Bayesian Games

Debraj Ray, November 2006

Unlike the previous notes, the material here is perfectly standard and can be found in the usual textbooks: see, e.g., Fudenberg-Tirole. For the examples in these notes (except for the very last section), I draw heavily on Martin Osborne's excellent recent text, *An Introduction to Game Theory*, Oxford University Press.

Obviously, incomplete information games — in which one or more players are privy to information that others don't have — has enormous applicability:

credit markets / auctions / regulation of firms / insurance / bargaining /lemons / public goods provision / signaling / . . . the list goes on and on.

1. A Definition

A Bayesian game consists of

1. A set of *players* N.

2. A set of *states* Ω, and a common prior μ on Ω.

3. For each player i a set of *actions* A_i and a set of *signals* or *types* T_i. (Can make actions sets depend on type realizations.)

4. For each player i, a mapping $\tau_i : \Omega \mapsto T_i$.

5. For each player i, a vN-M payoff function $f_i : A \times \Omega \mapsto \mathbb{R}$, where A is the product of the A_i's.

Remarks

A. Typically, the mapping τ_i tells us player i's *type*. The easiest case is just when Ω is the product of the T_i's and τ_i is just the appropriate projection mapping.

B. The prior on states need not be common, but one view is that there is no loss of generality in assuming it (simply put all differences in the prior into differences in signals and redefine the state space and s_i-mappings accordingly).

C. The prior and the (private) observation of one's own type yields a conditional posterior over the state space, given by

$$\mu_i(A, t_i) \equiv \frac{\mu(A \cap \tau_i^{-1}(t_i))}{\mu(\tau_i^{-1}(t_i))},$$

where $\tau_i^{-1}(t_i)$ is interpreted in the usual way as the *set* of all $\omega \in \Omega$ such that $\tau_i(\omega) = t_i$.

A (pure) strategy for player i is a map $s_i : T_i \mapsto A_i$. Define the induced mapping $\hat{s}_j$ on Ω by

$$\hat{s}_j(\omega) = s_j(\tau_j(\omega)).$$

Now we are in a position to define a type-dependent payoff function on the space of strategy profiles:

$$F_i(t_i, \mathbf{s}) = \int_\Omega f_i(s_i(t_i), \hat{s}_j(\omega), \omega) d\mu_i(\omega, t_i).$$

A *Bayesian Nash equilibrium* is a strategy profile $\mathbf{s}^*$ such that for every player i *and* every type t_i of that player,

$$F_i(t_i, \mathbf{s}^*) \geq F_i(t_i, a_i, \mathbf{s}^*_{-i})$$

for every action $a_i \in A_i$.

2. Some Applications

2.1. Introductory Examples. Here are some examples. In all of them, you should take care to go through the necessary formalities of the associated Bayesian game.

Example 1. An opponent of unknown strength. The row player's type is known, but the column player can be "strong" (with probability α) or "weak" (with probability $1 - \alpha$). If he is strong, the situation is

	F	Y
F	$-1, 1$	$1, 0$
Y	$0, 1$	$0, 0$

while if he is weak, the situation is

	F	Y
F	$1, -1$	$1, 0$
Y	$0, 1$	$0, 0$

Draw a tree to represent this game. Also, describe all the ingredients of the associated Bayesian game.

It is easy enough to solve for the Bayesian Nash equilibrium of this game. First note that if the opponent is strong, it is a dominant strategy for him to play F — fight. Now look at Row. If Row fights, he gets 1 if the opponent is weak and — by the dominance argument just made — he gets -1 if the opponent is strong. So his expected return from fighting is $-\alpha + (1 - \alpha)$, while his expected return from playing Y — yielding — is 0.

It follows that if $\alpha > 1/2$, Row will yield so that both types of Column fight. On the other hand, if $\alpha < 1/2$, Row will fight anyway so that the weak type of Column will not fight.

Example 2. (Classical) adverse selection. A firm A is taking over another firm B. The true value of B is known to B but unknown to A; it is only known by A to be uniformly distributed on $[0, 1]$. It is also known that B's value will flourish under A's ownership: it will rise to λx, where x is the pre-takeover value and $\lambda > 1$. All of this description is, in addition, common knowledge.

A strategy for A is a bid y. A strategy for B is a yes-no decision as a function of its value x. So if the type (B, x) accepts, A gets $\lambda x - y$, while (B, x) gets y. If (B, x) rejects, A gets 0 while (B, x) gets x.

It follows, therefore, that (B, x) will accept if $x < y$, while she will reject of $x > y$. Therefore the *expected value* of B — in A's eyes — conditional on B accepting the offer, is $y/2$. It follows that the overall expected payoff to A is

$$\lambda(y/2) - y = y\left[\frac{\lambda}{2} - 1\right].$$

Therefore the acquirer cannot buy unless the takeover target more than doubles in value after the takeover!

Example 3. Information. Sometimes more information can hurt both the informed and uninformed parties. Here is an example. There are two states and two corresponding payoff matrices, given as follows:

	L	M	R
T	1, 2	1,0	1, 3
B	4, 4	0, 0	0, 5

ω_1, prob. 1/2

	L	M	R
T	1, 2	1, 3	1, 0
B	4, 4	0, 5	0, 0

ω_2, prob. 1/2

Suppose, first, that both parties have the same 50-50 prior about the state. Then it is a dominant strategy for player 2 to play L. And player 1 will consequently play B, so that payoffs are $(4, 4)$.

Now suppose that player 2 is informed (player 1 is uninformed, as before). Then player 2 will *always* play M or R depending on the realization of the state. Knowing this, player 1 will play T and payoffs are always $(1, 3)$ regardless of state, a Pareto-deterioration.

Example 4. Infection. Consider the following two-player situation: Suppose that $T_1 = \{a, b\}$

	L	R
U	2, 2	0,0
D	3, 0	1, 1

α, prob. 9/13

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

β, prob. 3/13

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

γ, prob. 1/13

and $T_2 = \{a', b'\}$. Suppose that the type maps are given by

$$\tau_1(\alpha) = a, \quad \tau_1(\beta) = \tau_1(\gamma) = b,$$

and

$$\tau_2(\alpha) = \tau_2(\beta) = a', \quad \tau_2(\gamma) = b'.$$

Interpret this. Player 1 always knows the true game being played. (There is no difference between the payoff matrices at states β and γ.) Player 2 sometimes knows it (when the state is γ, that is).

So in state α, player 1 knows the game, player 2 does not know the game, and player 1 knows this latter fact.

In state β, player 1 knows the game, player 2 does not know the game, but player 1 does not know whether or not player 2 knows the game.

In state γ, players 1 and 2 *both* know the game, 2 knows that 1 knows the game (2 must, because 1 always knows), but 1 doesn't know if 2 knows the game.

So this is a good example to illustrate how "higher-order ignorance" can affect equilibrium in all states. Begin by studying player $(1,a)$. He will play B as a dominant strategy.

Now consider player $(2,a')$. Her posterior on (α,β) is $(3/4,1/4)$. So if $(2,a')$ plays L, her payoff is bounded above by

$$\frac{3}{4}\times 0+\frac{1}{4}\times 2=\frac{1}{2},$$

while if she plays R, her payoff is bounded *below* by

$$\frac{3}{4}\times 1+\frac{1}{4}\times 0=\frac{3}{4}.$$

So $(2,a')$ will definitely play R.

Now look at player $(1,b)$. He assigns a posterior of $(3/4,1/4)$ to (β,γ). In state β he encounters $(2,a')$ who we know will play R. So by playing T, our player $(1,b)$ gets at most

$$\frac{3}{4}\times 0+\frac{1}{4}\times 2=\frac{1}{2},$$

while if he plays B, her payoff is at least

$$\frac{3}{4}\times 1+\frac{1}{4}\times 0=\frac{3}{4}.$$

So player $(1,b)$ also plays B, just like his counterpart $(1,a)$. Now the infection is complete, and player $(2,b')$ must play R! The good outcome (T,L) is ruled out even at state γ, when both players know the game.

2.2. **Juries.** n jury members must decide whether or not acquit or convict a defendant. Everybody has the same payoff function, given as follows:

0 if you take the correct action (acquit when innocent, convict when guilty);

$-z$ if you convict an innocent person;

$-(1-z)$ if you acquit a guilty person.

The magnitude of z provides relative weightings on the jury's tolerance for "type 1" or "type 2" errors in this case.

A nice thing about the payoff magnitude z is that it can be cleanly interpreted as a (psychological) conviction threshold. If a juror feels that a person is guilty with probability r, then expected payoff from conviction is $-(1-r)z$, while the expected payoff from acquittal is $-(1-z)r$, so that she will convict if $r>z$. Thus z can be interpreted as the threshold probability of guilt above which the juror will vote for conviction.

Before a juror gets to vote, she receives a signal of guilt or innocence (this sums up the court evidence and all deliberations). Suppose that the signal is either g or i and the true state is G or I (for "guilty" or "innocent"). Assume the following signal structure:

$$\text{Prob}(g|G)=p>1/2;\quad \text{Prob}(i|I)=q>1/2,$$

where the signals are drawn in a (conditionally) independent way across jurors.

Finally, suppose that there is a prior probability of guilt (again common to all jurors); call it π.

Define a state by the element $\omega = (x; t_1, t_2, \ldots, t_n)$, where x is either G or B and t_k is either g or i for every k. Juror k simply gets the projection t_k. He can then form posteriors on the state space in the usual way.

Begin the analysis with just one juror. Conditional on receiving the signal g, her posterior probability of guilt is

$$\frac{p\pi}{p\pi + (1-\pi)(1-q)},$$

while conditional on receiving the signal i, her posterior probability of guilt is

$$\frac{(1-p)\pi}{(1-p)\pi + (1-\pi)q}.$$

If the juror's conviction threshold z exceeds both these numbers, she will acquit anyway, and if it is lower than both these numbers, she will convict anyway, so the interesting case, of course, is one in which she "acts according to her signal":

$$\frac{p\pi}{p\pi + (1-\pi)(1-q)} > z > \frac{(1-p)\pi}{(1-p)\pi + (1-\pi)q}.$$

What if there are n jurors, and conviction requires unanimity?

The key to thinking about this case is a concept called *pivotality.* A single juror's vote — convict or acquit — has to be placed in the context of what the other jurors are saying. By the unanimity rule, her vote has *no* effect if at least one other juror is voting to acquit, so the only case in which our juror is "pivotal" is one in which every other juror is voting to convict. *By the linearity of expected utility, her optimal action in the pivotal case must coincide with her overall optimal action.* This pivotality logic has bizarre consequences in a world of incomplete information.

Suppose that everybody continues to act according to their signal. Suppose our juror gets an innocent signal i. In the pivotal case, every other juror is voting guilty, so by our assumption they have *all* received signal g. Therefore in the pivotal case, our juror's posterior on guilt must be given by

$$\begin{aligned}
\text{Prob } (G|i; g, \ldots, g) &= \frac{\text{Prob } (i; g, \ldots, g|G)\pi}{\text{Prob } (i; g, \ldots, g|G)\pi + (1-\pi)\text{Prob } (i; g, \ldots, g|I)} \\
&= \frac{p^{n-1}(1-p)\pi}{p^{n-1}(1-p)\pi + q(1-a)^{n-1}(1-\pi)} \\
&= \frac{1}{1 + \frac{q}{1-p}\left(\frac{1-q}{p}\right)^{n-1}\frac{1-\pi}{\pi}}
\end{aligned}$$

which is very, very close to 1 for reasonable values of jury sizes. For instance, if $n = 12$, $\pi = 0.5$, and $p = q = 0.8$, this fraction is 0.999999046!

Indeed, one can show that no matter what the parameters, voting according to signal cannot be an equilibrium as n becomes very large.

How, then, do we characterize equilibrium? There is always an equilibrium in which everyone votes for acquittal, though this is weakly dominated. But the interesting equilibrium is the symmetric mixed strategy equilibrium in which jurors convict for sure if they get a g signal and with some probability β if they get an i signal. Now this dilutes the strength of guilt in the pivotal case; indeed, we have to dilute it just enough so that Prob $(G|i; g, \ldots, g)$ equals

$$\begin{aligned} & \frac{\text{Prob } (i|G)\text{Prob } (\text{vote } C|G)^{n-1}\pi}{\text{Prob } (i|G)\text{Prob } (\text{vote } C|G)^{n-1}\pi + \text{Prob } (i|I)\text{Prob } (\text{vote } C|G)^{n-1}(1-\pi)} \\ = \quad & \frac{(1-p)[p+(1-p)\beta]^{n-1}\pi}{(1-p)[p+(1-p)\beta]^{n-1}\pi + q[(1-q)+q\beta]^{n-1}(1-\pi)} = z! \end{aligned}$$

We may rewrite this as the condition

$$\frac{1}{1+\frac{q}{1-p}\frac{1-\pi}{\pi}Y} = z,$$

where

$$Y \equiv \left[\frac{(1-q)+q\beta}{p+(1-p)\beta}\right]^{n-1}.$$

Rearranging this equation, we see that

$$Y = \frac{(1-z)(1-p)\pi}{(1-\pi)qz} \equiv X,$$

so that — recalling the definition of Y —

$$\frac{(1-q)+q\beta}{p+(1-p)\beta} = X^{1/(n-1)},$$

or

$$\beta = \frac{pX^{1/(n-1)} - (1-q)}{q-(1-p)X^{1/(n-1)}},$$

which is easily seen to converge to 1 as n becomes large. For instance, if $n = 12$, $\pi = 0.5$, $p = q = 0.8$ and $z = 0.7$, $\beta \simeq 0.6$: each juror votes to convict with probability 0.6 *when they get an innocent signal*!

There are several reasons why this model should not be taken very seriously as a model of juries, but the pivotality argument is provocative and powerful and should be looked at carefully.

2.3. **Global Games.** A lot of refinements in game theory use the idea of a perturbation: small mistakes, payoff variations, and so on. The idea is to look at the limit of perturbed equilibria as the perturbations go to zero: this yields a refinement. These perturbations are usually taken to be independent, so that my beliefs regarding the other players' strategies *do not depend on my particular perturbation*.

A very interesting class of perturbations, pioneered by the paper of Carlsson and van Damme (1993), considers a somewhat different scenario. Imagine that the game itself is subject to a "public perturbation": for instance, in Cournot oligopoly, some shock affects the aggregate demand curve (as opposed to, say, the individual cost functions, which could plausibly taken to be independent). Then my (noisy) observation of the shock tells me something about the shock and consequently something about what *you* must have observed, and so my observations affect my beliefs about what you are going to do. But then — putting myself mentally in your shoes — I must also think about *your* beliefs about me (and about others). This creates a "belief drift" which can be controlled over some finite belief chain, but which will generally drift over the *entire* range of perturbations as we go far enough out in the belief chain. I am therefore forced to take account of the full, *global* possibilities in order to formulate a rational belief about what you are going to do. Hence the term global game.

Notice that a global game is an excellent example of a game with incomplete information. Our (correlated) observations form our *types*. The very fact that there is correlation, however, leads us to consider long, drifting chains of types.

In this section I will draw on an application of this idea to financial crises by Morris and Shin (1998). The application will be of interest in itself, but it will also illustrate the global games technique.

Consider a country which has pegged its exchange rate at some value e. (Think of e as the number of dollars required to buy one unit of the domestic currency.) We shall assume that the exchange rate is overvalued, in the following sense: suppose that there is some random variable θ (the state) on $[0,1]$ which determines the "true" exchange rate $f(\theta)$ were the currency to be allowed to float at θ. Then e always exceeds $f(\theta)$ for all $\theta \in [0,1]$.

But θ also influences the exchange rate: which is to say that $f(\theta)$ varies with θ. Arrange so that $f(\theta)$ is strictly increasing in θ. So the idea is that θ is some "fundamental" which influences the country's capacity to export or import, or to attract investment; the higher being θ, the more favorable the climate.

Now there is a bunch of speculators (of total measure 1), each of whom can sell one unit of the local currency. If they do, they pay a transactions cost t. If the government holds the peg, the exchange rate stays where it is, and the payoff to selling is $-t$. If the government abandons the peg, then the speculators make a profit of $e - f(\theta)$, so their net payoff is $e - f(\theta) - t$.

What about the government's decisions? It has only one decision to make: whether to abandon or to retain the peg. We assume that it will abandon the peg if the measure of speculators exceeds $a(\theta)$, where a is increasing in θ (that is, if the basic health of the economy is better, the government is more reluctant to abandon[1]).

We will assume that there is some positive value of θ, call it $\underline{\theta}$, such that below $\underline{\theta}$ the situation is so bad that the government will abandon the peg anyway. In other words we are assuming that $a(\theta) = 0$ for $\theta \in [0, \underline{\theta}]$. Then it rises but always stays less than one by assumption.

Consider, now, a second threshold for θ which we'll call $\bar{\theta}$: this is the point above which no one wants to sell the currency even though she feels that the government will abandon the

[1]See Morris and Shin (1998) for a very simple account of how to derive $a(\theta)$ from a somewhat more basic starting point.

peg for sure. In other words, $\bar{\theta}$ solves the equation

$$e - f(\bar{\theta}) - t = 0. \tag{1}$$

We will assume that such a $\bar{\theta}$, strictly less than one, indeed exists. But we also suppose that there is a gap between $\underline{\theta}$ and $\bar{\theta}$: that $\underline{\theta} < \bar{\theta}$.

[If there were no such gap, there wouldn't be a coordination problem to start with.]

Now we are ready to begin our discussion of this model. First assume that the realization of θ is perfectly observed by all agents, and that this information is common knowledge. Then there are obviously three cases to consider.

Case 1. $\theta \leq \underline{\theta}$. In this case, the government will abandon the peg for sure. The economy is not viable, all speculators must sell, and a currency crisis occurs.

Case 2. $\theta \geq \bar{\theta}$. In this case no speculator will attack the currency, and the peg will hold for sure.

Case 3. $\underline{\theta} < \theta < \bar{\theta}$. Obviously this is the interesting case, in which multiple equilibria obtain. There is an equilibrium in which no one attacks, and the government maintains the peg. There is another equilibrium in which everyone attacks and the government abandons the peg. This is a prototype of the so-called "second-generation" financial crises models, in which expectations — over and above fundamentals — play an important role.

So much for this standard model. Now we drop the assumption of common knowledge of realizations (but of course we maintain the assumption of common knowledge of the information *structure* that I am going to write down).

Suppose that θ is distributed uniformly on $[0, 1]$: its value will be known perfectly at the time the government decides whether or not to hold the peg or to abandon it. Initially, however, the realization of θ is noisy in the following sense: each individuals sees a signal x which is distributed uniformly on $[\theta - \epsilon, \theta + \epsilon]$, for some tiny $\epsilon > 0$ (where θ is the true realization). Conditional on the realization of θ, this additional noise is iid across agents.

Proposition **1.** *There is a unique value of the signal x such that an agent attacks the currency if $x < x^*$ and does not attack if $x > x^*$.*

This is an extraordinary result in the sense that a tiny amount of noise refines the equilibrium map considerably. Notice that as $\epsilon \to 0$, we are practically at the common knowledge limit (or are we? the question of what sort of convergence is taking place is delicate and important here), yet there is no "zone" of multiple equilibria! The equilibrium is unique.

What is central to the argument is the "infection" created by the lack of common knowledge (of realizations). To see this, we work through a proof of Proposition 1, with some informal discussion.

Start by looking at the point $\underline{\theta} - \epsilon$. Suppose that someone receives a signal x of this value or less. What is she to conclude? She doesn't know what everyone else has seen, but she *does* know that the signal is distributed around the truth with support of size 2ϵ. This means that the true realization *cannot* exceed $\underline{\theta}$, so that the government will abandon the peg for sure. So she will sell. That is, we've shown that for all

$$x \leq x_0 \equiv \underline{\theta} - \epsilon,$$

it is dominant to sell.

Now pick someone who has a signal *just* bigger than x_0. What does he conclude? Suppose, for now, he makes the assumption that *only* those with signals less than x_0 are selling; no one else is. Now what is the chance — given his own signal x — that someone else has received a signal not exceeding x_0? To see this, first note that the true θ must lie in $[x-\epsilon, x+\epsilon]$. For each such θ the chances that the signal for someone else is below x_0 is $(1/2\epsilon)[x_0 - (\theta - \epsilon)]$, so the overall chances are just these probabilities integrated over all conceivable values of θ, which yields $(1/2\epsilon)[x_0 - (x-\epsilon)]$. So the "infection" spreads: if x is close to x_0, these chances are close to 1/2. In this region, moreover, it is well known that the government's threshold is very low: close to zero sellers (and certainly well less than half the population) will cause an abandonment of the peg. Knowing this, such an x must sell. *Knowing that all with signals less than x_0 must sell, we have deduced something stronger: that some signals above x_0 must also generate an attack.*

So let us proceed recursively: Suppose we have satisfied ourselves that for some index n, everyone sells if the signal is no bigger than x_n (we already know this for x_0). We define x_{n+1} as the *largest* value of the signal for which people will want to sell, knowing that all below x_n are selling.

This is a simple matter to work out. First, fix any value of θ. Then everybody with a signal between $\theta - \epsilon$ and x_n (such an interval may be empty, of course) will attack, by the recursive assumption. Because these are the *only* attackers (also by the recursive assumption), the government will yield iff

$$\frac{1}{2\epsilon}[x_n - (\theta - \epsilon)] \geq a(\theta),$$

or

$$\theta + 2\epsilon a(\theta) \leq x_n + \epsilon$$

So we can define an implicit function $h(x, \epsilon)$ such that the above inequality translates into

$$\theta \leq h(x_n, \epsilon).$$

Put another way, the implicit function $h(x, \epsilon)$ solves the equation

$$h(x, \epsilon) + 2\epsilon a(h(x, \epsilon)) = x + \epsilon \tag{2}$$

Now consider a person with signal x. If she were to attack, her expected payoff would be given by

$$\frac{1}{2\epsilon}\int_{x-\epsilon}^{h(x_n,\epsilon)} [e - f(\theta)]d\theta - t. \tag{3}$$

The idea is to find the (obviously unique) value of x such that the expression in (3) just equals 0; call this x_{n+1}.

Trace this recursion starting all the way from $n = 0$: we have $x_0 = \underline{\theta} - \epsilon$. Then (remembering that $a(\theta) = 0$ for all $\theta \leq \underline{\theta}$) it is easy to see that (3) reduces to

$$\frac{1}{2\epsilon}\int_{x-\epsilon}^{\underline{\theta}} [e - f(\theta)]d\theta - t.$$

For $x \simeq x_0$, this is just

$$\frac{1}{2\epsilon}\int_{\underline{\theta}-2\epsilon}^{\underline{\theta}} [e - f(\theta)]d\theta - t,$$

which is certainly strictly positive. So x_1 is well-defined, and $x_1 > x_0$.

Now put x_1 in place of x_0, and repeat the process. Notice that h is increasing in x, so if we replace x_0 by x_1 in (3), then, evaluated at $x = x_1$, the payoff must turn strictly positive.[2] So the new x_2, which is the maximal signal for which people will sell under the belief that everyone less than x_1 sells, will be still higher than x_1. And so on: the recursion creates a strictly increasing sequence $\{x_n\}$, which converges from below to x^*, where x^* solves

$$\frac{1}{2\epsilon}\int_{x^*-\epsilon}^{h(x^*,\epsilon)} [e - f(\theta)]d\theta - t = 0. \tag{4}$$

It is very easy to see that there is a unique solution to x^* defined in this way. In fact, something stronger can be established:

Claim. If x^* is some solution to (4), and $x' > x^*$, then

$$\frac{1}{2\epsilon}\int_{x'-\epsilon}^{h(x',\epsilon)} [e - f(\theta)]d\theta - t < 0. \tag{5}$$

To prove this, consider any $x' > x^*$. Then two things happen: first, it is easy to see that

$$h(x', \epsilon) - x' < h(x^*, \epsilon) - x^*,$$

so that the support over which integration takes place in (4) is narrowed. Moreover, the stuff *inside* the integral is also smaller when we move from x^* to x', because $f(\theta)$ is increasing. So the LHS of (4) *unambiguously* falls when we move from x^* to x', and we are done with the Claim.

To learn a bit more about x^*, use (2) to see that $h(x, \epsilon) - x + \epsilon = 2\epsilon[1 - a(h(x, \epsilon))]$, so that

$$0 = \frac{1}{2\epsilon}\int_{x^*-\epsilon}^{h(x^*,\epsilon)} [e - f(\theta)]d\theta - t = [1 - a(h(x^*, \epsilon))]e - \frac{1}{2\epsilon}\int_{x^*-\epsilon}^{h(x^*,\epsilon)} f(\theta)d\theta - t,$$

or

$$e - \frac{1}{2\epsilon}\int_{x^*-\epsilon}^{h(x^*,\epsilon)} f(\theta)d\theta - t = a(h(x^*, \epsilon))e$$

A comparison of this equation with (1) categorically shows that x^* is bounded below $\bar{\theta}$ for small ϵ.

[2] This is on the assumption that the sequence $\{x_n\}$ stays bounded below $\bar{\theta}$. This will certainly be the case, see below, so it's not really an assumption at all.

So there is a unique solution to x^* and it is below $\bar{\theta}$, which justifies the previous recursive analysis (see in particular, footnote 2). Notice also that our analysis shows that *every* equilibrium must involve attack for signals less than x^*.

To complete the proof, we must show that no signal above x^* can ever attack. Suppose, on the contrary, that in some equilibrium some signal above x^* finds it profitable to attack. Take the supremum of all signals under which it is weakly profitable to attack: call this x'. Then at x' it is weakly profitable to attack. Suppose we now entertain a change in belief by supposing that everybody below x' attacks *for sure*; then this cannot change the weak profitability of attack at x'. But the profit is

$$\frac{1}{2\epsilon}\int_{x'-\epsilon}^{h(x',\epsilon)}[e-f(\theta)]d\theta - t,$$

which is nonnegative as we've just argued. But this contradicts the Claim.

So we have proved that there is a unique equilibrium to the "perturbed" game, in which a speculative attack is carried out by an individual if and only if $x \leq x^*$. As $\epsilon \to 0$, this has an effect of refining the equilibrium correspondence dramatically. To describe this, calculate the threshold x^* as $\epsilon \to 0$. The easiest way to do this is the "sandwich" inequality:

$$[e-f(h(x^*,\epsilon))][1-a(h(x^*,\epsilon))] \leq \frac{1}{2\epsilon}\int_{x^*-\epsilon}^{h(x^*,\epsilon)}[e-f(\theta)]d\theta \leq [e-f(x^*-\epsilon)][1-a(x^*-\epsilon)],$$

which is obtained by noting that $f(x^*-\epsilon) \leq f(\theta) \leq f(h(x^*,\epsilon))$ for all $\theta \in [x^*-\epsilon, h(x^*,\epsilon)]$. Both sides of the sandwich go to the same limit, because x^* and $h(x^*,\epsilon)$ — as well as the realization of the state — all go to a common limit, call it θ^*. This limit solves the condition

$$[e-f(\theta^*)][1-a(\theta^*)] = t. \tag{6}$$

It is obvious that there is a unique solution to (6).

Note: At this point be careful when reading Morris-Shin. There is an error in Theorem 2. See Heinemann (AER 2000) for a correction of this error which agrees with the calculations provided here.