1]$
(ii) $[s_A^1 = \{0,1\},\ s_A^2 = \{0,1\},\ s_B = \{0,\ s = [0.8,1],\ 1 - s\},\ m = 0.75]$
(iii) $[s_A^1 = \{0,1\},\ s_A^2 = \{0,1\},\ s_B = \{0, 1, 0\},\ 0.75 > m > 0.25]$
(iv) $[s_A^1 = \{0,1\},\ s_A^2 = \{0,1\},\ s_B = \{s = [0,1],\ 1 - s,\ 0\},\ m = 0.25]$
(v) $[s_A^1 = \{0,1\},\ s_A^2 = \{0,1\},\ s_B = \{s = 1, 0, 0\},\ m < 0.25]$

Question: Can we do better than this?

For (i), we established $m = 1$ using Bayes rule, so there is nothing really to argue with here, but for (ii) – (v), we just asserted $m \leq 0.75$ because it was consistent with everyone making a best response and it is off the equilibrium path so Bayes rule offers no guidance. Let us think a little harder about $m \leq 0.75$.

Question: Does $m \leq 0.75$ really make sense?

Regardless of *B*'s response, type 2 *A*'s payoff is always higher from choosing its equilibrium strategy *L*. So, there is little reason suspect a type 2 *A* would ever deviate from choosing *L*. Alternatively, a type 1 *A* could be better off choosing *R* instead of the (ii) - (v) equilibrium action *L* if *B* chooses *U*. Therefore, there is some reason for a type 1 *A* to deviate from the (ii) - (v) equilibrium strategy. If a type 1 *A* may deviate, but a type 2 *A* will never deviate, then Bayes rule would imply $m = 1 > 0.75$, which contradicts $m \leq 0.75$. Thus, we might suspect that $m \leq 0.75$ is an unreasonable belief for *B*, which rules out (ii) - (v) and leads us to the unique equilibrium prediction (i).

Note that this line of reasoning is the same as in Gibbons for signaling requirement 5. The message *R* for type 2 *A* is dominated by *L* because the minimum possible payoff from choosing message *L*, 105, is greater than the maximum possible payoff from choosing *R*, 100. This is not the case for type 1 *A*. Therefore, it seems reasonable to believe that any *R* message that is observed must come from a type 2 *A*, which contradicts $m \leq 0.75$.

Figure 8 demonstrates how we can modify this argument slightly to see the rationale for Gibbon's signaling requirement 6, which is more commonly referred to as the "Intuitive Criterion." The game in Figure 8 modifies the game in Figure 6 by giving *B* the chance to respond to *L* with either *N* or *S* if *A* is type 2. Sequential rationality implies *B* should respond to *L* by choosing *N*, which suggests the game in 8 is essentially identical to the game in 6. However, Gibbon's signaling requirement 5 is no longer enough to rule out equilibria with $m \leq 0.75$ because the minimum payoff from choosing message *L* is now 30 and less than the maximum payoff from choosing message *R*, 100. But the strategy choice that leads to this minimum payoff can never be part of a sequentially rational equilibrium for *B*. The "Intuitive Criterion" modifies Gibbon's signaling requirement 5 by saying that we only need to compare the equilibrium payoff from a message to the maximum possible equilibrium payoffs from some alternative message. In our example, the equilibrium payoff for type 2 *A* from choosing *L* is 105 not 30. Since 105 is greater than the maximum possible payoff from switching to message *R*, it seems reasonable to assume a type 2 *A* will never choose *R*. Again, this is not the case for a type 1 *A*, so it is unreasonable to suspect $m < 1$, which gets us back to our unique perfect Bayesian equilibrium.

Even the "Intuitive Criterion" can fail us however.

Example 2: $p^1 = 25$ *and* $p^2 = 75$

Plugging these values into *A*'s best response function above allows us to derive an equation for the circumstances under which each type is indifferent:

$p^1 = 25 = 50s_B(M) + 100s_B(D) \Rightarrow s_B(M) = 0.5 - 2s_B(D)$
$p^2 = 75 = 50s_B(M) + 100s_B(D) \Rightarrow s_B(M) = 1.5 - 2s_B(D)$

We can graph these equations in Figure 9 and identify *A*'s best response in each region of the simplex. To the left of $s_B(M) = 0.5 - 2s_B(D)$, type 1 *A* prefers *L*. To the right of $s_B(M) = 0.5 - 2s_B(D)$, type 1 *A* prefers *R*. To the left of $s_B(M) = 1.5 - 2s_B(D)$, type 2 *A* prefers *L*. To the right of $s_B(M) = 1.5 - 2s_B(D)$, type 2 *A* prefers *R*.

So, what are the perfect Bayesian equilibria for this game?

They are
(i) $[s_A^1 = \{1, 0\},\ s_A^2 = \{1/3, 2/3\},\ s_B = \{0,0.5,0.5\},\ m = 0.75]$
(ii) $[s_A^1 = \{0, 1\},\ s_A^2 = \{0, 1\},\ s_B = \{s = [0,0.5],\ 1 - s,\ 0\},\ m = 0.25]$
(iii) $[s_A^1 = \{0, 1\},\ s_A^2 = \{0, 1\},\ s_B = \{1, 0, 0\},\ m < 0.25]$

But, is $m \leq 0.25$ a reasonable belief? If we look at the simplex in Figure 9, any strategy for *B* that makes a type 2 *A* better off deviating from *L* in (ii) or (iii) also makes a type 1 *A* better off deviating from *L* in (ii) or (iii) (Note that the set of strategies for one player that make another player willing to deviate from the equilibrium strategy is referred to as the deviation set). But, there are some strategies for *B* that make a type 1 *A* better off deviating from *L* in (ii) or (iii), but not a type 2 *A*. Therefore, it would seem to make sense that a type 1 *A* would be more likely to deviate from *L* in (ii) or (iii) than a type 2 *A*. If this is the case, $m > 0.5$ which contradicts $m \leq 0.25$. So, why not dispense with this equilibrium?

Example 3: $p^1 = 60$ *and* $p^2 = 75$

Figure 10 reproduces Figure 9 for the payoffs in example 3. The primary difference between the two examples is that the pooling equilibrium strategies $s_A^1 = \{0, 1\}$ and $s_A^2 = \{0, 1\}$ that were supported by the off the equilibrium path belief $m \leq 0.25$ in example 2 must now be supported by the off the equilibrium path belief $m \leq 0.75$. So what is the big deal? The big deal is that it is no longer enough to just assume $m > P = 0.5$ in order to get rid of all these pooling equilibria. Instead, we must assume $m > 0.75 > P = 0.5$. Intuitively, in example 2, all we had to assume to get a unique equilibrium was that Type 1 *A*s are more likely to deviate from a pooling equilibrium than Type 2 As. For example 3, we must say how much more likely in order to get a unique equilibrium.

It gets even worse. The examples illustrated in Figures 11 and 12 show how simply looking at who is more likely to deviate from a pooling equilibrium may not help. In both these cases, which type of *A* is more likely to deviate depends on how *B* is likely to respond. If we only look at *B*'s best responses, the example in Figure 11 can then be resolved just as before. But in Figure 12, even limiting attention to *B*'s best responses is not enough.

Now, with all this said, the most common refinement that is used and the one that makes the most sense is the "Intuitive criterion" (Gibbons signaling requirement 6 and Cho and Kreps, 1987). The intuitive criterion is essentially what we used to eliminate the pooling equilibrium for Example 1, although we did not need its full power because we did not need to restrict deviation sets to only include best responses. The "Intuitive criterion" also does reasonably well in experimental test. The same cannot be said for refinements beyond the "Intuitive criterion."

Application: Environmental Conflict Game

Now let's look at a little richer game with continuous types. Here is the story. A developer (denoted by D) has bought a piece of land to build a mini mall. The value of development to the developer is V_D. A homeowner (denoted by H) lives across the street from the piece of land and would prefer to see the land remain open space. So, it would benefit by V_H if the developer is not able to follow through with its plans. The homeowner knows the developer stands to gain V_D from building a mini mall. The developer however does not know that the homeowner is against its development. It does however know that if a homeowner is against the development, the probability distribution for the value of the homeowner to stopping the development is $F(v)$ with density $f(v)$ where $v \in [v', v'']$. This distribution is common knowledge.

The developer really cannot do much to defend its right to develop until the homeowner indicates that it is against development. We will assume that the homeowner can do this in one of two ways. First, it can choose to come out fighting by making a costly investment in effort, x_H, that will influence the probability that it is successful stopping the development. The developer can then respond to this effort with its own effort, x_D. The probability that the homeowner is successful stopping the developer is $P(x_H, x_D) = x_H / (x_H + x_D)$, while $1 - P(x_H, x_D) = x_D / (x_H + x_D)$ is the probability that the developer is successful defending its right. Alternatively, the homeowner can come out whining, costlessly complaining about the developer's plans, in which case the developer has the opportunity to quiet the homeowner by investing effort x_D. The homeowner then has the opportunity to "put its money where its mouth is" by investing effort x_H. The probability the homeowner or developer is successful are the same as before. We will assume all players are risk neutral.

Essentially the game boils down to a Stackelberg contest with incomplete information. Who gets to lead in the Stackelberg contest depends on whether the homeowner chooses to fight or whine. If it chooses to fight, it gets to lead. If it chooses to whine, the developer gets to lead.

Question: What are the perfect Bayesian equilibria for this game?

To solve this game, we need to ask what are each player's sequentially rational strategies given their beliefs. We then need to check to make sure beliefs satisfy Bayes rule where applicable.

We can start by solving for the developer's sequentially rational strategy given the homeowner chooses to fight.

The firm's expected payoff given the homeowner comes out fighting is

$$U_D(F) = \int_{v'}^{v''}\left(\frac{x_D}{x_H + x_D}V_D - x_D\right)f_F(v)dv$$

where $f_F(v)$ and $F_F(v)$ are the density and distribution function of the developer's beliefs about the homeowners value (type) given it comes out fighting.

The first order condition for the developer is

$$\int_{v'}^{v''}\left(\frac{x_H}{(x_H + x_D)^2}V_D - 1\right)f_F(v)dv = 0,$$

which can be solved for the firm's best response:

$$x_D(x_H) = \sqrt{x_H V_D} - x_H$$

if we make things nice so we always get an interior solution.

Now you might be asking yourself: What about the integral? How come it doesn't appear in the developer's best response function?

It doesn't appear in the developer's best response function because the developer gets to see the homeowner's effort before it makes its choice. Since the developer's expected payoff depends only on the homeowner's effort and not the homeowner's value, this effort is all it needs to know to resolve its incomplete information:

$$\int_{v'}^{v''}\left(\frac{x_H}{(x_H + x_D)^2}V_D - 1\right)f_F(v)dv = \left(\frac{x_H}{(x_H + x_D)^2}V_D - 1\right)\int_{v'}^{v''} f_F(v)dv = \left(\frac{x_H}{(x_H + x_D)^2}V_D - 1\right).$$

Now, given the developer's best response, how hard should the homeowner come out fighting. The homeowner's expected payoff is

$$U_H(F) = \frac{x_H}{x_H + x_D(x_H)}V_H - x_H = \frac{x_H}{\sqrt{x_H V_D}}V_H - x_H .$$

Taking the first order condition and solving for x_H yields: $x_H(F) = V_H r_H^2/4$ where $r_i = \sqrt{\frac{V_i}{V_j}}$ for $i \neq j$. The developer's best response is then $x_D(F) = V_D r_H^2\left(\frac{1}{2} - \frac{r_H^2}{4}\right)$. We will also want to know the homeowner's payoff for these strategies, which is $U_H(F) = V_H r_H^2/4$.

Now that we have solved the game in the event that the homeowner chooses to fight, we can turn to the solving the game when the homeowner chooses to whine.

In this part of the game, the homeowner has the last move with an expected payoff equal to:

$$U_H(W) = \frac{x_H}{x_H + x_D} V_H - x_H \,,$$

which is now familiar enough for us to know its best response is:

$$x_H\left(x_D, V_H\right) = \sqrt{x_D V_H} - x_D \,.$$

Given the homeowner’s best response, the developer’s expected payoff when the homeowner whines is

$$U_D(W) = \int_{v'}^{v''} \left(\frac{x_D}{x_H\left(x_D, v\right) + x_D} V_D - x_D \right) f_W\left(v\right) dv = \int_{v'}^{v''} \left(\sqrt{\frac{x_D}{v}} V_D - x_D \right) f_W\left(v\right) dv$$

where $f_W(v)$ and $F_W(v)$ are the density and distribution function of the developer’s beliefs about the homeowners value (type) given it whines.

The first order condition is

$$\int_{v'}^{v''} \left(\frac{1}{2} \sqrt{\frac{1}{x_D v}} V_D - 1 \right) f_W\left(v\right) dv = 0 \,.$$

Solving this expression for x_D yields

$$x_D(W) = \frac{V_D}{4} \left(\int_{v'}^{v''} r_D \, f_W\left(v\right) dv \right)^2 = V_D \mathrm{E}_W(r_D)^2/4.$$

The homeowner’s sequentially rational response is then,

$$x_H\left(W\right) = V_H \, r_D E_W\left(r_D\right) \left(\frac{1}{2} - \frac{r_D E\left(r_D\right)}{4} \right).$$

Finally, to figure out how the homeowner should start the game, we need to know its payoff, which after a little bit of algebra can be written as $U_H(W) = V_H \left(1 - \frac{r_D E_W\left(r_D\right)}{2} \right)^2$.

The question we now have to ask is whether the homeowner should fight or whine. It should whine if $U_H(W) > U_H(F)$, which implies

$$V_H\left(1-\frac{r_D E_W(r_D)}{2}\right)^2 > \frac{V_H r_H^{\,2}}{4}$$

or

$$2r_H - E_W(r_D) - r_H^{\,2} > 0 .$$

Figures 13 and 14 show the implications of this equation. If beliefs are such that $E_W(r_D) < 1$, then the homeowner should whine if $1 + (1 - E_W(r_D))^{0.5} > r_H > 1 - (1 - E_W(r_D))^{0.5}$. Otherwise, the homeowner should fight. If beliefs are such that $E_W(r_D) > 1$, the homeowner should always fight.

What about beliefs? For there to be an equilibrium where some types of homeowners choose to whine, Bayes rule implies

$$f_W(v) = \frac{f(v)}{F\left(V_D\left(1+(1-E_W(r_D))^{0.5}\right)^2\right) - F\left(V_D\left(1-(1-E_W(r_D))^{0.5}\right)^2\right)},$$

so we need to see if we can find a $v_W' = V_D(1 - (1 - E_W(r_D))^{0.5})^2$ and $v_W'' = V_D(1 + (1 - E_W(r_D))^{0.5})^2$ such that $E_W(r_D) < 1$. There may be many. Ultimately, it will depend on how you specify the developer's initial beliefs.

Figure 15 illustrates how incomplete information changes the contest model with endogenous timing in the event that some homeowners choose to whine.

In the complete information contest with endogenous timing, if the homeowner's value exceeds the developer's ($r_H > 1$), it should whine and let the developer go first. If its value is less than the developer's ($r_H < 1$), it should choose to fight and take the lead.

In the incomplete information contest, this result changes in two ways. First, if the homeowner's value is much greater than the developers ($r_H > 1 + (1 - E_W(r_D))^{0.5}$), it should choose to fight and take the lead as oppose to whining and letting the developer take the lead. Second, if the homeowners value is less than the developers ($r_H < 1$), but not too much less ($r_H > 1 - (1 - E_W(r_D))^{0.5}$), it should choose to whine and let the developer go first instead of fighting.

Why do we get these results?

If the homeowner's value is really large, the developer tends to underestimate the homeowner's value and will tend to invest more effort than it should if it knew the homeowner's value. The only way the homeowner can clear up the developer's misperception is to strike first with its own effort.

If the homeowner's value is small, but not too small, the developer tends to overestimate this value and will tend to invest less effort than it should if it knew the homeowner's value. Now if the homeowner strikes first with its own effort, it lets the developer know that it should fight harder, something the homeowner would certainly like to avoid. So the homeowner keeps its type disguised by choosing to whine.

There are two fundamental effects in the model; a timing effect and an information effect. For some homeowner values the timing and information effect work together. For other values, they work against each other.

Above, we classified perfect Bayesian equilibria as either pooling, such that no new information is revealed during the play of the game, and separating, such that all information is revealed through the play of the game. For this game, we see an example of the third possibility which we referred to as partially separating/pooling. If the homeowner chooses to whine, it reveals some information, but does not completely resolve the developer's incomplete information.

Figure 1

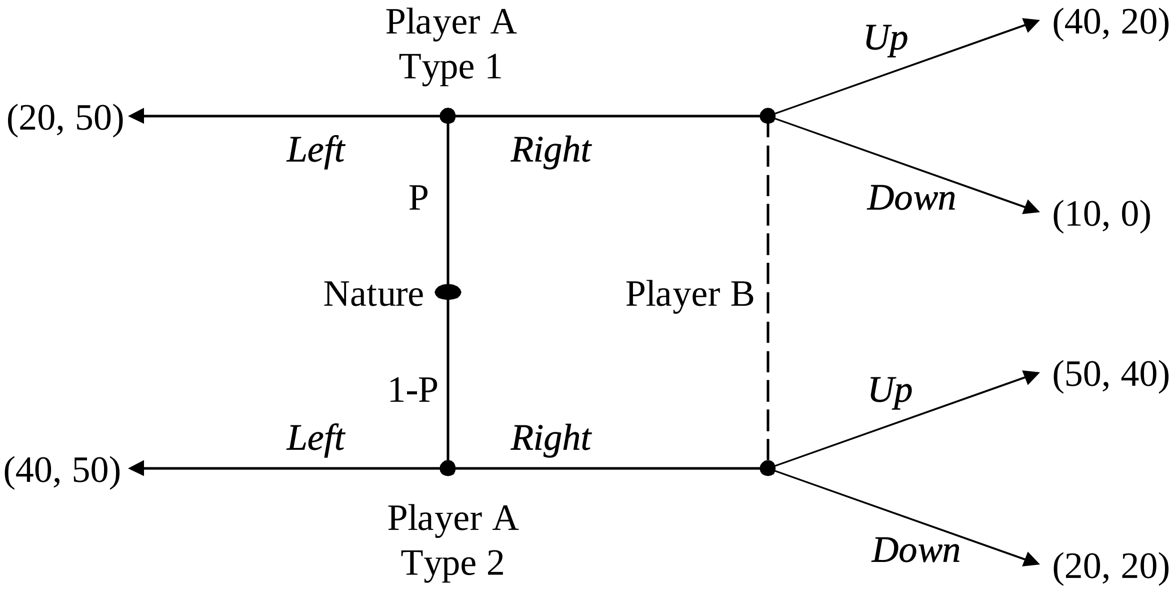
Player A
Type 1
(20, 50)
Left
Right
P
Nature
1-P
Left
Right
(40, 50)
Player A
Type 2
Player B
Up
(40, 20)
Down
(10, 0)
Up
(50, 40)
Down
(20, 20)

Figure 2

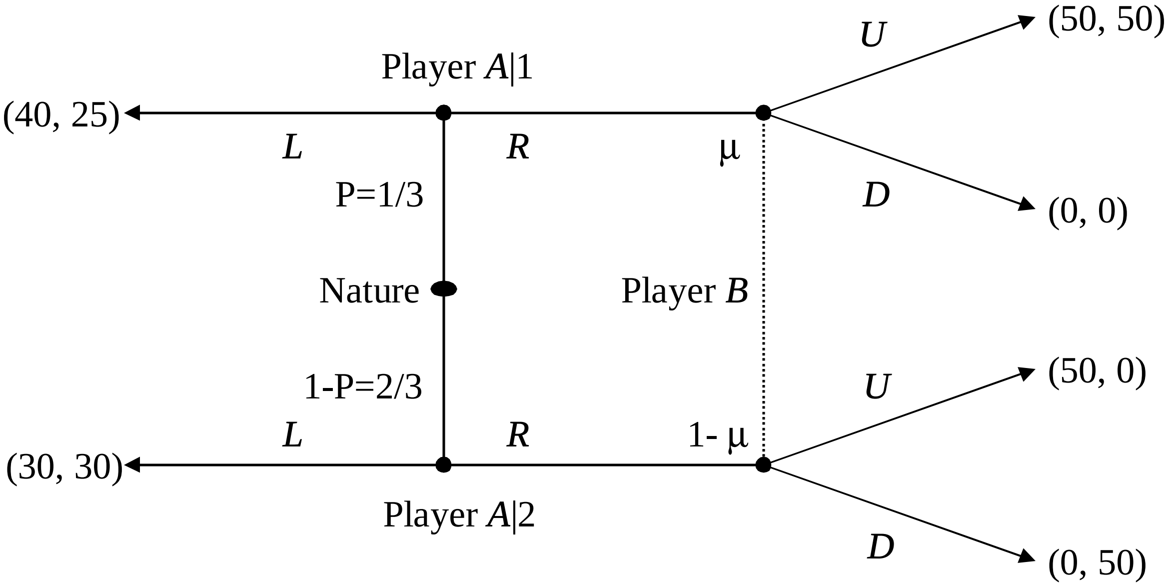

(50, 50)
U
Player A|1
(40, 25)
L
R
μ
D
(0, 0)
P=1/3
Nature
Player B
(50, 0)
U
1-P=2/3
L
R
1- μ
(30, 30)
Player A|2
D
(0, 50)

Figure 3

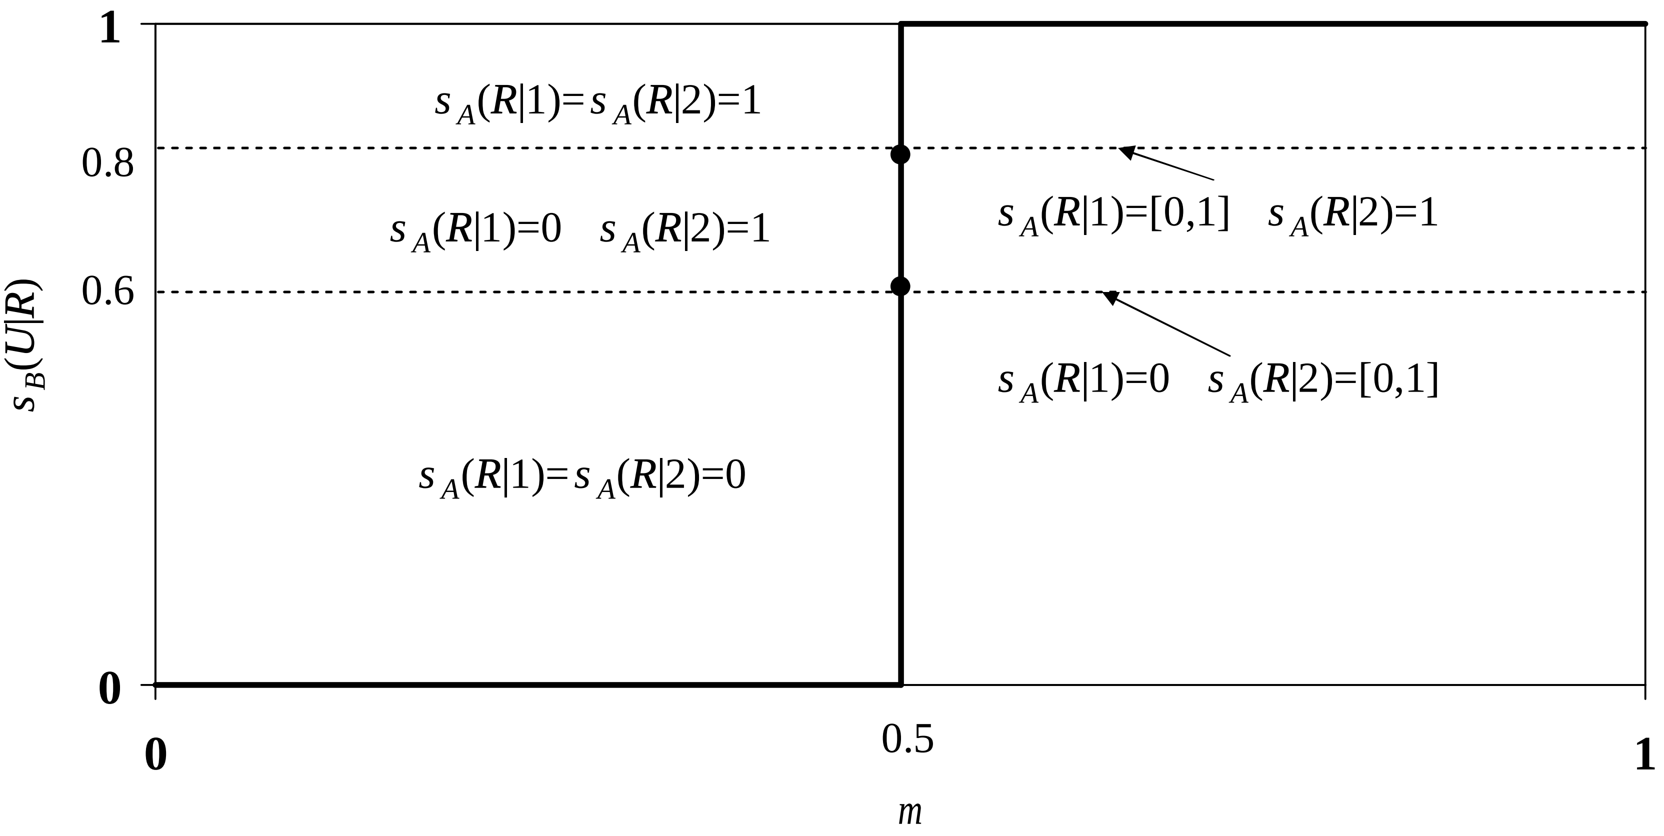
1
0.8
0.6
0
$s_B(U|R)$
$s_A(R|1)=s_A(R|2)=1$
$s_A(R|1)=0 \quad s_A(R|2)=1$
$s_A(R|1)=[0,1] \quad s_A(R|2)=1$
$s_A(R|1)=0 \quad s_A(R|2)=[0,1]$
$s_A(R|1)=s_A(R|2)=0$
0
0.5
1
m

Figure 4

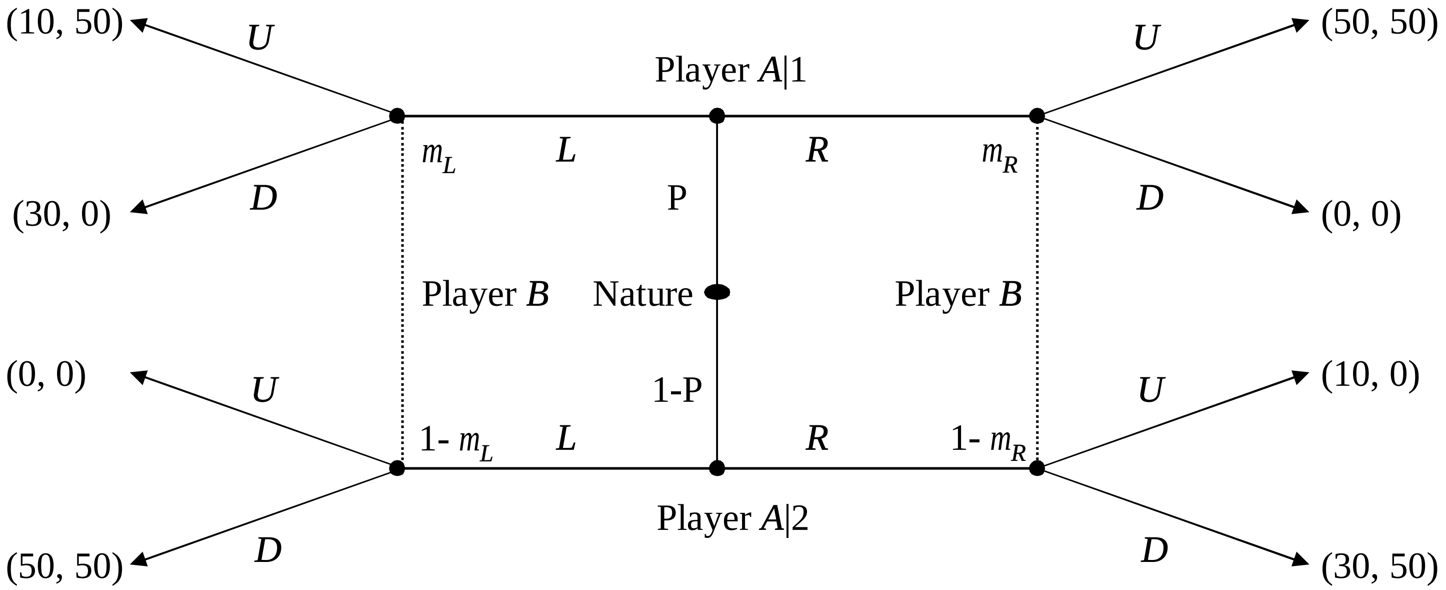

(10, 50)
U
Player A|1
U
(50, 50)
m_L
L
R
m_R
(30, 0)
D
P
D
(0, 0)
Player B
Nature
Player B
(0, 0)
U
1-P
U
(10, 0)
1- m_L
L
R
1- m_R
Player A|2
(50, 50)
D
D
(30, 50)

Figure 5

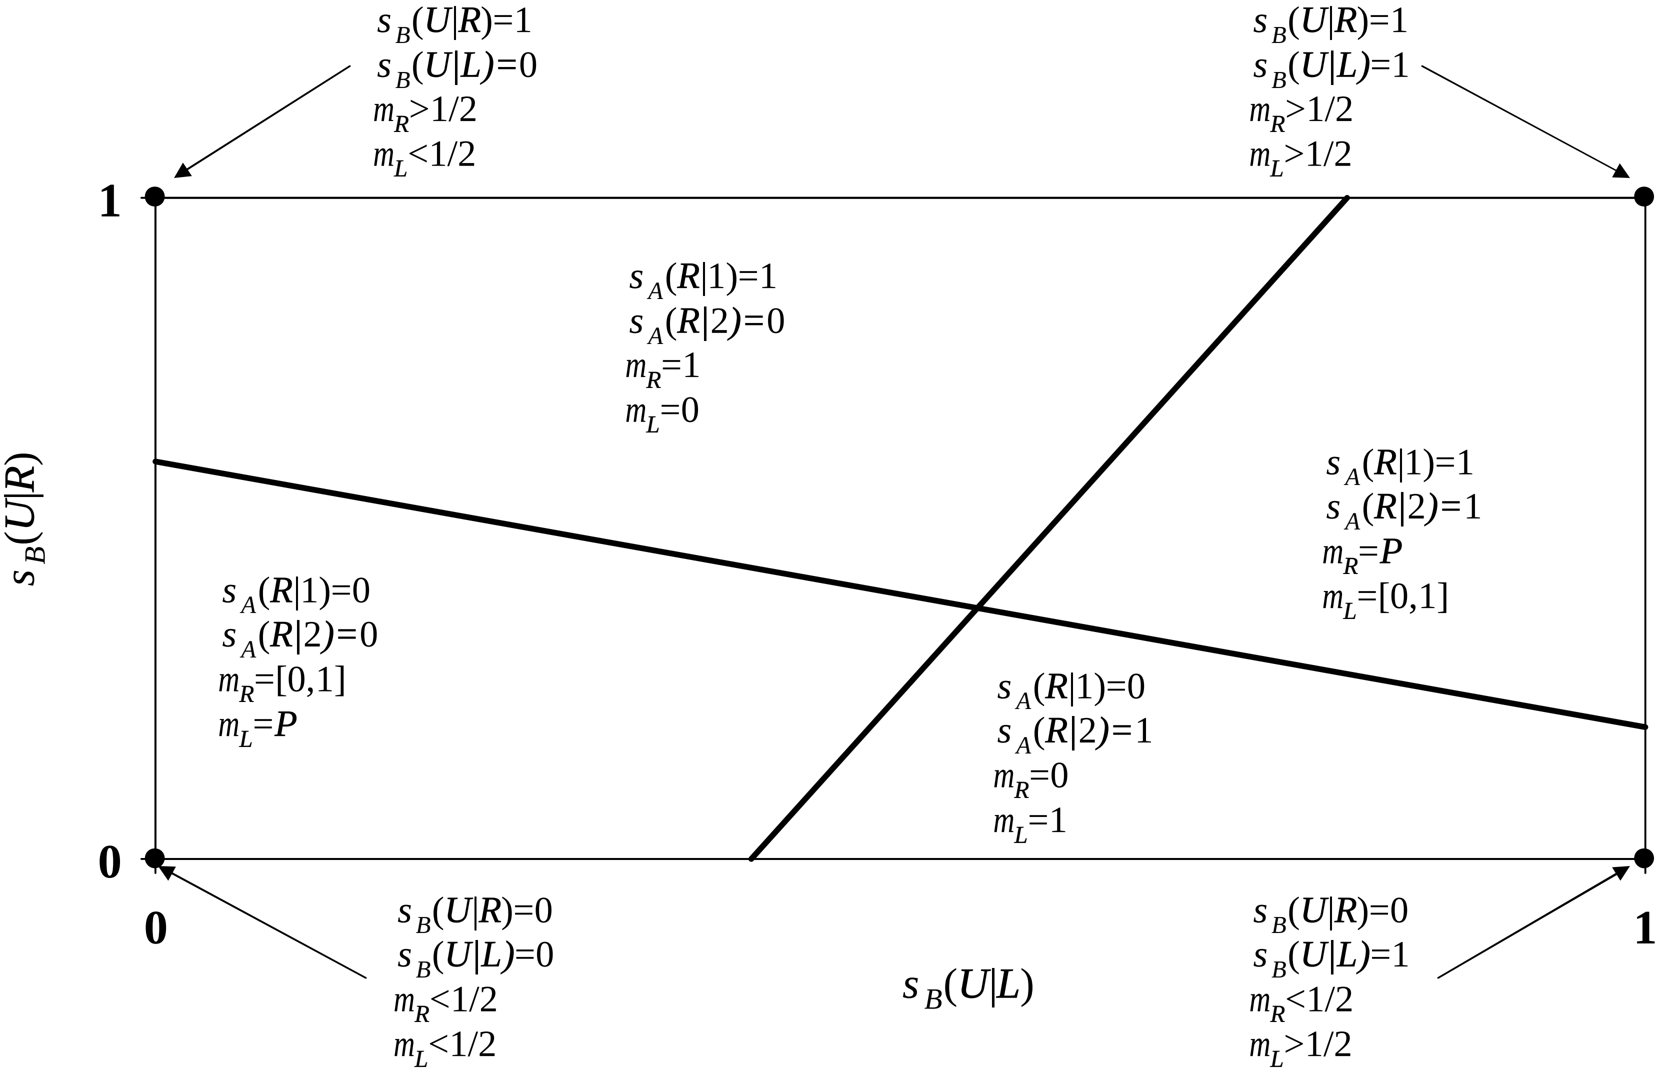

$s_B(U|R)=1$
$s_B(U|L)=0$
$m_R>1/2$
$m_L<1/2$
$s_B(U|R)=1$
$s_B(U|L)=1$
$m_R>1/2$
$m_L>1/2$
1
$s_A(R|1)=1$
$s_A(R|2)=0$
$m_R=1$
$m_L=0$
$s_B(U|R)$
$s_A(R|1)=1$
$s_A(R|2)=1$
$m_R=P$
$m_L=[0,1]$
$s_A(R|1)=0$
$s_A(R|2)=0$
$m_R=[0,1]$
$m_L=P$
$s_A(R|1)=0$
$s_A(R|2)=1$
$m_R=0$
$m_L=1$
0
0
1
$s_B(U|R)=0$
$s_B(U|L)=0$
$m_R<1/2$
$m_L<1/2$
$s_B(U|L)$
$s_B(U|R)=0$
$s_B(U|L)=1$
$m_R<1/2$
$m_L>1/2$

Figure 6

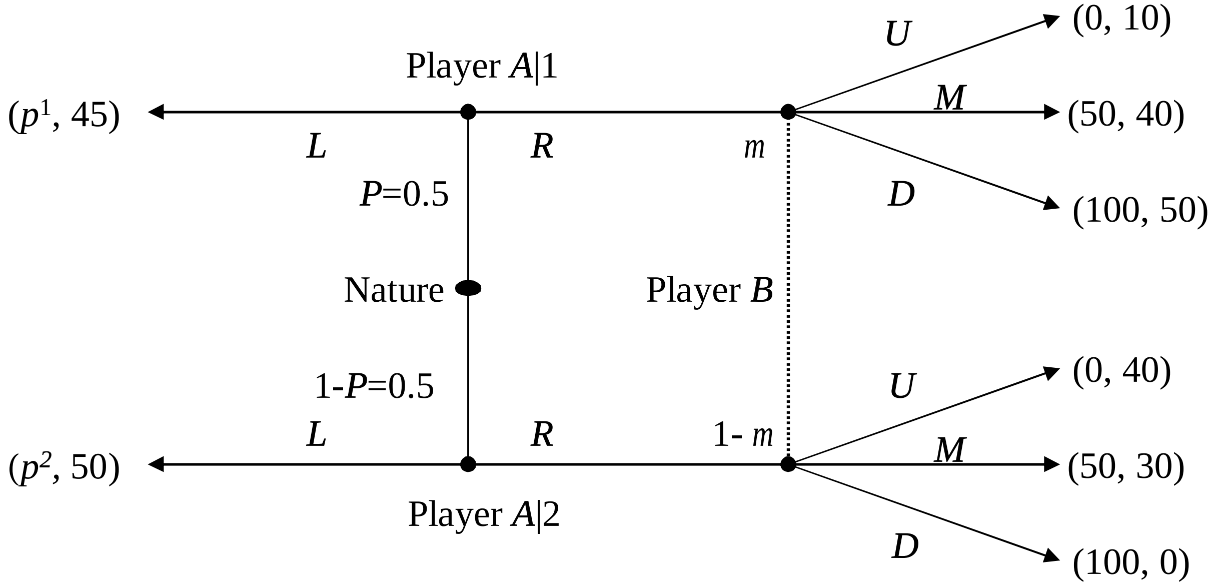

Example 1: $p^1 = 60$, $p^2 = 105$
Example 2: $p^1 = 25$, $p^2 = 75$
Example 3: $p^1 = 60$, $p^2 = 75$

Figure 7: Example 1

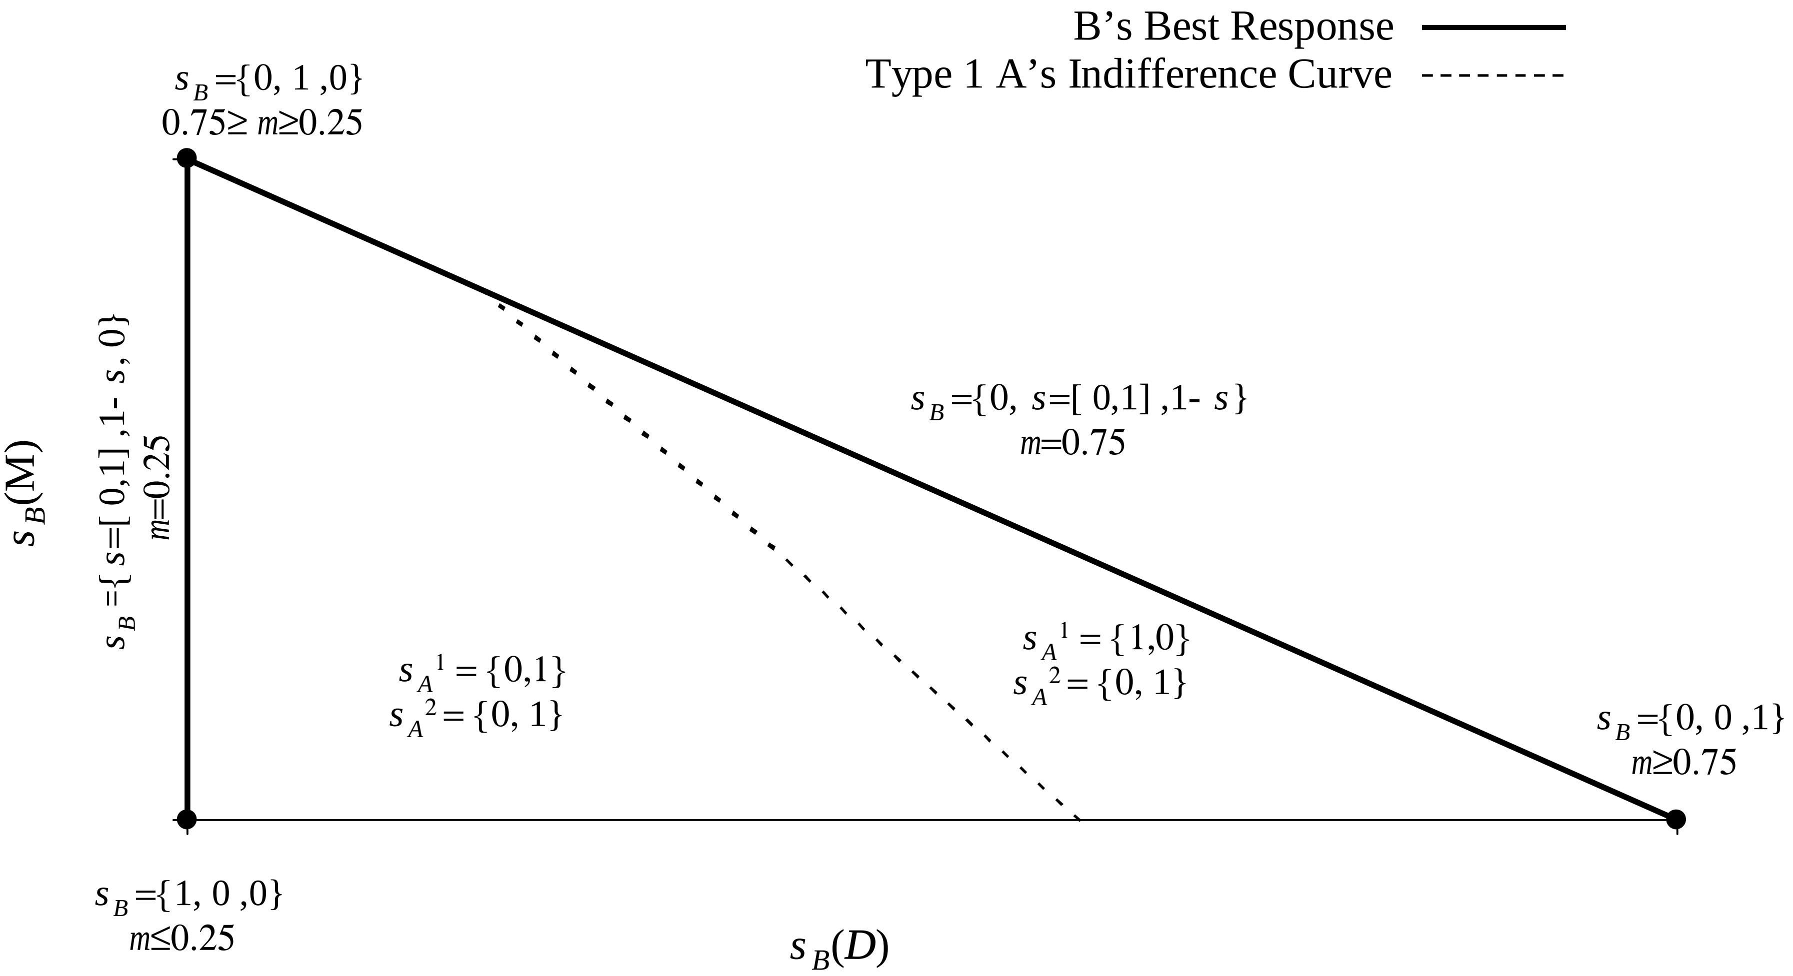

Figure 8

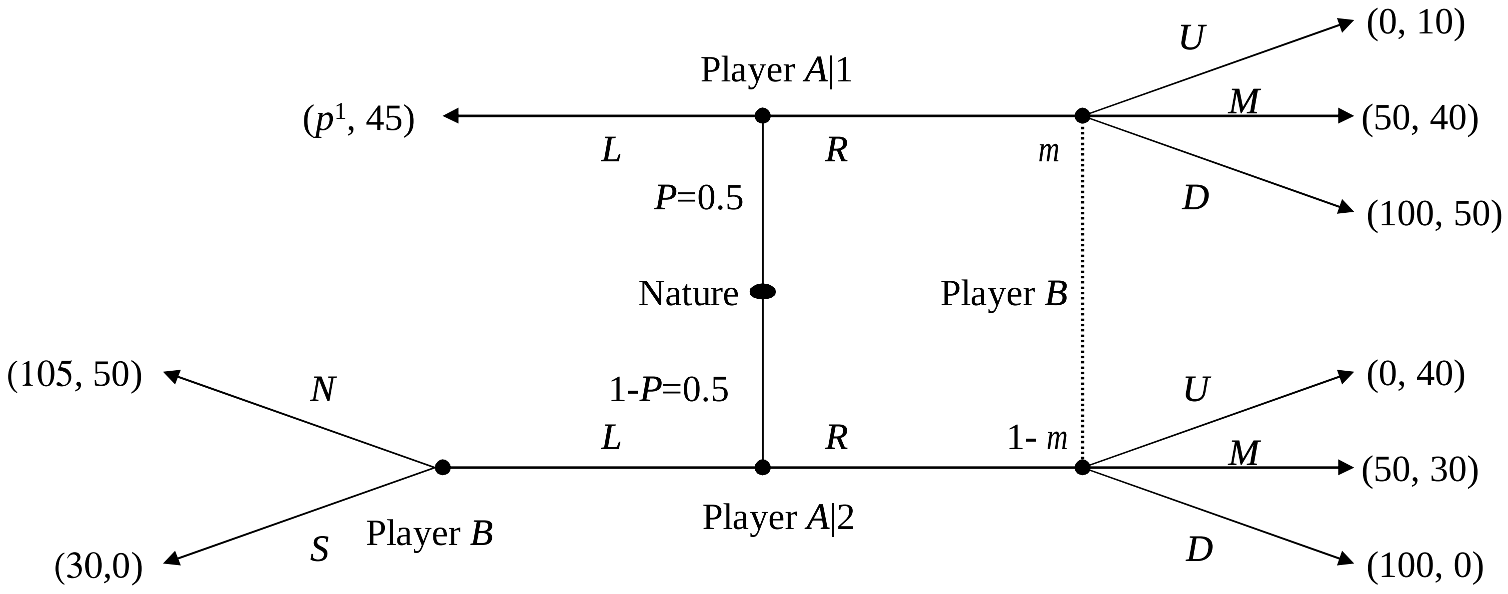

U
(0, 10)
Player A|1
M
(p^1, 45)
(50, 40)
L
R
m
D
P=0.5
(100, 50)
Nature
Player B
(105, 50)
N
U
(0, 40)
1-P=0.5
L
R
1- m
M
(50, 30)
Player A|2
S
Player B
D
(30,0)
(100, 0)

Figure 9: Example 2

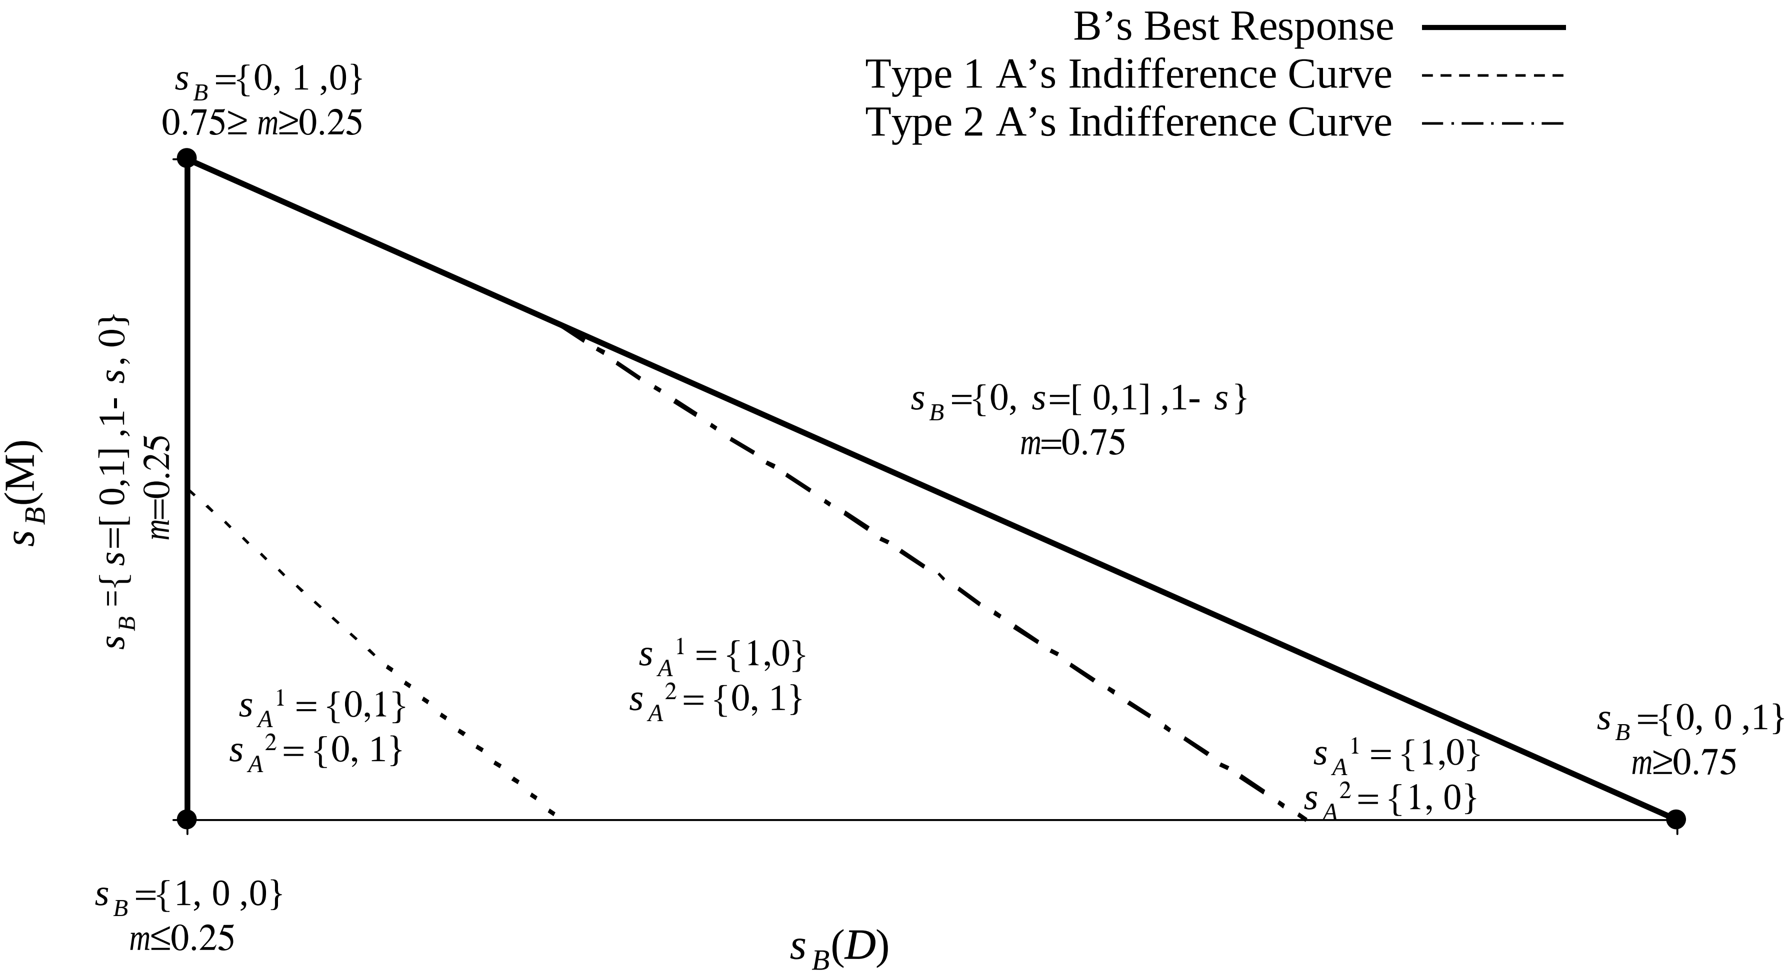

Figure 10: Example 3

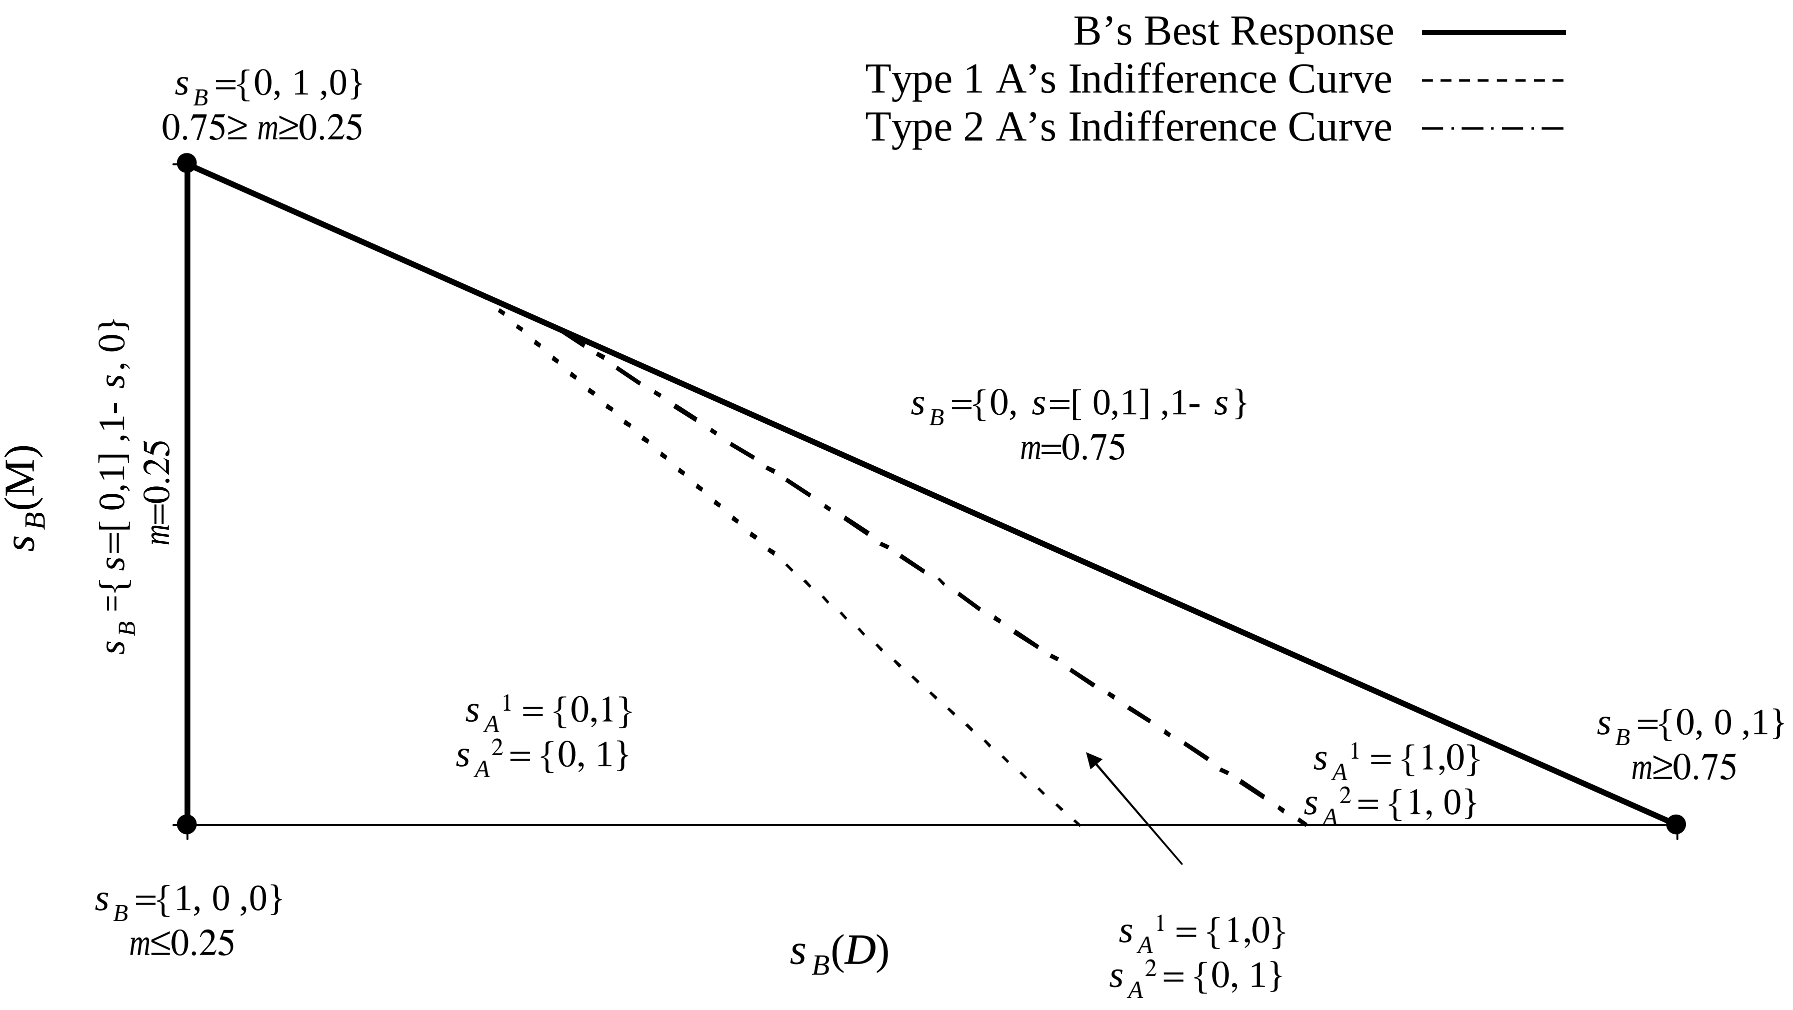

Figure 11: Example Of Why It May Be Useful To Only Look At *B*’s Best Responses

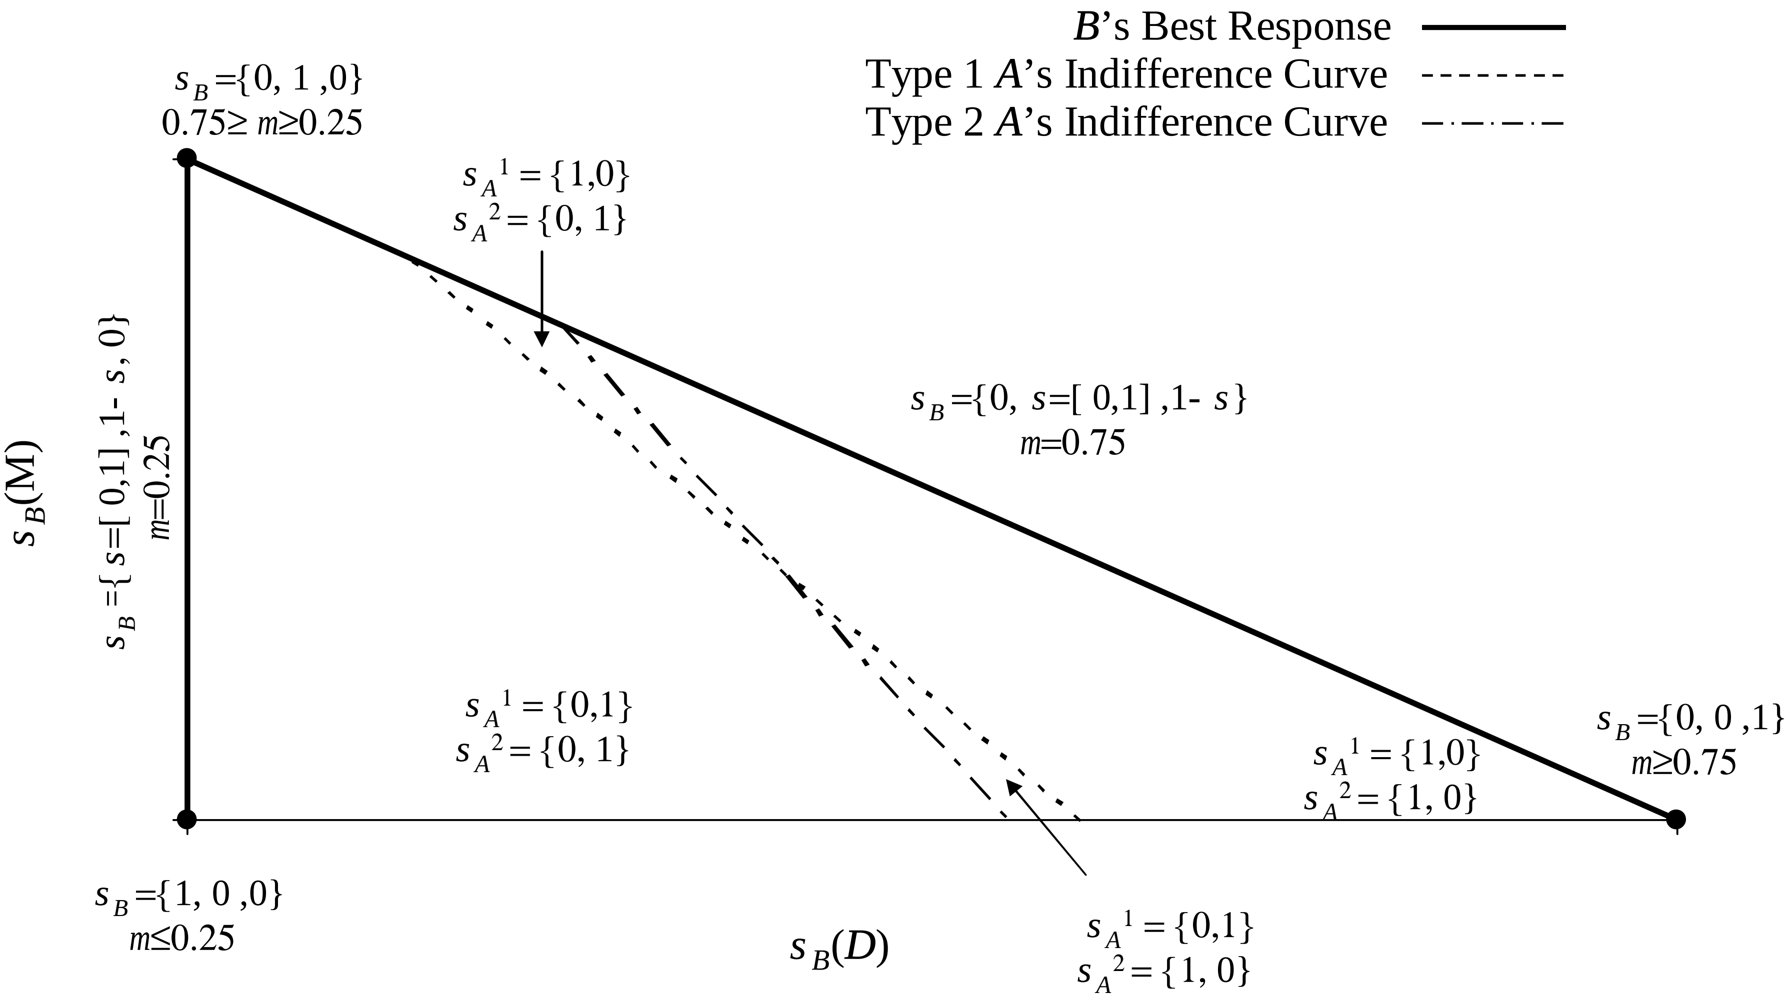

Figure 12: Example Where Deviation Sets Are Not So Useful

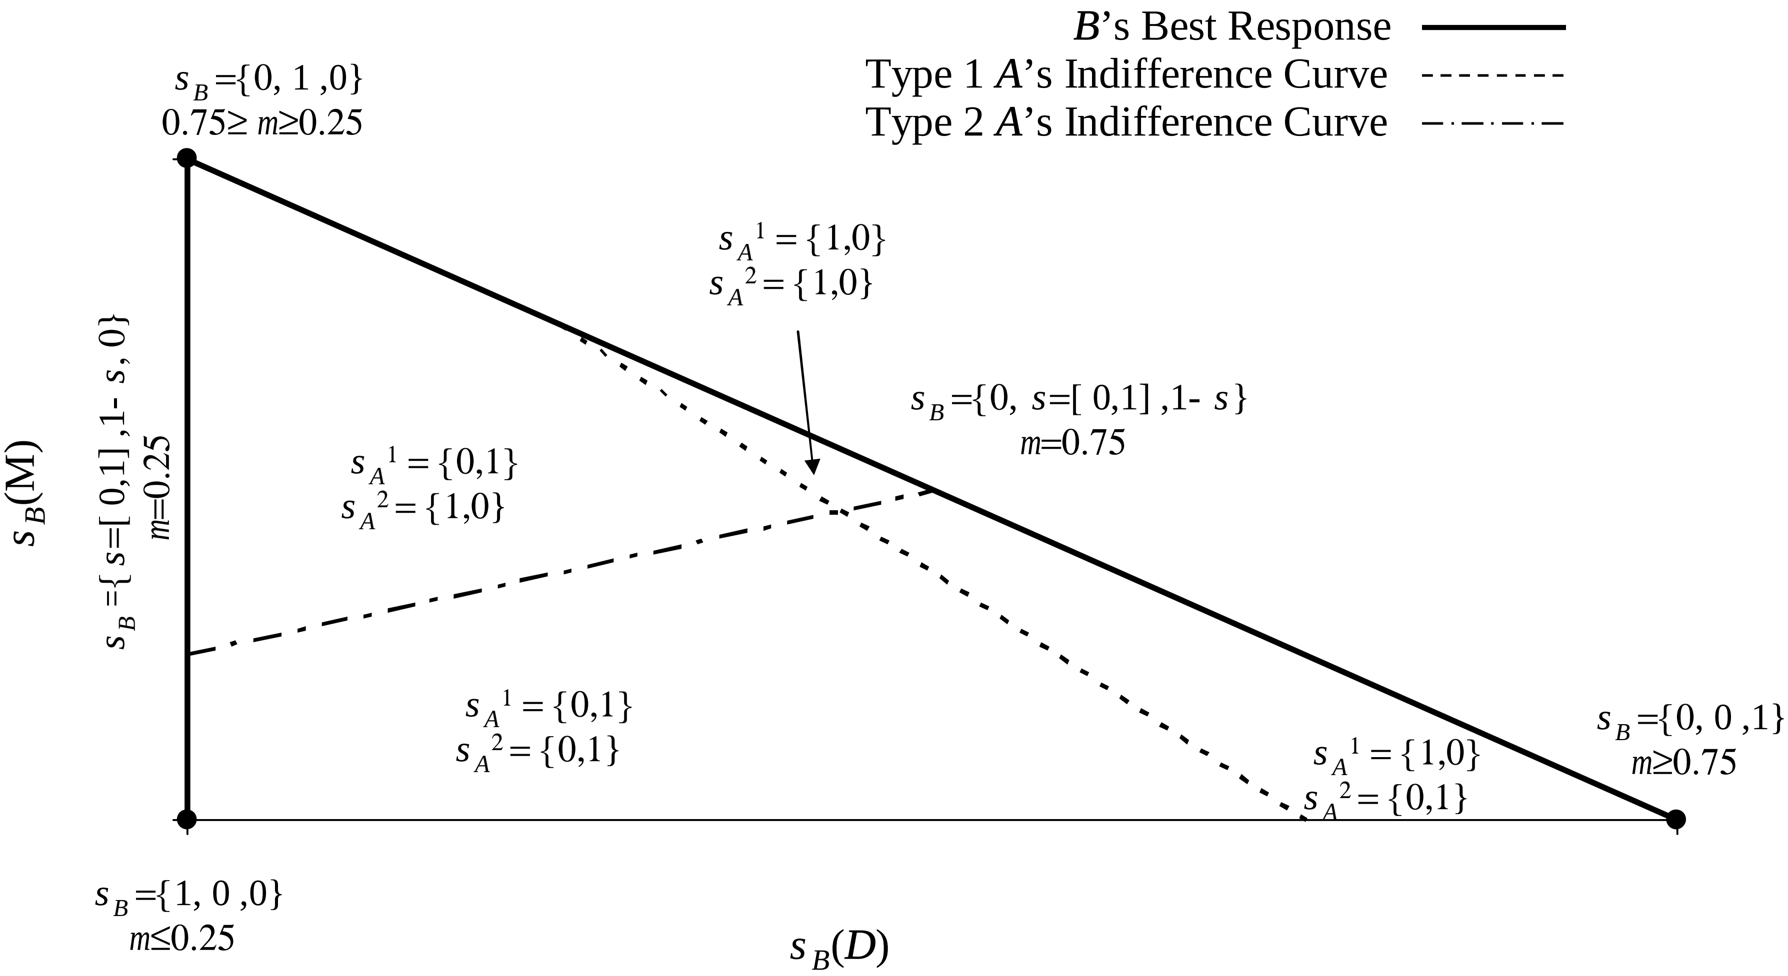

Figure 13: Case where $1>E_W(r_D)$

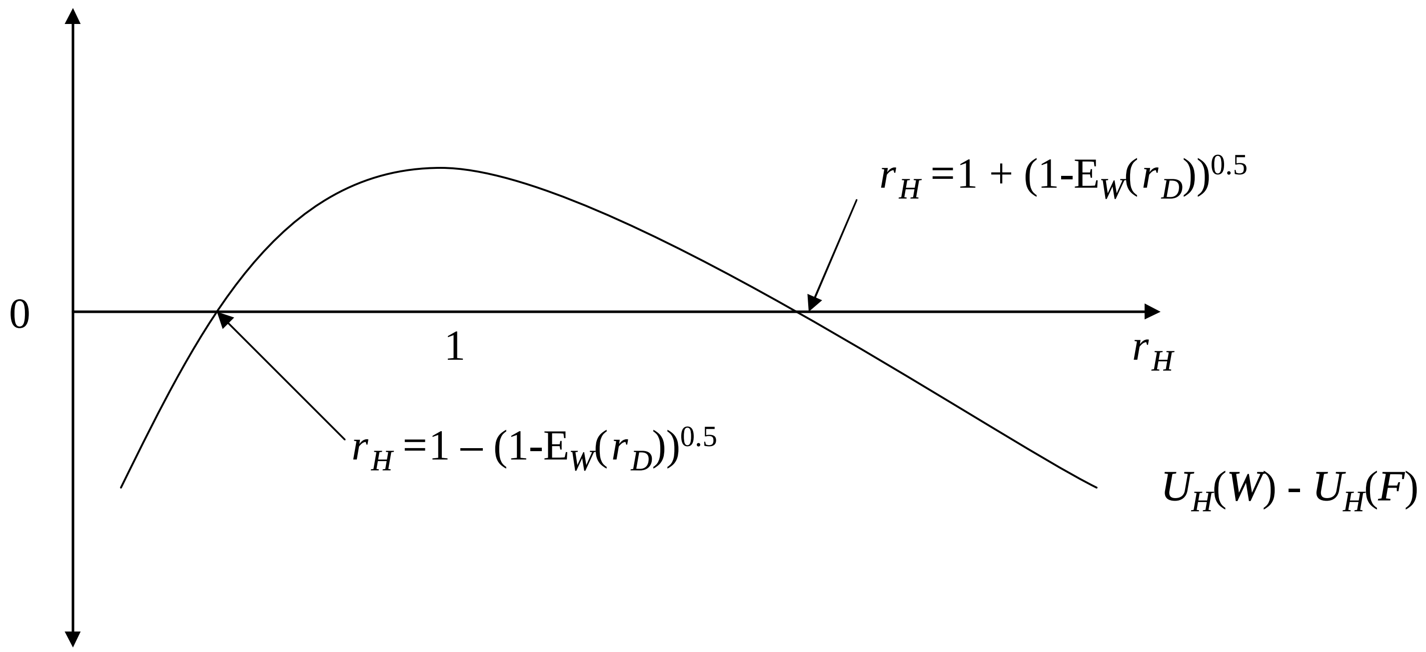

Figure 14: Case where $1<E_W(r_D)$

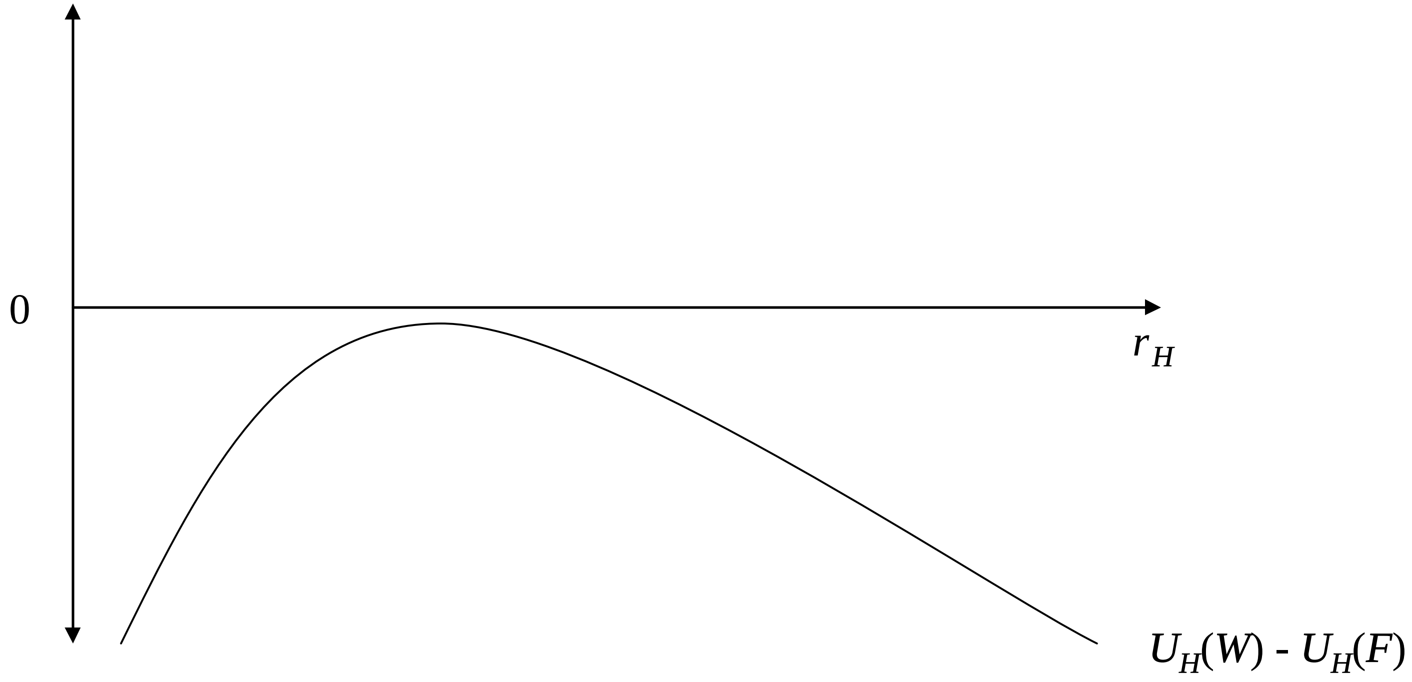

Figure 15: Comparison of Complete and Incomplete Information Games When $1 > E_W(r_D)$