

Perfect Bayesian Equilibrium in Sender-Receiver Games

<i>Introduction</i>	<u>1</u>
<i>Sender-receiver games</i>	<u>1</u>
<i>Strategies in sender-receiver games</i>	<u>4</u>
<i>Sender's best-response strategies</i>	<u>5</u>
<i>Receiver's best-response strategies</i>	<u>5</u>
Updating the Receiver's beliefs	<u>5</u>
Message-wise optimization	<u>6</u>
<i>Bayesian equilibrium</i>	<u>8</u>
<i>Perfect Bayesian equilibrium</i>	<u>11</u>
<i>The test of dominated messages</i>	<u>13</u>

Introduction

We previously studied static games of imperfect information.¹ Each player $i \in I$ has private information which was summarized by her type $\theta_i \in \Theta_i$. Each player knows her own type but does not in general know the types of her opponents. Each player i 's beliefs about the types $\theta_{-i} \in \Theta_{-i}$ of her opponents are derived from her knowledge of her own type θ_i and a common prior belief $p \in \Delta(\Theta)$ over the space of type profiles.

Nature moves first, picking a type profile $\theta \in \Theta$ according to the probability distribution $p \in \Delta(\Theta)$. Nature then privately informed each player $i \in I$ of her type: player 1 is type θ_1 , player 2 is type θ_2, \dots , player n is type θ_n . Then each player $i \in I$ of the n players simultaneously chooses an action $a_i \in A_i$ from her action space. A payoff $u_i(a, \theta)$ is then awarded to each player, which depends on the action profile $a \in A$ the players chose and the type profile θ Nature chose.

It was because the players simultaneously chose their actions that we called these games static. Now we want to generalize our analysis by considering dynamic games of incomplete information; i.e. we consider games in which some players take actions before others and these actions are observed to some extent by some other players.

Sender-receiver games

We consider here the simplest dynamic games of incomplete information: *sender-receiver* games. There

© 1996 by Jim Ratliff, <jim@virtualperfection.com>, <<http://virtualperfection.com/gametheory>>.

¹ See the "Static Games of Incomplete Information" handout.

are only two players: a Sender (S) and a Receiver (R). The Sender's action will be to send a *message* $m \in M$ chosen from a *message space* M to the Receiver. The Receiver will observe this message m and respond to it by choosing an action $a \in A$ from his action space A .

To make this game a simple but nontrivial game of incomplete information we endow the Sender with some private information which we describe by her type $\theta \in \Theta$. The Receiver has no private information, so he has but a single type, which we then have no need to mention further. The Receiver does have prior beliefs (i.e. prior to observing the Sender's message) about the Sender's type, which are described by the probability distribution $p \in \Delta(\Theta)$ over the Sender's type space Θ . In other words, before observing the Sender's message, the Receiver believes that the probability that the Sender is some particular type $\theta \in \Theta$ is $p(\theta)$.

We will typically assume that the type space Θ , the message space M , and the Receiver's action space A are finite sets.

After the Receiver takes an action $a \in A$, each player is awarded a payoff which can in general depend on the message m the Sender sent, the action a the Receiver took in response, and the type θ which Nature chose for the Sender. The payoffs to the Sender and Receiver to a (message, action, type) triple $(m, a, \theta) \in M \times A \times \Theta$ are $u(m, a, \theta)$ and $v(m, a, \theta)$, respectively. I.e. $u, v: M \times A \times \Theta \rightarrow \mathbb{R}$.

We can express this game of incomplete information as an extensive-form game of imperfect information by explicitly representing Nature, who chooses a type $\theta \in \Theta$ for the Sender. Because the Sender observes this choice of Nature, every Sender information set is a singleton, and the number of Sender information sets is equal to the number of possible Sender types, viz. $\#\Theta$. The Receiver observes only the message sent by the Sender. Therefore the number of Receiver information sets is equal to the number of possible messages the Sender can transmit, viz. $\#M$. Within each of his information sets the Receiver cannot distinguish between the Sender's possible types; so each Receiver information set has a number of nodes equal to the number of possible Sender types, viz. $\#\Theta$. Therefore the total number of Receiver nodes is the product of the cardinalities of the message and type spaces, viz. $\#M \cdot \#\Theta$.

Example: Life's a Beach?

For example, the Sender might be a job applicant and the Receiver an employer. The Sender's decision could be a choice between going to College or going to the Beach. The particular choice the Sender makes is her message.² The employer will observe this decision and decide to Hire or Reject the applicant. In this example the Sender's message space is $M = \{\text{College, Beach}\}$ and the Receiver's action space is $A = \{\text{Hire, Reject}\}$.

The Sender's private information could concern her aptitude: she knows whether she is Bright or

² Of course in this case the message seems more than a mere message. The terms "message" for the Sender and "action" for the Receiver both refer to actions taken by a player. The distinction between the two is only interpretational. We use "message" for the Sender's action to acknowledge that the Sender realizes that the Receiver will respond to the Sender's action, and therefore the Sender can attempt to influence the Receiver's response through her choice of message.

Dull. Therefore her type space would be $\Theta = \{\text{Bright, Dull}\}$. The Receiver's prior beliefs concerning the probability that the Sender is Bright or Dull can be described by a single number $\gamma \in [0, 1]$. With probability γ , the Sender is Bright; with probability $1 - \gamma$, she is Dull. I.e. the Receiver's prior beliefs $p \in \Delta(\Theta)$ are defined by $p(\text{Bright}) = \gamma$ and $p(\text{Dull}) = 1 - \gamma$.

Consider the simple sender-receiver game shown in Figure 1. Note that the Sender has two information sets, corresponding to her two types (viz. Bright and Dull). The Receiver also has two information sets, but these correspond to the Sender's two possible messages (viz. Beach and College) rather than to the Sender's possible types. (The Receiver's left-hand information set is his Beach information set and his right-hand information set is his College information set.)

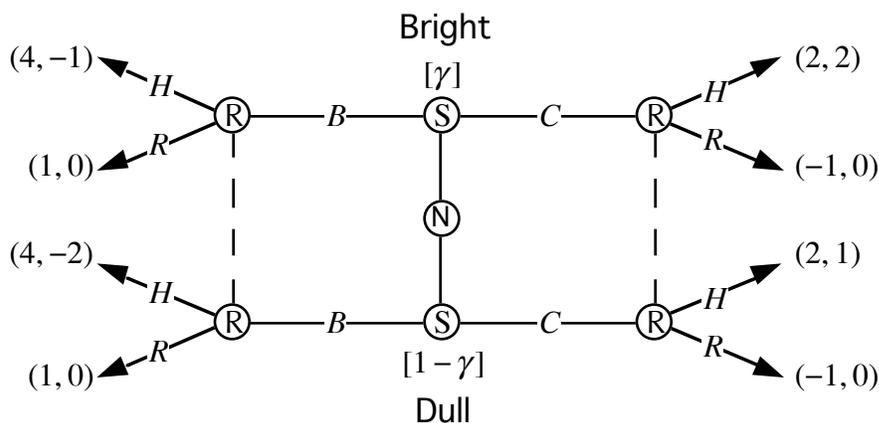


Figure 1: A simple sender-receiver game.

Let's interpret the payoffs shown in Figure 1. The first and second payoffs of each ordered pair are the Sender's and Receiver's payoffs, respectively, for a particular type/message/action triple. For a fixed type and Receiver action, the Sender's payoff to going to the Beach is always two greater than her payoff to going to College.³ For fixed educational and employment decisions, the Sender's payoff is independent of her type.⁴ For a fixed type and educational decision, the Sender receives a payoff from being Hired which is 3 greater than her payoff if she is Rejected.⁵ To summarize the Sender's payoffs (with the appropriate *ceteris paribus* qualifications): the Sender prefers the Beach over going to College, prefers being Hired over being Rejected, and is not discriminated against due to aptitude.

Whenever the Receiver Rejects an applicant, the Receiver gets a payoff of zero. Although the Sender's aptitude did not directly influence the Sender's payoffs, aptitude is payoff-relevant to the Receiver when he Hires: For a fixed educational decision, the Receiver's payoff to Hiring is 1 greater

³ For example, If the Bright applicant is Hired, she receives a payoff of 4 from the Beach but only 2 from College. If the Dull applicant is Rejected, she receives a payoff of 1 from the Beach but only -1 from College.

⁴ For example, if the applicant goes to the Beach and is Hired, she receives a payoff of 4 regardless of whether she is Bright or Dull.

⁵ For example, if the Bright applicant goes to College, she receives 2 if she is Hired and only -1 if she is Rejected. If the Dull applicant goes to the Beach, she receives 4 if she is Hired and only 1 if she is Rejected.

when the Sender is Bright than when she is Dull.⁶ Whereas the Sender disliked going to College, the Receiver appreciates a hired Sender's higher education: For a fixed type, the Receiver's payoff to Hiring is 3 greater when the Sender went to College rather than to the Beach.⁷ In fact, the Receiver's appreciation of a Sender's College education is so great that the Sender's educational decision is decisive in determining whether the Receiver should Hire or Reject: For both types of Sender, the Receiver's payoff to Hiring exceeds his payoff to Rejecting if and only if the Sender went to College.

Strategies in sender-receiver games

A pure strategy for a player in any extensive-form game is a mapping from her information sets to her available actions at the relevant information set. There is a one-to-one correspondence between the Sender's information sets and her type space Θ . Therefore a pure strategy for the Sender is a map $\bar{m}: \Theta \rightarrow M$ from her type space Θ to her message space M . There is a one-to-one correspondence between the Receiver's information sets and the Sender's message space. Therefore a pure strategy for the Receiver is a mapping $\bar{a}: M \rightarrow A$ from the Sender's message space M to the Receiver's action space A .

We can also define behavior strategies for the players. The Sender can send mixed messages. Let $\mathcal{M} \equiv \Delta(M)$ be the set of probability distributions over the Sender's message space M . A behavior strategy for the Sender is a map $\sigma: \Theta \rightarrow \mathcal{M}$ from her type space Θ to mixtures over her message space. Therefore, for all types $\theta \in \Theta$, $\sigma(\theta) \in \mathcal{M}$ is a mixture over messages. In particular, for any message $m \in M$, we denote by $\sigma(m|\theta)$ the probability, according to the Sender behavior strategy σ , that a type- θ Sender will send the message m .

For a given Sender strategy σ , a message m is *on the path* if, according to σ , there exists a type θ who sends m with positive probability. The set of on-the-path messages for Sender strategy σ is

$$M^+(\sigma) = \{m: \exists \theta \in \Theta, \sigma(m|\theta) > 0\} = \bigcup_{\theta \in \Theta} \text{supp } \sigma(\theta). \quad (1)$$

The Receiver can also randomize his actions in response to his message observation. Let $\mathcal{A} \equiv \Delta(A)$ be the set of probability distributions over the Receiver's action space. Then a behavior strategy for the Receiver is a map $\rho: M \rightarrow \mathcal{A}$ from the Sender's message space M to the mixed-action space \mathcal{A} .⁹ Therefore, for all messages $m \in M$, $\rho(m) \in \mathcal{A}$ is a mixture over Receiver actions. In particular, for any action $a \in A$, we denote by $\rho(a|m)$ the probability, according to the Receiver behavior strategy ρ , that the Receiver will choose the action a conditional on having observed the message $m \in M$.

⁶ If the Sender goes to College, the Receiver's payoff to Hire is 2 when the Sender is Bright but only 1 when the Sender is Dull. If the Sender goes to the Beach, the Receiver's payoff to Hire is -1 when the Sender is Bright and -2 when she is Dull.

⁷ For example, if the Receiver hires the Dull Sender, the Receiver gains a payoff of 1 if the Sender went to College compared to a payoff of -2 if the Sender had gone to the Beach instead.

⁸ Note that the support of $\sigma(\theta)$ is the set of messages which a type- θ Sender sends with positive probability when she is playing according to the behavior strategy σ .

⁹ The symbols σ and ρ were selected for these behavior strategies to be mnemonically friendly—i.e. in the hope that “sigma” and “rho” would suggest “Sender” and “Receiver,” respectively.

Sender's best-response strategies

We first ask: When is a Sender strategy $\bar{m} \in M^\Theta$ a best response to some Receiver behavior strategy $\rho \in \mathcal{A}^M$?¹⁰ Consider the case where a type- θ Sender chooses to send a message $m \in M$ knowing that the Receiver will respond according to his behavior strategy $\rho \in \mathcal{A}^M$. This Sender's expected utility will be a convex combination of her payoffs to particular pure actions by the Receiver, viz. $\sum_{a \in A} \rho(a|m) u(m, a, \theta)$. A pure-strategy $\bar{m} \in M^\Theta$ will be a best response for the Sender to a Receiver behavior strategy $\rho \in \mathcal{A}^M$ if, for every type $\theta \in \Theta$ of Sender, the message specified by $\bar{m}(\theta)$ maximizes the expected utility of a type- θ Sender given that the Receiver will respond to the message $\bar{m}(\theta)$ according to the strategy ρ . For a given Receiver mixed strategy $\rho \in \mathcal{A}^M$, the set of optimal messages for a type- θ Sender is

$$\tilde{M}(\rho, \theta) \equiv \arg \max_{m \in M} \sum_{a \in A} \rho(a|m) u(m, a, \theta). \quad (2)$$

Therefore a Sender strategy $\bar{m} \in M^\Theta$ is a best response to the Receiver strategy $\rho \in \mathcal{A}^M$ if and only if, for all $\theta \in \Theta$, $\bar{m}(\theta) \in \tilde{M}(\rho, \theta)$. A sender behavior strategy $\sigma \in \mathcal{M}^\Theta$ is a best response to the Receiver behavior strategy $\rho \in \mathcal{A}^M$ if and only if, for all $\theta \in \Theta$, $\text{supp } \sigma(\theta) \subset \tilde{M}(\rho, \theta)$.

Receiver's best-response strategies

Now we ask: When is a Receiver strategy $\bar{a} \in A^M$ a best response to a Sender behavior strategy $\sigma \in \mathcal{M}^\Theta$?

Updating the Receiver's beliefs

The Receiver chooses an action after he observes the Sender's message. He wants to choose the action which is optimal given the best beliefs he can have concerning the Sender's type. The Receiver enters the game with prior beliefs $p \in \Delta(\Theta)$ concerning the Sender's type θ . Because the Receiver knows the Sender's type-contingent message-sending strategy $\sigma \in \mathcal{M}^\Theta$, the Receiver might be able to infer something more about the Sender's type and thereby update his beliefs.

As long as the observed message is not totally unexpected given that the Sender is playing the behavior strategy σ —i.e. there is some type which, according to σ , sends that message with positive probability—we can use Bayes' Rule to update the Receiver's prior beliefs $p \in \Delta(\Theta)$. Specifically, for any observed on-the-path message $m \in M^+(\sigma)$, we denote the Receiver's posterior belief that the Sender is type θ by $p^B(\theta|m)$, which is given from Bayes' Rule by

$$p^B(\theta|m) \equiv \frac{p(\theta) \sigma(m|\theta)}{\sum_{\theta' \in \Theta} p(\theta') \sigma(m|\theta')}. \quad (3)$$

We see the justification for the restriction to on-the-path (not totally unexpected) messages: if some

¹⁰ For any sets A and B , A^B is the set of all functions from $B \rightarrow A$.

¹¹ The numerator of the right-hand side is the probability of the event "the Sender is type θ and sends message m ." The denominator is the probability that message m is sent.

message m is never sent regardless of which type the Sender is, then the denominator of the right-hand side will vanish.

In general we can define the Receiver's posterior beliefs even after observing off-the-path (and therefore totally unexpected) messages. For every message $m \in M$ we let $\tilde{p}(m) \in \Delta(\Theta)$ be the Receiver's posterior beliefs—after observing the message m —about the Sender's type. In other words, the Receiver attaches the probability $\tilde{p}(\theta | m)$ to the event that the Sender has type $\theta \in \Theta$ conditional upon the Receiver observing the message $m \in M$. So $\tilde{p}: M \rightarrow \Delta(\Theta)$ is a conditional posterior belief system.

Where does the Receiver's conditional posterior belief system $\tilde{p}: M \rightarrow \Delta(\Theta)$ come from? It is derived from the Receiver's prior beliefs $p \in \Delta(\Theta)$ and updated in response to his observation of the Sender's message m . We require that this updating be done according to Bayes' Rule *whenever possible*. This means that, for all on-the-path messages $m \in M^+(\sigma)$ and for all types $\theta \in \Theta$, $\tilde{p}(m | \theta) = p^B(m | \theta)$. Alternatively but equivalently, we can say that the conditional posterior belief system $\tilde{p}: M \rightarrow \Delta(\Theta)$ is consistent with Bayes' Rule if the restriction of \tilde{p} to the on-the-path messages $M^+(\sigma)$ is p^B .

Message-wise optimization

Consider a Receiver pure strategy $\hat{a} \in A^M$. If a type- θ Sender sends the message $m \in M$ and the Receiver responds according to his pure strategy \hat{a} , the Receiver's payoff will be $v(m, \hat{a}(m), \theta)$. The probability with which he receives this particular payoff is the probability of the event “the Sender is type θ and sends message m .” This probability is the probability that the Sender is type θ , viz. $p(\theta)$, multiplied by the probability that the Sender sends the message m conditional on the Sender being type θ , viz. $\sigma(m | \theta)$. Therefore the expected utility $V(\hat{a}, \sigma)$ to the Receiver who plays the strategy $\hat{a} \in A^M$ against the Sender behavior strategy $\sigma \in \mathcal{M}^\Theta$ is the sum of the probability-weighted payoffs $p(\theta)\sigma(m | \theta)v(m, \hat{a}(m), \theta)$ over all possible combinations of messages and types:

$$V(\hat{a}, \sigma) \equiv \prod_{m \in M} \prod_{\theta \in \Theta} p(\theta)\sigma(m | \theta)v(m, \hat{a}(m), \theta). \quad (4)$$

A Receiver strategy $\bar{a} \in A^M$ will be a best response to the Sender behavior strategy $\sigma \in \mathcal{M}^\Theta$ if and only if it maximizes the Receiver's expected utility over all possible Receiver pure strategies; i.e.

$$\bar{a} \in \arg \max_{\hat{a} \in A^M} V(\hat{a}, \sigma). \quad (5)$$

At first glance the optimization problem in (5) might appear problematic because it requires maximization over a function space. Fortunately the maximand, from (4), is additively separable in the various messages $m \in M$, so we'll be able to construct a best-response Receiver strategy $\bar{a} \in A^M$ via a message-by-message optimization to find individual best-response actions $\bar{a}(m)$ for each message $m \in M$. This simplification is justified by the following Lemma which you are invited to prove for yourself.

Lemma

Let A be a set and M be a finite set. Let f be a function $f: M \times A \rightarrow \mathbb{R}$. Then

$$\bar{a} \in \arg \max_{\hat{a} \in A^M} \prod_{m \in M} f(m, \hat{a}(m)) \quad (6)$$

if and only if, for all $m \in M$,

$$\bar{a}(m) \in \arg \max_{a \in A} f(m, a). \quad (7)$$

To apply this Lemma to the optimization problem (5) we define

$$f(m, a) \equiv \prod_{\theta \in \Theta} p(\theta) \sigma(m | \theta) v(m, a, \theta). \quad (8)$$

Now we have, from (4),

$$V(\hat{a}, \sigma) \equiv \prod_{m \in M} f(m, \hat{a}(m)). \quad (9)$$

Therefore, from (5), (9), and the Lemma, we see that the Receiver strategy $\bar{a} \in A^M$ is a best response to the Sender behavior strategy $\sigma \in \mathcal{M}^\Theta$ if and only if, for all $m \in M$,

$$\bar{a}(m) \in \arg \max_{a \in A} f(m, a). \quad (10)$$

If a message m is off the path, i.e. $m \in M \setminus M^+(\sigma)$, then it is sent by no type: for every type $\theta \in \Theta$, $\sigma(m | \theta) \equiv 0$. Therefore, $\forall m \in M \setminus M^+(\sigma)$, $\forall a \in A$, $f(m, a) \equiv 0$. Therefore all actions $a \in A$ are maximizers of $f(m, a)$ when m is an off-the-path message. I.e. $\forall m \in M \setminus M^+(\sigma)$,

$$A = \arg \max_{a \in A} f(m, a). \quad (11)$$

When $m \in M^+(\sigma)$ is an on-the-path message, it is useful to divide the maximand of (10) by the guaranteed-nonzero probability that m is sent, viz. $\sum_{\theta' \in \Theta} p(\theta') \sigma(m | \theta')$. This does not change the set of maximizers of (10). This division allows us, using (8) and (3), to express the condition (10) in terms of the Receiver's Bayes-updated beliefs: $\forall m \in M^+(\sigma)$,

$$\bar{a}(m) \in \arg \max_{a \in A} \prod_{\theta \in \Theta} p^B(\theta | m) v(m, a, \theta). \quad (12)$$

But this maximand is simply the Receiver's expected utility, given her Bayes-updated beliefs about the Sender's type, when she chooses the action $a \in A$ after observing the on-the-path message $m \in M^+(\sigma)$. Therefore condition (12) states that it is necessary and sufficient—in order that the Receiver strategy $\bar{a} \in A^M$ to be a best response to the Sender behavior strategy $\sigma \in \mathcal{M}^\Theta$ —that it specify for each on-the-path message an action which is a best response to that message given the Receiver's Bayes-updated beliefs.

For a given conditional posterior belief system $\tilde{p}: M \rightarrow \Delta(\Theta)$, the Receiver's expected utility to the action $a \in A$ conditional upon having observed the message $m \in M$ is $\sum_{\theta \in \Theta} \tilde{p}(\theta | m) v(m, a, \theta)$. Therefore for

a given conditional posterior-belief system \tilde{p} , the set of Receiver best-response actions to some message m is given by

$$\tilde{A}(\tilde{p}, m) = \arg \max_{a \in A} \prod_{\theta \in \Theta} \tilde{p}(\theta | m) v(m, a, \theta). \quad (13)$$

A Receiver pure strategy $\bar{a} \in A^M$ is a best response to the Sender behavior strategy $\sigma \in \mathcal{M}^\Theta$ if and only if, for all $m \in M$, $\bar{a}(m) \in \tilde{A}(p^B, m)$. A Receiver behavior strategy $\rho \in \mathcal{A}^M$ is a best response to the Sender behavior strategy $\sigma \in \mathcal{M}^\Theta$ if and only if, for all $m \in M$, $\text{supp } \rho(m) \subset \tilde{A}(p^B, m)$.

Bayesian equilibrium

Definition A Bayesian equilibrium of the sender-receiver game is a triple $(\sigma, \rho, \tilde{p}) \in \mathcal{M}^\Theta \times \mathcal{A}^M \times (\Delta(\Theta))^M$ satisfying the following three conditions:

1 For all types $\theta \in \Theta$,

$$\text{supp } \sigma(\theta) \subset \tilde{M}(\rho, \theta), \quad (14)$$

2 For all on-the-path messages $m \in M^+(\sigma)$,

$$\text{supp } \rho(m) \subset \tilde{A}(\tilde{p}, m), \quad (15)$$

3 The conditional posterior belief system \tilde{p} is consistent with Bayes' Rule whenever possible in the sense that the restriction of \tilde{p} to the on-the-path messages $M^+(\sigma)$ is p^B .

Note that optimality from the Receiver is required only at on-the-path information sets. Therefore the only Receiver information sets at which the specification of the conditional posterior belief system \tilde{p} enters into the definition of Bayesian equilibrium is at on-the-path-message information sets, where these beliefs are just the ones derived from Bayes' Rule from (3).

Example: Bayesian equilibrium in Life is a Beach?

We consider further the game of Figure 1. In order to make the Receiver's posterior belief system explicit on the extensive form, we indicate by bracketed probabilities at each Receiver node the Receiver's belief at each of his information sets that he is located at that node conditional on having reached that information set. See Figure 2. For example, if the Receiver observes Beach, then $s \in [0, 1]$ is the probability the Receiver attaches to the event that the Sender is Bright. If the Receiver observes College, then $1 - t$ is the probability the Receiver attaches to the event that the Sender is Dull.

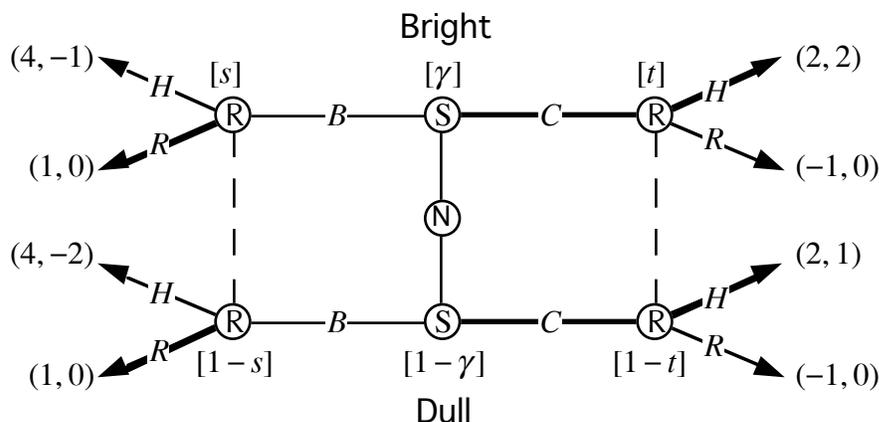


Figure 2: Life is a Beach? with the Receiver’s conditional posterior beliefs indicated.

We can represent a strategy profile by the ordered sextuple $(X, Y; L, R; s, t)$, where

- $X =$ Sender’s action if Bright,
- $Y =$ Sender’s action if Dull,
- $L =$ Receiver’s action if Beach is observed,
- $R =$ Receiver’s action if College is observed,
- $s =$ Receiver’s belief (probability), given that Beach is observed, that the Sender is Bright,
- $t =$ Receiver’s belief (probability), given that College is observed, that the Sender is Bright.

Consider the following strategy profile: $(C, C; R, H; *, \gamma)$.¹² This strategy profile is depicted in Figure 2 by the thick line segments. (We note that according to this strategy profile the Beach message is never sent by any type of Sender and is therefore off-the-path. Therefore in order to evaluate whether this strategy profile is a Bayesian equilibrium we need not specify conditional posterior beliefs for the Receiver at this information set. Hence the “*” in the above specification.)

Let’s verify that this strategy profile is a Bayesian equilibrium of this game. First we check whether any type of Sender wishes to deviate away from going to College in favor of going to the Beach instead, given the hiring policies of the Receiver. Each type of Sender receives a payoff of 2 from conforming to College. Each would receive a lower payoff of 1 instead if she went to the Beach. Therefore neither type of Sender would deviate.

To check whether the Receiver would prefer to change his hiring policy given the Sender’s type-contingent strategy we need only check the only on-the-path information set, viz. the College information set. The easy way to see that Hiring is optimal at the College information set is to notice that

¹² I.e. both types of Sender go to College. The Receiver Rejects any Beachgoers and Hires any College graduates. If the Receiver observes College, he believes that the probability is γ that the Sender is Bright.

Hiring is better for the Receiver than Rejecting for each type of Sender separately. Therefore regardless of the Receiver’s belief t , the corresponding convex combination of Hiring payoffs will exceed the zero he would get if he Rejects. More formally... For any Receiver beliefs $\gamma \in [0, 1]$ that the Sender is Bright conditional on observing College, the Receiver’s expected payoff to Hiring at the College information set is $2\gamma + (1 - \gamma) > 0$. Therefore for any $\gamma \in [0, 1]$ the specified strategy profile is a Bayesian equilibrium.

The specification of posterior beliefs at the College information set, viz. $t = \gamma$, implies that, even after observing the Sender’s message, the Receiver’s beliefs about the Sender’s type are unchanged from her prior beliefs. This no-updating result occurs because this is a *pooling* strategy profile—i.e. all types of the Sender send the same message. We can also use (5) to see formally that this specification $t = \gamma$ is consistent with Bayes’ Rule. (This is the last step in verifying that the strategy profile is a Bayesian equilibrium.) Letting $m = \text{College}$ and $\theta = \text{Bright}$,

$$\begin{aligned}
 p^B(\text{Bright} | \text{College}) &= \frac{p(\text{Bright}) \sigma(\text{College} | \text{Bright})}{p(\text{Bright}) \sigma(\text{College} | \text{Bright}) + p(\text{Dull}) \sigma(\text{College} | \text{Dull})} \\
 &= \frac{\gamma \cdot 1}{\gamma \cdot 1 + (1 - \gamma) \cdot 1} = \gamma.
 \end{aligned}
 \tag{16}$$

I.e. $\tilde{p}(\text{Bright} | \text{College}) = t = \gamma = p^B(\text{Bright} | \text{College})$, exactly as required by condition 3 for Bayesian equilibrium.

Now let’s look at another strategy profile: $(B, B; R, R; \gamma, *)$. This strategy profile is indicated below in Figure 3. Each type of Sender is sending the optimal message, given the Receiver’s hiring policy, by choosing Beach. (Because each type of Sender will be Rejected whatever message she sends, she’ll choose the most pleasant message, viz. go to the Beach.) To check the optimality of the Receiver’s hiring plans we need to check only the single on-the-path information set, viz. Beach. Uneducated Senders aren’t worth hiring, so Rejection at this information set is optimal for the Receiver. You can also verify, similarly to the demonstration for the strategy profile of Figure 2, that the specification $s = \gamma$ is consistent with Bayes’ Rule.

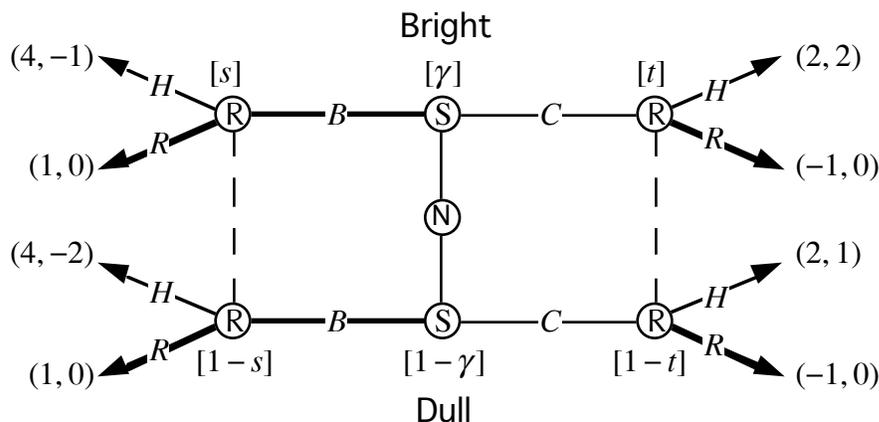


Figure 3: A less credible Bayesian equilibrium.

However, note why the above strategy profile’s specification for the Sender is a best response to the

Receiver's hiring plans: Each type of Sender eschews College because the Receiver plans to Reject college-educated applicants. But regardless of the type of Sender the Receiver would be better off Hiring, rather than Rejecting, a college-educated applicant. No matter what off-the-path posterior belief $t \in [0, 1]$ we specified, Hiring would be the unique best response for the Receiver at his College information set. This equilibrium is undesirable because it relies on an incredible off-the-path action by the Receiver.

Perfect Bayesian equilibrium

We saw in the above example that the strategy profile depicted in Figure 3 was a Bayesian equilibrium of the game but was suspect because it relied on a nonoptimal action by the Receiver at an off-the-path Receiver information set. We can eliminate this strategy profile by a simple strengthening of our solution concept.

Definition A perfect Bayesian equilibrium of the sender-receiver game is a triple $(\sigma, \rho, \tilde{p}) \in \mathcal{M}^\Theta \times \mathcal{A}^M \times (\Delta(\Theta))^M$ satisfying the following three conditions:

1 For all types $\theta \in \Theta$,

$$\text{supp } \sigma(\theta) \subset \tilde{M}(\rho, \theta), \quad (17)$$

2 For all messages $m \in M$,

$$\text{supp } \rho(m) \subset \tilde{A}(\tilde{p}, m), \quad (18)$$

3 The conditional posterior belief system \tilde{p} is consistent with Bayes' Rule whenever possible in the sense that the restriction of \tilde{p} to the on-the-path messages $M^+(\sigma)$ is p^B .

Note that the only difference between this definition of perfect Bayesian and the earlier definition of Bayesian equilibrium is in the strengthening of the original Receiver-optimality condition (15)—which imposed optimality only at on-the-path-message information sets—resulting in (18), which requires optimality of the Receiver's strategy at all message information sets. Note from (13) that this also implies that now the Receiver's posterior beliefs are important even at off-the-path-message information sets. However, we aren't constrained by Bayes' Rule in the specification of these off-the-path beliefs.

The strategy profile from Figure 3 would fail to be a perfect Bayesian equilibrium regardless of how we specified $t \in [0, 1]$ because, as we saw in the analysis of the example of Figure 2, for any beliefs Hiring is better for the Receiver at the College information set is better than Rejecting there.

Also note that if all messages are on the path then if the strategy profile is a Bayesian equilibrium it is also a perfect Bayesian equilibrium.

Example: Perfect Bayesian equilibria can still be undesirable

Consider the same basic game we've been considering but with the different payoffs shown in Figure 4. Note that for a fixed hiring decision each type of Sender prefers going to Beach over going to College, but the Bright Sender finds College less onerous than the Dull Sender does. In fact this difference is extreme in the following sense: A Bright Sender is willing to incur the cost of College if it means that it makes the difference between being Hired and being Rejected.¹³ However, the Dull Sender finds College such a drag that she's unwilling to skip the Beach regardless of the effect her action has on the hiring decision of the Receiver.

For a fixed education decision the Receiver prefers to Hire the Bright Sender but prefers to Reject the Dull Sender. For a fixed type of Sender, the Receiver is indifferent between hiring a College-educated vs. a Beach-tanned Sender. Note that with this payoff structure education is unproductive. But because going to College has a higher cost for the lower-ability type of Sender, education might provide a costly signal of the Sender's type to the Receiver.

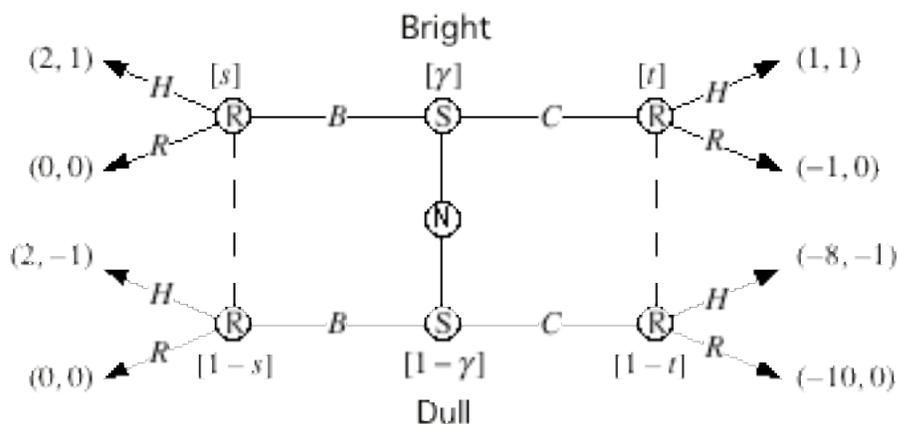


Figure 4: Education is unproductive but an effective signal of ability.

Consider the strategy profile $(C, B; R, H; 0, 1)$. This is not only a Bayesian equilibrium but also a perfect Bayesian equilibrium (because every message is on the path). This is a *separating* equilibrium because each type of Sender chooses a different action. (When each type of Sender sends a distinct message, the Receiver can deduce with certainty the identity of the Sender from her observed message.) You can use (3) to verify that the posterior-belief assignments $s = 0$ and $t = 1$ are those determined by Bayes' Rule.

Let $\gamma \in [0, 1/2)$. Consider the strategy profile $(B, B; R, R; \gamma, t)$, where $t \in [0, 1/2)$. This is a *pooling* strategy profile. This is a perfect Bayesian equilibrium. The off-the-path posterior beliefs imply that if a defection to College is observed, the defector is more likely to be Dull than Bright. However, such off-the-path beliefs are objectionable for the following reason: No matter what influence a deviation to

¹³ If going to the Beach implies that the Bright Sender will be Rejected, then going to the Beach implies a payoff of zero. If going to College is necessary to be Hired, then College implies a payoff of 1. Therefore the Bright Sender will go to College if that is necessary for being Hired.

College might have on the Receiver's hiring decision, the Dull Sender would never find going to College worthwhile. However, the Bright Sender would be willing to go to College if that convinced the Receiver that the Sender was indeed Bright and therefore should be hired.

The test of dominated messages

We saw in the above example that the pooling perfect Bayesian equilibrium profile was undesirable because it relied on the Receiver interpreting a deviation as coming from a type who would never find it optimal to deviate. The College message was *dominated* for the Dull type in the following sense: No matter how badly-for-the-Sender the Receiver might respond to the prescribed message Beach and no matter how favorably-for-the-Sender the Receiver might respond to the deviation message College, the Dull Sender would still prefer to send the prescribed message.

Denote the set of Receiver actions which are best responses, conditional on the message m , for *some* conditional posterior beliefs by

$$\hat{A}(m) = \bigcap_{\tilde{p} \in (\Delta(\Theta))^M} \tilde{A}(\tilde{p}, m).$$

(A Sender who sends the message $m \in M$ would never have to worry about a Receiver response which fell outside of the set $\hat{A}(m)$, because such an action would not be a best-response by the Receiver to any posterior belief she could possibly hold.)

Definition Message $m \in M$ is dominated for type $\theta \in \Theta$ if there exists a message $m' \in M$ such that

$$\min_{a \in \hat{A}(m')} u(m', a, \theta) > \max_{a \in \hat{A}(m)} u(m, a, \theta). \quad (19)$$

Definition Let $\psi \equiv (\sigma, \rho, \tilde{p})$ be a perfect Bayesian equilibrium. The equilibrium ψ fails the test of dominated messages if there exist types $\theta', \theta'' \in \Theta$ and an off-the-equilibrium-path message $m \in M \setminus M^+(\sigma)$ such that¹⁴

- 1 The receiver puts positive weight, conditional on m being observed, that the message was sent by type θ' , i.e. $\tilde{p}(\theta' | m) > 0$,
- 2 m is dominated for type θ' , and
- 3 m is not dominated for type θ'' .

Before we can reject an equilibrium because it puts positive weight on a deviant message originating

¹⁴ It is more common to allow m to be any message in M . The test stated here is equivalent because no on-the-equilibrium-path message m could possibly be dominated for a type θ' for whom $\tilde{p}(\theta' | m) > 0$, because, along the equilibrium path, \tilde{p} is derived by Bayes' rule. (I.e. this would imply that $\sigma(m | \theta') > 0$, and thus that, in equilibrium, type θ' were sending a dominated message.) The statement given here simplifies the proof of the theorem to come.

from a type for whom the message is dominated, we must be able to identify a type of Sender for whom this message is *not* dominated. Otherwise this logic would force us to put zero weight on *all* types at this information set, and this would not be a legitimate conditional probability distribution.

We see that the pooling perfect Bayesian equilibrium fail the test of dominated messages.

Example: The separating equilibrium disappears and the pooling becomes reasonable.

Consider the example in Figure 5. Now College, though more costly for the Dull than for the Bright Sender, is not as costly for the Dull Sender as it was in the example of Figure 4. Going to College is no longer dominated for the Dull Sender; she would be willing to go to College if that made the difference between being Hired and being Rejected.

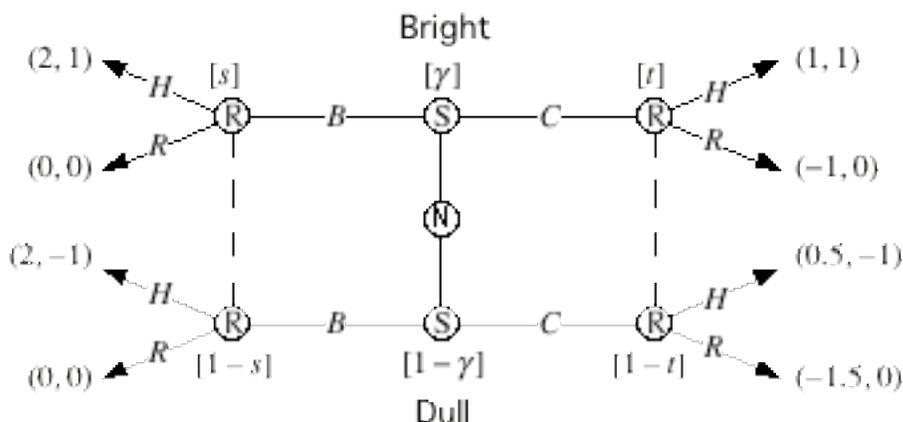


Figure 5: Pooling is now reasonable and separation is not.

The separating strategy profile $(C, B; R, H; 0, 1)$, which was a perfect Bayesian equilibrium in the game of Figure 4, is not an equilibrium of the present game, because the Dull Senders would now deviate to going to College.

The pooling equilibrium $(B, B; R, R; \gamma, t)$, where $\gamma, t \in [0, 1/2)$ of Figure 4 is not only still a perfect Bayesian equilibrium in this game, it is no longer rejected by the test of dominated messages.

Example: The test of dominated messages is not strong enough

Consider the game of Figure 6. The Receiver prefers that the Dull type be uneducated. The Bright Sender actually likes College, while the Dull Sender still finds it degrading. The Receiver strictly prefers to Hire, rather than Reject, a Bright Sender, and finds that College is unproductive when the Sender is Bright. The Receiver is indifferent between Hiring and Rejecting a Dull Beachbum Sender and strictly prefers to Reject a College-educated Dull Sender.

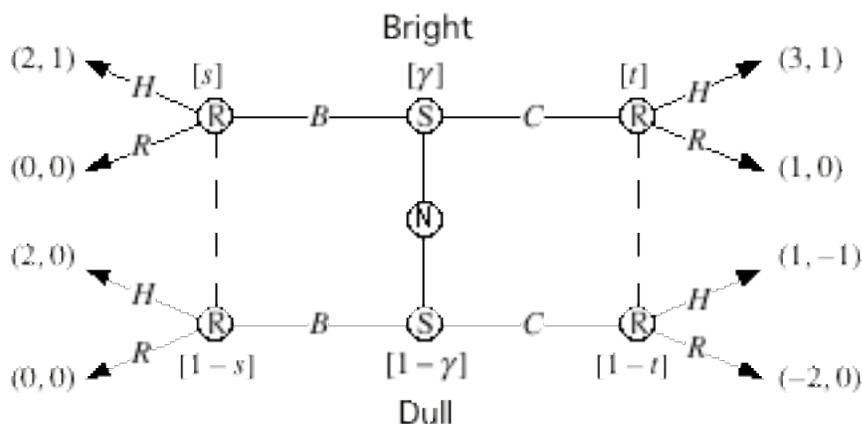


Figure 6: The test of dominated messages is not strong enough.

Consider the following equilibrium: $(B, B; H, R; \gamma, t)$, for $t \in [0, 1/2)$. I.e. if a deviation to College is observed, it is more likely that the deviator is a Dull Sender. This equilibrium passes the test of dominated messages because the Dull Sender could do worse by going to Beach (getting a zero) than by the most optimistic hopes for going to College, where she could get a 1. However, the Bright Sender could hope to gain by deviation relative to his equilibrium potential, but the Dull type cannot hope this. Therefore we shouldn't attribute positive probability to the Dull Sender deviating.

Definition Let $\psi \equiv (\sigma, \rho, \tilde{p})$ be a perfect Bayesian equilibrium. Let $\bar{u}(\theta)$ be the type- θ Sender's expected payoff in this equilibrium. Message $m \in M$ is equilibrium dominated, with respect to ψ , for type $\theta \in \Theta$ if

$$\bar{u}(\theta) > \max_{a \in \hat{A}(m)} u(m, a, \theta). \tag{20}$$

I quickly verify that domination implies equilibrium domination:

Fact If $m \in M$ is dominated for type $\theta \in \Theta$ then, for every perfect Bayesian equilibrium ψ , m is equilibrium dominated with respect to ψ for type θ .

Proof Let $m' \in M$ be a message which dominates m for type θ . For any equilibrium Receiver strategy ρ , the Sender's expected payoff to the message m' is

$$\int_{a \in \hat{A}(m')} \rho(a | m') u(m', a, \theta) \geq \min_{a \in \hat{A}(m')} u(m', a, \theta), \tag{21}$$

which is derived from (18) and $\tilde{A}(\tilde{p}, m') \subset \hat{A}(m')$. For any $m'' \in \text{supp } \sigma(\theta)$,

$$\bar{u} = \int_{a \in A} \rho(a | m'') u(m'', a, \theta). \tag{22}$$

Assume that m is not equilibrium dominated with respect to the equilibrium ψ . Then from (19), the converse of (20), (21), and (22),

$$\prod_{a \in A(m')} \rho(a | m') u(m', a, \theta) > \prod_{a \in A} \rho(a | m'') u(m'', a, \theta). \quad (23)$$

Therefore $m'' \notin \tilde{M}(\rho, \theta)$ —note in (2) that $\rho(a | m') = 0$ for $m' \in A \setminus \hat{A}(m')$ —which contradicts (17). ☺

Definition Let $\psi \equiv (\sigma, \rho, \tilde{p})$ be a perfect Bayesian equilibrium. The equilibrium ψ fails the refinement \mathcal{I} —the Intuitive Criterion—if there exist types $\theta', \theta'' \in \Theta$ and an off-the-equilibrium-path message $m \in M \setminus M^+(\sigma)$ such that^{15,16}

1 The receiver puts positive weight, conditional on m being observed, that the message was sent by type θ' , i.e. $\tilde{p}(\theta' | m) > 0$,

2 m is equilibrium dominated with respect to ψ for type θ' , and

3 m is not equilibrium dominated with respect to ψ for type θ'' .

It is often asserted (or at least strongly suggested) that \mathcal{I} is an equilibrium refinement of \mathcal{D} .¹⁷ However, a perfect Bayesian equilibrium strategy profile can pass the Intuitive Criterion yet fail the test of dominated messages. Yet, if a perfect Bayesian equilibrium survives the Intuitive Criterion, then there exists a perfect Bayesian equilibrium which yields the same outcome (i.e. probability distribution over terminal nodes) and which survives both the test of dominated messages and the Intuitive Criterion. (See Ratliff [1993].)

Reference

- Cho, In-Koo and David M. Kreps [1987] “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics* **102** 2 (May): 179–221.
- Fudenberg, Drew and Jean Tirole [1989] “Noncooperative Game Theory for Industrial Organization: An Introduction and Overview,” in *Handbook of Industrial Organization*, eds. Richard Schmalensee and Robert D. Willig, Vol. 1, North-Holland, pp. 259–327.
- Fudenberg, Drew and Jean Tirole [1991] *Game Theory*, MIT Press.
- Gibbons, Robert [1992] *Game Theory for Applied Economists*, Princeton University Press.
- Kreps, David M. [1990] *A Course in Microeconomic Theory*, Princeton University Press.
- Ratliff, James D. [1993] “A Note on the Test of Dominated Messages and the Intuitive Criterion in

¹⁵ As in the definition of refinement \mathcal{D} , the restriction of m to off-the-equilibrium-path messages is without loss of generality.

¹⁶ For more on these refinements see Cho and Kreps [1987] and Kreps [1990: 436].

¹⁷ For example, Fudenberg and Tirole [1989: 312] say “The idea is roughly to extend the elimination of weakly dominated strategies to strategies which are dominated relative to equilibrium payoffs. So doing eliminates more strategies and thus refines the equilibrium concept further.” Fudenberg and Tirole [1991: 446–447] suggest that replacing the equilibrium path by its payoff results in an equilibrium refinement whose rejection requirements are weaker and easier to apply. Kreps [1990: 436] says that the Intuitive Criterion is a “stronger test” than the test of dominated messages. Gibbons [1992: 213] says that, because equilibrium dominance is easier to satisfy than dominance, the Intuitive Criterion makes the test of dominated messages redundant.

Sender-Receiver Games,” mimeo.