BAYESIAN EQUILIBRIUM

MICHAEL PETERS

The games like matching pennies and prisoner's dilemma that form the core of most undergrad game theory courses are games in which players know each others' preferences. Notions like iterated deletion of dominated strategies, and rationalizability actually go further in that they exploit the idea that each player puts him or herself in the shoes of other players and imagines that the others do the same thing. Games and reasoning like this apply to situations in which the preferences of the players are *common knowledge.* When we want to refer to situations like this, we usually say that we are interested in games of *complete information.*

Most of the situations we study in economics aren't really like this since we are never really sure of the motivation of the players we are dealing with. A bidder on eBay, for example, doesn't know whether other bidders are interested in a good they would like to bid on. Once a bidder sees that another has submit a bid, he isn't sure exactly how much the good is worth to the other bidder. When players in a game don't know the preferences of other players, we say a game has *incomplete information.* If you have dealt with these things before, you may have learned the term asymmetric information. This term is less descriptive of the problem and should probably be avoided.

The way we deal with incomplete information in economics is to use the approach originally described by Savage in 1954. If you have taken a decision theory course, you probably learned this approach by studying Anscombe and Aumann (last session), while the approach we use in game theory is usually attributed to Harsanyi. What it says is that we start to think about these problems by trying to formulate a set of different possibilities, then proceed by assigning *subjective probabilities* to these possibilities, and evaluating them using expected utility. So, for example, when bidding against another player in an eBay auction, we would assign probabilities to the different possible valuations a player might have. If we can figure out what each each of these different *types* of players are doing, we could calculate the probability of winning an auction by finding the subjective probability with which we believe our competitor's bid is below our bid.

Bayesian Nash Equilibrium. There is a generalization of Nash equilibrium that applies to games of incomplete information. This solution concept is called Bayes Nash equilibrium (i.e., an equilibrium is called a Bayesian Nash equilibrium or Bayes Nash Equilibrium).

A *Bayesian* game is a collection

$$\Gamma = \{N, \{A_i\}, \{u_i\}, \{\Theta_i\}, \{F_i\}\}$$

where N is the set of players, A_i is a set of actions available to player i, Θ_i is a set of possible *types* (in the sense of Savage or Harsanyi) player i could have, and

$$u_i : \prod_{i=1}^{N} A_i \times \prod_{i=1}^{N} \Theta_i \to \mathbf{R}$$

is a utility function for player i that describes the payoff that i gets for each array of actions used by the players and for each array of types possessed by the players, F_i is the probability distribution that player i believes describes the joint distribution of players' types.

In a complete information game, all players do is to choose an action (or a mixture over actions). In a Bayesian game, we need to proceed as if players choose a *strategy rule,* which for player i is a mapping $\sigma_i : \Theta_i \to A_i$ that describes the action that player i will use for each of his types. If you want to think about mixed strategies, the player i would have a set of mixtures over his actions, given by $\triangle(A_i)$ and his strategy would be a mapping from Θ_i into $\triangle(A_i)$.

In a Bayes Nash equilibrium, the strategies must all be best replies to one another, just as they are in a Nash Equilibrium. Though that bit is straightforward, an immediate complication probably occurs to you - best replies maximize payoff against another player's action. We can't really guess the other players' actions here, because we don't know their payoffs.

The reason the approach I described above works is that players are assumed to know the strategy *rules* that each of the other players are using. Once the player thinks he knows what action (or mixed strategy) each type of the other player will use, he or she can use their subjective beliefs about the probabilities of the various types of the other players in order to find their own best reply.

Formally, a Bayesian Nash equilibrium is a profile of N strategy rules $\{\sigma_i^*(\cdot)\}$ such that

$$\mathbf{E}_{F_i}\{u_i(\sigma_1^*(\theta_1), \ldots \sigma_N^*(\theta_N), \theta_1, \ldots \theta_N)\} \geq$$
$$\mathbf{E}_{F_i}\{u_i(\sigma_1^*(\theta_1), \ldots \sigma_i'(\theta_i), \ldots, \sigma_N^*(\theta_N), \theta_1, \ldots \theta_N)\}$$

for every $i \in N$ and every feasible strategy rule σ_i'.

This seems an odd way to express it, because each player takes the expected payoff he or she receives across all of his or her possible types in evaluating a strategy. It would seem that a player knows her own type, if nothing else, so a strategy should be evaluated relative to the type she has rather than all her possible types. Part of the resolution of this confusion is that this equality is equivalent to the following one. Let θ_{-i} denote the list of types for players other than i. A set of strategy rules $\{\sigma_i^*(\cdot)\}$ is a Bayesian Nash equilibrium if and only if for each $i \in N$ and each $\theta_i \in \Theta_i$

$$\mathbf{E}_{F_i}\{u_i(\sigma_1^*(\theta_1), \ldots \sigma_i^*(\theta_i), \ldots, \sigma_N^*(\theta_N), \theta_1, \ldots \theta_N)|\theta_i\} \geq$$
$$\mathbf{E}_{F_i}\{u_i(\sigma_1^*(\theta_1), \ldots a', \ldots, \sigma_N^*(\theta_N), \theta_1, \ldots \theta_N)|\theta_i\}$$

for each $a' \in A_i$. The expectation means to take the expected payoff by integrating across the types of the other players using beliefs that are conditional on the type that player i knows she has. (**Problem 1:** Prove these two conditions are equivalent - use the law of iterated expectations.)

When you compute the conditional expectation of a profile θ_{-i} of types for the other players, the way you do it is to use Bayes rule - you divide the probability of the profile of types given by θ_{-i} and i's own type θ_i by the probability that i has type θ_i. For example, if θ_i is finite for each player i, then each profile of actions $\{\theta_i, \theta_{-i}\}$ has probability $f_i(\theta_i, \theta_{-i})$ for player i, with

$$F_i(\theta_i, \theta_{-i}) = \sum_{\theta_i' \leq \theta_i} \sum_{\theta_{-i}' \leq \theta_{-i}} f_i(\theta_i', \theta_{-i}').$$

With this restriction, the conditional probability of a profile of types θ_{-i} of the others is given by

$$\Pr(\theta_{-i}|\theta_i) = \frac{f_i(\theta_i, \theta_{-i}),}{\sum_{\theta_{-i}} f_i(\theta_{-i}', \theta_i)}.$$

This is how a Bayesian equilibrium is usually defined, but this explains the reason why it is referred to as a Bayesian equilibrium - you use Bayes rule when you evaluate the strategies.

0.1. **Example.** Probably the best known example of a simple Bayesian equilibrium with a common prior is the First Price Independent Private Values Auction. In this auction there are N bidders who are going to submit bids. The highest bidder wins the auction and pays his bid. The action space for each player is a continuum, and we'll assume for the example it is $A_i = B = [0,1]$ for all i. In other words, each player i's action is a bid b_i (it seems sensible to use b_i for bid here instead of a_i which we used above to describe a generic action). Payoff types all come from a continuum as well $\Theta = [0,1]^N$. The common prior is that each player's type is independently drawn using a uniform distribution on $[0,1]$. This means that and $F = U_{[0,1]} \times \cdots \times U_{[0,1]}$.

Now we need to describe the payoffs in this game by showing what each player gets given every profile of actions and for each of his or her possible type. This is

$$u_i(b_1, \ldots b_N, \theta_i) = \begin{cases} \frac{\theta_i - b_i}{r+1} \text{ if } b_i \geq b_j \forall j \neq i \text{ and } r = \#\{j \neq i : b_j = b_i\} \\ 0 \text{ otherwise} \end{cases}$$

A Bayesian equilibrium is a collection of strategy rules $\{\sigma_i^*\}$. To find this collection for an auction, a nice way to start is just to guess that each will use *the same* monotonic bidding rule $b^*(\theta)$. In the general notation, we are just guessing that $\{\sigma_i^*\} = \{b^*, b^*, \ldots, b^*\}$. If we can find this rule, then we will have a fully separating equilibrium in which bidders with higher values submit higher bids.

Formally this gives

$$\begin{gathered}\mathbf{E}_F \{u_i(b_1^*(\theta_1), \ldots b_i^*(\theta_i), \ldots, b_N^*(\theta_N), \theta_1, \ldots \theta_N) | \theta_i\} \\ = \int_0^{b^{*-1}(b)} \cdots \int_0^{b^{*-1}(b)} (\theta - b) \, d\theta_2 \ldots d\theta_N = \\ (\theta - b)\left[b^{*-1}(b)\right]^{N-1}.\end{gathered}$$

Here, the reason we know that b^* has an inverse is because it is monotonic.

We can make a second jump by noticing that one way to think about a strategy rule that is part of a Bayesian Nash equilibrium is to acknowledge that if deviating in the bid is unprofitable, then it should also be unprofitable to submit the bid that would have been made by a person with a different valuation. This is the core argument in the revelation principle, which we'll discuss later. Then $(\theta - b(\theta'))[\theta']^{N-1}$ should reach its maximum at $\theta' = \theta$ or

$$b'(\theta) = \frac{(\theta - b(\theta))}{\theta}(N-1).$$

This is a well known ordinary differential equation which has a linear solution $b = \frac{N-1}{N}\theta$ (make sure to check this by computing the derivative in the payoff function above explicitly).

0.2. **Higher Order Beliefs.** Notice a few things that you may not be familiar with. First, the general formulation, unlike the auction example above, players can believe that types of different players are correlated. For example, suppose player 1 has two types L and H, as does player 2. Player i's belief about the joint distribution of these types (F) is given by the following box:

	L	H
L	$\frac{1}{8}$	$\frac{3}{8}$
H	$\frac{3}{8}$	$\frac{1}{8}$

In this box, 1's type is listed at the beginning of each row, while 2's type is listed at the top of each column. The cells give the probabilities of the corresponding pairs. Applying Bayes rule, when player 1 has type L, she thinks that 2 is a type L with probability $\frac{1}{4}$ while when she has type H, she believes that 2 is a type L with probability $\frac{3}{4}$ (Verify this by applying Bayes rule explicitly so you can do the calculation). In other words, different types of the same player can have different beliefs about the types of the other players.

Notice secondly, that unlike this example with correlation, different players could have different beliefs about the joint probability of types. When F_i and F_j are the same for all players, then we usually refer to that as a *common prior* assumption. There is a sense in which assuming a common prior is without loss of generality if you define types broadly enough. However, for our purposes we'll treat common prior beliefs as a special case.

Now the *type* in this formulation actually gives us a lot of information. For example, we know from the box above that a type L player believes that the other player has type L with probability $\frac{1}{4}$ and type H with probability $\frac{3}{4}$. If we stick with the common prior assumption, a type L player also has some beliefs about what the other player believes. For example, he thinks that there is a $\frac{1}{4}$ probability that the other player believes that he has type L with probability $\frac{1}{4}$ and that he has type H with probability $\frac{3}{4}$, and a $\frac{3}{4}$ probability that the other player has the reverse beliefs - i.e., the other player believes that he has type L with probability $\frac{3}{4}$ and type H with probability $\frac{1}{4}$. Similarly, we could describe a player's beliefs about another player's beliefs, and so on.

So when we write down a *type* like L or H, we are in a way just using a shorthand that describes a unending description of the player's beliefs about different events.

Furthermore, there are different ways to describe these events. For example, suppose we add the presumption that players are both using the strategy rule where type L plays action a while type H plays action b. Then the player's type describes his beliefs about what the other player will do. For example, a player of type L believes that the other player will choose action a with probability $\frac{1}{4}$ and action b with probability $\frac{3}{4}$. At the same time, he believes that there is a $\frac{1}{4}$ chance that the other player believes he will take action a with probability $\frac{1}{4}$ and a $\frac{3}{4}$ chance that the other player believes he will take action a with probability $\frac{3}{4}$. You can see that in this simple problem, neither player is entirely sure of the other player's beliefs. This is quite different from the examples we looked at above in which each worker in the directed search story was sure of the other player's beliefs about his type.

There is something else you should notice at this point about this formulation. If you have ever used this stuff before you probably thought about auctions, where types are willingness to pay - players aren't sure how much each other is willing to pay. For instance in the example above, you probably thought that L meant low willingness to pay. However there is another thing that players might not know about one another - their beliefs. In the example above, player 1 could be absolutely sure that player 2 is willing to pay \$100 for a camera on eBay. What he might not know is whether player 2 believes that he has low willingness to pay (equal to \$100 say), or a high willingness to pay equal to \$200. Player 1 might believe that player 2 won't bid against him in an auction, not because he believes that 2 has a low willingness to pay, but because he believes that 2 believes that 1 has a high willingness to pay. Player 1 might simply be unsure about this, even if he is sure of 2's willingness to pay.

Once you open the door to this kind of incomplete information, Bayesian equilibrium can start to lose some of its predictive power. Here is an example that you have probably seen - matching pennies.

	H	T
H	$1,-1$	$-1,1$
T	$-1,1$	$1,-1$

As you know, this game has a single Nash equilibrium where each player mixes with equal probability across his two available actions. Matching pennies, as it is usually described, is a game of complete information. That is really just a special case of Bayesian equilibrium in which each player believes he knows the payoffs of the other player, believes that the other player knows he knows etc, etc. Yet one might imagine that the numbers in the cells describe players' payoff types which each player knows, while at the same time, players have incomplete information about other players' *belief* types. Since there are different ways to describe belief types, lets use beliefs about the actions and beliefs about actions of the other players and use this to show how the outcome HH can be part of a Bayesian Nash equilibrium.

What it means for an outcome to be part of a Bayesian Nash equilibrium is that there is a profile of types for the players such that HH will be the outcome (for sure) if players have those types. To find this, we need to find a belief type for Player 1 that would cause him to choose H for sure. Since there are many ways to express this, lets describe belief types as beliefs about beliefs about .... actions.

In the following script $\mathcal{B}$ means 'believes' while $\Longleftarrow$ means 'because' .

$(a_1 = H) \Longleftarrow$
$1\mathcal{B}\,(a_2 = H) \Longleftarrow$
$1\mathcal{B}2\mathcal{B}\,(a_1 = T) \Longleftarrow$
$1\mathcal{B}2\mathcal{B}1\mathcal{B}\,(a_2 = T) \Longleftarrow$
$1\mathcal{B}2\mathcal{B}1\mathcal{B}2\mathcal{B}\,(a_1 = H)$
$\cdots$ *repeating.*

So the first line in the table above says that the action chosen by player 1 is H because .... In the second row, $1\mathcal{B}\,(a_2 = H)$ means that 1 believes that player 2 will play H. So the two lines together say that player 1 plays H because he believes that player 2 will play H. The second line might be referred to as player 1's *first order belief* (about player 2). Normally, we think of the first order as being a belief about an action or a parameter of another player's utility function u_i.

All the other lines in the table describe beliefs about beliefs. For example, the third line in the table gives a description of player 1's belief about player 2's first order belief. The rows in this table go on for ever describing 1's beliefs the beliefs of the other player to some order. The reason I wrote "repeating" in the last row of the table is because we would just replace $(a_1 = H)$ with $1\mathcal{B}(a_2 = H)$, which is what we put in the second line. So we would just go on adding the strings in the first four rows of the table to the strings we already had. This would describe player 1's beliefs of every order.

Instead of describing 1 by writing out this infinitely long string of symbols, lets just refer to this sequence of statements by saying that 1 has belief type t_H^1. We could describe a similar string of statements for a belief type for player 1 who would want to play T, refer to this string as t_T^1, and do the same thing for two possible types for player 2 - t_H^2 and t_T^2.

Lets just stick with the simple type spaces $\Theta_1 = \{t_H^1, t_T^1\}$ and $\Theta_2 = \{t_H^2, t_T^2\}$. As we have defined it so far, we can construct a Bayesian equilibrium in which both players play H for sure. This is supported with beliefs F_1 and F_2 given respectively by

	t_H^2	t_T^2
t_H^1	1	0
t_T^1	0	0

and

	t_H^2	t_T^2
t_H^1	0	0
t_T^1	1	0

If you are familiar with game theory, you will recognize the logic behind the type construction is sometimes called rationalizability. Every profile of actions is rationalizable in matching pennies. Yet you can see that every rationalizable profile of actions can also be thought of as a Bayesian equilibrium in which players have incomplete information about each others' belief types.

What this means here is that player 1 has a subjective prior belief in which he believes with probability 1 that player 2 has a belief type that leads him to play H. Notice that if player 2 shared the same prior, his belief type would assign probability 1 to player 1 being type t_H^1. Then if we tried to construct the Bayesian equilibrium from this common prior, player 2 should realize that 1's best reply is to play H, not T as he believes. So we could not support this outcome with a common prior.

This isn't very helpful in practise, as any profile of actions for which every action survives iterated deletion of strictly dominated strategies is rationalizable, so can be understood as a Bayesian equilibrium. In fact, if we just take one very minor step in the interpretation of the common prior above, we can say that there is a common prior belief for which the outcome HH is the unique outcome consistent with Bayesian equilibrium. Game theory doesn't have much predictive content if all you can say is players will use rationalizable actions. Nonetheless, higher order beliefs seem - at least at very low levels or reasoning - realistic, so we need to have models in which higher order beliefs are somewhat richer than we assumed in the directed search example. At the same time we want these models to have some predictive content. To get this, we need to make assumptions about the type space *and* the prior belief. So far, the most useful game in which higher order beliefs are important, but also tractable, is something called a *global game.*

0.3. **Example: Email game.** The traditional story has two generals on opposite sides of a valley. General 2 has a big platoon of soldiers he can use to attack, but he doesn't have enough soldiers to defeat the enemy down in the valley by himself. General 1 is waiting for her platoon to arrive. If they do, and the Generals both attack the valley simultaneously they will defeat the enemy, but if General 1 goes it alone she will lose. The Generals are planning a coodinated attack at midday. General 1 has promised to send a text message to General 2 when her platoon arrives.

It is wartime, and text messages don't always work, which both generals know. So each general has installed an app on their phone that automatically replies to text messages saying that they have been received. Of course, these replies might not get through either. If the reply does get through the other general's phone will send an automatic reply to say they have received the reply, etc.

The noon deadline has almost arrived. General 2 keeps looking at his phone waiting for the text. Soon he will have to decide whether to attack or not.

This strange interaction can be modelled as a Bayesian game, and its solution is as strange as the game itself.

The 'payoff types' of the players are actually common knowledge. We can model the payoffs of each player as follows: when the state is 1, the game the players play is given by

	Attack	Don't
Attack	$-1,-1$	$-1,0$
Don't	$0,-1$	$0,0$

In state 2, the game is

	Attack	Don't
Attack	$1,1$	$-1,0$
Don't	$0,-1$	$0,0$

Then they jointly want to attack in state 2, but not attack in state 1.

The belief types of the players aren't common knowledge. We'll try to show that incomplete information about belief types will cause players to be uncertain about each other's actions even when they know the state.

We'll describe belief types with a pair of indices. This will allows us to show how the belief types can be drawn from a common prior distribution. For General 1, one belief type for the general is indexed by the belief type that occurs when she sees that the platoon hasn't arrived. There are a bunch of other belief types for General 1 in which she sees that the platoon *has* arrived and has received n replies to her text message.

For General 2, all his belief types are indexed by the number of text messages he has received. We can put it into a table where each cell gives a probability assigned to the outcome where General 1's belief type is given by the column index of the cell while General 2's belief type is given by the row index.

	No Platoon	$(2,0)$	$(2,1)$	$(2,2)$	$(2,3)$	$(2,4)$
0	$\frac{1}{2}$	$\frac{1}{2}\epsilon$	0	0	0	$\ldots$
1	0	$\frac{1}{2}(1-\epsilon)\epsilon$	$\frac{1}{2}(1-\epsilon)^2\epsilon$	0	0	$\ldots$
2	0	0	$\frac{1}{2}(1-\epsilon)^3\epsilon$	$\frac{1}{2}(1-\epsilon)^4\epsilon$	0	$\ldots$
3	0	0	0	$\frac{1}{2}(1-\epsilon)^5\epsilon$	$\frac{1}{2}(1-\epsilon)^6\epsilon$	$\ldots$
4	0	0	0	0	$\frac{1}{2}(1-\epsilon)^7\epsilon$	$\ldots$
5	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

For example, the probabilty that General 2 has belief type 0 while General 1 has belief type $(2,0)$ is equal to the probability that the state is 2 times the probability that General 1's message is lost. On the other hand the probability that 2 has type 3 while 1 has type $(2,0)$ is zero because General 1 will never send a reply to any message.

If General 2 receives any email messages at all, he must believe that state is 2 with probability 1. That is the only belief that is consistent with Bayes rule. To see this, suppose General 2 gets 3 text messages. Then

$$\Pr(s=2|n=3) = \frac{\Pr(n=3|s=2)\Pr(s=2)}{\Pr(n=3)}$$

$$= \frac{\left((1-\epsilon)^5\epsilon + (1-\epsilon)^6\epsilon\right)\frac{1}{2}}{\frac{1}{2}(1-\epsilon)^5\epsilon + \frac{1}{2}(1-\epsilon)^6\epsilon}$$

To find the equilibrium, we need to figure out what each player would do for each of their possible belief types. One belief type is easy. If General 1 learns that her platoon hasn't arrived by midday, she knows the state is 1. In state 1 General 1 has a strictly dominant strategy which is Don't Attack.

If General 1 sees her platoon arrive, she sends the text. If she doesn't get a reply, there are two possible belief types for General 2, 0 or 1. The probability with which General 1 believes that General 2 is belief type 0 conditional on General 1 having type $(2,0)$ is, by Bayes rule

$$\Pr(0|(2,0)) = \frac{\Pr(0\cap(2,0))}{\Pr((2,0))}$$

To find this, divide the number in the cell $(0,(2,0))$ by the sum of the numbers down the column headed $(2,0)$. This gives

$$\frac{\frac{1}{2}\epsilon}{\frac{1}{2}\epsilon + \frac{1}{2}(1-\epsilon)\epsilon} = \frac{1}{2-\epsilon}$$

General 1 has to decide whether to attack. She believes the state is 2 with probability 1. Yet she doesn't know the belief type of General 2, though she thinks the belief type 0 is more likely than 1.

So we should start by trying to figure out what General 2 would do if his belief type were 0. He knows that if General 1 has belief type "No Troops", then she won't attack. So let suppose we make him as optimistic as possible in the sense that he believes that if the troops do arrive, then General 1 will attack.

Lets use Bayes rule to do the calculation using the joint distribution table above. The probability with which General 2 believes there are no troops conditional on

$n = 0$ (no text message) is

$$\Pr\{\text{"No troops"}|n = 0\} =$$
$$\frac{\Pr\{\text{"No troops"} \cap n = 0\}}{\Pr(n = 0)} =$$
$$\frac{\frac{1}{2}}{\frac{1}{2} + \frac{1}{2}\epsilon} = \frac{1}{1 + \epsilon}.$$

Not surprisingly, if General two doesn't get a message, he thinks it is very likely that the platoon didn't arrive. Then given the optimistic assumption that General 1 will surely attack when the troops arrive, the expected payoff for General 2 if he attack with no text messages is

$$\frac{1}{1 + \epsilon}(-1) + \frac{\epsilon}{1 + \epsilon}1 = \frac{\epsilon - 1}{1 + \epsilon} < 0.$$

Notice that General 2 can't really be sure what General 1 will do, and what we have just shown is that General 2 won't attack without a text message no matter what he thinks General 1 will do. Don't attack is a dominant strategy for General 2 once we realize that don't attack is a dominant strategy for General 1 if there are no troops.

Now we are in 'higher order' land for General 1. She realizes that General 2 realizes that Attack is a dominated strategy for General 1 when there are no troops. This reasoning by will lead General 1 to conclude that "Attack" is a dominated strategy for General 2 when he receives no text message. I'll try not to write a sentence like that again. Notice that if we think of the cells in the beliefs table above as representing possible profiles of belief types for the two players, all the cells in the first column of the table will lead to outcomes where neither player attacks.

At this point, we might as well try to figure out what General 1 will do when she sees the troops have arrived, but hasn't received a message (in other words, $(2, 0)$ is her belief type). She knows the state is two, and she has figured out that General 1 won't attack if he doesn't receive a text. So lets suppose she is really optimistic and believes that General 2 will attack as long as he has received at message.

Again, use the table to compute conditional probabilities. Conditional on her belief type $(2, 0)$ she believes General 2 has belief type 0 with probability

$$\frac{\frac{1}{2}\epsilon}{\frac{1}{2}\epsilon + \frac{1}{2}(1 - \epsilon)\epsilon} = \frac{1}{1 + (1 - \epsilon)} > \frac{1}{2}$$

Then given her optimistic prediction that General 2 will attack as long as he gets a text message, her payoff from attacking is

$$\frac{1}{1 + (1 - \epsilon)}(-1) + \frac{1 - \epsilon}{1 + (1 - \epsilon)} < 0.$$

Now maybe you find this slightly more surprizing. General 1 knows the troops have arrived but hasn't received a reply to her text. She will not attack no matter what she thinks General 1 will do when he does get a message. "Attack" is dominated strategy once we delete dominated strategies for some of General 1's belief types.

At this point we could go on and discuss what General 2 would choose to do if he has received the message that the troops arrived but hasn't received a response to this reply (in other words General 2 has belief type 1. Instead, lets jump right ahead and suppose that General 2 has belief type n (has received n messages from General 1 - use n=3 if you find it helpful to refer to the table above). When General 2 has n messages, he is in somewhat the same predicament as he is when he has only 1 message. He knows that General 1 either has $n-1$ messages or n messages and he believes that conditional on n General 1 has $n-1$ messages with probability

$$\frac{\frac{1}{2}(1-\epsilon)^{2n-1}\epsilon}{\frac{1}{2}(1-\epsilon)^{2n-1}\epsilon+\frac{1}{2}(1-\epsilon)^{2n}\epsilon}=\frac{1}{2-\epsilon}.$$

What is different at this point is that he knows the troops have arrived, because he did get the initial text from General 2.

Notice that *if* n were equal to 1, he would know that if his reply to General 1's initial message is lost, then "Attack" for General 1 would be a dominated strategy. So lets try to extend the approach above. Suppose that he thinks that if General 1's belief type were $(2, n-1)$ then the strategy "Attack" would be dominated for General 1 (we haven't established this yet). At the same time, suppose he holds out the belief that if his message does get through to General 1, that General 1 will surely attack. Then if General 2 attacks, his payoff is

$$(-1)\left(\frac{1}{2-\epsilon}\right)+(1)\left(\frac{1-\epsilon}{2-\epsilon}\right)<0.$$

Now this is starting to get messy, so here is the beliefs table again:

	No Platoon	$(2,0)$	$(2,1)$	$(2,2)$	$(2,3)$	$(2,4)$
0	$\frac{1}{2}$**	$\frac{1}{2}\epsilon$**	0	0	0	$\ldots$
1	0	$\frac{1}{2}(1-\epsilon)\epsilon$**	$\frac{1}{2}(1-\epsilon)^2\epsilon$	0	0	$\ldots$
2	0	0	$\frac{1}{2}(1-\epsilon)^3\epsilon$	$\frac{1}{2}(1-\epsilon)^4\epsilon$	0	$\ldots$
3	0	0	0	$\frac{1}{2}(1-\epsilon)^5\epsilon$	$\frac{1}{2}(1-\epsilon)^6\epsilon$	$\ldots$
4	0	0	0	0	$\frac{1}{2}(1-\epsilon)^7\epsilon$	$\ldots$
5	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

The difference here is that I have added asterisks (**) to each cell in which we have established neither General will attack. Each cell represents a possible pair of belief types. The asterisks that occur in the cell $\{1,(2,0)\}$ follow just by replacing n with 1 in the argument above and noting that "Attack" actually is a dominated strategy for General 1 if her belief type is $(2,0)$.

So lets take the same approach with General 1 and assume that she has received $n>0$ messages. Notice that this only happens if the troops arrive, and General 2 actually knows that they have arrived. Lets also suppose that General 1 knows that if General 2 also has n messages, then the strategy "Attack" for General 2 will be dominated. This is true when, for example, General 1 sees 1 message.

When General 1 has n messages, she believes that General 1 has n messages with probability

$$\frac{\frac{1}{2}(1-\epsilon)^{2n}\epsilon}{\frac{1}{2}(1-\epsilon)^{2n}\epsilon+\frac{1}{2}(1-\epsilon)^{2n+1}\epsilon}=\frac{1}{2-\epsilon}.$$

Given her belief that General 2 won't attack if he also has n messages, the best she can do by attacking is

$$(-1)\frac{1}{2-\epsilon}+(1)\frac{1-\epsilon}{2-\epsilon}<0$$

That calculation is based on the presumption that 2 will attack is he has $n+1$ messages. Hopefully you can see that means that General 1 with type $(2,1)$ won't attack no matter what she thinks 1 will do if the has $n+1$ messages.

Now we are finished. We can put an asterisk in the cell $\{1,(2,1)\}$. Once we do that, we can put an asterisk in the cell $\{2,(2,1)\}$ using the argument we made for General 1. That lets us put an asterisk in the cell $\{2,(2,2)\}$, and so on. Proof by induction.

1. How deeply Do players think

- here are a pair of games invented by Terri Kneeland now at UCL - called ring games
-
  Player 1

	c	d
a	15	0
b	5	10

  Player 2

	e	f
c	15	0
d	5	10

  Player 3

	a	b
e	10	5
f	5	0

-
  Player 1

	c	d
a	15	0
b	5	10

  Player 2

	e	f
c	15	0
d	5	10

  Player 3

	a	b
e	5	0
f	10	5

- each player plays in each role in each game. A fully rational player in game 1 should play e as player 3, c as player 2 and a as player 1. in game two he should play f,d,
- there are at least two other interesting kinds of players - level 1's play the same way as player 1 in both games, but switch actions as player 2.
- level 0's switch actions as player 2, but play the same way in both games as player 1 or 2