MA300.2 Game Theory II, LSE

Lecture 10: Sequential Games with Imperfect Information

1. The Spence Signaling Model

Or: a model of education in which you don't really learn anything ...

[But that's not why this model is famous. It's because this is one of the first signaling models, and indeed, it is one of the models that motivated the definition of sequential equilibrium.]

An employer faces a worker of unknown ability θ. The ability of the worker is known to the worker though, and is either $\theta = H$ or $\theta = L$, where $H > L > 0$. Interpret these numbers as the money value of what the worker would produce working in the firm.

The worker would like to transmit the knowledge of her ability to the firm; the problem is how to do so in a credible way. Think of education as just such a device.

1.1. **The Game.** Specifically, suppose that a worker can choose to acquire e units of education, where e is any nonnegative number. Of course, the worker will have to study hard to obtain her education, and this creates disutility (studying for exams, doing homework, etc.). Assume that a worker of true ability θ expends e/θ in disutility. The point is, then, that H-types can acquire education easier than L-types (there is a bit more going on in the particular specification that I've adopted that I will clarify presently).

The game proceeds as follows:

1. Nature moves and chooses a worker type, H or L. The type is revealed to the worker but not to the employer.

2. The worker then chooses e units of education. This is perfectly observed by the employer.

3. The employer observes e and forms an estimate of θ. He then pays the worker a salary equal to this estimate, which is just the conditional expectation of θ given e, written $\mathbb{E}(\theta/e)$.

4. The H-worker's payoff is $\mathbb{E}(\theta/e) - (e/H)$, and the L-worker's payoff is $\mathbb{E}(\theta/e) - (e/L)$.

The game is set up so simply that the employer's expected payoff is zero. Essentially, we assume that the worker's choice of education is visible to the world at large so that perfect competition must push her wage to $\mathbb{E}(\theta/e)$, the conditional expectation of θ given e. [If you want to drop this assumption, you can do so costlessly by assuming that the worker gets paid some fraction of the expected θ; nothing of substance will change.]

Very Important. Note well that $\mathbb{E}(\theta/e)$ is not just a given. How it is computed will depend on worker strategies. See below for more detail.

1.2. **Single Crossing.** Suppose that a worker of type θ uses a probability distribution μ_θ over different education levels. First observe that if e is a possible choice of the high worker and e' a possible choice of the low worker, then it must be that $e \geq e'$. This follows from the following important *single-crossing* argument:

The H-type *could* have chosen e' instead of e, so

$$\mathbb{E}(\theta/e) - \frac{e}{H} \geq \mathbb{E}(\theta/e') - \frac{e'}{H}, \tag{1}$$

while the L-type *could* have chosen e instead of e', so

$$\mathbb{E}(\theta/e') - \frac{e'}{L} \geq \mathbb{E}(\theta/e) - \frac{e}{L}. \tag{2}$$

Adding both sides in (1) and (2), we see that

$$(e - e')\left(\frac{1}{L} - \frac{1}{H}\right) \geq 0.$$

Because $(1/L) > (1/H)$, it follows that $e \geq e'$.

Essentially, if the low type *weakly* prefers a higher education to a lower one, the high type would *strictly* prefer it. So a high type can never take strictly less education than a low type in equilibrium.

This sort of result typically follows from the assumption that being a high type reduces not just the *total* cost from taking an action but also the *marginal* cost of that action; in this case, of acquiring one more unit of education. As long as this feature is present, we could replace the cost function e/θ by any cost function and the same analysis goes through.

1.3. **Equilibrium.** Now that we know that the high type will not invest any less than the low type, we are ready to describe the equilibria of this model. There are three kinds of equilibria here; the concepts are general and apply in many other situations.

1. *Separating Equilibrium.* Each type takes a different action, and so the equilibrium action reveals the type perfectly. It is obvious that in this case, L must choose $e = 0$, for there is nothing to be gained in making a positive effort choice.

What about H? Note: she cannot play a mixed strategy because each of her actions fully reveals her type, so she might as well choose the least costly of those actions. So she chooses a single action: call it e^*, and obtains a wage equal to H. Now these are the crucial incentive constraints; we must have

$$H - \frac{e^*}{L} \leq L, \tag{3}$$

otherwise the low person will try to imitate the high type, and

$$H - \frac{e^*}{H} \geq L, \tag{4}$$

otherwise the high person will try to imitate the low type.

Look at the smallest value of e^* that just about satisfies (3); call it e_1. And look at the largest value of e^* that just about satisfies (4); call it e_2. It is very easy to see that e_1 is smaller than e_2, so the two restrictions above are not inconsistent with each other.

Now it is easy to see that any outcome in which the low type chooses 0 and the high type chooses some $e^* \in [e_1, e_2]$ is supportable as a separating equilibrium. To show this we must also specify the beliefs of the employer. There is a lot of leeway in doing this. Here is one set of beliefs that works: the employer believes that any $e < e^*$ (if observed) comes from the low type, while any $e > e^*$ (if observed) comes from the high type. These beliefs are consistent because sequential equilibrium in this model imposes no restrictions on off-the-equilibrium beliefs.

Given these beliefs and equations (3) and (4), it is very easy to see that no type will want to deviate. We are done.

2. *Pooling Equilibrium.* There is also a family of pooling equilibria in which only one signal is received in equilibrium. It is sent by both types, so the employer learns nothing new about the types. So if it sees that signal — call it e^* — it simply pays out the expected value calculated using the prior beliefs: $pH + (1-p)L$.

Of course, for this to be an equilibrium two conditions are needed. First, we need to specify employer beliefs off the equilibrium path. Again, a wide variety of such beliefs are compatible; here is one: the employer believes that any action $e \neq e^*$ is taken by the low type. [It does not have to be this drastic.[1]] Given these beliefs, the employer will "reward" any signal not equal to e^* with a payment of L. So for the types not to deviate, it must be that

$$pH + (1-p)L - \frac{e^*}{\theta} \geq L,$$

but the binding constraint is clearly for $\theta = L$, so rewrite as

$$pH + (1-p)L - \frac{e^*}{L} \geq L.$$

This places an upper bound on how big e^* can be in any pooling equilibrium. Any e^* between 0 and this bound will do.

3. *Hybrid Equilibria.* There is also a class of "hybrid equilibria" in which one or both types randomize. For instance, here is one in which the low type chooses 0 while the high type randomizes between 0 (with probability q) and some e with probability $1-q$. If the employer sees e he knows the type is high. If he sees 0 the posterior probability of the high type there is — by Bayes' Rule — equal to

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

and so the employer must pay out a wage of precisely

$$\frac{qp}{qp + (1-p)}H + \frac{1-p}{qp + (1-p)}L.$$

But the high type must be *indifferent* between the announcement of 0 and that of e, because he willingly randomizes. It follows that

$$\frac{qp}{qp + (1-p)}H + \frac{1-p}{qp + (1-p)}L = H - \frac{e}{H}.$$

To complete the argument we need to specify beliefs everywhere else. This is easy as we've seen more than once (just believe that all other e-choices come from low types). We therefore have a hybrid equilibrium that is "semi-separating".

[1] For instance, the employer might believe that any action $e < e^*$ is taken by the low type, while any action $e > e^*$ is taken by types in proportion to their likelihood: $p : 1-p$.

In the Spence model all three types of equilibria coexist. Part of the reason for this is that beliefs can be so freely assigned off the equilibrium path, thereby turning lots of outcomes into equilibria. What we turn to next is a way of narrowing down these beliefs. To be sure, to get there we have to go further than just sequential equilibrium.

2. The Intuitive Criterion

Consider a sequential equilibrium and a non-equilibrium announcement (such as an nonequilibrium choice of education in the example above). What is the other recipient of such a signal (the employer in the example above) to believe when she sees that signal?

Sequential equilibrium imposes little or no restrictions on such beliefs in signalling models. [We have seen, of course, that in other situations — such as those involving moves by Nature — that it does impose several restrictions, but not in the signalling games that we have been studying.] The purpose of the Intuitive Criterion is to try and narrow beliefs further. In this way we eliminate some equilibria and in so doing sharpen the predictive power of the model.

Consider some non-equilibrium signal e. Consider some type of a player, and suppose even if she were to be treated in the best possible way following the emission of the signal e, she *still* would prefer to stick to her equilibrium action. Then we will say that signal e is *equilibrium-dominated* for the type in question. She would never want to emit that signal, except purely by error. Not strategically.

The Intuitive Criterion (IC) may now be stated.

If, under some ongoing equilibrium, a non-equilibrium signal is received which is equilibrium-dominated for some types but not others, then beliefs cannot place positive probability weight on the former set of types.

Notice that IC places no restrictions on beliefs over the types that are *not* equilibrium dominated, and in addition it also places no restrictions if *every* type is equilibrium-dominated. For then the deviation signal is surely an error, and once that possibility is admitted, all bets about who is emitting that signal are off.

The idea behind IC is the following "speech" that a sender (of signals) might make to a recipient:

Look, I am sending you this signal which is equilibrium-dominated for types A, B or C. But it is not so for types D and E. Therefore you *cannot* believe that I am types A, B or C.

Let us apply this idea to the Spence model.

Proposition **1.** *In the Spence signaling model, a single equilibrium outcome survives the IC, and it is the separating equilibrium in which L plays 0 while H plays* e_1*, where* e_1 *solves (3).*

Proof. First we rule out all equilibria in which types H and L play the same value of e with positive probability. [This deals with all the pooling and all the hybrid equilibria.]

At such an e, the payoff to each type θ is

$$\lambda H + (1-\lambda)L - \frac{e}{\theta},$$

where λ represents the employer's posterior belief after seeing e. Now, there always exists an $e' > e$ such that

$$\lambda H + (1-\lambda)L - \frac{e}{L} = H - \frac{e'}{L},$$

while at the same time,

$$\lambda H + (1-\lambda)L - \frac{e}{H} < H - \frac{e'}{H}.$$

It is easy to see that if we choose e'' very close to e' but slightly bigger than it, it will be equilibrium-dominated for the low type —

$$\lambda H + (1-\lambda)L - \frac{e}{L} > H - \frac{e''}{L},$$

while it is not equilibrium-dominated for the high type:

$$\lambda H + (1-\lambda)L - \frac{e}{H} < H - \frac{e''}{H}.$$

But now the equilibrium is broken by having the high type deviate to e''. By IC, the employer must believe that the type there is high for sure and so must pay out H. But then the high type benefits from this deviation relative to playing e.

Next, consider all separating equilibria in which L plays 0 while H plays some $e > e_1$. Then a value of e' which is still bigger than e_1 but smaller than e can easily be seen to be equilibrium-dominated for the low type but not for the high type. So such values of e' must be rewarded with a payment of H, by IC. But then the high type will indeed deviate, breaking the equilibrium.

This proves that the only equilibrium that can survive the IC (in the Spence model) is the one in which the low type plays 0 and the high type chooses e_1. □

The heart of the intuitive criterion is an argument which is more general: it is called a *forward induction* argument. The basic idea is that an off-equilibrium signal can be due to one of two things: an error, or strategic play. If at all strategic play can be suspected, the error theory must play second fiddle: that's what a forward induction argument would have us believe.