

Analytic Theory II: Dynamic Games with Incomplete Information

Branislav L. Slantchev

Department of Political Science, University of California – San Diego

March 11, 2020

Contents

1	The Building Blocks	3
1.1	Sequential Rationality	3
1.2	Consistent Beliefs	5
1.3	Weak Perfect Bayesian Equilibrium	6
2	Perfect Bayesian Equilibrium	11
2.1	Myerson’s Card Game	11
2.2	Classifying PBE by Strategy Types	13
2.3	A Two-Period Reputation Game	14
2.4	Spence’s Education Game	17
3	Computing Perfect Bayesian Equilibria	22
3.1	The Yildiz Game	22
3.2	The Myerson-Rosenthal Game	24
3.3	One-Period Sequential Bargaining	27
3.4	A Three-Player Game	28
3.5	Rationalist Explanation for War	29
3.6	The Perjury Trap Game	30

The Nash equilibrium definition of rationality is too weak to pick out moves that fail to satisfy its requirements when these moves occur off the path of play. This is not a problem in strategic-form games, where the dynamic element is lost and as a result players cannot condition their behavior on their knowledge about what has happened in the game by the time they have to act. In extensive-form games that preserve the idea of sequences of actions, this weakness of Nash equilibrium is a serious problem because it sometimes rationalizes behaviors along the path of play by expectations that involve non-rational behavior off the path of play. Since the contingencies where such behaviors occur do not arise with positive probability if the equilibrium strategies are followed, the Nash requirement of best-responding while holding the other players' strategies fixed cannot eliminate such strategies (in a sense, Nash equilibrium cannot evaluate the optimality of a player's behavior in contingencies it cannot "see" under the strategies of the other players). We resolved this by strengthening the definition of rationality to require that strategies remain best-responses in these contingencies as well (subgame-perfection).

Since static games of incomplete information also ignore the notion of sequencing, Bayesian Nash equilibrium (which is really just Nash equilibrium in the Harsanyi-transformed game that represents the incomplete information) will similarly fail to identify the analogous suboptimal behavior if the dynamic element is introduced. This means that we need to provide the analogue to subgame perfection for games of incomplete information: a definition of rationality that takes into account what players know when it is their turn to move. The problem now is that after the Harsanyi transformation that converts the game into one of imperfect information, we are usually left with very few, if any, proper subgames to analyze because of the non-singleton information sets after Nature's "move". In this setting, subgame perfection will usually have no bite because it does not allow us to split information sets. The strengthening of our definition of rationality will do just that: enable us to evaluate the expected payoffs of various actions at a given information set using the knowledge that (a) this information set has been reached, and (b) players must form some beliefs about which node at that information set could have been reached. As before, this definition will continue to be Nash equilibrium — we are just eliminating Nash equilibria that involve unreasonable behavior from the set of admissible solutions.

As a motivating example of how subgame perfection might not help us, consider the complete-information game in Fig. 1 (p. 2). Player 1 gets to choose between U , M , or D . If he chooses D , the game ends. If he chooses either U or M , player 2 gets to choose between L and R without knowing what action player 1 has taken except that it was not D .




Figure 1: The SPE Example Game.

What are the Nash equilibria? By inspection of the strategic form, we see that there are two Nash equilibria in pure strategies, $\langle U, L \rangle$ and $\langle D, R \rangle$. However, there is something unsatisfying about the second one. In the $\langle D, R \rangle$ equilibrium, player 2 seems to behave irrationally. If her information set is ever reached, playing L strictly dominates R . So player 1 should not be induced to choose D by player 2's incredible threat to choose L . However since player 2's information set is off the equilibrium path, Nash equilibrium does not evaluate the optimality of play there. Moreover, neither does SPE: player 2's information set is not a singleton. The game has only one subgame—the entire game itself—and as a result, all Nash

equilibria, including $\langle D, R \rangle$, are subgame perfect. We need an approach that evaluates the optimality of player 2's response at her information set.

1 The Building Blocks

The approach we are going to take involves two requirements that, taken jointly, will allow us to evaluate the optimality of behavior at any information set, including those that are not singletons. The first step is to observe that in a dynamic setting, each player must have some beliefs about what has transpired in the game at the time that it is their turn to choose an action. We shall evaluate the consequences of various actions in terms of these beliefs. The second step is to note that such beliefs cannot be arbitrary—they must be consistent with what the player thinks the strategies of the other players are. In other words, all players must be best-responding to the beliefs about the other players' strategies. In a sense, this is exactly what Nash equilibrium does because we can interpret the strategy profiles as sets of expectations (beliefs) about behaviors: all players are then best-responding to what they believe the other players are going to do, and these beliefs are predicated on the assumption that the other players are best-responding to their beliefs. In subgame perfection, we required players to ask themselves what they would do in every possible contingency. For contingencies that could not be reached in some strategy profile, players performed a thought experiment by pretending that the contingency has arisen. In effect, they conditioned their behavior on the knowledge of what had happened in the game so far. Conceptually, then, we are not introducing anything new to our definition of rationality. We merely need the notational apparatus that allows us to extend the implications of our concept of rationality to situations where the existing approaches leave us without a way to apply it.

1.1 Sequential Rationality

Consider any dynamic game and an arbitrary information set for player i , h_t . When it is player i 's turn to move, he *knows* that the game has reached h_t (as a fact). If this information set is a singleton, then he also knows precisely what all preceding actions for all players must have been. That is, he also *knows* the history of the game leading to h_t . It is this knowledge that the definition of subgame perfection operates on because it allows one to start a new game from that information set and apply the Nash equilibrium requirements to it. If the information set is not a singleton, however, then this technique does not work. Instead of *knowing* the history of the game that has led to h_t , player i can only have some idea about how his information set was reached; that is, he will have some **belief** about the past actions of all players. For now, we are not going to examine where these beliefs come from; we are just asserting that they must exist. Thus, by analogue with subgame perfection, we are going to imagine that the contingency has arisen, but unlike subgame perfection, we are not going to require that players know precisely how this has happened.

Going back to the game in Fig. 1 (p. 2), player 2 may not know whether player 1 has played U or M by the time her information set is reached, but she will have some belief about which of these two he could have played. That is, the fact that her information set has been reached forces player 2 to consider it a fact that player 1 must have played either U or M because these are the only two actions that could have possibly led to this contingency arising. As usual, we represent beliefs about events with probability distributions. In this case, let $\mu \in [0, 1]$ be the probability that player 1 chose U , and $1 - \mu$ be the probability that he chose M . Think of this as the probability distribution over the nodes in her information set: each node is the result of a different history of play, and so beliefs about these histories can be represented simply by probabilities over the nodes. Since each node represents a unique way of reaching the information set, the nodes are mutually exclusive. Since the information set can only be reached by one of these histories, the nodes are also exhaustive. The fact that the nodes are exhaustive and mutually

exclusive means that the probabilities associated with them must sum up to 1; that is, they must form a probability distribution.

To define these beliefs generally, recall from our formal definition of extensive-form games that \mathcal{X} denotes the set of nodes, \mathcal{H} denotes the set of histories (information sets), and $h(x) \in \mathcal{H}$ is the information set that contains node $x \in \mathcal{X}$.

DEFINITION 1. A **system of beliefs** is a mapping, $\mu : \mathcal{X} \rightarrow [0, 1]$, such that for all $h \in \mathcal{H}$, the following holds: $\sum_{x \in h} \mu(x) = 1$.

Intuitively, a system of beliefs μ specifies the relative probabilities of being at each node for any information set. Obviously, if the information set is a singleton, $h(x) = \{x\}$, then $\mu(x) = 1$ (beliefs are degenerate). Think of beliefs as probability distributions over the nodes in an information set; there must be as many of these distributions as there are information sets.

These systems of beliefs allow us to define the analogue to a subgame, which we shall call a **continuation game**: every information set begins a new continuation game, which includes the remaining game tree that emanates from that information set, and nothing else. That is, this is basically a subgame without the requirement that it starts with a singleton information set.

Just as we required strategies to be best responses in a subgame, we are now going to require that they are best responses in a continuation game. Each action at a given information set is going to have different consequences depending on which node in the information set the play has actually reached. The system of beliefs then allows us to compute the expected payoff of each action as a function of the probability distribution over the nodes. For instance, if player 2 chooses L , she must expect that the outcome will be (U, L) with probability μ (her belief that player 1 has chosen U), and that the outcome will be (M, L) with probability $1 - \mu$ (her belief that player 1 has chosen M). Correspondingly, her expected payoff from choosing L is:

$$U_2(L) = \mu u_2(U, L) + (1 - \mu)u_2(M, L) = \mu(1) + (1 - \mu)(2) = 2 - \mu$$

Analogously, since choosing R leads to the outcome (U, R) with probability μ and the outcome (M, R) with probability $1 - \mu$ under her beliefs, her expected payoff from choosing R is:

$$U_2(R) = \mu u_2(U, R) + (1 - \mu)u_2(M, R) = \mu(0) + (1 - \mu)(1) = 1 - \mu.$$

The best-response is, as usual, the action that yields the highest expected payoff. In this case, we note that $2 - \mu > 1 - \mu$ for all $\mu \in [0, 1]$. This means that player 2's best response at her information set must be L irrespective of her beliefs about player 1 choosing U or M . Since our definition of rationality requires best responses, merely having beliefs in this game is sufficient to eliminate any solution that involves player 2 choosing R at her information set. This is so because the strategy R is strictly dominated: player 2 is better off choosing L for any action that player 1 might take that reaches her information set. Most games, of course, are not going to have such convenient strictly dominated strategies, and merely having beliefs will not be sufficient to determine behavior in the continuation game.

The notion of beliefs allows us to define best responses for any continuation game as the strategies that yield the highest expected payoff under the system of beliefs. Formally, let $U_i(\sigma_i(h), \sigma_{-i}(h) | h, \mu)$ denote player i 's expected payoff starting at information set h , with beliefs $\mu(x)$ about being at any node $x \in h$, if he follows strategy $\sigma_i(h)$ while the other players follow $\sigma_{-i}(h)$, where $\sigma(h)$ denotes the restriction of the strategies to the continuation game starting at h . We then define best responses as follows:

DEFINITION 2. A strategy σ_i is **sequentially rational** at information set h given a system of beliefs μ if

$$U_i(\sigma_i(h), \sigma_{-i}(h) | h, \mu) \geq U_i(\tilde{\sigma}_i(h), \sigma_{-i}(h) | h, \mu)$$

for all $\tilde{\sigma}_i(h) \in \Sigma_i(h)$.

That is, a strategy is sequentially rational for a player at a given information set if, given his beliefs, it yields the highest expected utility in the continuation game starting at that information set while holding the other players' strategies fixed. This is easier to understand if you think about behavior strategies in extensive form games. Recall that behavior strategies specify probability distributions over the nodes in an information set, just like systems of beliefs do. The calculation of the expected payoff in a continuation game is exactly the same as the computation of expected payoffs from a behavior strategy.

Sequential rationality at an information set implies that a player would not want to deviate from the strategy at that information set. This means that if a strategy, σ_i , is sequentially rational at all information sets where player i has to move given his beliefs μ and the strategies σ_{-i} for the other players, then it is player i 's *best response* to σ_{-i} under μ . In our example, player 2 has a unique sequentially rational strategy, L .

Having specified a player's best responses as a function of his beliefs, we must now ask where these beliefs come from.

1.2 Consistent Beliefs

Let us pause for a second to think what it is that we are after here: a definition of rational behavior as a best response to what a player thinks the other players are doing. Since this definition is going to apply to all players, each of them must expect the others to be best-responding as well. In the present context, this means that players must all expect everyone to be choosing sequentially rational strategies given their beliefs. This means that each player's beliefs about the history of play must be restricted to everyone using sequentially rational strategies. In other words, all beliefs must be somehow derived from the optimal strategies.

Consider the game in Fig. 1 (p. 2) and some strategy σ_1 for player 1 such that $\sigma_1(U) > 0$ or $\sigma_1(M) > 0$. If player 2 thinks that player 1 is using σ_1 and her information set is reached, what belief must she hold about which node she is at? That is, what is the likelihood that player 1 chose, say U , conditional on him having chosen either U or M and knowing that the probabilities with which he chooses them are $\sigma_1(U)$ and $\sigma_1(M)$, respectively? Since this involves conditional probabilities, it is perhaps not surprising that we are going to use Bayes rule to compute them:

$$\mu = \Pr(U|U \text{ or } M) = \frac{\Pr(U)}{\Pr(U \text{ or } M)} = \frac{\sigma_1(U)}{\sigma_1(U) + \sigma_1(M)}.$$

You can see now why we wanted one of these two probabilities for player 1 to be non-zero: Bayes rule is undefined for zero-probability events (you cannot condition on stuff that cannot happen). Intuitively, since there are only two histories consistent with player 2's information set, the total probability of reaching that set is the probability that one of them is chosen. Since histories are exhaustive and mutually exclusive, this probability is the sum of the probabilities associated with the different histories: $\sigma_1(U) + \sigma_1(M)$. Since we are interested in the likelihood that one particular history has occurred, U in this case, we divide the unconditional probability of that history occurring, $\sigma_1(U)$, by the total probability.

The notion that beliefs are consistent with the strategies when they are derived from them by Bayes rule whenever possible is defined formally as follows:

DEFINITION 3. A system of beliefs μ is **consistent** with the strategies σ if, for any information set h such that $\Pr(h|\sigma) > 0$ and for any $x \in h$,

$$\mu(x) = \frac{\Pr(x|\sigma)}{\Pr(h|\sigma)}.$$

To see that this is indeed Bayes rule that conditions on the fact that the information set is reached, write the full formula:

$$\Pr(x|h, \sigma) = \frac{\Pr(h|x, \sigma) \Pr(x|\sigma)}{\Pr(h|\sigma)},$$

and observe $\Pr(h|x, \sigma) = 1$ because knowing that some $x \in h$ has been reached means that h has been reached.

With these two concepts in mind, we can now revise our definition of rationality.

1.3 Weak Perfect Bayesian Equilibrium

The idea is simple: strategies must be sequentially rational given beliefs (best responses) and beliefs must be consistent with the strategies (derived by Bayes rule whenever the strategies allow it). We shall call a strategy profile σ and a system of beliefs μ an **assessment** (σ, μ) , and use this to define the solution concept formally as follows,

DEFINITION 4. An assessment (σ, μ) is a **Weak Perfect Bayesian Equilibrium (WPBE)** if σ is sequentially rational given μ , and μ is consistent with σ .

It is important to realize that this definition places absolutely no restrictions on beliefs at information sets that do not occur if the strategy profile σ is followed; i.e., on information sets *off the path of play*. We only require that they are *some* probability distributions but there is no consistency restriction on *what* those distributions must be. This allows the analyst to assign arbitrary probability distributions (beliefs) at such information sets. The definition still requires that the strategies are sequentially rational everywhere, which includes these off-the-path sets with arbitrary beliefs.

As it turns out, for some games this definition is sufficient to rule out implausible NE and SPE. Going back to our example from Fig. 1 (p. 2), note that any WPBE requires that player 2 chooses L at her information set no matter what beliefs she might have. Player 1's best response to this is to choose U , which leaves us with the assessment $(\langle U, L \rangle, \mu = 1)$ as the unique WPBE.

We can infer from this result that not all Nash equilibria are WPBE, and that not all SPE are WPBE. We shall see soon that not all WPBE are SPE either. For now, we can state an immediate result that follows directly from the definition of WPBE:

PROPOSITION 1. *Every WPBE is a Nash equilibrium, but not every Nash equilibrium is a WPBE.* □

Proof. Consider an assessment (σ, μ) that is WPBE. Since the strategies are sequentially rational at all information sets, they must be sequentially rational at all information sets that are reached with positive probability under σ (and μ is derived by Bayes rule at all such sets). This means that they are best responses to each other, and so σ is a Nash equilibrium.

The second part of the claim—that there are Nash equilibria that are not WPBE—is established with the example above. In particular, Nash equilibrium allows strictly dominated strategies to be played at information sets that do not occur under σ , but WPBE will not allow it. ■

To see how one would solve for WPBE in slightly more involved situations, consider a modification of our example: the game shown in Fig. 2 (p. 7). One key difference is that now player 2 no longer has a strictly dominant strategy at her information set: her best response depends on what player 1 does. The other differences involve player 1's payoffs.




Figure 2: The WPBE Example Game.

Let σ_i denote a strategy for player i and start with player 2, whose expected payoffs are:

$$U_2(L) = \mu(1) + (1 - \mu)(1) = 1$$

$$U_2(R) = \mu(0) + (1 - \mu)(2) = 2 - 2\mu,$$

which implies that her sequentially rational strategy is L whenever $1 \geq 2 - 2\mu \Leftrightarrow \mu \geq 1/2$, and R otherwise. We can write her best response as:

$$BR_2(\mu) = \begin{cases} \sigma_2(L) = 1 & \text{if } \mu > 1/2 \\ \sigma_2(L) = 0 & \text{if } \mu < 1/2 \\ \sigma_2(L) \in [0, 1] & \text{if } \mu = 1/2. \end{cases}$$

Turning now to player 1, observe that U strictly dominates D . This means that no matter what player 1 believes at the outset of the game, choosing D is never sequentially rational. Therefore, $\sigma_1(D) = 0$ in any WPBE, which in turn means that $\sigma_1(U) + \sigma_1(M) = 1$ in any WPBE. This tells us that player 2's information set must occur on the path of play in any WPBE, and so Bayes rule will pin down her beliefs there. Since her best response involves a critical value of μ , let us examine all possibilities:

- Suppose that $\mu > 1/2$, in which case player 2 must choose L . Player 1's best response is to choose M . But this implies that $\mu = 0$ because this is the only belief consistent with his strategy, a contradiction. Therefore, there is no such WPBE.
- Suppose that $\mu < 1/2$, in which case player 2 must choose R . Player 1's best response is to choose U . But this implies that $\mu = 1$ because this is the only belief consistent with his strategy, a contradiction. Therefore, there is no such WPBE.
- Suppose that $\mu = 1/2$, in which case player 2 is indifferent, and so any response is sequentially rational. For this belief to be consistent with player 1's strategy, it must be the case that

$$\mu = \frac{\sigma_1(U)}{\sigma_1(U) + \sigma_1(M)} = \sigma_1(U) = \frac{1}{2},$$

where we used the fact that $\sigma_1(U) + \sigma_1(M) = 1$ in any WPBE. Thus, we conclude that player 1 must be mixing in this equilibrium with $\sigma_1(U) = \sigma_1(M) = 1/2$. He would only do so if both U and M are sequentially rational; i.e., if he is indifferent among them. Since

$$U_1(U) = \sigma_2(L)(2) + (1 - \sigma_2(L))(1) = 1 + \sigma_2(L)$$

$$U_1(M) = \sigma_2(L)(3) + (1 - \sigma_2(L))(0) = 3\sigma_2(L),$$

this can only happen when $\sigma_2(L) = 1/2$. This yields the WPBE.

We conclude that this game has a unique WPBE: $(\langle \sigma_1^*(U) = \sigma_1^*(M) = 1/2, \sigma_2^*(L) = 1/2 \rangle, \mu = 1/2)$. Generally, when all information sets occur with positive probability under the equilibrium strategy profile, as it is the case here, we will omit the system of belief when writing the solutions because these are uniquely defined by Bayes rule. For instance, in this example we would say that the game has a unique WPBE, in which player 1 chooses between U and M with equal probabilities, and player 2 chooses between L and R with equal probabilities. Since all information sets occur under σ^* , all Nash equilibria must also be SPE and WPBE. A quick check of the strategic form reveals that this is indeed the case: the unique MSNE specifies the strategies we just found.

Let us now make the game a bit more interesting so that D is no longer strictly dominated, as shown in Fig. 3 (p. 8). We can now no longer assert that player 1 cannot play D in WPBE, and as a result cannot assume that μ will be defined by Bayes rule.




Figure 3: The Interesting WPBE Example Game.

Since player 2's payoffs are the same as in Fig. 2 (p. 7), her sequentially rational strategies remain the same. We shall examine the possible assessments by looking at the various beliefs she might hold:

- Suppose that $\mu > 1/2$, in which case player 2 must choose L . Player 1's best response is to choose M , which puts player 2's information set on the path of play. The only consistent belief is $\mu = 0$, a contradiction. There is no WPBE with beliefs $\mu > 1/2$.
- Suppose that $\mu < 1/2$, in which case player 2 must choose R . Player 1's best response is to choose D , which leaves player 2's information set off the path of play. Bayes rule is undefined, and WPBE places no restrictions on her beliefs. In particular, we could pick any probability distribution that could rationalize her response. Thus, there is a continuum of WPBE where the assessments take the form $(\langle D, R \rangle, \mu < 1/2)$.
- Suppose that $\mu = 1/2$, in which case player 2 can mix. Player 1's expected payoffs are

$$U_1(U) = \sigma_2(L)(2) + (1 - \sigma_2(L))(-1) = 3\sigma_2(L) - 1$$

$$U_1(M) = \sigma_2(L)(3) + (1 - \sigma_2(L))(0) = 3\sigma_2(L).$$

and so M strictly dominates U ; i.e., U is never sequentially rational. Thus, $\sigma_1(U) = 0$ in any WPBE, which implies that if $\sigma_1(M) > 0$ in WPBE, Bayes rule will pin down the only consistent belief to be $\mu = 0$, and player 2 would not be indifferent. We conclude that $\sigma_1(M) = 0$ must obtain as well in any such WPBE, which in turn means that player 2's information set must be off the path of play, and we can assign whatever assessment we need to rationalize her strategy. Since player 1 must be willing to choose D over M , it must be that $U_1(D) \geq U_1(M) \Leftrightarrow \sigma_2(L) \leq 1/3$. Thus, there is a continuum of WPBE where the assessments take the form $(\langle D, \sigma_2(L) \leq 1/3 \rangle, \mu = 1/2)$.

This game has infinitely many WPBE characterized by different off-the-path beliefs we assign for player 2 (and, in the second set of solutions, by her different mixing probabilities). All these WPBE are payoff

equivalent (the players obtain (1, 3)), and observationally equivalent as well (the game ends with player 1 choosing D). In that sense, the multiplicity of WPBE is not particularly troubling. On the other hand, since U is strictly dominated by M for player 1, it seems quite odd that any beliefs other than $\mu = 0$ should be admissible. If we were to eliminate these, the sole surviving assessment would be $((D, R), \mu = 0)$. We shall soon see this stronger consistency requirement in action.

How do these WBNE relate to Nash equilibria (and SPE, since all Nash equilibria in this game are subgame-perfect)? The strategic form in Fig. 3 (p. 8) shows that M strictly dominates U , so $\sigma_1(U) = 0$ in any NE. Moreover, if $\sigma_1(M) > 0$, then R strictly dominates L for player 2. But if player 2 chooses R , then player 1 would choose D . Therefore, $\sigma_1(M) = 0$ in any NE as well. We conclude that player 1 must play D , which makes player 2 indifferent. She can mix with any probability as long as $U_1(D) = 1 \geq U_1(M) = 3\sigma_2(L)$, or $\sigma_2(L) \leq 1/3$. Thus, there is a continuum of MSNE: $\langle D, \sigma_2(L) \leq 1/3 \rangle$. All of them are subgame-perfect. All WPBE that we found are in this set of MSNE as well.

We saw earlier that not all SPE are WPBE (indeed, this is how we eliminated the unreasonable solution in the game from Fig. 1 (p. 2)). We now show that not all WPBE are SPE either; that is, there exist WPBE that are not subgame-perfect. We shall do this by way of the example game in Fig. 4 (p. 9).




Figure 4: The WPBE $\not\subseteq$ SPE Game.

The subgame that starts with player 2's move has a unique Nash equilibrium, in which she chooses U and player 3 chooses R . With that expectations, player 1 has a unique best response to choose B . Therefore, this game has a unique SPE: $\langle B, U, R \rangle$.

Consider now sequentially rational strategies for player 3. His expected payoffs are:

$$\begin{aligned} U_3(L) &= \mu(1) + (1 - \mu)(2) = 2 - \mu \\ U_3(R) &= \mu(3) + (1 - \mu)(1) = 1 + 2\mu, \end{aligned}$$

which means that his best response is:

$$BR_3(\mu) = \begin{cases} \sigma_3(L) = 1 & \text{if } \mu < 1/3 \\ \sigma_3(L) = 0 & \text{if } \mu > 1/3 \\ \sigma_3(L) \in [0, 1] & \text{if } \mu = 1/3. \end{cases}$$

Since player 2's payoffs are:

$$\begin{aligned} U_2(U) &= \sigma_3(L)(2) + (1 - \sigma_3(L))(3) = 3 - \sigma_3(L) \\ U_2(D) &= \sigma_3(L)(1) + (1 - \sigma_3(L))(1) = 1, \end{aligned}$$

choosing U strictly dominates choosing D . Her best response in any WPBE must be U , so $\sigma_2(U) = 1$. This makes it easy to compute player 1's expected payoff from choosing B as:

$$U_1(B) = \sigma_3(L)(1) + (1 - \sigma_3(L))(3) = 3 - 2\sigma_3(L),$$

so his best response is:

$$BR_1(\sigma_2(U) = 1, \sigma_3) = \begin{cases} \sigma_1(T) = 1 & \text{if } \sigma_3(L) > 1/2 \\ \sigma_1(T) = 0 & \text{if } \sigma_3(L) < 1/2 \\ \sigma_1(T) \in [0, 1] & \text{if } \sigma_3(L) = 1/2. \end{cases}$$

Suppose that $\sigma_1(B) > 0$, which puts player 2's information set on the path of play. In this case, $\sigma_2(U) = 1$ pins down $\mu = 1$, so player 3 must choose R , to which player 1's best response is B . Therefore, the assessment $(\langle B, U, R \rangle, \mu = 1)$ is WPBE. Not surprisingly, when all information sets are on the path of play, there is no difference between SPE and WPBE (and Nash).

Suppose now that $\sigma(B) = 0$, which leaves player 2's information set off the path of play. This, in turn, puts player 3's information set off the path of play as well, and so Bayes rule is undefined.¹ We are free to assign whatever beliefs we want to rationalize player 3's strategy. Since player 2 has a unique sequentially rational strategy in U , and we wish to ensure that player 1 chooses T , getting player 3 to choose L will be sufficient. We can get this by assigning some $\mu < 1/3$. Therefore, any assessment $(\langle T, U, L \rangle, \mu < 1/3)$ is WPBE. None of these are subgame perfect.²

Why did this happen? Subgame perfection places strong restrictions on strategies in all subgames. In this instance, it requires player 3 to best-respond in the subgame that starts with player 2's move. In effect, this forces player 3 to behave as if player 2's information set has been reached. Since player 2 has a strictly dominant strategy to choose U , this restricts player 3 to choosing R . WPBE, on the other hand, does not place any restrictions on information sets off the path of play. It only requires that players choose strategies that are sequentially rational given *some* beliefs. In this instance, this means that player 3 does not have to take into account the fact that U is strictly dominant for player 2, which allows him to maintain the (strange) belief that she played D probability that exceeds $2/3$. In fact, there is a WPBE in which he believes that she plays D with certainty! This is clearly undesirable: since player 2 has a strictly dominant strategy, any reasonable belief for player 2 should require $\mu = 1$, which would make the WPBE subgame perfect. But no such restriction exists in the current definition.

Before moving on, note that our examples have collectively established the following result:

PROPOSITION 2. *Not every WPBE is a SPE, and not every SPE is a WPBE in games of imperfect information.* □

As the examples suggest, the discrepancies arise when information sets that are not singletons are left off the path of play. This suggests that games of perfect information might not have this problem, as indeed turns out to be the case.

PROPOSITION 3. *Every WPBE is a SPE in games of perfect information.* □

¹The formula is:

$$\mu = \frac{\sigma_1(B)\sigma_2(U)}{\sigma_1(B)\sigma_2(U) + \sigma_1(B)\sigma_2(D)},$$

and it is clearly undefined when $\sigma_1(B) = 0$. Make sure you understand that you cannot just divide both the numerator and the denominator by $\sigma_1(B)$ to cancel that term.

²There is also a continuum of assessments where player 3 mixes, which are also not subgame perfect but WPBE. I leave these to you as an exercise.

Proof. WPBE requires that strategies are sequentially rational given *some* belief at each information set. Since each information set is a singleton in games of perfect information, the game admits only a single belief system. But then sequential rationality guarantees that the strategies are Nash in all subgames; i.e., that they are subgame perfect. ■

The intuition for this result is that if all information sets are singletons, then there is only one belief that we must assign for each of them irrespective of whether they occur on or off the path of play: the probability of being at the single node in an information set must be 1. Sequential rationality then boils down to our standard Nash best response for the subgame (and continuation game) that starts at that node, which in turn ensures that the sequentially rational strategies are subgame perfect.

The possibility that WPBE is not subgame perfect is troubling. It is especially so when you recognize why this is the case: by assigning arbitrary beliefs for off-the-path information sets, we are permitting players to threaten each other with implausible beliefs. This means that we must make the solution concept more demanding.

2 Perfect Bayesian Equilibrium

If our only goal were to ensure that all WPBE are subgame perfect, then there is a straightforward strengthening of the equilibrium definition that would satisfy this: require that the strategies are WPBE in all subgames. This will force recalculation of beliefs from the start of each subgame irrespective of whether it is reached by the strategies, and avoid problems of the sort we saw in Fig. 4 (p. 9).

DEFINITION 5. An assessment (σ, μ) is a **Perfect Bayesian Equilibrium (PBE)** if it is a WPBE in all subgames.

Since this requirement ensures that PBE are subgame perfect and because the original example from Fig. 1 (p. 2) shows that there are SPE that are not PBE, the following is immediate:

PROPOSITION 4. *Every PBE is a SPE but not every SPE is a PBE.* □

In practice, applied work with games of incomplete information almost always uses PBE.³ Let us look at several example games.

2.1 Myerson's Card Game

Recall the card that we have now seen a couple of times, reproduced here in Fig. 5 (p. 12). Previously, we solved this by converting it to strategic form and finding the Nash equilibrium. Let us now find the PBE.

We start with player 2's sequentially rational strategy. Her expected payoffs are:

$$\begin{aligned} U_2(m) &= \mu(2) + (1 - \mu)(-2) = 4\mu - 2 \\ U_2(p) &= \mu(-1) + (1 - \mu)(-1) = -1, \end{aligned}$$

so her best response is

$$BR_2(\mu) = \begin{cases} \sigma_2(m) = 1 & \text{if } \mu > 1/4 \\ \sigma_2(m) = 0 & \text{if } \mu < 1/4 \\ \sigma_2(m) \in [0, 1] & \text{if } \mu = 1/4. \end{cases}$$

³The definition of PBE in Fudenberg and Tirole is actually a bit more demanding than the one I have given here, but this one is sufficient for this course.




Figure 5: Myerson's Card Game.

Player 2's belief is defined as

$$\mu = \frac{(1/2)\sigma_1(R)}{(1/2)\sigma_1(R) + (1/2)\sigma_1(r)} = \frac{\sigma_1(R)}{\sigma_1(R) + \sigma_1(r)}$$

whenever $\sigma_1(R) > 0$ or $\sigma_1(r) > 0$. It is undefined, of course, if both are zero. Since player 1's payoffs are:

$$\begin{aligned} U_1(F) &= -1 & U_1(R) &= 1 - 3\sigma_2(m) \\ U_1(f) &= 1 & U_1(r) &= 1 + \sigma_2(m), \end{aligned}$$

his sequentially rational strategies are:

$$BR_1(\sigma_2|\text{black}) = \begin{cases} \sigma_1(R) = 1 & \text{if } \sigma_2(m) < 2/3 \\ \sigma_1(R) = 0 & \text{if } \sigma_2(m) > 2/3 \\ \sigma_1(R) \in [0, 1] & \text{if } \sigma_2(m) = 2/3, \end{cases}$$

and

$$BR_1(\sigma_2|\text{red}) = \begin{cases} \sigma_1(r) = 1 & \text{if } \sigma_2(m) > 0 \\ \sigma_1(r) \in [0, 1] & \text{if } \sigma_2(m) = 0. \end{cases}$$

We now look for PBE using the critical value of player 2's posterior beliefs:

- Suppose that $\mu > 1/4$, so $\sigma_2(m) = 1$. Player 1's best responses are $\sigma_1(R) = 0$ and $\sigma_1(r) = 1$, which imply that $\mu = 0$, a contradiction. There is no such PBE.
- Suppose that $\mu < 1/4$, so $\sigma_2(m) = 0$. Player 1's best response if his card is black is $\sigma_1(R) = 1$. There are two possibilities if his card is red:
 - Suppose that $\mu \in (0, 1/4)$, so his best response is $\sigma_1(r) = 1$ as well. Then $\mu = 1/2$, a contradiction. There is no such PBE.
 - Suppose that $\mu = 0$, in which case $\mu = 1/(1 + \sigma_1(r)) > 0$, a contradiction. There is no such PBE.
- Suppose that $\mu = 1/4$, so player 2 is indifferent and can mix. There are two possibilities:

- Suppose that $\sigma_2(m) > 0$, so that player 1's best response when the card is red is $\sigma_1(r) = 1$. Bayes rule then requires that

$$\mu = \frac{\sigma_1(R)}{\sigma_1(R) + 1} = \frac{1}{4} \quad \Rightarrow \quad \sigma_1(R) = \frac{1}{3},$$

which implies that player 1 must be willing to mix when the card is black. He would only do so when $\sigma_2(m) = 2/3$. The assessment $((\sigma_1(R) = 1/3, \sigma_1(r) = 1), \sigma_2(m) = 2/3), \mu = 1/4$ is PBE.

- Suppose that $\sigma_2(m) = 0$, so that player 1 is indifferent when the card is red, so that $\sigma_1(r) \in [0, 1]$. Since this also implies that $\sigma_1(R) = 1$, Bayes rule requires that

$$\mu = \frac{1}{1 + \sigma_1(r)} = \frac{1}{4} \quad \Rightarrow \quad \sigma_1(r) = 3,$$

which is clearly impossible. There is no such PBE.

We conclude that this game has a unique PBE, in which player 1 always raises when the card is red, and bluffs by raising with probability $1/3$ when the card is black. Player 2 meets with probability $2/3$. In this PBE, when player 2 observes that player 1 has raised, her updated belief that the card is black is $\mu = 1/4$, which is a reduction from her prior belief of $1/2$. This, of course, is the MSNE we originally found. Since player 2's information set is reached with positive probability under these strategies, it should not be surprising that this MSNE is a PBE as well (and a SPE).

You might wonder if it is possible to construct a PBE exploiting the idea that player 2's information set could be left off the path of play and then assigning some arbitrary beliefs for her to act upon. Since all PBE are Nash equilibria and the game has a unique MSNE, the answer is negative. But let's try it anyway. To leave player 2's information set off the path of play, it is necessary that $\sigma_1(R) = \sigma_1(r) = 0$ in that putative PBE. To rationalize player 1's choice when the card is red, $\sigma_1(r) = 0$, it is necessary that $\sigma_2(m) = 0$, but in that case $\sigma_1(R) = 1$, a contradiction. We cannot construct a PBE with strategies that do not reach player 2's information set.

Observe now that you can think of this situation as a simple signaling game. Player 1's type is whether he holds a winning (red) or losing (black) card, and the common prior is that the two types are equally likely. Since player 1 can condition his choice on the card he holds, he has type-contingent strategies that can potentially reveal something to player 2. That is, after observing player 1's choice, player 2 can make inferences about his type. She begins the game thinking that there is a 50-50 chance that he holds the winning card, but after observing him raising, she updates to believe that there is a 75-25 chance that this is so. These are the odds that leave her indifferent between meeting and passing, which rationalizes her willingness to mix. Her strategy, in turn, rationalizes player 1's willingness to bluff when he holds the losing card (he always raises with a winning card). Thus, player 1 raising is a noisy signal that he has a winning card.

The signal works because player 1 does not always raise when he has a losing card. If he were to do that, $\sigma_1(R) = \sigma_1(r) = 1$, then player 2 would not learn anything from observing him raise, so her posterior belief will remain the same as her prior: $\mu = 1/2$. But in that case, she will meet for sure, $\sigma_2(m) = 1$, which induces an unacceptable outcome for the player with the losing card.

2.2 Classifying PBE by Strategy Types

The discussion of transmission of information by strategies leads to a useful classification of PBE in incomplete information games depending on how much information they reveal about the privately known types of the players:

- If all types of a player choose their actions with the same probabilities, then observing any particular action conveys no new information about their types, and the posterior beliefs remain the same as the priors. These strategies are called **pooling** because all types pool on the same probability distribution.
- If all types of a player choose different actions with certainty, then observing any particular action fully reveals the information about their type, and the posterior beliefs assign probability 1 to that type irrespective of the priors. This is why these strategies are called **separating**: the actions they prescribe “separate” the types fully.
- If the types of a player choose different actions with different non-degenerate probabilities, then observing any particular action will convey *some* information about their types, which will lead to the posterior beliefs getting updated but without causing them to degenerate. This is why these strategies are called **semi-separating** (or semi-pooling, or hybrid): the actions they prescribe reveal something about their types but the other players remain uncertain. The PBE we found in Myerson’s Card Game involves a strategy for player 1 that is of this type.

In games with one-sided incomplete information (where one player is fully informed but the other is not), the type of strategy the informed player uses in PBE is often used to label the PBE itself. Thus, in Myerson’s Card Game, we would normally refer to the PBE as semi-separating. This classification of strategies (and PBE) sometimes makes it easier to organize one’s analysis, as the following example shows.

2.3 A Two-Period Reputation Game

There are two firms, $i \in \{1, 2\}$, in the market, and their interaction unfolds over two periods. In the first period, firm 1 can accommodate, A , or fight (by cutting prices), F , and in the second period, firm 2 chooses whether to stay, S , or exit, X , and after that firm 1 chooses whether to cut prices or not. All choices are observable. If firm 1 accommodates and both firms are in the market, it is a duopoly, and each firm i receives a payoff $d_i > 0$. If firm 2 exits, then it receives a payoff of 0, and if firm 1 does nothing else, it receives the monopoly price, $m_1 > d_1$. If firm 1 fights, firm 2 receives $w_2 < 0$. Firm 1’s payoff depends on whether it prefers a monopoly to cutting prices. The type that prefers monopoly has a price-cutting payoff of $w_1 < 0$, whereas the type that prefers to cut prices has a price-cutting payoff $\hat{w}_1 > m_1$. For simplicity, we shall refer to the type that prefers monopoly as “sane” and the type that prefers to cut-prices as “crazy”. This type is privately known to firm 1, and firm 2’s (common knowledge prior) belief that firm 1 is sane is $p \in (0, 1)$. The total payoff for each firm is the discounted sum of its per-period payoffs. The firms have a common discount factor $\delta \in (0, 1)$.

Since all PBE are subgame perfect, let us simplify the game a bit by looking at the final information sets for firm 1 in the second period. Since the firm knows its own type and can observe all prior actions, these information sets are singletons. This means that we can derive the sequentially rational (subgame perfect) strategies by backward induction. Since there is no future action by firm 2, firm 1’s behavior at all its final information sets simply depends on its type: the sane type does nothing, and the crazy type cuts prices. If firm 2 exits, the payoffs are $(m_1, 0)$ if firm 1 is sane, and $(\hat{w}_1, 0)$ if it is crazy. If firm 2 stays, the payoffs are (d_1, d_2) if firm 1 is sane, and (\hat{w}_1, w_2) if it is crazy. We can now represent the resulting situation with the extensive-form game in Fig. 6 (p. 15).

Let μ_1 denote firm 2’s belief that firm 1 is sane after being fought in the first period, and let μ_2 denote that belief after being accommodated in the first period. We begin by deriving firm 2’s sequentially rational




Figure 6: The Reputation Game.

strategies at each of her two information sets. Firm 2's expected payoffs are:

$$\begin{aligned}
 U_2(S|F) &= \mu_1 d_2 + (1 - \mu_1) w_2 & U_2(S|A) &= \mu_2 d_2 + (1 - \mu_2) w_2 \\
 U_2(X|F) &= 0 & U_2(X|A) &= 0.
 \end{aligned}$$

These define a critical threshold value,

$$\mu^* = \frac{-w_2}{d_2 - w_2} \in (0, 1),$$

where the fact that μ^* is a valid probability follows from $w_2 < 0 < d_2$. The sequentially rational strategies at both information sets are the same:

$$BR_2(\mu) = \begin{cases} \sigma_2(S|\mu) = 1 & \text{if } \mu > \mu^* \\ \sigma_2(S|\mu) = 0 & \text{if } \mu < \mu^* \\ \sigma_2(S|\mu) \in [0, 1] & \text{if } \mu = \mu^*. \end{cases}$$

That is, firm 2 will stay only if it sufficiently convinced that firm 1 is sane.

Turning now to firm 1, observe that F strictly dominates A for the crazy type, which means that $\sigma_1(F|\text{crazy}) = 1$ in any PBE. This ensures that the corresponding information set must be on the path of play in any PBE, which means that μ_1 must be defined by Bayes rule as:

$$\mu_1 = \frac{p\sigma_1(F|\text{sane})}{p\sigma_1(F|\text{sane}) + (1-p)(1)},$$

whereas μ_2 might or might not be on the path of play, depending on what the sane type does. If $\sigma_1(A|\text{sane}) > 0$, then Bayes rule requires that it be:

$$\mu_2 = \frac{p\sigma_1(A|\text{sane})}{p\sigma_1(A|\text{sane}) + (1-p)(0)} = 1.$$

The expected payoffs for the sane firm 1 are:

$$\begin{aligned} U_1(F|\text{sane}) &= w_1 + \delta m_1 + \delta \sigma_2(S|\mu_1)(d_1 - m_1) \\ U_1(A|\text{sane}) &= d_1 + \delta m_1 + \delta \sigma_2(S|\mu_2)(d_1 - m_1). \end{aligned}$$

This defines a condition,

$$\sigma_2(S|\mu_2) - \sigma_2(S|\mu_1) \geq \frac{d_1 - w_1}{\delta(m_1 - d_1)} \equiv \Delta > 0, \quad (1)$$

which we can use to specify the sequentially rational strategy as follows:

$$BR_{\text{sane}} = \begin{cases} \sigma_1(F|\text{sane}) = 1 & \text{if } \sigma_2(S|\mu_2) - \sigma_2(S|\mu_1) > \Delta \\ \sigma_1(F|\text{sane}) = 0 & \text{if } \sigma_2(S|\mu_2) - \sigma_2(S|\mu_1) < \Delta \\ \sigma_1(F|\text{sane}) \in [0, 1] & \text{if } \sigma_2(S|\mu_2) - \sigma_2(S|\mu_1) = \Delta. \end{cases}$$

The left-hand side of (1) is the difference in firm 2's probability of staying following accommodation, $\sigma_2(S|\mu_2)$ and a fight, $\sigma_2(S|\mu_1)$. The sane firm 1 will fight in the first period only if this difference is sufficiently large; that is, it will fight if firm 2 is much more likely to stay after being accommodated than after being fought. Another way of saying this is that the sane firm 1 will fight in period 1 only if this makes firm 2's exit quite a bit more likely in the second period. This it can achieve by increasing firm 2's belief that he is the crazy type (which always fights); that is, by establishing a reputation for toughness (or, better, by bluffing).⁴

Observe now that if $\Delta \geq 1$, then (1) can never be satisfied. This means that the sane firm 1 must accommodate in the first period: $\sigma_1(F|\text{sane}) = 0$. Since the crazy type always fights, both of firm 2's information sets are on the path of play: $\mu_1 = 0 < \mu^* < 1 = \mu_2$. This means that firm 2 will exit after observing a fight, $\sigma_2(S|\mu_1) = 0$, but stay after being accommodated, $\sigma_2(S|\mu_2) = 1$. Since $1 \leq \Delta$ by supposition, the sane type accommodates. Therefore, if $\Delta \geq 1$, the game has a unique PBE, in which the sane type accommodates in both periods, the crazy type fights in both periods, and firm 2 stays whenever accommodated, and exits otherwise. Since the strategies fully reveal the type of firm 1, this is a **separating** PBE.

Assume now that $\Delta < 1$, so it is possible to satisfy (1). Since the crazy type always fights, let us find the PBE using the sane type's strategy:

- **Pooling:** $\sigma_1(F|\text{sane}) = 1$, which implies that $\mu_1 = p$ and μ_2 is undefined. There are now two possibilities to consider:
 - If $p > \mu^*$, then $\sigma_2(S|\mu_1) = 1$, which means that $\sigma_2(S|\mu_2) - 1 < \Delta$ irrespective of how μ_2 is defined, and thus $\sigma_1(F|\text{sane}) = 0$, a contradiction. There is no such PBE.
 - If $p \leq \mu^*$, then $\sigma_2(S|\mu_1) = 0$, and all that is necessary to rationalize the sane type's strategy is $\sigma_2(S|\mu_2) > \Delta$. Any $\mu_2 > \mu^*$ will accomplish this (causing firm 2 to stay after accommodation), which means that there is a continuum of PBE in which both types of firm 1 fight in the first period, and firm 2 exists if, and only if, it is accommodated. Intuitively, this strategy works because firm 2 is already quite pessimistic: her prior belief that firm 1 is sane is very

⁴I don't like this way of thinking about reputation, for reasons we discussed in class. Briefly, this method conceptualizes reputation as someone's belief that you are a type that you are not (one, perhaps, that you would like to be), rather than their belief that you are the type that you are. In this model, this conceptualization means that the sane type wants firm 2 to believe that it is the crazy type, rather than that it is the sane type. To me, this sounds like bluffing about being someone you are not rather than establishing a reputation for who you are.

low, $p \leq \mu^*$. With such a belief, staying is too risky, so firm 2 exits even though there is a positive probability that firm 1 was bluffing and that it is, in fact, sane, and so would accommodate in the second period if firm 2 stayed. All these PBE are both observationally and payoff equivalent: on the path of play firm 1 fights, and firm 2 exits. Since they only differ in the off-the-path beliefs that support them, they are *generically unique*, and so we do not have a multiple equilibria problem.

- **Separating:** $\sigma_1(F|\text{sane}) = 0$, which implies that $\mu_1 = 0$ and $\mu_2 = 1$. This implies that $\sigma_2(S|\mu_2) - \sigma_2(S|\mu_1) = 1 > \Delta$, which is satisfied. But then $\sigma_1(F|\text{sane}) = 1$, a contradiction. There is no such PBE. (In other words, $\Delta \geq 1$, which we found above to be sufficient for a separating PBE to exist, turns out to be necessary as well.)
- **Semi-separating:** $\sigma_1(F|\text{sane}) \in (0, 1)$, which implies $\mu_2 = 1$, and thus $\sigma_2(S|\mu_2) = 1$, so firm 2 must stay after accommodation. Since the sane type must be willing to mix, this now means that $\sigma_2(S|\mu_1) = 1 - \Delta \in (0, 1)$, and so firm 2 must be mixing after a fight. This requires that $\mu_1 = \mu^*$, or, using the definition of μ_1 , that

$$\frac{p\sigma_1(F|\text{sane})}{p\sigma_1(F|\text{sane}) + 1 - p} = \mu^* \quad \Rightarrow \quad \sigma_1(F|\text{sane}) = \frac{(1-p)\mu^*}{p(1-\mu^*)}.$$

Since this expression is clearly positive, all that is required for it to be a valid probability is that it is less than one:

$$\frac{(1-p)\mu^*}{p(1-\mu^*)} < 1 \quad \Leftrightarrow \quad p > \mu^*.$$

Thus, we conclude that the game has a semi-separating PBE but only when $p > \mu^*$.

We have thus found three solutions: a unique separating PBE when $\Delta \geq 1$ that does not depend on the priors, and, when $\Delta < 1$, a generically unique pooling PBE when $p \leq \mu^*$, and a unique semi-separating PBE when $p > \mu^*$. Observe that these PBE are mutually exclusive (if a configuration of the exogenous parameters is associated with one of them, then it is not associated with another) and exhaustive (all possible configurations of the exogenous parameters are associated with an equilibrium). This is great for predictions since we do not have to deal with indeterminacies that arise when the model produces more than one solution for some specification of the exogenous variables. This, however, is a result of the restrictive assumption that the crazy type always fights, which implies that (a) fighting is never a zero-probability event, and (b) accommodation fully reveals firm 1's type. In more common applied models, the assumptions will not be that stark (and convenient), and so the analysis will become more involved, as the next example shows.

2.4 Spence's Education Game

Now that you are in graduate school, you probably have a good reason to think education is important.⁵ 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

⁵Or 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?

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.⁶

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 worker knows his own ability but the potential employer does not. The employer thinks that the prior probability of the candidate having high ability is $p \in (0, 1)$, and this belief is common knowledge (perhaps there is a study about average productivity in the industry). The worker chooses 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. Since the job market is competitive, the employer must offer a competitive wage such that the expected profit is zero. Let $\mu(e)$ denote the employer's posterior belief that the worker has high ability given that he observed e level of education. The employer's expected payoff is

$$U_E(e) = \mu(e)(2 - w(e)) + (1 - \mu(e))(1 - w(e)) = 1 - w(e) + \mu(e).$$

Intuitively, the wage starts at the compensation for the low ability worker, $1 - w(e)$, and increases in the probability that the worker is high ability, $\mu(e)$. Since in a competitive environment the expected payoff for the employer is zero, it follows that

$$w(e) = 1 + \mu(e).$$

The worker's payoffs are:

$$\begin{aligned} U_H(e) &= w(e) - \frac{e}{2} = 1 + \mu(e) - \frac{e}{2} && \text{if he is the } H \text{ type;} \\ U_L(e) &= w(e) - e = 1 + \mu(e) - e && \text{if he is the } L \text{ type.} \end{aligned}$$

Let $\sigma_H(e)$ be the probability that the H type chooses the level of education e , and let $\sigma_L(e)$ be the corresponding probability for the L type. We can write the employer's posterior belief (the probability that the worker is the H type conditional on an observed level of education e) as:

$$\mu(e) = \frac{\sigma_H(e)p}{\sigma_H(e)p + \sigma_L(e)(1 - p)}.$$

With just two types for the worker, there is no *a priori* reason to expect that any equilibrium would involve more than two different levels of education. Let e_H denote the level chosen by the H type, and e_L denote the level for the L type. We now wish to find the set of PBE.

⁶Here, 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... Or maybe not if your grandfather went to said elite school and the stadium is named after your family.

Separating Equilibria. In these PBE, $e_H \neq e_L$. From Bayes' rule, $\mu(e_H) = 1$ and $\mu(e_L) = 0$, and so we have $w(e_H) = 2$ and $w(e_L) = 1$. With this wage, the L worker's payoff is $U_L(e_L) = 1 - e_L$, so he will choose $e_L^* = 0$ because anything else would make him worse off. What about the H worker? His payoff is $U_H(e_H) = 2 - e_H/2$. Observe now that this should be at least as good as mimicking L 's behavior: if he does that and chooses $e = 0$, then the employer will conclude that he is the low-ability type and offer him the wage $w(e_L)$, so his payoff would be $U_H(0) = 1$. His payoff from $e_H > 0$ must be at least as good,

$$U_H(e_H^*) \geq U_H(e_L^*) \Leftrightarrow 2 - e_H^*/2 \geq 1 \Leftrightarrow e_H^* \leq 2.$$

We conclude that H 's equilibrium level of education cannot be too high or else he would just get no education and stick with the low wage. This of $e = 2$ as obtaining a Master's Degree: going for a Ph.D. will just hurt your bottom line.

On the other hand, H 's education level cannot be too low or else the L type will try to mimic it. To see that, observe that in a separating equilibrium, the low type also must have no incentive to imitate the behavior of the high type. This means that

$$U_L(e_L^*) \geq U_L(e_H^*) \Leftrightarrow 1 \geq 2 - e_H^* \Leftrightarrow e_H^* \geq 1.$$

We conclude that H 's equilibrium level of education cannot be too low or else the low-ability worker would be able to acquire it if doing so would convince the employer that he has high ability. For the educational level to be separating, it must be so high that L cannot profit from imitating H 's behavior. Think of $e = 1$ as obtaining a Bachelor's Degree: if you do not at least get that, then your education cannot possibly reveal to your employer that you are the high ability type.

We conclude that $e_H^* \in [1, 2]$ and $e_L^* = 0$. Although we have pinned down L 's type, we have not actually done so for the H type, we have just narrowed the possibilities. In fact, any $e_H^* \in [1, 2]$ can be sustained in equilibrium with appropriate beliefs.

To see what I mean, pick some e_H^* in that range and note that the employer only expects to see e_H^* or no education at all in equilibrium, any $e \notin \{e_H^*, 0\}$ is off the equilibrium path of play. We cannot use Bayes rule to ensure consistency of beliefs after such educational levels: $\mu(e)$ is undefined. This means that we can assign any beliefs we want. Consider the following beliefs:

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

These are the simplest beliefs (on and off the path) that will sustain the choice of e_H^* in equilibrium. Deviating to a higher level does not benefit the high-ability worker because he's already getting the highest wage and any additional education represents an unnecessary cost. Obviously, since the low-ability type cannot profit from e_H^* , he certainly cannot profit from a higher level either. If, on the other hand, the high-ability type were to attempt a lower level of education, the employer will infer that he is the L type and offer the minimum wage. This leaves him strictly worse off, so he has no incentive to deviate. Clearly, any deviation to a positive level of education that still leaves the employer convinced that he is the L type cannot be profitable for the low-ability type either.

We conclude that any $e_H^* \in [1, 2]$ can be sustained in PBE using the belief system specified above.⁷ Although the solution is indeterminate in the sense that it does not predict the precise level, it does give us

⁷Other beliefs can work too. For instance, we can assign any beliefs after $e > e_H^*$, and as long as $\mu(e)$ is sufficiently low for $e \in [1, e_H^*)$, the high ability type will not deviate. This means that for any e_H^* in the range there are multiple PBE that can work. Since there are also infinite e_H^* that can work, we have a serious multiplicity problem. However, for any given e_H^* , the beliefs induce the same probability distribution over the outcomes, so these PBE are essentially equivalent.

the important substantive conclusion: in any of these equilibria the high type chooses an education level that is (a) sufficiently high to prevent the low type from profiting by acquiring it, and (b) not so high as to make it unprofitable for himself.

We can now employ some forward induction logic to eliminate all but one of these separating PBE. Think about it this way: suppose in equilibrium $e_H^* > 1$ but the high type deviates to $1 \leq e < e_H^*$. With the off-the-path beliefs we assigned above, he will be punished when the employer infers that he is the low-ability type. However, this deviation would not profit the low-ability type even if the employer were to infer he has high ability. Therefore, the only type who can profit from is the high ability type. But this means, the employer should believe that $\mu(e) = 1$. We shall see more of this logic in just a minute, for now it suffices to say that our original system of beliefs appears unreasonable for such deviations. The only reasonable system will be:

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

With these beliefs, only $e_H^* = 1$ can be supported in equilibrium. Intuitively, the high-ability worker would pick the lowest possible level of education that can separate him from the low-ability worker. The low-ability worker cannot profit from choosing this level: $U_L(e_H^*) = 2 - e_H^* = 1 = U_L(e_L^*)$, so he has no incentive to deviate even if doing so would convince the employer that he's the high ability type. We therefore have a unique separating PBE with $e_H^* = 1$ and $e_L^* = 0$, with the beliefs specified above.

This refinement allows us to make a sharp prediction: the high ability worker will pick the lowest possible education level that will still deter the low-ability worker. Education will have instrumental value because it will reveal to the employer the type of worker he is considering hiring.

Pooling Equilibria. In these PBE, $e_H = e_L = e^*$, and Bayes' rule gives $\mu(e^*) = p$ because the employer learns nothing. The wage offered then is $w(e^*) = 1 + p$. With this wage, the L worker's type payoff is $U_L(e^*) = 1 + p - e^*$ and the H worker's type payoff is $U_H(e^*) = 1 + p - e^*/2$. Observe now that the worst that can happen to a worker is for the employer to conclude that he is the low-ability type, in which case the wage would be $w(e) = 1$. In equilibrium, even the low-ability type should be able to do at least as well by choosing some education as by getting no education at all and facing (possibly) the worst-case scenario. That is, $U_L(e^*) \geq 1$, which implies $e^* \leq p$. In other words, only sufficiently low levels of education can be supported in a pooling equilibrium. Any level above that will make the low-ability worker prefer to get the low wage without investing in education.

Observe that this level is lower than the one required to sustain separation: $e^* < 1 \leq e_H^*$. Why was it necessary to deter the low-ability type from deviations all the way up to $e = 1$ in the separating case but only up to p in the pooling case? In other words, in the separating case $e \in (p, 1)$ could potentially be profitable so it was necessary to make sure that the high-ability type did not pick that level. The reason is the wage being offered: in the pooling equilibrium the wage is $w(e_L^*) = 1 < w(e^*) = 1 + p < w(e_H^*) = 2$. That is, because the employer is uncertain which type he's considering, so he will offer less than what he would have offered to a worker who is known to be of high ability. On the other hand, the wage has to be higher than what he would offer a worker of known low ability or else it would be impossible to get the high-ability worker to invest in any education at all.

At any rate, any $e^* \leq p$ can be sustained in a pooling equilibrium. The system of beliefs that can do that must (a) prevent the low-ability type from investing in no education—mimicking the high-ability type must be profitable, and (b) prevent the high-ability type from investing in more education. Of course, now every $e \neq e^*$ is off the path, and as a zero-probability event does not allow us to use Bayes rule to derive the posterior beliefs. However, the following system can support the pooling equilibrium:

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

This is a very simple system (on and off the path): if the employer observes any unexpected level of education, he decides that the worker is the low-ability type. Any deviation from e^* will then cause him to offer $w(e) = 1 < w(e^*) = 1 + p$. Clearly, these beliefs will prevent any deviation to $e > e^*$ for either type. We also know that $e^* \leq p$ ensures that the low-ability worker would not want to stay uneducated. It is easy to see that this means that the high-ability worker would not want to say uneducated either: $U_L(e^*) \geq U_L(0) \Rightarrow U_H(e^*) > U_H(0)$. To see that, observe that $U_L(e^*) \geq U_L(0) \Rightarrow e^* \leq p$. This now means $U_H(e^*) = 1 + p - e^*/2 \geq 1 + p - p/2 = 1 + p/2 > 1 = U_H(0)$. In other words, no type will want to deviate. Therefore, any $e^* \leq p$ can be supported in a pooling equilibrium with these beliefs.⁸

By a logic analogous to the one above, there is no reason to expect workers to invest in any education if doing so would reveal nothing about their abilities. This argument would eliminate all pooling PBE except the one with $e^* = 0$ where no worker invests anything in education and the employer offers a wage is higher than the minimum but less than the maximum. How much higher depends on his prior belief that the worker is of high ability: the stronger this belief, the higher the wage he will offer. Note that uncertainty here hurts the high-ability worker (who gets a salary lower than what the employer would have offered if he knew his type) and benefits the low-ability one (who gets a salary higher than what the employer would have offered if he knew his type). This will generally be the case in these signaling games.

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

Observe now that the model seems to be making two completely opposite predictions about the instrumental value of education. If we predict the separating PBE, the our conclusion would be that education is a very useful signaling device. If, on the other hand, we predict the pooling PBE, then we conclude that education is completely useless as a signaling device. So which is it?

It would be nice if we could eliminate one of the equilibrium types. The obvious candidate for that are the pooling PBE because in them the employer's beliefs are suspect: to prevent the high-ability worker from attempting to reveal his type by choosing a higher level of education, the employer threatens that whenever he sees very high education levels, he will infer that the worker is the low-ability type. This inference seems implausible.

Let me make this a bit more intuitive. Suppose e^* is a high-school diploma, and $e = 3/2$ is a Master's Degree. The employer's belief essentially says, "I expect to see resumes where high-school diploma is the educational level and will believe that the candidate is of high ability with probability p ; if I see a master's degree, I will conclude that the candidate is of low ability for sure." This seems incredible: since acquiring education is costly, the only type who could potentially profit by getting more is the high-ability type. The employer is "threatening with incredible beliefs" in much the same way players could threaten with incredible actions off the equilibrium path. What beliefs would be more reasonable?

Suppose that $p = 1/5$ and you are the employer and you get a resume with a Master's degree and a cover letter that states:

I know that you now think I have low ability because I acquired a Master's degree. However, suppose you believed that I am the high-ability type instead. You would offer me the high wage $w = 2$. If I am really the high-ability type, my payoff will be $2 - (3/2)(1/2) = 5/4$. If, on the other hand, I am the low-ability type, my payoff would be $2 - 3/2 = 1/2$. If I had invested in a high-school diploma only, you would have offered me the wage $w = 1 + p = 6/5$. Observe now that I can potentially profit from a Master's degree only if I really am the high-ability type because $5/4 > 6/5 > 1/2$. In other words, if I were the low-ability type, then I would not profit from getting a Master's degree even if doing so were to convince you that I was the high-ability type. I would never acquire a Master's degree if I am the low-ability type.

⁸Again, there are other beliefs that can work here but they all induce the same probability distribution over the outcomes to all resulting PBE are essentially equivalent.

Therefore, I could only have acquired it because I am the high-ability worker. Therefore, my Master's degree is a credible signal that I am the high-ability worker. Therefore, you should update to believe that I am the high-ability worker. (Which means you should offer me the high wage.)

In fact, for any $p < 1$, we can always find $e \in [1, 2]$ which the high-type would prefer to choose if doing so would convince the employer that he is the high type but that the low-type would not pick. This argument (known as the Intuitive Criterion) would eliminate the pooling equilibria. We say that the pooling PBE are *unintuitive*, and therefore should not be considered as substantive predictors of play. This leaves the separating PBE. Of course, we already saw that a similar forward induction logic eliminates all but one of these as well. We are then left with a very sharp prediction: education is useful and high-ability workers will acquire the minimum level that is sufficient to deter the low-ability types from getting it.

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.

Finally, this result may be normatively troubling for it suggests that low-ability workers will be doomed to lower wages and education is the institution that enforces this inequality. If it is true that there is no intrinsic value-added to education, then universities are simply perpetuating the wealth inequalities associated with abilities. Now, you may think that it's ok that high-ability workers always earn more than low-ability ones. I, on the other hand, prefer to think that the University can educate a previously low-ability person and turn him/her into a high-ability worker that a firm can hire at a higher wage. Whatever you believe, it is obvious that we should both support strict standards for University education: if standards lapse, even low-ability types will be able to acquire it, and this will force high-ability types to aim even higher to get employers to offer better wages. We have, in fact, seen this already. With the decline of high-school education, more and more employers started to require bachelor's degrees before they would offer better salaries. As more students without adequate background flooded universities, many of them reacted by watering down the requirements so they can keep graduation rates up. The predictable result is that now you need to go for a Master's degree because a Bachelor's is no longer informative of quality. But obtaining a Master's degree is very expensive and it may not really improve your skills much beyond what a good Bachelor's degree can. It seems it will be in our common interest to strengthen the University requirements. Now this is how you can stretch a formal model way beyond its capacity! But hey, don't let anyone tell you that we are not interested in policy recommendations!

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. 7 (p. 23) from notes by Muhamet Yildiz. 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. This is very convenient for it ensures that player 2's information set will always be reached with positive probability in any equilibrium, so we do not have to worry about off-the-path beliefs. We conclude that 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.

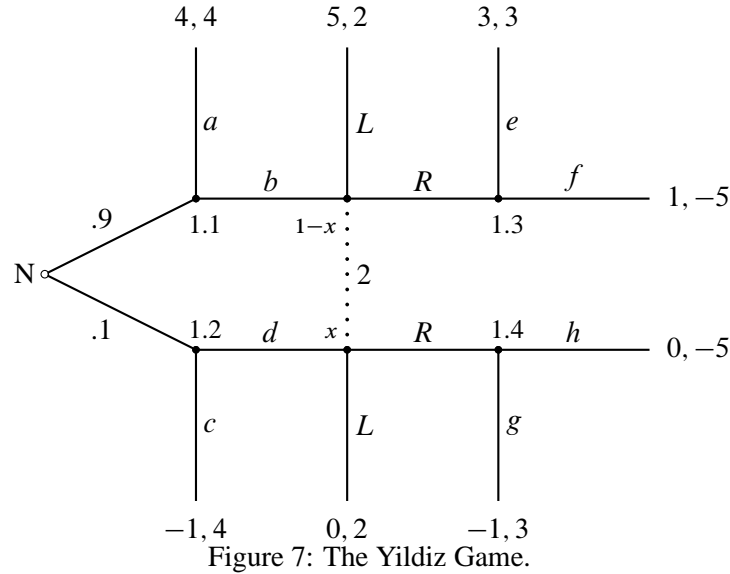


Figure 7: The Yildiz Game.

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.

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 1 is the informed party (knows the outcome of the chance move by Nature) and player 2 is the one who has incomplete information. Player 2 will attempt to infer information from player 1's actions and because the players have somewhat conflicting interests, player 1 obfuscates the inference by playing a mixed strategy (which prevents player 2 from learning with certainty what he knows). Since the informed player moves first, this is an instance of a signaling game. The game in this section reverses this: the first mover is the uninformed player now and he must take an action that would induce the other player to reveal some information. Since the preferences are again somewhat conflicting, player 2 will have incentives to obfuscate this inferences in her turn, making the screening process harder for player 1.




Figure 8: The Myerson-Rosenthal Game.

The game is depicted in Fig. 8 (p. 24). The interpretation is as follows. Players take turns being generous or selfish until someone is selfish or both have been generous twice. Each player loses \$1 by being generous, but gains \$5 each time the other player is generous. (So actions s , S , s' , and S' are selfish, but g , G , and g' , and G' are generous.) The catch is that player 1 is unsure whether player 2 is capable of being selfish: he estimates that with probability $19/20$ she can be selfish but with (small) probability $1/20$ she is the virtuous kind whose integrity compels her to be generous regardless of player 1's behavior. That is, she always chooses to be generous whenever she has to move. Of course, player 2 knows her own type.

At his first information set, player 1 believes that player 2 is virtuous with probability $1/20$. Let y denote his (posterior) belief that she is virtuous after they have taken two generous actions. Observe now that at her last information set 2.2, the selfish player 2's only sequentially rational choice is S' , which means that in any PBE she will always be selfish there. We now have to find the rest of the strategies and beliefs.

Suppose player 1 chose g' with certainty in equilibrium. The only way this would be sequentially rational is if the expected payoff from s' did not exceed the expected payoff from g' given 2's sequentially rational strategy, or if $4 \leq 8y + 3(1 - y)$, which requires $y \geq 1/5$. Because 1 is choosing g' for sure, player 2's expected payoff from choosing G at 2.1 is 9, which is strictly better than getting 5 by playing S , and so she would certainly choose G . Given that she would choose G , player 1's expected payoff from choosing g at his first information set would be $(1/20)(8) + (19/20)(3) = 3.25$, which is strictly greater than 0, which is what he would get by playing s . Therefore, he would choose g for sure. But this means that player 1's second information set is now along the path of play, and Bayes rule gives

$$y = \frac{(1/20)(1)}{(1/20)(1) + (19/20)(1)} = 1/20 < 1/5,$$

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

Suppose player 1 chose s' 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 s' for sure, player 2 would certainly choose S at 2.1 because the expected payoff is strictly greater. Given her sequentially rational strategy, choosing g would yield player 1 the expected payoff of $4(1/20) + (-1)(19/20) = -3/4$. Hence, the sequentially rational choice at this information set is s . 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 s and s' 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 S at 2.1 and S' at 2.2.

We have a continuum of PBE 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 g instead of s . Then, using Bayes rule his posterior belief would have to be:

$$y = \frac{(1/20)\epsilon}{(1/20)\epsilon + (19/20)\epsilon\sigma_2(G)} = 1$$

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 S , and so $\sigma_2(G) = 0$). 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. However, let's see how the idea of full consistency would eliminate the PBE we just found. The posterior belief y is given by:

$$y = \frac{(1/20)\sigma_1(g)}{(1/20)\sigma_1(g) + (19/20)\sigma_1(g)\sigma_2(G)} = \frac{1}{1 + 19\sigma_2(G)},$$

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

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

However, as we have seen, player 2's only sequentially rational strategy is to play S with certainty, and so $\sigma_2(G) = 0$, 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,

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

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

$$\begin{aligned} & (19/20)\left[(-1)(1 - \sigma_2(G)) + (4\sigma_1(s') + 3(1 - \sigma_1(s')))\sigma_2(G)\right] + (1/20)\left[4\sigma_1(s') + 8(1 - \sigma_1(s'))\right] \\ &= (19/20)\left[(-1)(15/19) + (4(4/5) + 3(1/5))(4/19)\right] + (1/20)\left[4(4/5) + 8(1/5)\right] = 1/4. \end{aligned}$$

Because this expected payoff is strictly greater than 0, which is what player 1 would get if he chose s , sequential rationality requires that he chooses g with certainty. We conclude that the following strategies and beliefs constitute a perfect Bayesian (and the unique sequential) equilibrium with $y = 1/5$:

- Player 1 chooses g with probability 1, and s' with probability $4/5$;
- Player 2 chooses G with probability $4/19$, and S' with probability 1.

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 may be 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 selfish player 2 would reciprocate with

higher probability because she wants player 1 to update his beliefs to an even higher probability that she is virtuous, which would induce him to be generous the second time around, at which point she would defect and reap her highest payoff of 9. Notice how in this PBE player 1's posterior belief went from $1/20$ up to $y = 1/5$. Of course, the selfish 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 has a pot of money worth 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. 9 (p. 27).




Figure 9: 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. 10 (p. 28). This is a slightly modified version of a game in notes by David Myatt.




Figure 10: The Three Player Game.

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

$$3(1 - \sigma_2(L)) + \sigma_2(L) \left[5(1/3) + 2(2/3) \right] = 3.$$

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.

3.6 The Perjury Trap Game

This one is from notes by Jean-Pierre Langlois. All similarities to any people, living or dead, or any events, in Washington D.C. or elsewhere, are purely coincidental. A prosecutor, whom we shall call (randomly) Ken, is investigating a high-ranking government official, whom we shall call (just as randomly) Bill. A young woman, Monica, has worked for Bill and is suspected of lying earlier to protect him. Ken is considering indicting Monica but he is really after the bigger fish: he has reason to believe that Monica holds some evidence concerning Bill and is hoping to get her to cooperate fully by offering her immunity. The problem is that he cannot be sure that she will, in fact, cooperate once granted immunity and even if she does cooperate, the evidence she has may be trivial. However, since her testimony will force Bill to take a public stand, Ken hopes to trap him into perjury or at least into admitting his guilt. Monica is most afraid of being discredited and, all else equal, would rather not lie. She really wants to be vindicated if she tells the truth or else to see Bill admit to all the facts. Bill, of course, wants to avoid getting trapped or

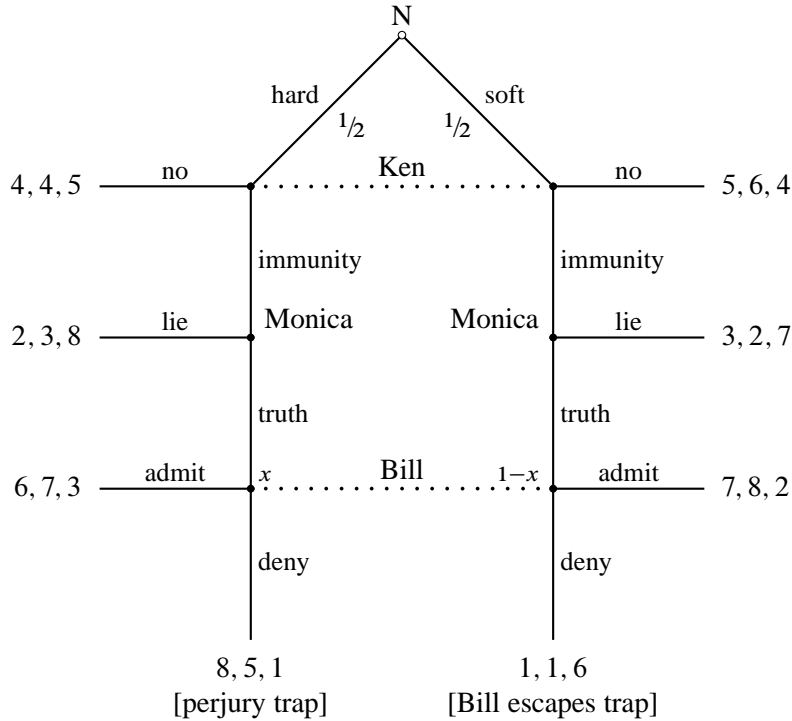


Figure 11: The Langlois Perjury Game.

admitting to any transgressions. Assuming that both Ken and Bill estimate that there's a 50 : 50 chance of Monica's evidence being hard, Fig. 11 (p. 31) shows one possible specification of this game.

We begin by finding the sequentially rational strategies for the players. Bill will deny whenever the expected payoff from denying, $U_B(D)$, exceeds his expected payoff from admitting, $U_B(A)$. Let x denote Bill's belief that the evidence is hard when he takes the stand. Then,

$$U_B(D) = x(1) + (1 - x)(6) = 6 - 5x$$

$$U_B(A) = x(3) + (1 - x)(2) = 2 + x,$$

so he will deny whenever $6 - 5x > 2 + x \Rightarrow x < 2/3$. That is, Bill will deny if he believes that the evidence is hard with probability less than $2/3$; otherwise, he will admit guilt. He is, of course, indifferent if $x = 2/3$, so he can randomize.

Turning now to Monica. Although she knows the quality of the evidence she has, she is not sure what Bill will do if she tells the truth. Let p denote the probability that Bill will deny if he is called to testify. If the evidence is hard, Monica will therefore expect to get $5p + 7(1 - p)$ if she tells the truth and 3 if she lies. Observe that her payoff from telling the truth is at least 5, and as such is always strictly better than her payoff from lying. That is, telling the truth strictly dominates lying here. This means that in any equilibrium Monica will always tell the truth if the evidence is hard.

What if the evidence is soft? Lying gives her a payoff of 2, whereas telling the truth gives her an expected payoff of $1p + 8(1 - p) = 8 - 7p$. Therefore, she will tell the truth if $8 - 7p > 2 \Rightarrow p < 6/7$. That is, if Monica knows the evidence is soft, she will tell the truth if she expects Bill to deny it with probability less than $6/7$; otherwise she will lie. (If $p = 6/7$, she is, of course, indifferent and can randomize.)

We can now inspect the various candidate equilibrium profiles by type:

- **Pooling Equilibrium.** Since Monica always tells the truth when the evidence is hard, the only possible pooling equilibrium is when she also tells the truth if the evidence is soft. Suppose that in

equilibrium Monica tells the truth when the evidence is soft. To make this sequentially rational, it has to be the case that $p \leq 6/7$. If Ken offers immunity, he will expect her to tell the truth no matter what, so his expected payoff from doing so is:

$$U_K(I) = (1/2)[8p + 6(1 - p)] + (1/2)[1p + 7(1 - p)] = (1/2)(13 - 4p).$$

If he decides not to offer immunity, then his expected payoff is $U_K(N) = (1/2)(4) + (1/2)(5) = (1/2)(9)$. Hence, he will offer immunity whenever $U_K(I) \geq U_K(N)$, or when $13 - 4p \geq 9 \Rightarrow p \leq 1$. That is, no matter what Bill does, Ken will always offer immunity. If that's the case, Bill cannot update his beliefs: Ken offers immunity and Monica tells the truth regardless of the quality of evidence. Therefore, $x = 1/2$, which implies that Bill will, in fact, deny for sure (recall that he does so for any $x < 2/3$). Hence, $p = 1$, which contradicts the requirement $p \leq 6/7$, which is necessary to get Monica to tell the truth when the evidence is soft. This is a contradiction, so such an equilibrium cannot exist.

- **Separating Equilibrium.** Since Monica always tells the truth when the evidence is hard, the only such equilibrium involves her lying when it is soft. Suppose that in equilibrium Monica lies when the evidence is soft. To make this sequentially rational, it has to be the case that $p \geq 6/7$. If Ken offers immunity, he expects a payoff of:

$$U_K(I) = (1/2)[8p + 6(1 - p)] + (1/2)(3) = (1/2)(9 + 2p).$$

We already know that his expected payoff from not making an offer is $(1/2)(9)$, so he will prefer to offer immunity whenever $9 + 2p \geq 9 \Rightarrow p \geq 0$. That is, no matter what Bill does, Ken will always offer immunity. This now enables Bill to infer the quality of the evidence with certainty: since Ken offers immunity no matter what but Monica only tells the truth if the evidence is hard, if Bill ever finds himself on the witness stand, he will know that the evidence must be hard for sure; that is $x = 1$. In this case, his sequentially rational response is to admit guilt (recall that he does so whenever $p > 2/3$), which means $p = 0$. But this contradicts the requirement that $p \geq 6/7$, which is necessary to get Monica to lie when the evidence is soft. This is a contradiction, so such an equilibrium cannot exist.

- **Semi-separating Equilibrium.** Since Monica always tells the truth when the evidence is hard, the only such equilibrium involves her mixing when the evidence is soft. Suppose that in equilibrium Monica mixes when the evidence is soft. To make this sequentially rational, it has to be the case that $p = 6/7$, which means that Bill must be mixing as well, which implies $x = 2/3$. Let q denote the probability that Monica tells the truth when the evidence is soft. If Ken offers immunity, he expects a payoff of:

$$\begin{aligned} U_K(I) &= (1/2)[8p + 6(1 - p)] + (1/2)[q(1p + 7(1 - p)) + (1 - q)(3)] \\ &= (1/2)[9 + 2p + 2q(2 - 3p)]. \end{aligned}$$

As before, his expected payoff from making no offer is $(1/2)(9)$, which means that he will prefer to offer immunity whenever $9 + 2p + 2q(2 - 3p) \geq 9 \Rightarrow p + q(2 - 3p) \geq 0$. Using $p = 6/7$, this reduces to $q \leq 3/2$. In other words, he will offer immunity no matter what probability q Monica uses. This now pins down Bill's posterior belief by Bayes' rule:

$$x = \frac{(1/2)(1)(1)}{(1/2)(1)(1) + (1/2)(1)q} = \frac{1}{1 + q}.$$

Because Bill is willing to mix, we know that $x = 2/3$. Substituting this in the equation above and solving for q yields: $q = 1/2$. This is the unique PBE.

Therefore the following strategies constitute the unique perfect Bayesian equilibrium of the Perjury Game:

- Ken always offers immunity;
- Monica tells the truth if the evidence is hard, and tells the truth with probability $\frac{1}{2}$ if the evidence is soft;
- Bill denies with probability $\frac{6}{7}$, believes that the evidence is hard with probability $\frac{2}{3}$.

The gamble is worth Ken's while: the probability of catching Bill in the perjury trap equals the likelihood of Monica having hard evidence, $\frac{1}{2}$, times the likelihood that Bill denies the allegations, $\frac{6}{7}$, for an overall probability of $\frac{3}{7}$, or approximately 43%. Bill is going to have a hard time in this game.