Game Theory:
Static and Dynamic Games of Incomplete Information

Branislav L. Slantchev
Department of Political Science, University of California – San Diego

June 3, 2005

Contents.

1 Static Bayesian Games **2**
1.1 Building a Plant . . . 2
1.1.1 Solution: The Strategic Form . . . 4
1.1.2 Solution: Best Responses . . . 5
1.2 Interim vs. Ex Ante Predictions . . . 6
1.3 Bayesian Nash Equilibrium . . . 6
1.4 The Battle of the Sexes . . . 10
1.5 An Exceedingly Simple Game . . . 12
1.6 The Lover-Hater Game . . . 13
1.7 The Market for Lemons . . . 16

2 Dynamic Bayesian Games **18**
2.1 A Two-Period Reputation Game . . . 19
2.2 Spence's Education Game . . . 22

3 Computing Perfect Bayesian Equilibria **24**
3.1 The Yildiz Game . . . 24
3.2 The Myerson-Rosenthal Game . . . 25
3.3 One-Period Sequential Bargaining . . . 29
3.4 A Three-Player Game . . . 29
3.5 Rationalist Explanation for War . . . 31

Thus far, we have only discussed games where players knew about each other's utility functions. These games of complete information can be usefully viewed as rough approximations in a limited number of cases. Generally, players may not possess full information about their opponents. In particular, players may possess *private information* that others should take into account when forming expectations about how a player would behave.

To analyze these interesting situations, we begin with a class of games of incomplete information (i.e. games where at least one player is uncertain about another player's payoff function) that are the analogue of the normal form games with complete information: Bayesian games (static games of incomplete information). Although most interesting incomplete information games are dynamic (because these allow players to lie, signal, and learn about each other), the static formulation allows us to focus on several modeling issues that will come handy later.

The goods news is that you already know how to solve these games! Why? Because you know how to solve games of imperfect information. As we shall see, Harsanyi showed how we can transform games of incomplete information into ones of imperfect information, and so we can make heavy use of our perfect Bayesian equilibrium.

Before studying dynamic (extensive form) games of incomplete information, let's take a look at static (normal form) ones.

1 Static Bayesian Games

1.1 Building a Plant

Consider the following simple example. There are two firms in some industry: an incumbent (player 1) and a potential entrant (player 2). Player 1 decides whether to build a plant, and simultaneously player 2 decides whether to enter. Suppose that player 2 is uncertain whether player 1's building cost is 1.5 or 0, while player 1 knows his own cost. The payoffs are shown in Fig. 1 (p. 2).

	Enter	*Don't*
Build	$0, -1$	$2, 0$
Don't	$2, 1$	$3, 0$

high-cost

	Enter	*Don't*
Build	$1.5, -1$	$3.5, 0$
Don't	$2, 1$	$3, 0$

low-cost

Figure 1: The Two Firm Game.

Player 2's payoff depends on whether player 1 builds or not (but is not directly influenced by player 1's cost). Entering for player 2 is profitable only if player 1 does not build. Note that "don't build" is a dominant strategy for player 1 when his cost is high. However, player 1's optimal strategy when his cost is low depends on his prediction about whether player 2 will enter. Denote the probability that player 2 enters with y. Building is better than not building if

$$\begin{aligned} 1.5y + 3.5(1-y) &\geq 2y + 3(1-y) \\ y &\leq 1/2. \end{aligned}$$

In other words, a low-cost player 1 will prefer to build if the probability that player 2 enters is less than $1/2$. Thus, player 1 has to predict player 2's action in order to choose his own action, while player 2, in turn, has to take into account the fact that player 1 will be conditioning his action on these expectations.

For a long time, game theory was stuck because people could not figure out a way to solve such games. However, in a couple of papers in 1967-68, John C. Harsanyi proposed a method that allowed one to transform the game of incomplete information into a game of imperfect information, which could then be analyzed with standard techniques. Briefly, the **Harsanyi transformation** involves introducing a prior move by Nature that determines player 1's "type" (in our example, his cost), transforming player 2's incomplete information about player 1's cost into imperfect information about the move by Nature.

Letting p denote the prior probability of player 1's cost being high, Fig. 2 (p. 3) depicts the Harsanyi transformation of the original game into one of imperfect information.

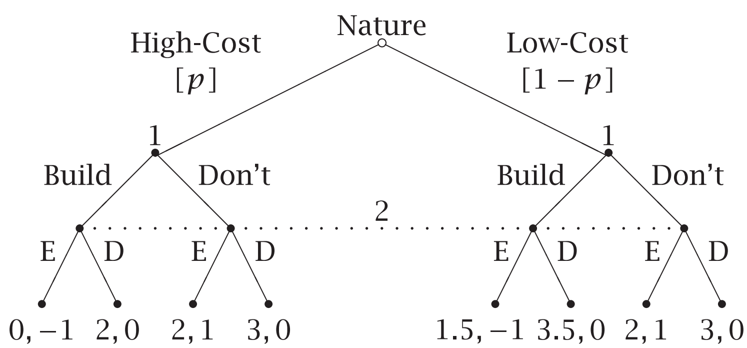

Figure 2: The Harsanyi-Transformed Game from Fig. 1 (p. 2).

Nature moves first and chooses player 1's "type": with probability p the type is "high-cost" and with probability $1 - p$, the type is "low-cost." It is standard to assume that both players have the same prior beliefs about the probability distribution on nature's moves. Player 1 observes his own type (i.e. he learns what the move by Nature is) but player 2 does not. We can now apply the Nash equilibrium solution concept to this new game. Harsanyi's **Bayesian Nash Equilibrium** (or simply Bayesian Equilibrium) is precisely the Nash equilibrium of this imperfect-information representation of the game.

Before defining all these things formally, let's skip ahead and actually solve the game in Fig. 2 (p. 3). Player 2 has one (big) information set, so his strategy will only have one component: what to do at this information set. Note now that player 1 has two information sets, so his strategy must specify what to do if his type is high-cost and what to do if his type is low-cost. One might wonder why player 1's strategy has to specify what to do in both cases, after all, once player 1 learns his type, he does not care what he would have done if he is of another type.

The reason the strategy has to specify actions for both types is roughly analogous for the reason the strategy has to specify a complete plan for action in extensive-form games with complete information: player 1's optimal action depends on what player 2 will do, which in turn depends on what player 1 would have done at information sets even if these are never reached in equilibrium. The Bayesian Nash equilibrium will be a *triple* of strategies: one for player 1 of the high-cost type, another for player 1 of the low-cost type, and one for player 2.

1.1.1 Solution: The Strategic Form

Let's write down the strategic form representation of the game in Fig. 2 (p. 3). Player 1's pure strategies are $S_1 = \{BB, BD, DB, DD\}$, where the first component of each pair tells his what to do if he is the high-cost type, and the second component if he is the low-cost type. Player 2 has only two pure strategies, $S_2 = \{E, D\}$. The resulting payoff matrix is shown in Fig. 3 (p. 4).

Player 2

Player 1		E	D
	BB	$1.5 - 1.5p, -1$	$3.5 - 1.5p, 0$
	BD	$2 - 2p, 1$	$3 - p, 0$
	DB	$1.5 + 0.5p, 2p - 1$	$3.5 - 0.5p, 0$
	DD	$2, 1$	$3, 0$

Figure 3: The Strategic Form of the Game in Fig. 2 (p. 3).

DB strictly dominates BB and DD strictly dominates BD. Eliminating the two strictly dominated strategies reduces the game to the one shown in Fig. 4 (p. 4).

Player 2

Player 1		E	D
	DB	$1.5 + 0.5p, 2p - 1$	$3.5 - 0.5p, 0$
	DD	$2, 1$	$3, 0$

Figure 4: The Reduced Strategic Form of the Game in Fig. 2 (p. 3).

If player 2 chooses E, then player 1's unique best response is to choose DD regardless of the value of $p < 1$. Hence $\langle DD, E\rangle$ is a Nash equilibrium for all values of $p \in (0, 1)$. Note that E strictly dominates D whenever $2p - 1 > 0 \Rightarrow p > 1/2$, and so player 2 will never mix in equilibrium in this case.

Let's then consider the cases when $p \leq 1/2$. We now also have $\langle DB, D\rangle$ as a Nash equilibrium. Suppose now that player 2 mixes in equilibrium. Since she is willing to randomize,

$$\begin{aligned} U_2(E) &= U_2(D) \\ \sigma_1(DB)(2p - 1) + (1 - \sigma_1(DB))(1) &= 0 \\ \sigma_1(DB) &= \frac{1}{2(1 - p)}. \end{aligned}$$

Since player 1 must be willing to randomize as well, it follows that

$$\begin{aligned} U_1(DB) &= U_1(DD) \\ \sigma_2(E)(1.5 + 0.5p) + (1 - \sigma_2(E))(3.5 - 0.5p) &= 2\sigma_2(E) + 3(1 - \sigma_2(E)) \\ \sigma_2(E) &= 1/2. \end{aligned}$$

Hence, we have a mixed-strategy Nash equilibrium with $\sigma_1(DB) = 1/[2(1 - p)]$, and $\sigma_2(E) = 1/2$ whenever $p \leq 1/2$.

If $p = 1/2$, then player 2 will be indifferent between her two pure strategies if player 1 chooses DB for sure, so she can randomize. Suppose she mixes in equilibrium. Then player

1's expected payoff from DB would be $\sigma_2(E)(1.5 + 0.25) + (1 - \sigma_2(E))(3.5 - 0.25) = 3.25 - 1.5\sigma_2(E)$. Player 1's expected payoff from DD is then $2\sigma_2(E) + 3(1 - \sigma_2(E)) = 3 - \sigma_2(E)$. He would choose DB whenever $3.25 - 1.5\sigma_2(E) \geq 3 - \sigma_2(E)$, that is, whenever $\sigma_2(E) \leq 1/2$. Hence, there is a continuum of mixed strategy Nash equilibria when $p = 1/2$: Player 1 chooses DB and player 2 randomizes with $\sigma_2(E) \leq 1/2$.

Summarizing the results, we have the following Nash equilibria:

- Neither the high nor low cost types build, and player 2 enters;
- If $p \leq 1/2$, there are two types of equilibria:
  - the high-cost type does not build, but the low-cost type does, and player 2 enters;
  - the high-cost type does not build, but the low-cost type builds with probability $1/[2(1-p)]$, and player 2 enters with probability $1/2$.
- If $p = 1/2$, there is a continuum of equilibria: the high cost player 1 does not build, but the low-cost does, and player 2 enters with probability less than $1/2$.

1.1.2 Solution: Best Responses

Noting that the high-cost player 1 never builds, let x denote the probability that the low-cost player 1 builds. As before, let y denote player 2's probability of entry. The best-responses for the low-cost player 1 and player 2 are

$$BR_1(y|L) = \begin{cases} 1 & \text{if } y < 1/2 \\ [0,1] & \text{if } y = 1/2 \\ 0 & \text{if } y > 1/2 \end{cases} \qquad BR_2(x) = \begin{cases} 1 & \text{if } x < \overline{x} \\ [0,1] & \text{if } x = \overline{x} \\ 0 & \text{if } x > \overline{x} \end{cases}$$

$$\text{where} \quad \overline{x} = \frac{1}{2(1-p)}.$$

To see how the best-responses were obtained, note that the low-cost player 1 strictly prefers building to not building when the expected utility of building exceeds the expected utility of not building:

$$\begin{aligned} U_1(B|L) &\geq U_1(D|L) \\ 1.5y + 3.5(1-y) &\geq 2y + 3(1-y) \\ y &\leq 1/2 \end{aligned}$$

Similarly, player 2 prefers to enter when the expected utility of doing so exceeds the expected utility of not entering:

$$\begin{aligned} U_2(E) &\geq U_2(D) \\ pU_2(E|H) + (1-p)U_2(E|L) &\geq pU_2(D|H) + (1-p)U_2(D|L) \\ p(1) + (1-p)(-x + 1 - x) &\geq 0 \\ 1 - 2x + 2px &\geq 0 \\ x &\leq \frac{1}{2(1-p)} \equiv \overline{x} \end{aligned}$$

Given these best-responses, the search for a Bayesian Nash Equilibrium boils down to finding a pair (x, y), such that x is optimal for player 1 with low cost against player 2 and y is optimal for player 2 against player 1 given beliefs p and player 1's strategy (x for the low cost and "don't build" for the high cost).

For instance, $(x = 0, y = 1)$ is an equilibrium for any p (here, player 1 does not build regardless of type and player 2 enters). This is the first Nash equilibrium we found above. Also, $(x = 1, y = 0)$ is an equilibrium if $1 \geq 1/[2(1-p)] \Rightarrow p \leq 1/2$ (here, the low-cost player 1 builds, the high-cost player 1 does not, and player 2 does not enter). This is the other pure-strategy equilibrium we found.

You should verify that the mixed-strategy Nash equilibria from the previous method can also be recovered with this one. This yields the full set of equilibria.

1.2 Interim vs. Ex Ante Predictions

Suppose in the two-firm example player 2 also had private information and could be of two types, "aggressive" and "accommodating." If she must predict player 1's type-contingent strategies, she must be concerned with how an aggressive player 2 might think player 1 would play for each of the possible types for player 1 and also how an accommodating player 2 might think player 1 would play for each of his possible types. (Of course, player 1 must also estimate both the aggressive and accommodating type's beliefs about himself in order to predict the distribution of strategies he should expect to face.)

This brings up the following important question: How should the different types of player 2 be viewed? On one hand, they can be viewed as a way of describing different information sets of a *single* player 2 who makes a type-contingent decision at the *ex ante* stage. This is natural in Harsanyi's formulation, which implies that the move by Nature reveals some information known only to player 2 which affects her payoffs. Player 2 makes a type-contingent plan expecting to carry out one of the strategies after learning her type. On the other hand, we can view the two types as two different "agents," one of whom is selected by Nature to "appear" when the game is played.

In the first case, the single *ex ante* player 2 predicts her opponent's play at the *ex ante* stage, so all types of player 2 would make the same prediction about the play of player 1. Under the second interpretation, the different "agents" would each make her prediction at the *interim* stage after learning her type, and thus different "agents" can make different predictions.

It is worth emphasizing that in our setup, players plan their actions *before* they receive their signals, and so we treat player 2 as a single *ex ante* player who makes type-contingent plans. Both the aggressive and accommodating types will form the same beliefs about player 1. (For more on the different interpretations, see Fudenberg & Tirole, section 6.6.1.)

1.3 Bayesian Nash Equilibrium

A static game of imperfect information is called a **Bayesian game**, and it consists of the following elements:

- a set of players, $N = \{1, \ldots, n\}$,

  and, for each player $i \in N$,

- a set of actions, A_i, with the usual $A = \times_{i \in N} A_i$,
- a set of types, Θ_i, with the usual $\Theta = \times_{i \in N} \Theta_i$,

- a probability function specifying i's belief about the type of other players given his own type, $p_i : \Theta_i \to \triangle(\Theta_{-i})$,
- a payoff function, $u_i : A \times \Theta \to \mathbb{R}$.

Let's explore these definitions. We want to represent the idea that each player knows his own payoff function but may be uncertain about the other players' payoff functions. Let $\theta_i \in \Theta_i$ be some type of player i (and so Θ_i is the set of all player i types). Each type corresponds to a different payoff function that player i might have.

We specify the pure-strategy space A_i (with elements a_i and mixed strategies $\alpha_i \in \mathcal{A}_i$) and the payoff function $u_i(a_1, \ldots, a_n | \theta_1, \ldots, \theta_n)$. Since each player's choice of strategy can depend on his type, we let $s_i(\theta_i)$ denote the pure strategy player i chooses when his type is θ_i ($\sigma_i(\theta_i)$ is the mixed strategy). Note that in a Bayesian game, pure strategy spaces are constructed from the type and action spaces: Player i's set of possible (pure) strategies S_i is the set of all possible functions with domain Θ_i and range A_i. That is, S_i is a collection of functions $s_i : \Theta_i \to A_i$.

If player i has k possible payoff functions, then the type space has k elements, $\#(\Theta_i) = k$, and we say that player i has k types. Given this terminology, saying that player i knows his own payoff function is equivalent to saying that he knows his type. Similarly, saying that player i may be uncertain about other players' payoff functions is equivalent to saying that he may be uncertain about their types, denoted by θ_{-i}. We use Θ_{-i} to denote the set of all possible types of the other players and use the probability distribution $p_i(\theta_{-i}|\theta_i)$ to denote player i's *belief* about the other players' types θ_{-i}, given his knowledge of his own type, θ_i.[1] For simplicity, we shall assume that Θ_i has a finite number of elements.

If player i knew the strategies of the other players as a function of their type, that is, he knew $\{\sigma_j(\cdot)\}_{j \neq i}$, player i could use his beliefs $p_i(\theta_{-i}|\theta_i)$ to compute the expected utility to each choice and thus find his optimal response $\sigma_i(\theta_i)$.[2]

Following Harsanyi, we shall assume that the timing of the static Bayesian game is as follows: (1) Nature draws a type vector $\theta = (\theta_1, \ldots, \theta_n)$, where θ_i is drawn from the set of possible types Θ_i using some objective distribution p that is common knowledge; (2) Nature reveals θ_i to player i but not to any other player; (3) the players simultaneously choose actions, player i chooses from the feasible set A_i; and then (4) payoffs $u_i(a_1, \ldots, a_n | \theta)$ are received.

Since we assumed in step (1) above that it is common knowledge that Nature draws the vector θ from the prior distribution $p(\theta)$, player i can use Bayes' Rule to compute his posterior belief $p_i(\theta_{-i}|\theta_i)$ as follows:

$$p_i(\theta_{-i}|\theta_i) = \frac{p(\theta_{-i}, \theta_i)}{p(\theta_i)} = \frac{p(\theta_{-i}, \theta_i)}{\sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}, \theta_i)}.$$

Furthermore, the other players can compute the various beliefs that player i might hold depending on i's type. We shall frequently assume that the players' types are independent, in which case $p_i(\theta_{-i})$ does not depend on θ_i although it is still derived from the prior distribution $p(\theta)$.

[1] In practice, the players' types are usually assumed to be independent, in which case $p_i(\theta_{-i}|\theta_i)$ does not depend on θ_i, and so we can write the beliefs simply as $p_i(\theta_{-i})$.

[2] This is where the assumption that Θ_i is finite is important. If there is a continuum of types, we may run into measure-theoretic problems.

Now that we have the formal description of a static Bayesian game, we want to define the equilibrium concept for it. The notation is somewhat cumbersome but the intuition is not: each player's (type-contingent) strategy must be the best response to the other players' strategies. That is, a **Bayesian Nash equilibrium** is simply a Nash equilibrium in a Bayesian game.

Given a strategy profile $s(\cdot)$ and a strategy $s_i'(\cdot) \in S_i$ (recall that this is a type-contingent strategy, with $s_i' \in S_i$, where S_i is the collection of functions $s_i : \Theta_i \to A_i$), let $(s_i'(\cdot), s_{-i}(\cdot))$ denote the profile where player i plays $s_i'(\cdot)$ and the other players follow $s_{-i}(\cdot)$, and let

$$\left(s_i'(\theta_i), s_{-i}(\theta_{-i})\right) = \left(s_1(\theta_1), \ldots, s_{i-1}(\theta_{i-1}), s_i'(\theta_i), s_{i+1}(\theta_{i+1}), \ldots, s_N(\theta_N)\right)$$

denote the value of this profile at $\theta = (\theta_i, \theta_{-i})$.

Definition 1. Let G be a Bayesian game with a finite number of types Θ_i for each player i, a prior distribution p, and strategy spaces S_i. The profile $s(\cdot)$ is a (pure-strategy) **Bayesian equilibrium** of G if, for each player i and every $\theta_i \in \Theta_i$,

$$s_i(\theta_i) \in \arg\max_{s_i' \in S_i} \sum_{\theta_{-i}} u_i\left(s_i', s_{-i}(\theta_{-i}) | \theta_i, \theta_{-i}\right) p(\theta_{-i}|\theta_i),$$

that is, no player wants to change his strategy, even if the change involves only one action by one type.[3]

Simply stated, each type-contingent strategy is a best response to the type-contingent strategies of the other players. Player i calculates the expected utility of playing every possible type-contingent strategy $s_i(\theta_i)$ given his type θ_i. To do this, he sums over all possible combinations of types for his opponents, θ_{-i}, and for each combination, he calculates the expected utility of playing against this particular set of opponents: The utility, $u_i(s_i', s_{-i}(\theta_{-i})|\theta_i, \theta_{-i})$, is multiplied by the probability that this set of opponents θ_{-i} is selected by Nature: $p(\theta_{-i}|\theta_i)$. This yields the optimal behavior of player i when of type θ_i. We then repeat the process for all possible $\theta_i \in \Theta_i$ and all players.

Example 1. Consider a simple example. You are player 1 and you are playing with two opponents, A, and B. Each of them has two types. Player A can be either t_A^1 with probability p_A or t_A^2 with probability $1 - p_A$, and player be can be either t_B^1 with probability p_B or t_B^2 with probability $1 - p_B$. Each of these types has two actions at his disposal. Player A can choose either a_1 or a_2, and player B can chooses either b_1 or b_2. You can choose from actions c_1, c_2, and c_3 and you can be one of two types, θ^1 or θ^2.

We let player 1 be player i and use Definition 1. First define θ_{-i}, the set of all possible combination of opponent types. Since there are two opponents with two types each, there are four combinations to consider:

$$\Theta_{-1} = \left\{(t_A^1, t_B^1), (t_A^1, t_B^2), (t_A^2, t_B^1), (t_A^2, t_B^2)\right\}$$

Of course, $\Theta_1 = \{\theta^1, \theta^2\}$. For each $\theta_1 \in \Theta_1$, we have to define $s_1(\theta_1)$ as the strategy that maximizes player 1's payoff given what the opponents do when we consider all possible combinations of opponents types, Θ_{-i}.

[3]The general definition is a bit more complicated but we here have used the assumption that each type has a positive probability, and so instead of maximizing the *ex ante* expected utility over all types, player i maximizes his expected utility conditional on his type θ_i for each θ_i.

Note that the probabilities associated with each type of opponent allow player 1 to calculate the probability of a particular combination being realized. Thus we have the following probabilities $p(\theta_{-1}|\theta_1)$:

$$p(t_A^1, t_B^1) = p_A p_B \qquad p(t_A^2, t_B^1) = (1 - p_A)p_B$$
$$p(t_A^1, t_B^2) = p_A(1 - p_B) \qquad p(t_A^2, t_B^2) = (1 - p_A)(1 - p_B)$$

where we suppressed the conditioning on θ^1 because the realizations are independent from player 1's own type.

We now fix a strategy profile for the other two players to check player 1's optimal strategy for that profile. The players are using type-contingent strategies themselves. Given the number of available actions, the possible (pure) strategies are $s_A(t_A^1) = s_A(t_A^2) \in \{a_1, a_2\}$, and $s_B(t_B^1) = s_B(t_B^2) \in \{b_1, b_2\}$. So, suppose we want to find player 1's best strategy against the profile where both types of player A choose the same action, $s_A(t_A^1) = s_A(t_A^2) = a_1$, but the two types of player B choose different actions, $s_B(t_B^1) = b_1$, and $s_B(t_B^2) = b_2$.

We have to calculate the summation over all θ_{-1}, of which there are four. For each of these, we calculate the probability of this combination of opponents occurring (we did this above) and then multiply it by the payoff player 1 expects to get from his strategy if he is matched with these particular types of opponents. This gives the expected payoff of player 1 from following his strategy against opponents of the particular type. Once we add all the terms, we have player 1's expected payoff from his strategy.

So, suppose we want to calculate player 1's expected payoff from playing $s_1(\theta^1) = c_1$:

$$\begin{aligned}
&u_1\left(c_1, s_A(t_A^1), s_B(t_B^1)\right) p(t_A^1, t_B^1) + u_1\left(c_1, s_A(t_A^1), s_B(t_B^2)\right) p(t_A^1, t_B^2)\\
+&u_1\left(c_1, s_A(t_A^2), s_B(t_B^1)\right) p(t_A^2, t_B^1) + u_1\left(c_1, s_A(t_A^2), s_B(t_B^2)\right) p(t_A^2, t_B^2)\\
=&u_1(c_1, a_1, b_1)\, p_A p_B + u_1(c_1, a_1, b_2)\, p_A(1 - p_B)\\
+&u_1(c_1, a_1, b_1)\,(1 - p_A)p_B + u_1(c_1, a_1, b_2)\, p(1 - p_A)(1 - p_B)
\end{aligned}$$

We would then do this for actions c_2 and c_3, and then pick the action that yields the highest payoff from the three calculations. This is the arg max strategy. That is, it is the strategy that maximizes the expected utility.[4] This yields type θ^1 the best response to the strategy profile specified above.

We shall have to find the optimal response to this strategy profile if player 1 is of type θ^2. We then have to find player A's and player B's optimal strategies given what they know about the other players. Once all of these best responses are found, we can match them to see which constitute profiles with strategies that are mutual best responses. That is, we then proceed as before, when we found best responses and equilibria in normal form games.

Here's an example with formal notation. Suppose there are two players, player 1 and player 2 and for each player there are two possible types. Player i's possible types are θ_i and θ_i'. Furthermore, suppose that the types are independently distributed with the probability of θ_1 being p and the probability of θ_2 being q. For a given strategy profile (s_1^*, s_2^*), the expected payoff of player 1 of type θ_1 is

$$qu_1(s_1^*(\theta_1), s_2^*(\theta_2)|\theta_1, \theta_2) + (1 - q)u_1(s_1^*(\theta_1), s_2^*(\theta_2')|\theta_1, \theta_2'),$$

[4]Unlike the max of an expression which denotes the expression's maximum when the expression is evaluated, the arg max operator finds the *parameter(s)*, for which the expression attains its maximum value. In our case, the arg max simply instructs us to pick the strategy that yields the highest expected payoff when matched against the profile of opponent's strategies.

and for a given mixed strategy profile (σ_1^*, σ_2^*), the expected payoff of player 1 of type θ_1 is

$$q \sum_{a \in A} \sigma_1^*(a_1|\theta_1)\sigma_2^*(a_2|\theta_2)u_1(a_1, a_2|\theta_1, \theta_2) + (1-q) \sum_{a \in A} \sigma_1^*(a_1|\theta_1)\sigma_2^*(a_2|\theta_2')u_1(a_1, a_2|\theta_1, \theta_2').$$

A Bayesian equilibrium will consist of *four* type-contingent strategies, one for each type of each player. Some equilibria may depend on particular values of p and q, and others may not.

The existence of a Bayesian equilibrium is an immediate consequence of the existence of Nash equilibrium.

1.4 The Battle of the Sexes

Consider the following modification of the Battle of the Sexes: player 1 is unsure about player 2's payoff if they coordinate on going to the ballet, and player 2 is unsure about player 1's payoff if they coordinate on going to the fight. That is, the player 1's payoff in the (F, F) outcome is $2 + \theta_1$, where θ_1 is privately known to the man; and the woman's payoff in the (B, B) outcome is $2 + \theta_2$, where θ_2 is privately known to the woman. Assume that both θ_1 and θ_2 are independent draws from a uniform distribution $[0, x]$.[5]

	F	B
F	$2+\theta_1, 1$	$0, 0$
B	$0, 0$	$1, 2+\theta_2$

Figure 5: Battle of the Sexes with Two-Sided Incomplete Information.

In terms of the formal description,

- Players: $N = \{1, 2\}$
- Actions: $A_1 = A_2 = \{F, B\}$
- Types: $\Theta_1 = \Theta_2 = [0, x]$
- Beliefs: $p_1(\theta_2) = p_2(\theta_1) = 1/x$ (where we used the fact that the uniform probability density function is $f(x) = 1/x$ when specified for the interval $[0, x]$)
- Payoffs: u_1, u_2 as described in Fig. 5 (p. 10).

Note that in this game, each player has a continuum of types, and so Θ_i is infinite. We shall look for a Bayesian equilibrium in which player 1 goes to the fight if θ_1 exceeds some critical value x_1 and goes to the ballet otherwise, and player 2 goes to the ballet if θ_2 exceeds some critical value x_2 and goes to the fight otherwise. Formally, let $\sigma_1(\theta_1)$ denote the probability that player 1 goes to the fight, that is:

$$\sigma_1(\theta_1) = \Pr[\theta_1 > x_1] = 1 - \Pr[\theta_1 \le x_1] = 1 - \frac{x_1}{x}.$$

[5]The choice of the distribution is not important but we do have in mind that these privately known values only slightly perturb the payoffs.

Similarly, the probability that player 2 goes to the ballet is

$$\sigma_2(\theta_2) = \Pr[\theta_2 > x_2] = 1 - \Pr[\theta_2 \le x_2] = 1 - \frac{x_2}{x}.$$

Suppose the players play the strategies just specified. We now want to find x_1, x_2 that make these strategies a Bayesian equilibrium. Given player 2's strategy, player 1's expected payoffs from going to the fight and going to the ballet are:

$$\begin{aligned}
\mathrm{E}\left[u_1(F|\theta_1,\theta_2)\right] &= (2+\theta_1)(1-\sigma_2(\theta_2)) + (0)\sigma_2(\theta_2) = \frac{x_2}{x}(2+\theta_1) \\
\mathrm{E}\left[u_1(B|\theta_1,\theta_2)\right] &= (0)(1-\sigma_2(\theta_2)) + (1)\sigma_2(\theta_2) = 1 - \frac{x_2}{x}
\end{aligned}$$

Going to the fight is optimal if and only if the expected utility of doing so exceeds the expected utility of going to the ballet:

$$\begin{aligned}
\mathrm{E}\left[u_1(F|\theta_1,\theta_2)\right] &\ge \mathrm{E}\left[u_1(B|\theta_1,\theta_2)\right] \\
\frac{x_2}{x}(2+\theta_1) &\ge 1 - \frac{x_2}{x} \\
\theta_1 &\ge \frac{x}{x_2} - 3.
\end{aligned}$$

Let $x_1 = x/x_2 - 3$ denote the critical value for player 1. Player 2's expected payoffs from going to the ballet and going to the fight given player 1's strategy are:

$$\begin{aligned}
\mathrm{E}\left[u_2(B|\theta_1,\theta_2)\right] &= (0)\sigma_1(\theta_1) + (2+\theta_2)(1-\sigma_1(\theta_1)) = \frac{x_1}{x}(2+\theta_2) \\
\mathrm{E}\left[u_2(F|\theta_1,\theta_2)\right] &= (1)\sigma_1(\theta_1) + (0)(1-\sigma_1(\theta_1)) = 1 - \frac{x_1}{x}
\end{aligned}$$

and so going to the ballet is optimal if and only if:

$$\begin{aligned}
\mathrm{E}\left[u_2(B|\theta_1,\theta_2)\right] &\ge \mathrm{E}\left[u_2(F|\theta_1,\theta_2)\right] \\
\frac{x_1}{x}(2+\theta_2) &\ge 1 - \frac{x_1}{x} \\
\theta_2 &\ge \frac{x}{x_1} - 3.
\end{aligned}$$

Let $x_2 = x/x_1 - 3$ denote the critical value for player 2. We now have the two critical values, so we solve the following system of equations

$$\begin{aligned}
x_1 &= x/x_2 - 3 \\
x_2 &= x/x_1 - 3.
\end{aligned}$$

The solution is $x_1 = x_2$ and $x_2^2 + 3x_2 - x = 0$. We now solve the quadratic, whose discriminant is $D = 9 + 4x$, for x_2. The critical values are:

$$x_1 = x_2 = \frac{-3 + \sqrt{9+4x}}{2}.$$

The Bayesian equilibrium is thus the pair of strategies:[6]

$$s_1(\theta_1) = \begin{cases} F & \text{if } \theta_1 > x_1 \\ B & \text{if } \theta_1 \le x_1 \end{cases} \qquad\qquad s_2(\theta_2) = \begin{cases} F & \text{if } \theta_2 \le x_2 \\ B & \text{if } \theta_2 > x_2 \end{cases}$$

[6]There are other Bayesian equilibria in this game. For example, all types θ_1 choose F (B) and all types θ_2 choose F (B) are both Bayesian equilibria.

where

$$x_1 = x_2 = \frac{-3 + \sqrt{9 + 4x}}{2}.$$

Note that the strategies do not specify what to do for $\theta_i = x_i$ because the probability of this occurring is 0 (the probability of any particular number drawn from a continuous distribution is zero). It is customary that one of the inequalities, it does not matter which, is weak in order to handle the case. By convention, it is '$<$' that is specified as '$\leq$'.

In the Bayesian equilibrium, the probability that player 1 goes to the fight equals the probability that player 2 goes to the ballet, and they are:

$$1 - \frac{x_1}{x} = 1 - \frac{x_2}{x} = 1 - \frac{-3 + \sqrt{9 + 4x}}{2x}. \tag{1}$$

It is interesting to see what happens as uncertainty disappears (i.e. x goes to 0). Taking the limit of the expression in Equation 1 requires an application of the L'Hôpital rule:

$$\lim_{x \to 0} \left[1 - \frac{-3 + \sqrt{9 + 4x}}{2x} \right] = 1 - \lim_{x \to 0} \left[\frac{\frac{d}{dx}(-3 + \sqrt{9 + 4x})}{\frac{d}{dx}(2x)} \right] = 1 - \lim_{x \to 0} \frac{2(9 + 4x)^{-1/2}}{2} = \frac{2}{3}$$

In other words, as uncertainty disappears, the probabilities of player 1 playing F and player 2 playing B both converge to $2/3$. But these are *exactly the probabilities of the mixed strategy Nash equilibrium of the complete information case!* That is, we have just shown that as incomplete information disappears, the players' behavior in the pure-strategy Bayesian equilibrium of the incomplete-information game approaches their behavior in the mixed-strategy Nash equilibrium in the original game of complete information.

Harsanyi (1973) suggested that player j's mixed strategy represents player i's uncertainty about j's choice of a pure strategy, and that player j's choice in turn depends on a small amount of private information. As we have just shown (and as can be proven for the general case), a mixed-strategy Nash equilibrium can almost always be interpreted as a pure-strategy Bayesian equilibrium in a closely related game with a little bit of incomplete information. The crucial feature of a mixed-strategy Nash equilibrium is not that player j chooses a strategy randomly, but rather that player i is uncertain about player j's choice. This uncertainty can arise either because of randomization or (more plausibly) because of a little incomplete information, as in the example above. This is called **purification of mixed strategies**.

1.5 An Exceedingly Simple Game

Let's start with a simple example (exercise 3.5 in Myerson, p. 149). Player 1 can be of type α with probability .9 and type β with probability .1 (from player 2's perspective). The payoff matrices are in Fig. 6 (p. 12).

	L	R
U	$2,2$	$-2,0$
D	$0,-2$	$0,0$

$t_1 = \alpha(.9)$

	L	R
U	$0,2$	$1,0$
D	$1,-2$	$2,0$

$t_1 = \beta(.1)$

Figure 6: The Two Type Game.

The easiest way to solve this is to construct the strategic form of the Bayesian game. The expected payoffs are as follows:

- $U_1(U, L) = 2(.9) + 0(.1) = 1.8$ and $U_2(U, L) = 2(.9) + 2(.1) = 2$
- $U_1(U, R) = -2(.9) + 1(.1) = -1.7$ and $U_2(U, R) = 0$
- $U_1(D, L) = 0(.9) + 1(.1) = .1$ and $U_2(D, L) = -2$
- $U_1(D, R) = 0(.9) + 2(.1) = .2$ and $U_2(D, R) = 0$

The resulting payoff matrix is shown in Fig. 7 (p. 13).

	L	R
U	$1.8, 2$	$-1.7, 0$
D	$.1, -2$	$.2, 0$

Figure 7: The Strategic Form of the Game from Fig. 6 (p. 12).

There are two Nash equilibria in pure strategies, $\langle U, L\rangle$ and $\langle D, R\rangle$, and a mixed strategy equilibrium $\langle {}^1\!/_2[U], {}^{19}\!/_{36}[L]\rangle$. Note, in particular, that in two of these equilibria, player 2 chooses L with positive probability. If you look back at the original payoff matrices in Fig. 6 (p. 12), this result may surprise you because $\langle D, R\rangle$ is a Nash equilibrium in the separate games against both types of player 1. In fact, it is the unique equilibrium when $t_1 = \beta$. On the other hand, the result is perhaps not surprising because $\langle U, L\rangle$ is obviously focal in the sense that it is best for both players. However, if we relax the common knowledge assumption (about the probabilities associated with player 1's type), then there will be no Bayesian equilibrium where player 2 would choose L! (This is a variant of Rubinstein's electronic mail game.)

1.6 The Lover-Hater Game

Suppose that player 2 has perfect information and two types, l and h. Type l loves going out with player 1 whereas type h hates it. Player 1 has only one type and is uncertain about player 2's type and believes the two types are equally likely. We can describe this formally as a Bayesian game:

- Players: $N = \{1, 2\}$
- Actions: $A_1 = A_2 = \{F, B\}$
- Types: $\Theta_1 = \{x\}$, $\Theta_2 = \{l, h\}$
- Beliefs: $p_1(l|x) = p_1(h|x) = 1/2$, $p_2(x|l) = p_2(x|h) = 1$
- Payoffs: u_1, u_2 as described in Fig. 8 (p. 14).

Since player 1 has only one type, we suppress all references to his type from now on. Let's begin by analyzing player 2's optimal behavior for each of the two types. Fix an arbitrary strategy for player 1 and denote it by σ_1. When listing best responses, we shall denote the probability that players choose F (and so they choose B with complementary probability).

	F	B
F	$2,1$	$0,0$
B	$0,0$	$1,2$

type l

	F	B
F	$2,0$	$0,2$
B	$0,1$	$1,0$

type h

Figure 8: The Lover-Hater Battle of the Sexes.

- Player 2 of type l. The expected utilities from playing F and B are

$$\begin{aligned}\mathrm{E}\left[u_2(\sigma_1, F|l)\right] &= \sigma_1(F)\\ \mathrm{E}\left[u_2(\sigma_1, B|l)\right] &= 2(1-\sigma_1(F))\end{aligned}$$

  and type l's best response is

$$BR_2(\sigma_1|l) = \begin{cases} 1 & \text{if } \sigma_1(F) > 2/3 \\ [0,1] & \text{if } \sigma_1(F) = 2/3 \\ 0 & \text{if } \sigma_1(F) < 2/3 \end{cases}$$

  That is, type l will go to the fight whenever $\sigma_1(F) > 2/3$.

- Player 2 of type h. The expected utilities from playing F and B are

$$\begin{aligned}\mathrm{E}\left[u_2(\sigma_1, F|h)\right] &= 1-\sigma_1(F)\\ \mathrm{E}\left[u_2(\sigma_1, B|h)\right] &= 2\sigma_1(F)\end{aligned}$$

  and so type h's best response is

$$BR_2(\sigma_1|h) = \begin{cases} 1 & \text{if } \sigma_1(F) < 1/3 \\ [0,1] & \text{if } \sigma_1(F) = 1/3 \\ 0 & \text{if } \sigma_1(F) > 1/3 \end{cases}$$

  That is, type h will go to the fight whenever $\sigma_1(F) < 1/3$.

- Player 1. Given player 2's strategy $\sigma_2 \equiv (\sigma_2(\cdot|l), \sigma_2(\cdot|h)$, player 1's expected payoffs to playing F and B are

$$\begin{aligned}\mathrm{E}\left[u_1(F, \sigma_2)\right] &= \frac{1}{2}\sigma_2(F|l)(2) + \frac{1}{2}\sigma_2(F|h)(2) = \sigma_2(F|l) + \sigma_2(F|h)\\ \mathrm{E}\left[u_1(B, \sigma_2)\right] &= \frac{1}{2}(1-\sigma_2(F|l)) + \frac{1}{2}(1-\sigma_2(F|h)) = 1 - \frac{\sigma_2(F|l)+\sigma_2(F|h)}{2}\end{aligned}$$

  and so player 1's best response is

$$BR_1(\sigma_2) = \begin{cases} 1 & \text{if } \sigma_2(F|l) + \sigma_2(F|h) > 2/3 \\ [0,1] & \text{if } \sigma_2(F|l) + \sigma_2(F|h) = 2/3 \\ 0 & \text{if } \sigma_2(F|l) + \sigma_2(F|h) < 2/3 \end{cases}$$

  That is, player 1 will go to the fight whenever $\sigma_2(F|l) + \sigma_2(F|h) > 2/3$.

We now must find a triple $(\sigma_1, \sigma_2(\cdot|l), \sigma_2(\cdot|h))$ such that player 1's strategy is a best response to the strategies of both types of player 2 and each type of player 2's strategy is a best response to player 1's strategy. Let's check for various types of equilibria.

- **Pooling Equilibrium.**[7] Is there a pure strategy equilibrium where both h and l choose F, that is, $\sigma_2(F|l) = \sigma_2(F|h) = 1$? Player 1's best response to these strategies is to play F with probability 1, but in that case type h's optimal response is to play B with probability 1. Therefore, there is no such equilibrium.

  Is there a pure strategy pooling equilibrium where $\sigma_2(F|l) = \sigma_2(F|h) = 0$? In this case player 1's best response is to play B with probability 1, to which type h's optimal response is to play F with probability 1. Therefore, there is no such equilibrium.

- **Separating Equilibrium.**[8] Is there a pure strategy equilibrium where type l chooses F and type h chooses B, that is, $\sigma_2(F|l) = 1$, and $\sigma_2(F|h) = 0$? Player 1's best response is to choose F with probability 1, to which the specified strategies for both types of player 2 are best responses. Therefore,

  $$\left(\sigma_1(F), \sigma_2(F|l), \sigma_2(F|h)\right) = (1, 1, 0)$$

  is a pure-strategy separating Bayesian equilibrium.

  Is there a pure strategy separating equilibrium where $\sigma_2(F|l) = 0$ and $\sigma_2(F|h) = 1$? Player 1's best response is to play F with probability 1, to which type h's strategy is not a best response. Therefore, there is no such equilibrium.

- **Semi-separating Equilibrium.**[9] Is there a semi-separating equilibrium, in which both types of player 2 mix? No, because this would require $\sigma_1(F) = 2/3$ and $\sigma_1(F) = 1/3$ at the same time. Therefore, there is no such equilibrium.

  Is there an equilibrium where only type l mixes? Suppose that type l mixes in equilibrium. This implies that player 1 mixes with $\sigma_1(F) = 2/3$, which in turn implies $\sigma_2(F|h) = 0$. Since we also require that $\sigma_2(F|l) + \sigma_2(F|h) = 2/3$ (or else player 1 would not be mixing), it follows that $\sigma_2(F|l) = 2/3$. Therefore,

  $$\left(\sigma_1(F), \sigma_2(F|l), \sigma_2(F|h)\right) = (2/3, 2/3, 0)$$

  is a semi-separating Bayesian equilibrium.

  Is there an equilibrium where only type h mixes? Suppose that type h mixes in equilibrium. This implies that player 1 must also be mixing with $\sigma_1(F) = 1/3$. This in turn means that type l's optimal action is to play B, or $\sigma_2(F|l) = 0$. Since we also require $\sigma_2(F|l) + \sigma_2(F|h) = 2/3$ for player 1 to mix, we conclude that $\sigma_2(F|h) = 2/3$. Therefore,

  $$\left(\sigma_1(F), \sigma_2(F|l), \sigma_2(F|h)\right) = (1/3, 0, 2/3)$$

  is a semi-separating Bayesian equilibrium.

Voilà! The game has one separating and two semi-separating Bayesian equilibria.

[7] In a pooling equilibrium, all types of player 2 choose the same action.
[8] In a separating equilibrium, each type of player 2 chooses a different action.
[9] In a semi-separating equilibrium, at least one type of player 2 plays a mixed strategy.

1.7 The Market for Lemons

If you have bought or sold a used car, you know something about markets with asymmetric information. Typically, the seller knows far more about the car he is offering than the buyer.[10] Generally, buyers face a significant informational disadvantage. As a result, you might expect that buyers will tend not to do very well in the market, making them cautious and loath to buy used cars, in turn making sellers worse off when the market fails due to lack of demand.

Let's model this! Suppose you, the buyer, are in the market for a used car. You meet me, the seller, through an add in the *Penny Pincher* (never a good place to look for a good car deal), and I offer you an attractive 15-year old *Firebird* for sale. You love the car, it has big fat tires, it peels rubber when you hit the gas, and it's souped up with a powerful 6-cylinder engine. It also sounds cool and has a red light in the interior. You take a couple of rides around the block and it handles like a dream.

Then you suddenly have visions of the souped up engine exploding and blowing you up to smithereens, or perhaps a tire getting loose just as you screech around that particularly dangerous turn on US-1. In any case, watching the fire-fighters dowse your vehicle while you cry at the curb or watching said fire-fighters scrape you and your car off the rocks, is not likely to be especially amusing. So you tell me, "The car looks great, but how do I know it's not a lemon?"

I, being completely truthful and honest as far as used car dealers can be, naturally respond with "Oh! I've taken such good care of it. Here're are all the receipts from the regular oil changes. See? No receipts for repairs to the engine because I have *never had problems* with it! It's a peach, trust me."

You have, of course, taken my own course on repeated games and so say, "A-ha! But you will not deal with me in the future again after the sale is complete, and so you have no interest in cooperating today because I cannot punish you tomorrow for not cooperating today! You will say whatever you think will get me to buy the car."

I sigh (Blasted game theory! It was so much easier to cheat people before.) and tell you, "Fair enough. The *Blue Book* value of the car is v dollars. Take a look at the car, take a couple of more rides around the block if you wish and then decide whether you are willing to pay the *Blue Book* price and I will decide whether to offer you the car at that price."

We shall assume that if a car is peach, it is worth B (you, the buyer) \$3,000 and worth S (me, the seller) \$2,000. If it is a lemon, then it is worth \$1,000 to B and \$0 to S. In each case, your valuation is higher than mine, so under complete information trade should occur with the surplus of \$1,000 divided between us. However, there is asymmetric information. I know the car's condition, while you only estimate that the likelihood of it being a peach is q. Each of us has two actions, to trade or not trade at the market price v. We simultaneously announce what we are going to do. If we both elect to trade, then the trade takes place. Otherwise, I keep the car and you go home to deplore the evils of simultaneous-moves games. The situation is depicted in Fig. 9 (p. 17) (rows for the buyer and columns for the seller).

As before, let's derive the best responses. Here S can be thought of as having two types, l if his car is a lemon, and p if his car is a peach. Fix an arbitrary strategy for the buyer, let σ_b denote the probability that he elects to trade, and calculate seller's best responses (the

[10] For example, I once sold in Texas an old 1989 *Firebird* knowing full well that the engine tended to overheat if run at highway speeds (e.g. 85-90 mph) for over 20 minutes. I was so busy selling it, I forgot to mention this little detail to the buyer. Once I had the money, I left the state.

			S					S	
		T	N				T	N	
B	T	$3000 - v, v$	$0, 2000$		B	T	$1000 - v, v$	$0, 0$	
	N	$0, 2000$	$0, 2000$			N	$0, 0$	$0, 0$	
		peach (q)					lemon ($1 - q$)		

Figure 9: The Market for Lemons.

probability that he will choose trade) as a function of this strategy.

- Seller with a peach:

$$
\begin{aligned}
\mathrm{E}\left[u_s(T, \sigma_b|\text{peach})\right] &= \sigma_b v + (1 - \sigma_b)(2000) \\
\mathrm{E}\left[u_s(N, \sigma_b|\text{peach})\right] &= \sigma_b(2000) + (1 - \sigma_b)(2000) = 2000
\end{aligned}
$$

  and so the best-response for this seller is

$$
BR_s(\sigma_b|\text{peach}) = \begin{cases} 1 & \text{if } \sigma_b > 0 \text{ and } v > 2000 \\ [0,1] & \text{if } \sigma_b = 0 \text{ or } v = 2000 \\ 0 & \text{if } \sigma_b > 0 \text{ and } v < 2000 \end{cases}
$$

- Seller with a lemon:

$$
\begin{aligned}
\mathrm{E}\left[u_s(T, \sigma_b|\text{lemon})\right] &= \sigma_b v + (1 - \sigma_b)(0) = \sigma_b v \\
\mathrm{E}\left[u_s(N, \sigma_b|\text{lemon})\right] &= \sigma_b(0) + (1 - \sigma_b)(0) = 0
\end{aligned}
$$

  and so the best-response for this seller is

$$
BR_s(\sigma_b|\text{lemon}) = \begin{cases} 1 & \text{if } \sigma_b > 0 \text{ and } v > 0 \\ [0,1] & \text{if } \sigma_b = 0 \text{ or } v = 0 \\ 0 & \text{if } \sigma_b > 0 \text{ and } v < 0 \end{cases}
$$

- The buyer's expected utilities from trading and not trading are (letting σ_p and σ_l denote the probability of trading for the peach seller and the lemon seller respectively):

$$
\begin{aligned}
\mathrm{E}\left[u_b(T, \sigma_p, \sigma_l)\right] &= q(3000 - v)\sigma_p + (1 - q)(1000 - v)\sigma_l \\
\mathrm{E}\left[u_b(N, \sigma_p, \sigma_l)\right] &= 0
\end{aligned}
$$

  and so the best-response for the buyer is

$$
BR_b(\sigma_p, \sigma_l) = \begin{cases} 1 & \text{if } q\sigma_p(3000 - v) > (1 - q)\sigma_l(v - 1000) \\ [0,1] & \text{if } q\sigma_p(3000 - v) = (1 - q)\sigma_l(v - 1000) \\ 0 & \text{if } q\sigma_p(3000 - v) < (1 - q)\sigma_l(v - 1000) \end{cases}
$$

As before, we must find a profile $(\sigma_b, \sigma_p, \sigma_l)$, along with some possible restrictions on q, such that the buyer strategy is a best-response to the strategies of the two types of sellers, and each seller type's strategy is a best response to the buyer's strategy.[11]

- **Pooling Equilibrium.** Both the lemon and the peach are sold with probability 1. For this equilibrium to exist, it must be the case that $v \geq 2000$ for otherwise S would not sell the car. Given this condition, B is only willing to buy if

$$q(3000 - v) \geq (1 - q)(v - 1000)$$
$$2000q + 1000 \geq v \geq 2000$$

  and so there is a price v that works if and only if $2000q + 1000 \geq 2000$, which implies $q \geq 1/2$. We conclude that

$$\left(\sigma_b, \sigma_p, \sigma_l\right) = (1, 1, 1)$$

  is a pooling equilibrium if $v \geq 2000$ and $q \geq (v - 1000)/2000 \geq 1/2$. In other words, if the probability of a peach is sufficiently high, then trade occurs and in equilibrium both the peach and the lemon will be traded.

- **Separating Equilibrium.** Is there an equilibrium where the peach is traded but the lemon is not? No because this requires both $\sigma_b > 0$ and $2000 < v < 0$, which is clearly impossible.

  Is there an equilibrium where only the lemon is traded but the peach is not? We need $0 \leq v \leq 2000$ and $\sigma_b > 0$. B will be willing to buy with positive probability if and only if

$$0 \geq (1 - q)(v - 1000)$$
$$0 \leq (1 - q)(1000 - v)$$

  which only holds for $v \leq 1000$. In other words, for any $v \in [0, 1000]$, there exists an equilibrium where only the lemon is traded. The reason for this is intuitive, if the valuation of the car is so low, then S would only bring the lemon to the market. Given that only the lemon is going to be for sale, B is not willing to spend more than \$1,000.

The results are not very encouraging for the buyer: It is not possible to obtain an equilibrium where only peaches are traded and lemons are not. In the pooling equilibrium both types of cars are sold, and in the separating equilibrium only the lemon is sold. Furthermore, if $q < 1/2$, then only lemons are traded in equilibrium (because the pooling equilibrium disappears). Thus, with asymmetric information markets can sometimes fail.

2 Dynamic Bayesian Games

Recall now the different types of Bayesian equilibria that we studied. In a **pooling equilibrium**, different types of some player take the same actions; in a **separating equilibrium**,

[11] We shall only look at equilibria where trade occurs with positive probability. There is an equilibrium where B buys with probability 0 and both types of S sell with probability 0, but it is not terribly interesting because it relies on knife-edge indifference conditions. We also do not look at semi-separating equilibria because these (if they existed) would require either $v = 0$ or $v = 2000$, which are also too knife-edge to be interesting.

different types take different actions; in a **semi-separating equilibrium**,[12] some subset of types take the same actions, and another subset take different actions.[13]

Among the most frequently studied games are the signaling games, in which an informed player gets to move first (perhaps signaling some of the information it has) and the uninformed player gets to move second, making use of the information revealed from the first stage.[14]

2.1 A Two-Period Reputation Game

There are two firms, ($i = 1, 2$) and in period 1 both are in the market but only firm 1 takes an action a_1 from the action space {Prey, Accommodate}. Firm 2 has profit D_2 if firm 1 accommodates and $P_2 < 0 < D_2$ if firm 1 preys. Firm 1 can be of two types, "sane" which makes D_1 when it accommodates and $P_1 < D_1$ when it preys, but which prefers to be a monopolist with profit $M_1 > D_1$ per period; and "crazy" which prefers predation to everything else (for simplicity, we assume it gets $\hat{P}_1 > M_1$). Let p be the probability that film 1 is sane.

In period 2, only firm 2 chooses an action a_2 from the action space {Stay, Exit}. If firm 2 stays, it obtains D_2 if firm 1 is actually sane and P_2 if firm 1 is actually crazy; if firm 2 exits, it obtains a payoff of 0. The sane firm gets D_1 if firm 2 stays, and M_1 if firm 2 exits. Let δ denote the (common) discount factor between the two periods. The extensive form game is illustrated in Fig. 10 (p. 19).

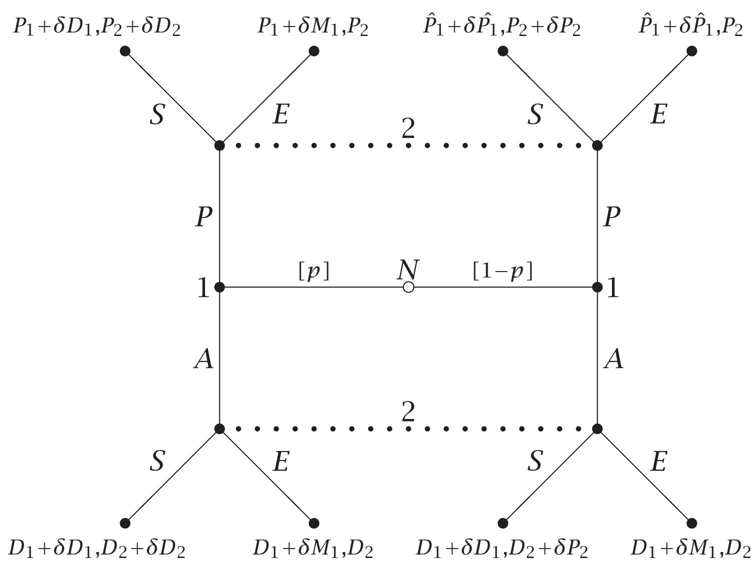

Figure 10: The Reputation Game.

[12] Also sometimes called *partially separating* or *partially pooling*.

[13] Some authors also distinguish *hybrid* strategies, where some types play pure strategies, and other types play mixed strategies. This, of course, results in players sometimes separating and sometimes pooling in equilibrium. Hence, our old terminology that called such equilibria *semi-separating*. When there are more than two types, it may be useful to be a bit more precise.

[14] In contrast, screening games are those where the uninformed player moves first and takes an action that is designed to screen its opponent's type by causing certain types to take one action, and different types to take another. Generally, screening games are much easier to analyze than signaling games. There are also interesting models where both screening and signaling take place (one such model of wartime negotiations is analyzed by yours truly).

Since we have assumed that the crazy type always preys, it is interesting to study the behavior of the sane type. The idea is that although it prefers to accommodate in any period if that period is played by itself, it might want to prey the first period if doing so would induce firm 2 to exit in the second period. That is, can the sane firm 1 behave predatorily to build reputation for a crazy one and increase its payoff in the second period?

Let $q(a_1)$ denote firm 2's posterior (updated) belief that firm 1 is sane after it observes firm 1's action a_1 in the first period. That is, $q(P)$ is the updated belief given predation, and $q(A)$ is the updated belief given accommodation. As before, we study potential perfect Bayesian equilibria by their class.

- **Pooling Equilibria.** In a pooling PBE both types of firm 1 prey, and so firm 2's posterior belief after observing the action is the same as the prior after preying $q(P) = p$, and $q(A) = 1$ after accommodating.[15] Firm 2's payoff if it stays is $p(P_2 + \delta D_2) + (1-p)(P_2 + \delta P_2)$ and if it exits, the payoff is $pP_2 + (1-p)P_2 = P_2$. Firm 2 will exit in the second period if

$$\begin{aligned} P_2 &\geq p(P_2 + \delta D_2) + (1-p)(P_2 + \delta P_2) \\ 0 &\geq \delta[pD_2 + (1-p)P_2] \\ 0 &\geq pD_2 + (1-p)P_2 \end{aligned} \tag{2}$$

That is, if condition (2) is met, then firm 2 prefers to exit. If it exits, firm 1's payoff is $P_1 + \delta M_1$. If firm 1 accommodates, its payoff is $(1+\delta)D_1$ because firm 2 updates to believe it is facing the sane type with certainty, and stays. Hence, the sane firm 1 will prefer to prey in the first period if the following condition is satisfied:

$$\begin{aligned} P_1 + \delta M_1 &\geq (1+\delta)D_1 \\ \delta(M_1 - D_1) &\geq D_1 - P_1 \end{aligned} \tag{3}$$

Therefore, when conditions (2) and (3) are satisfied, the pooling PBE exists.

If condition (2) is satisfied with equality, then firm 2 is indifferent between staying and exiting. Let r be the probability that it exits, and so the sane firm's payoff from preying is $(1-r)(P_1 + \delta D_1) + r(P_1 + \delta M_1) = P_1 + \delta D_1 + r\delta(M_1 - D_1)$. To induce the sane entrant to pray, the probability of exit must be such that

$$\begin{aligned} P_1 + \delta D_1 + r\delta(M_1 - D_1) &\geq (1+\delta)D_1 \\ r\delta(M_1 - D_1) &\geq D_1 - P_1. \end{aligned}$$

And therefore, for any

$$r \geq \frac{D_1 - P_1}{\delta(M_1 - D_1)}$$

the sane firm will prefer to prey, and since condition (2) is satisfied with equality, firm 2 is indifferent between staying and exiting, and so can play the mixed strategy r. There exists a continuum of pooling PBE in this case. However, satisfying (2) with equality is a knife-edge condition and even the slightest perturbation in the parameters would violate it. For this reason, such PBE are usually ignored in applied work.

[15] There is no pooling PBE where both types accommodate because we assumed that the crazy type always preys.

- **Separating Equilibria.** In a separating equilibrium, the sane firm accommodates and the crazy firm preys. Firm 2's updated beliefs are $q(P) = 0$ and $q(A) = 1$, and so it has complete information in the second period. Firm 2's best response is to stay after accommodation and exit after predation. The sane firm's payoff from accommodating is then $(1 + \delta)D_1$ and the payoff from preying is $P_1 + \delta M_1$ given firm 2's strategy. The sane firm will prefer to separate and accommodate if

$$\begin{aligned} (1 + \delta)D_1 &\geq P_1 + \delta M_1 \\ D_1 - P_1 &\geq \delta(M_1 - D_1) \end{aligned} \tag{4}$$

  This condition is, of course, the reverse of (3), which ensured that it would prefer to pool. We conclude that if (4) is satisfied, then the sane firm accommodates, the crazy one preys, firm 2 updates to believe as described, exits if it observes predation, and stays if it observes accommodation in the first period. This is the separating PBE, and condition (4) is both necessary and sufficient for its existence.

- **Semi-separating Equilibria.** In a semi-separating equilibrium, the sane type randomizes and the crazy type preys. Let r_1 denote the probability that the sane type preys. How does firm 2 update its beliefs? Since the crazy type never accommodates, $q(A) = 1$, and so firm 2 updates to believe that firm 1 is sane with probability 1 following accommodation. In this case, firm 2's best response is to stay. On the other hand, if firm 2 observes predation, then its updated belief is $q(P) = pr_1/(pr_1 + 1 - p)$ from Bayes' rule.

  Since the sane type is randomizing, it must be indifferent between preying and accommodating. If it accommodates then, given firm 2's best response, the payoff is $(1+\delta)D_1$. If it preys, then its payoff is $P_1 + \delta D_1 + r\delta(M_1 - D_1)$, as we found above (recall that r is the probability that firm 2 exits after it observes predation). Hence, the sane firm will randomize whenever

$$r = \frac{D_1 - P_1}{\delta(M_1 - D_1)} \tag{5}$$

  For firm 2 to randomize after predation, it must be indifferent between staying and exiting given its updated belief:

$$\begin{aligned} q(P)(P_2 + \delta D_2) + (1 - q(P))(P_2 + \delta P_2) &= P_2 \\ q(P)D_2 + (1 - q(P))P_2 &= 0 \end{aligned} \tag{6}$$

  Thus, if the sane firm preys with probability r_1 such that condition (6) is satisfied, then firm 2 will be indifferent between staying and exiting after predation, and in particular can play the mixed strategy where it exits with probability r defined in condition (5), which in turn makes the sane firm 1 indifferent. Therefore, if both conditions (5) and (6) are satisfied, the strategies and beliefs described above constitute a semi-separating PBE.

Although there are different types of PBE in this game, it is *not* the case that we are dealing with multiple equilibria (except the special knife-edge case of continuum of pooling PBE). For each set of different values of the exogenously specified parameters, the model has a (generically) unique PBE. This is due to the restrictive assumption that the crazy firm always preys, which results in two facts: (i) predation is never a zero-probability event, and (ii) accommodation reveals firm 1's type with certainty. In more interesting (and more common) models, this will not be the case, as we see in the next example.

2.2 Spence's Education Game

Now that you are in graduate school, you probably have a good reason to think education is important.[16] Although I firmly believe that education has intrinsic value, it would be stupid to deny that it also has economic, or instrumental, value as well. As a matter of fact, I am willing to bet that the majority of students go to college not for the sake of knowledge and bettering themselves, but because they think that without the skills, or at least the little piece of paper they get at the end of four years, they will not have good chances of finding a decent job. The idea is that potential employers do not know you, and will therefore look for some signals about your potential to be a productive worker. A university diploma, acquired after meeting rigorous formal requirements, is such a signal and may tell the employer that you are intelligent and well-trained. Employers will not only be more willing to hire such a person, but will probably pay premium to get him/her. According to this view, instead of making people smart, education exists to help smart people prove that they are smart by forcing the stupid ones to drop out.[17]

The following simple model is based on Spence's (1973) seminal contribution that preceded the literature on signaling games and even the definition of equilibrium concepts like PBE. There are two types of workers, a high ability (H) and a low ability (L) type. The prior probability of having high ability is $p \in (0,1)$ and after Nature selects the type, the worker learns it and can choose a level of education $e \geq 0$ before applying for a job. The cost of obtaining an educational level e is e for the low ability worker, and $e/2$ for the high ability worker. (In other words high ability workers find education much less costly.)

The only thing the employer observes is the level of education. The employer offers a wage $w(e)$ as a function of the educational level, and the employers' payoff is $2 - w(e)$ if the worker turns out to have high ability, and $1 - w(e)$ if he turns out to have low ability. The worker's payoffs are $w(e) - e/2$ if he is the H type, and $w(e) - e$ if he is the L type. Since the job market is competitive, the employer must offer a competitive wage such that the expected profit is zero.

Let $\mu(H|e)$ denote the belief of the employer that the worker has high ability given it observes a level of education e. Because of the competitiveness requirement, $(2-w(e))\mu(H|e) + (1-w(e))(1-\mu(H|e)) = 0$, which implies that the wage schedule must satisfy $w(e) = 2\mu(H|e) + (1-\mu(H|e))$. Let e_H be the level of education chosen by the H type, and e_L be the level of education chosen by the L type. We want to find the set of PBE of this game.

- **Pooling Equilibria.** In these PBE, $e_H = e_L = e^*$, and Bayes' rule gives $\mu(H|e^*) = p$, and $\mu(L|e^*) = 1 - p$ because the employer learns nothing. The wage offered then is

[16] Or maybe not. I went to graduate school because I really did not want to work a regular job from 8:00a to 5:00p, did not want to be paid for writing programs (my B.S. is in Computer Science) even if meant making over 100k, and did not want to have a boss telling me what to do. I had no training in Political Science whatsoever, and so (naturally) decided it would be worth a try. Here I am now, several years later, working a job from 7:00a to 11:00p including weekends, making significantly less money, and although without a boss, having to deal with a huge government bureaucracy. Was this economically stupid? Sure. Am I happy? You betcha. Where else do you get paid to read books, think great thoughts, and corrupt the youth?

[17] Here, perhaps, is one reason why Universities that are generally regarded better academically tend to attract smart students, who then go on to earn big bucks. They make the screening process more difficult, and so the ones that survive it are truly exceptional... maybe.

$w(e^*) = 2p + 1 - p = 1 + p$, and so the worker's payoffs are

$$u(w, e^*|H) = 1 + p - e^*/2$$
$$u(w, e^*|L) = 1 + p - e^*.$$

Observing a level of education $e \neq e^*$ is a zero-probability event, and so we must assign some beliefs to the employer for this case before we can proceed. The easiest thing to do is to assign pessimistic beliefs such that $\mu(H|e) = 0$ whenever $e \neq e^*$, and so the employer updates to believe the worker has a low ability if it observes any level of eduction other than e^*. This minimizes the temptation to deviate for both players. Of course, if the employer believes it is hiring the L type, $w(e) = 1$, and the worker's payoffs are $1 - e/2$ if type H, and $1 - e$ if type L. Since choosing e^* must be at least as good as choosing any other e for both types, we have

$$1 + p - e^*/2 \geq 1 - e/2$$
$$1 + p - e^* \geq 1 - e$$

for all $e \geq 0$. These inequalities are satisfied whenever $e^* \leq p$. Therefore, any $e^* \leq p$ can be supported in equilibrium by using a belief system similar to the one above:

$$\mu(H|e) = \begin{cases} p & \text{if } e = e^* \\ 0 & \text{if } e \neq e^* \end{cases}$$

Therefore, there is a continuum of pooling equilibria.

- **Separating Equilibria.** In these PBE, $e_H \neq e_L$. From Bayes' rule, $\mu(H|e_H) = 1$ and $\mu(H|e_L) = 0$, and so we have $w(e_H) = 2$ and $w(e_L) = 1$. Given that, type L worker chooses $e_L = 0$ because anything else would leave him strictly worse off. In equilibrium, it must not have an incentive to mimic type H's action, so type H must choose $e_H > 0$ to prevent that (and vice versa, it must also not be the case that type H wants to mimic type L's behavior). We have two conditions:

  $$2 - e_H/2 \geq 1$$
  $$1 \geq 2 - e_H$$

  The first prevents the H type from choosing education level of 0 (the L type's action), and the second prevents the L type from choosing e_H (the H type's action). We therefore have $e_H \in [1, 2]$. The following belief system can support any e_H in this interval:

  $$\mu(H|e) = \begin{cases} 0 & \text{if } e < e_H \\ 1 & \text{if } e \geq e_H \end{cases}$$

  That is, if the employer observes any level of education less than e_H, it updates to believe that the worker's type is L, and if it observes any level greater than or equal to e_H, it updates to believe that the worker's type is H. Note again that we have assigned beliefs for probability zero events ourselves because in the separating PBE, any level of education except 0 and e_H is off the equilibrium path.

  Therefore, there is a continuum of separating equilibria in which the low ability worker chooses educational level 0, and the high ability worker chooses level $e_H \in [1, 2]$, and the employer uses the belief system specified above. In this case, as the case with pooling PBE, the leeway in specifying off-the-path beliefs produces a multiplicity of equilibria.

- **Semi-separating Equilibria.** I leave these to you as an exercise.

There is a continuum of separating equilibria, and so we cannot say anything about the level of education a worker with high ability will choose other than it will be high enough to deter a low-ability worker from getting it. The punch line of the model, however, is clear and captures the idea of education as a signaling device.

There are two substantive insights you should take away from the result. First, the only way for a high ability worker to get the high-paying job he deserves is to signal his type by investing in costly education. Otherwise, the employer will treat him as a low-ability worker. This corresponds quite well to the empirical observation that workers with more years of schooling on the average tend to earn higher wages.

Second, the value of education as a signaling device depends not on the skills that workers receive through it, but on the costs they have to pay to acquire it. The critical insight here is that for education to be useful as a signalling device, it is sufficient that education is costlier for the low ability type to acquire. It does not matter if education really has any value added as long as it is less costly for the high-ability type.

3 Computing Perfect Bayesian Equilibria

We now look at several examples of how we can characterize PBE in extensive form games.

3.1 The Yildiz Game

Consider the game in Fig. 11 (p. 25). Backward induction on player 1's actions at his two penultimate information sets 1.3 and 1.4 tells us that in any PBE he must be choosing e and h respectively. Furthermore, at 1.2 he must be choosing d because doing so would yield a strictly higher payoff (of 0) no matter what player 2 does. Hence, in any PBE, player 1's strategy must specify playing d at 1.2, e at 1.3, and h at 1.4, each with probability 1.

Let x denote player 2's posterior belief that she is at the lower node in her information set. Suppose that player 1 also chooses a with certainty. In this case, Bayes rule would pin down player 2's belief to $x = 1$, in which case she would certainly choose L. But if she chooses L at her information set, then player 1 could do strictly better by choosing b instead of a at his information set 1.1, and therefore it cannot be the case that in PBE he would choose a with certainty.

Suppose now that player 1 chose b with certainty at 1.1. In this case, Bayes rule pins down player 2's belief to $x = .1/(.1 + .9) = .1$ (intuitively, she can learn nothing new for player 1's action). Given player 1's sequentially rational strategy at his last information sets, the expected payoff from choosing L then is $(.1)(2) + (.9)(2) = 2$, and the expected payoff from choosing R then is $(.1)(-5) + (.9)(3) = 2.2$. Hence, the only sequentially rational strategy for player 2 would be to choose R with certainty. However, if she chooses R for sure, then player 1 can do better by playing a at the information set 1.1 because this would give him a payoff of 4, which is strictly better than the payoff of 3 he would get from playing b for sure. Therefore, it cannot be the case that in PBE he would choose b with certainty.

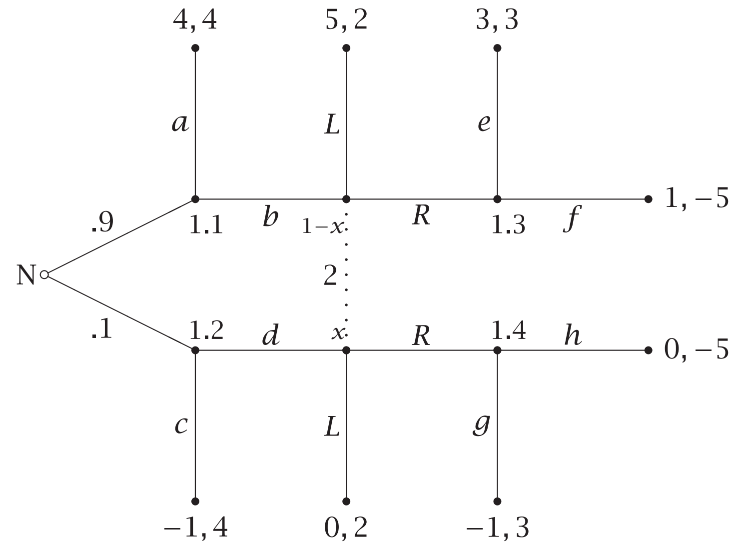

Figure 11: A Game from Muhamet Yildiz's Notes.

We conclude that in equilibrium player 1 must be mixing at information set 1.1. Let p denote the probability with which he chooses b, and let q denote the probability with which player 2 chooses R. Because player 1 is willing to mix, it follows that the expected payoff from choosing a must be the same as the expected payoff from choosing b, or $4 = q(3)+(1-q)(5)$, which gives $q = .5$. That is, because player 1 is mixing in equilibrium, it must be the case that player 2 is mixing as well.

But for player 2 to be willing to mix, it must be the case that she is indifferent between choosing L and R at her information set. That is, the expected payoff from L must equal the expected payoff from R, or $x(2) + (1 - x)(2) = x(-5) + (1 - x)(3)$, which gives $x = 1/8$. Only if her posterior belief is exactly $1/8$ would she be willing to mix.

From Bayes rule, $x = (.1)(1)/[(.1)(1) + (.9)p]$, and hence player 1 must choose p such that $x = 1/8$. Solving the equation yields the correct value for $p = 7/9$, and so this must be the equilibrium mixing probability for player 1 at 1.1. We conclude that the game has a unique perfect Bayesian equilibrium in the following strategies:

- Player 1 chooses b with probability $7/9$ at 1.1, and chooses with certainty d at 1.2, e at 1.3, and h at 1.4;
- Player 2 chooses R with probability $1/2$.

Player 2's beliefs at her information set are updated according to Bayes rule to $x = 1/8$. The strategies are sequentially rational given the beliefs and beliefs are consistent with the strategies. Hence, we have a PBE.

3.2 The Myerson-Rosenthal Game

This makes the previous example a bit more complicated. In the Yildiz Game, player knew the move by Nature and could disclose some of his information to player 2 by his actions. Because they have somewhat conflicting interests, player 1 obfuscates a bit by playing a mixed strategy, which prevents player 2 from inferring with certainty what he knows. In the next game, player 1 would have to make his choices without knowing the move by Nature.

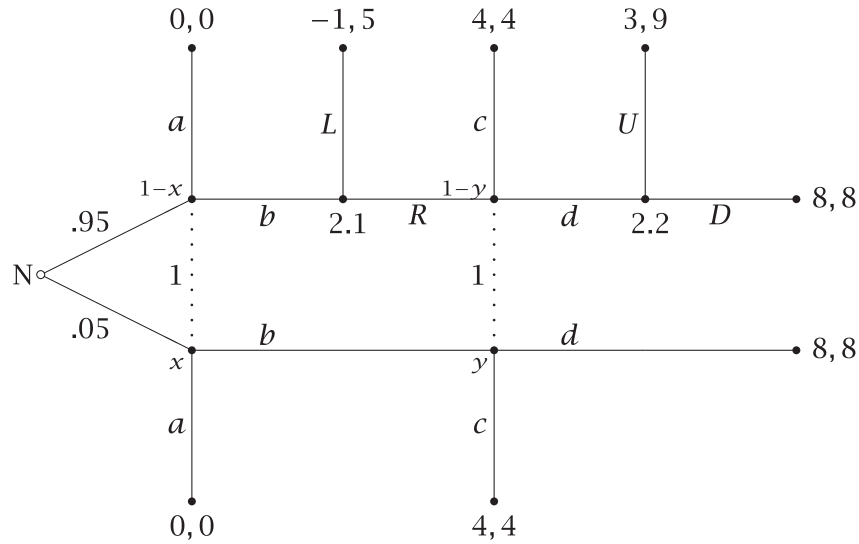

Figure 12: The Myerson-Rosenthal Game.

The game is depicted in Fig. 12 (p. 26). The interpretation is as follows. After one chance event (probability .95), players take turns being either generous or selfish until someone is selfish or until both have been generous twice. Each player loses \$1 by being generous, but gains \$5 each time the other player is generous. (So actions a, L, c, and U are selfish, but b, R, and d, and D are generous.) After another chance event (probability .05), player 2 is incapable of being selfish. Player 1 does not observe the chance outcome. This game could be interpreted as player 1 being uncertain about whether player 2 is the virtuous kind whose integrity would compel her reciprocate any act of generosity or whether she is capable of being selfish as well.

As before, we label player 1's posterior beliefs in some convenient manner. Here, let x denote the probability with which he believes that he is at the lower node in his first information set, and let y denote the probability with which he believe he is at the lower node in his second information set.

Clearly, $x = .05$ is the only belief consistent with the probability distribution of the chance events. Moreover, at her last information set 2.2, player 2's only sequentially rational choice is U. We now have to find the rest of the strategies and beliefs.

Suppose player 1 chose d with certainty in equilibrium. The only way this would be sequentially rational is if the expected payoff from c did not exceed the expected payoff from d given 2's sequentially rational strategy, or if $4 \le 8y + 3(1-y)$, which requires $y \ge 1/5$. Because 1 is choosing d for sure, player 2's expected payoff from choosing R at 2.1 is 9, which is strictly better than getting 5 by playing L, and so she would certainly choose R. Given that she would choose R, player 1's expected payoff from choosing b at his first information set would be $(.05)(8) + (.95)(3) = 3.25$, which is strictly greater than 0, which is what he would get by playing a. Therefore, he would choose b for sure. But this means that player 1's second information set is now along the path of play, and Bayes rule gives

$$y = \frac{(.05)(1)}{(.05)(1) + (.95)(1)} = .05 < 1/5,$$

which contradicts the necessary condition that makes playing d with certainty sequentially rational. Therefore, there cannot be a PBE where player 1 chooses d with certainty.

Suppose player 1 chose c with certainty in equilibrium. The only way this could be sequentially rational is (by reversing the inequality in the previous paragraph) if $y \leq 1/5$. Because 1 is playing c for sure, player 2 would certainly choose L at 2.1 because the expected payoff is strictly greater. Given her sequentially rational strategy, choosing b would yield player 1 the expected payoff of $4x + (-1)(1 - x) = 5x - 1 = -0.75$. Hence, the sequentially rational choice at this information set is a. This leaves player 1's second information set off the path of play, so Bayes rule cannot pin down the beliefs there. In this case, we are free to assign any beliefs, and in particular we can assign some $y \leq 1/5$. We have therefore found a continuum of PBE in this game:

- Player 1 chooses a and c with certainty at the respective information sets; if he ever finds himself at his second information set, his belief is $y \leq 1/5$;
- Player 2 chooses L at 2.1 and U at 2.2.

The reason we have a continuum of PBE is because there is an infinite number of beliefs that satisfy the requirement. However, all these PBE are equivalent in a very important sense: they predict the same equilibrium path of play, and they only differ in beliefs following zero-probability events.

This may be a bit disconcerting in the sense that this equilibrium seems to require unreasonable beliefs by player 1. Here's why. Suppose there is an extremely small probability $\epsilon > 0$ that player 1 makes a mistake at his first information set and plays b instead of a. Then, using Bayes rule his posterior belief would have to be $y = (.05)\epsilon/[(.05)\epsilon + (.95)\epsilon\sigma_2(R)] = 1$ because the only way to get to player 1's second information set would be from the lower node at his first information set (recall that player 2 chooses L, and so $\sigma_2(R) = 0$). Note that this is true regardless of how small ϵ we take. But $y = 1$ contradicts the requirement that $y \leq 1/5$. In other words, it does not seem reasonable for player 1 to hold such beliefs because even the slightest error would require $y = 1$.

The PBE solution concept is too weak to pick out this problem. The stronger solution concept of **sequential equilibrium** will eliminate all of the above PBE that require these unreasonable beliefs. Intuitively, sequential equilibrium simply formalizes the argument from the previous paragraph. Instead of requiring that beliefs are consistent along the equilibrium path only, it requires that they are *fully consistent*: that is, that they are consistent for slightly perturbed behavior strategies that reach all information sets with positive probability (and so Bayes rule would pin beliefs down everywhere). A belief vector π is fully consistent with a strategy σ if, and only if, there exist behavior strategy profiles that are arbitrarily close to σ and that visit all information sets with positive probability, such that the beliefs vectors that satisfy Bayes rule for these profiles are arbitrarily close to π.

Sequential equilibria are therefore a subset of the perfect Bayesian equilibria and, more importantly, always exist. Unfortunately, they can be quite difficult to compute because checking full consistency requires finding the limits of systems of beliefs in sequences of games in which the perturbed behavior strategies converge to the strategies under consideration. We will not cover sequential equilibria in this class (but you should definitely invest the time to study them because they are quite important). However, let's see how the idea of full consistency would eliminate the PBE we just found.

The posterior belief y is given by:

$$y = \frac{(.05)\sigma_1(b)}{(.05)\sigma_1(b) + (.95)\sigma_1(b)\sigma_2(R)} = \frac{1}{1 + 19\sigma_2(R)},$$

where the latter inequality would have to hold even when $\sigma_1(b) = 0$ because it would hold for any slightly perturbed behavior strategies with $\sigma_1(b) > 0$. Returning to our solution, the requirement that $y \leq 1/5$ then translates into:

$$\frac{1}{1 + 19\sigma_2(R)} \leq \frac{1}{5} \Leftrightarrow \sigma_2(R) \geq \frac{4}{19}.$$

However, as we have seen, player 2's only sequentially rational strategy is to play L with certainty, and so $\sigma_2(R) = 0$, which contradicts this requirement. Hence, no beliefs $y \leq 1/5$ are fully consistent, and therefore none of these PBE are sequential equilibria.

Finally, we turn to the possibility that player 1 mixes at his second information set in equilibrium. Since he is willing to randomize, he must be indifferent between his two actions, or in other words, $8y + 3(1 - y) = 4$ which yields $y = 1/5$. As we have seen already,

$$y = \frac{1}{1 + 19\sigma_2(R)} = \frac{1}{5} \Leftrightarrow \sigma_2(R) = \frac{4}{19}.$$

This is the full consistency requirement that must also hold in PBE for any $\sigma_1(b) > 0$. If player 2 is willing to randomize, she must be indifferent between her two actions: $5 = 4\sigma_1(c) + 9(1 - \sigma_1(c))$, which implies that $\sigma_1(c) = 4/5$. Turning now to player 1's move at his first information set, choosing b would yield an expected payoff of

$$\begin{aligned}(.95)\Big[(-1)(1 - \sigma_2(R)) + (4\sigma_1(c) + 3(1 - \sigma_1(c)))\sigma_2(R)\Big] + (.05)\Big[4\sigma_1(c) + 8(1 - \sigma_1(c))\Big] \\ = (.95)\Big[(-1)(15/19) + (4(4/5) + 3(1/5))(4/19)\Big] + (.05)\Big[4(4/5) + 8(1/5)\Big] = 0.25.\end{aligned}$$

Because this expected payoff is strictly greater than 0, which is what player 1 would get if he chose a, sequential rationality requires that he chooses b with certainty.

We conclude that the following strategies and beliefs constitute a perfect Bayesian (and the unique sequential) equilibrium:

- Player 1 chooses b with probability 1, and c with probability $4/5$;
- Player 2 chooses R with probability $4/19$, and U with probability 1;
- Player 1 updates to believe $x = .05$, and $y = 1/5$.

Substantively, this solution tells us that player 1 must begin the game by being generous. Small amounts of doubt can have significant impacts on how rational players behave. If player 1 were sure about 2's capacity for being selfish, then perpetual selfishness would be the only equilibrium outcome. If, however, it is common knowledge that player 2 is generous by disposition, the result is different. Even when player 1 attaches a very small probability to this event, he must be generous at least once because this would encourage 2 to reciprocate even if she can be selfish. The reason player 2 would reciprocate with higher probability is because she wants player 1 to update his beliefs to an even higher probability that she is incapable of being selfish, which would induce him to be generous the second time around, at which point player 2 would defect and reap her highest payoff of 9. Notice how in this PBE player 1's posterior belief went from .05 up to $y = .20$. Of course, player 2 would not want to try to manipulate player 1's beliefs unless there was an initial small amount of uncertainty that would cause player 1 to doubt her capacity for being selfish.

3.3 One-Period Sequential Bargaining

There are two players, a seller S and a buyer B. The buyer receives a gift v, but the seller does not know its exact amount. He believes that it is $v = \$20$ with probability π, and $v = \$10$ with probability $1 - \pi$. The seller sets the price $p \geq 0$ for a product that the buyer wants to get at the cheapest price possible. After observing the price, B either buys, yielding the payoff vector $(p, v - p)$, or does not, yielding $(0, 0)$. The game is shown in Fig. 13 (p. 29).

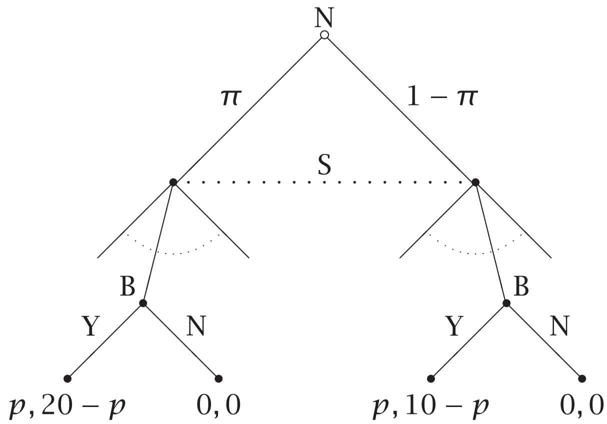

Figure 13: The One-Period Bargaining Game.

Player B would accept any $p \leq 20$ at her left information set (that is, if she received \$20) and would accept any $p \leq 10$ at her right information set (that is, if she received \$10). In other words, B buys iff $v \geq p$. This means that if S sets the price at $p = \$10$, then he is sure to sell the product and get a payoff of 10. If he sets the price at $10 < p \leq 20$, then B would only buy if she had \$20, in which case the seller's expected payoff is πp. Finally, the seller's payoff for any $p > 20$ is zero because B would never buy.

Consequently, the seller would never ask for more than \$20 or less than \$10 in equilibrium. What is he going to ask for then? The choice is between offering \$10 (which is the maximum a poor B would accept) and something the rich B would accept. Because any $p > 10$ will be rejected by the poor B, the seller would not ask for less than \$20, which is the maximum that the rich B would accept. Hence, the seller's choice is really between offering \$10 and \$20. When would he offer \$20?

The expected payoff from this offer is 20π, and the expected payoff from \$10 is 10 (because it is always accepted). Therefore, the seller would ask for \$20 whenever $20\pi \geq 10$, or $\pi \geq 1/2$. In other words, if S is sufficiently optimistic about the amount of money the buyer has, he will set the price at the ceiling. If, on the other hand, he is pessimistic about the prospect, he would set the price at its lowest. The seller is indifferent at $\pi = 1/2$.

3.4 A Three-Player Game

Let's try the game with three players shown in Fig. 14 (p. 30). This is a slightly modified version of a game in notes by David Myatt.

Player 3's expected payoff from choosing a is $4x + 0(1 - x) = 4x$, and his expected payoff

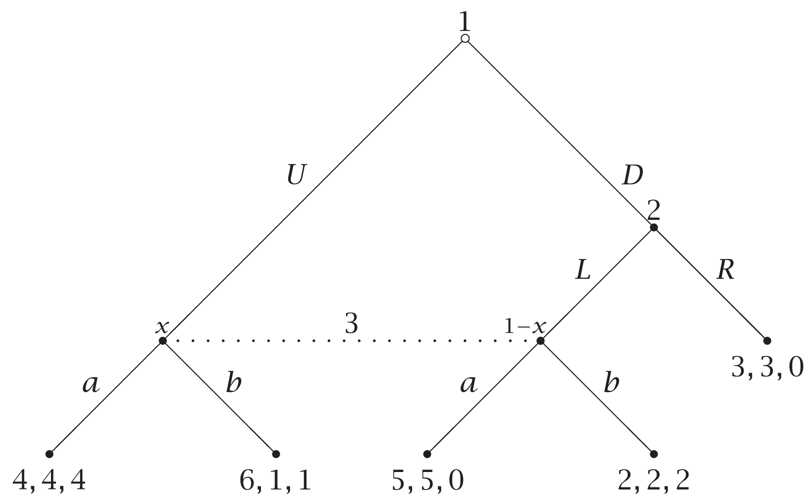

Figure 14: The Three Player Game.

from choosing b is $x + 2(1 - x) = 2 - x$. The sequentially rational best response is:

$$\sigma_3(a) = \begin{cases} 1 & \text{if } x > 2/5 \\ 0 & \text{if } x < 2/5 \\ [0,1] & \text{otherwise.} \end{cases}$$

Suppose then that $x > 2/5$, and so player 3 is sure to choose a at his information set. In this case, player 2 would strictly prefer to choose L, and given this strategy, player 1's optimal choice is D. Given these strategies, Bayes rule pins down $x = 0$, which contradicts the requirement that $x > 2/5$. Hence, there is no such PBE.

Suppose now that $x < 2/5$, and so player 3 is sure to choose b at his information set. In this case, player 2 strictly prefers to choose R. Given her strategy, player 1's best response would be U. In this case, Bayes rule pins down $x = 1$, which contradicts the requirement that $x < 2/5$. Hence, there is no such PBE.

We conclude that in PBE, $x = 2/5$, and so player 3 would be willing to mix. Player 2's expected payoff from L would then be $5\sigma_3(a) + 2(1 - \sigma_3(a)) = 3\sigma_3(a) + 2$, and her payoff from R is 3. Hence, her best response would be:

$$\sigma_2(L) = \begin{cases} 1 & \text{if } \sigma_3(a) > 1/3 \\ 0 & \text{if } \sigma_3(a) < 1/3 \\ [0,1] & \text{otherwise.} \end{cases}$$

Suppose then that $\sigma_3(a) > 1/3$, and so she would choose L for sure. In this case, player 1's expected payoff from U is $4\sigma_3(a) + 6(1 - \sigma_3(a)) = 6 - 2\sigma_3(a)$. His expected payoff from D would be $5\sigma_3(a) + 2(1 - \sigma_3(a)) = 2 + 3\sigma_3(a)$. He would therefore choose U if $\sigma_3(a) < 4/5$, would choose D otherwise, and would be indifferent when $\sigma_3(a) = 4/5$. However if he chooses D for sure, then Bayes rule pins down $x = 0$, which contradicts $x = 2/5$. Similarly, if he chooses U for sure, Bayes rule pins down $x = 1$, which is also a contradiction. Therefore, he must be mixing, which implies that $\sigma_3(a) = 4/5 > 1/3$, and so player 2's strategy is sequentially rational. What is the mixing probability? It must be such that $x = 2/5$, which implies that $\sigma_1(U) = {}^2/_5$. We conclude that the following strategies and beliefs constitute a perfect Bayesian equilibrium:

- Player 1 chooses U with probability $^2\!/_5$
- Player 2 chooses L with probability 1
- Player 3 chooses a with probability $^4\!/_5$, and updates to believe $x = {}^2/_5$.

Suppose now that $\sigma_3(a) < {}^1/_3$, and so player 2 would choose R for sure. In this case, player 1's expected payoff from D is 3, which means that he would choose U if $6 - 2\sigma_3(a) > 3$. But since $\sigma_3(a)$ can at most equal 1, this condition is always satisfied, and therefore player 1 would always choose U. In this case, Bayes rule pins down $x = 1$, which contradicts the requirement that $x = {}^2/_5$. Hence, there can be no such PBE.

Finally, suppose that $\sigma_3(a) = {}^1/_3$, and so player 2 is indifferent between her two actions. Player 1's expected payoff from D in this case would be:

$$3(1 - \sigma_2(L)) + \sigma_2(L)\Big[5({}^1/_3) + 2({}^2/_3)\Big] = 3.$$

As we have seen already, in this case he would strictly prefer to choose U. But in this case, Bayes rule pins down $x = 1$, which contradicts the requirement that $x = {}^2/_5$. Hence, no such PBE exists. We conclude that the PBE identified in the preceding paragraph is the unique solution to this game.

3.5 Rationalist Explanation for War

Two players bargain over the division of territory represented by the interval $[0, 1]$. Think of 0 as player 1's capital and 1 as player 2's capital. Each player prefers to get a larger share of territory measured in terms of distance from his capital. Assume that players are risk-neutral, and so the utilities of a division $x \in [0, 1]$ are $u_1(x) = x$ and $u_2(x) = 1 - x$, respectively.

The structure of the game is as follows. Nature draws the the war costs of player 2, c_2, from a uniform distribution over the interval $[0, 1]$. Player 2 observes her costs but player 1 does not. The war costs of player 1, $c_1 \in [0, 1]$, are common knowledge. Player 1 makes a demand $x \in [0, 1]$, which player 2 can either accept or reject by going to war. If she goes to war, player 1 will prevail with probability $p \in (0, 1)$. The player who wins the war, gets his most preferred outcome.

We begin by calculating the expected utility of war for both players:

$$\begin{aligned} U_1(\text{War}) &= pu_1(1) + (1 - p)u_1(0) - c_1 = p - c_1 \\ U_2(\text{War}) &= pu_2(1) + (1 - p)u_2(0) - c_2 = 1 - p - c_2. \end{aligned}$$

Before we find the PBE of this game, let's see what would happen under complete information. Player 1 will never offer anything less than what he expects to get with fighting, and hence any offer that he would agree to must be $x \geq p - c_1$. Similarly, player 2 will never accept anything less than what she expects to get with fighting, and hence any offer that she would agree to must be $1 - x \geq 1 - p - c_2$, or $x \leq p + c_2$. Hence, the set of offers that *both* prefer to war is $[p - c_1, p + c_2]$. Because costs of war are non-negative, this interval always exists. In other words, *there always exists a negotiated settlement that both players strictly prefer to going to war.* With complete information, war will never occur in equilibrium in this model.

What happens with asymmetric information? Since player 2 knows her cost when the offer is made, we can figure out what offers she will accept and what offers she will reject. Accepting an offer x yields her a payoff of $1 - x$, while rejecting it yields her a payoff of $1 - p - c_2$. She will therefore accept an offer if, and only if, $1 - x \geq 1 - p - c_2$, or, in terms of the costs, if

$$c_2 \geq x - p.$$

Player 1 does not know what c_2 is, but knows the distribution from which it is drawn. From his perspective, the choice boils down to making an offer and risk getting it rejected. Given player 2's sequentially rational strategy, from player 1's perspective the probability that an offer x is accepted is the probability that $c_2 \geq x - p$, or, given the uniform assumption,

$$\Pr(c_2 \geq x - p) = 1 - \Pr(c_2 < x - p) = 1 - x + p.$$

Hence, if player 1 makes an offer x, it will be accepted with probability $1 - x + p$, in which case he would obtain a payoff of x, and it will be rejected with probability $1 - 1 + x - p = x - p$, in which case he would obtain an expected payoff of $p - c_1$. The expected utility from offering x is therefore:

$$U_1(x) = (1 - x + p)(x) + (x - p)(p - c_1).$$

Player 1 will choose x that maximizes his expected utility:

$$\frac{\partial U_1(x)}{\partial x} = 1 - 2x + 2p - c_1 = 0 \Leftrightarrow x^* = \frac{1 + 2p - c_1}{2}.$$

The perfect Bayesian equilibrium is as follows:

- Player 1 offers $\min\{\max\{0, x^*\}, 1\}$.
- Player 2 accepts all offers $x \leq c_2 - p$, and rejects all others.

In the PBE, the *ex ante* risk of war is $x^* - p = \frac{1-c_1}{2} > 0$ as long as $c_1 < 1$. In other words, the risk of war is always strictly positive. This contrasts the complete information case where the equilibrium probability of war is zero. Hence, this model provides an explanation of how rational players can end up in a costly war. This is the well-known *risk-return* trade off: player 1 balances the risk of having an offer rejected against the benefits of offering to keep for himself slightly more. This result persists in models with richer bargaining protocols, where pre-play communication is allowed, and even where players can intermittently fight.