


MA300.2 Game Theory II, LSE

Lecture 9: Sequential Games with Imperfect Information

1. SOME EXAMPLES

For the last two lectures we return to extensive-form games, but this time we focus on imperfect information and especially on signalling games. Our first task, however, is to formulate an appropriate refinement of subgame perfection which will be central to all that follows. We shall develop the notion of a *sequential equilibrium*, due to David Kreps and Robert Wilson.

Example 1. Consider the following extensive game:




Notice that this game has no subgames. The information set where 2 moves might *look* like a subgame, but if you go back to the definition you will see that it isn't. A subgame starts at a singleton information set (and there are other restrictions as well).

Consider the following strategy profile, in which 1 plays a , and 2 plays L . This is a Nash equilibrium. Because there are no subgames, this is also a subgame-perfect Nash equilibrium. But it is a silly equilibrium, because if 2 ever found herself in a situation where she has to move, she would want to play R no matter what her beliefs regarding where she is "inside" that information set. [There are lots of other equilibria that make sense. In all of them 2 plays R and 1 plays either b or c or mixes between the two.]

In this example life was simple because we did not even have to figure out what 2's beliefs needed to be. So why can't we simply extend subgame perfection to include this sort of case, and be done with it? We could, except that in some situations there could be plausible restrictions on beliefs, and such beliefs could be useful in determining the course of play.


Example 2. Consider the following sequential game in which Nature moves first, then 1, then (possibly) 2. Nature's move chooses one of two types of states with 50-50 probability. Notice that neither 1 nor 2 knows the true state when they play.



In this game observe that if 2 were to have to move she would not have an unambiguously “correct” course of action: it would all depend on her beliefs. Now consider the strategy profile in which 1 plays *Out*, and 2 plays *R* (if she finds herself having to move). This is Nash, it is subgame perfect. It also satisfies the criterion in Example 1: there is *some* belief for 2 which justifies her action.

But 2 cannot hold any arbitrary belief. If the structure of the game is common knowledge (as it assuredly is by assumption), then 2 knows that 1 does not know the true state. Consequently, if 1 plays *In* instead of *Out*, then this play *cannot* be predicated on the state (unless 1 gets some secret vibe about the state which affects his action without him knowing about it, which is possible but ruled out). Consequently, 2 must believe that it is equally likely that she is at one node or the other in her information set. But then she prefers to play *L*. Consequently, this strategy profile does not constitute a “credible” equilibrium. [An equilibrium that is reasonable is one in which 1 plays *In* and 2 plays *L*.]

Well then, why don't we simply extend subgame perfection to these information sets and ask beliefs there to be consistent with whatever it is that Nature is doing (as in Example 2), and otherwise allow for arbitrary beliefs (as in Example 1)? We could, of course, but other concerns remain. *Example 3*. In the following sequential game, a player replaces Nature:



Consider the strategy profile in which 3 plays b , 1 plays Out , and 2 plays R (if she finds herself having to move). As before, this is Nash and subgame perfect, and satisfies both the criteria at the end of the last example. But what if this strategy profile is commonly believed by all: in that case 2, must form a theory about how she came to move. Clearly, 1 did not play Out , so one theory is just that 1 deviated. There are other theories as well: such as 1 deviated *and* 3 also deviated. But these alternatives are not “minimal” explanations of how 2 came to be in the position that she is in now (of having to move). Indeed, the only minimalist theory that she can have is that only 1 deviated, so that she must then believe that she is at the right node. In that case, however, she will want to play L , and the proposed equilibrium is destroyed. If you didn’t find this very convincing, look at it another way. What must it take for the above strategy profile to go through? 1 must really fear that if he deviates, 2 will then layer on all sorts of extra explanations, justifying a move of R for her. [After all, she must believe that she is at the left node with at least probability $2/3$, even when she is initially believing that 3 is playing b , according to the proposed equilibrium.] We rule out equilibria that need to be “supported” in this convoluted way. This example is a core example, because it tells us that at unreached information sets (according to the proposed strategy) that are unexpectedly reached, the player there must have a belief that is consistent not just with Nature (if Nature moves at all) but also with a minimal theory of deviations from the proposed equilibrium. For those of you who are philosophers, this is another instance of Occam’s Razor.

2. FORMAL DEFINITION

The notion of a sequential equilibrium is meant to capture these ideas (and more). One way to develop the idea is to think of an equilibrium as both a *strategy profile* and a *belief system*. [Indeed, this is always true of Nash equilibria but the concept we now develop makes explicit use of this double representation.] A strategy profile σ simply prescribes actions at every information set, or probability distributions over actions. A belief system μ assigns beliefs (or probability distributions) over nodes at every information set.

Notice that we didn’t even bring in the question of who takes the action or who holds these beliefs. The answer is that the person who is to move at any information set is the one who moves according to the “piece” of the strategy profile that acts there. As for beliefs, this is a bit more complicated. certainly the beliefs at any information set belong to the person who is to move there, but given that everything about an equilibrium is common knowledge, everyone else must believe that the person there holds these beliefs, everyone must believe that everyone believes that the person there holds these beliefs, etc.

Now we have to relate σ and μ to each other. There are two aspects to this:

First, strategy profile σ must be “rational” given beliefs μ ; this is the idea of *sequential rationality*.

Second, beliefs μ must be “rational” given strategy profile σ ; this is the notion of *consistency*.

Take sequential rationality first. This is simple, or at least we’ve been down this road before. It simply states that at every information set, *given the beliefs prescribed by μ* , no individual must want to deviate from the piece of the strategy assigned to her there.

Now consider consistency. This restricts the sort of beliefs one can have. There are two kinds of restrictions. One is for information sets that are “reached” under the strategy profile: these are the information sets which will become active with positive probability under σ . For these sets μ must be given by Bayes’ Rule, given σ . The second is for information sets which are not reached at all; indeed, these sets were the focus of all our examples. Bayes’ Rule cannot be applied to sets which are reached with probability zero; yet, as we’ve seen from the examples above, we do need some specification. This is where the full definition of consistency is needed:

Say that μ is *consistent* given σ if there exists a sequence (σ^m, μ^m) of strategy profiles and beliefs such that (a) $(\sigma^m, \mu^m) \rightarrow (\sigma, \mu)$ as $m \rightarrow \infty$; (b) σ^m is *completely mixed* for every m in that it assigns strictly positive probability to every action at every information set; and (c) for each m , μ^m is derived from σ^m by applying Bayes’ Rule to every information set.

Thus we require μ to be the “limit of a system of consistent beliefs” given a sequence of completely mixed strategy profiles that converge to the strategy profile in question.

To complete the definition, say that (σ, μ) is a *sequential equilibrium* if σ is sequentially rational given μ and μ is consistent given σ .

3. DISCUSSION

Just sequential rationality and that part of the definition which insists that μ must come from Bayes’ rule for information sets that are reached with positive probability takes us all the way to subgame perfection. This is because at unreached information sets that are singletons there can be only one conceivable value for μ , Bayes’ Rule or no Bayes’ Rule! So we are all the way up to subgame perfection. The remaining part of the definition dealing with limits of beliefs handles other aspects of consistency, such as Examples 1–3 above and additional situations besides.

Let’s take a quick look at the examples through the eyes of this definition. In Example 1, we would intuitively like to impose on restrictions on beliefs: after all, if player 1 is supposed to play a it is hard to assign relative likelihoods to the two deviations b and c .¹ And indeed, consistency imposes no further restrictions. To see this, fix the strategy profile examined in that example, any number α between 0 and 1 and consider the completely mixed strategy that assigns probability $1 - \epsilon$ to a , $\alpha\epsilon$ to b and $(1 - \alpha)\epsilon$ to c . Bayes’ Rule will give us probabilities α and $1 - \alpha$ at the information set where 2 was supposed to move, and note that α is arbitrary, so that any beliefs can be supported.²

In Example 2, notice that the limit of any completely mixed consistent sequence (relative to the strategy profile under consideration) must assign probabilities 50-50 to the two nodes

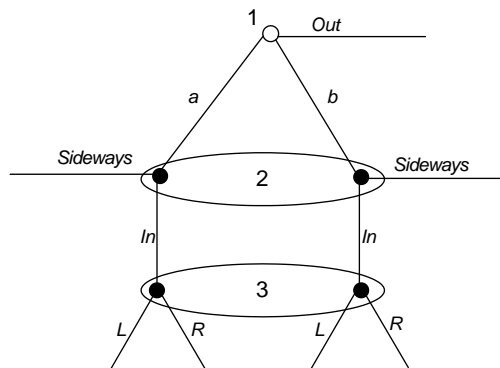
¹Or is it? This takes us all the way to the philosophy of what trembles really mean: are they completely random mistakes or (partly or fully controlled) “errors”. If the former, we can stick to the interpretation in the text. If the latter, we might say that more dangerous trembles are less likely than less dangerous trembles (for the player concerned). This leads us into refinements of sequential equilibrium, such as the notion of *proper equilibrium*, due to Roger Myerson. But this discussion takes us out of the scope of the lectures.

²A little more subtle are the “edge beliefs” $\alpha = 0$ and $\alpha = 1$; the former can be supported, for instance, using the completely mixed sequence $(1 - \epsilon, \epsilon^2, \epsilon - \epsilon^2)$. [Why?]

at 2's unreached information set. Similarly, in Example 3, any totally mixed sequence must generate the consistent belief $(0, 1)$ for the strategy profile under consideration (do this formally).


There are other restrictions imposed by the notion of sequential equilibrium.

Example 4. Consider this game tree (the payoffs are not needed):



Consider the strategy profile in which player 1 chooses *Out* and player 2 chooses *Sideways*. It is fine to have the beliefs $(1, 0)$ at the first unreached information set and it is also fine to have the beliefs $(0, 1)$ at the second unreached information set, but sequential equilibrium will not allow both these beliefs to be part of *one* consistent belief system. indeed, any consistent belief system will require the same probabilities on the left and right nodes at each of the two information sets. If the player at the two sets is the same one this makes a lot of sense. If the players are different (as in the diagram), however, then a real restriction has been imposed: that the two players must have beliefs which are not just individually consistent but *consistent with each other*.

Example 5. We have already alluded to the “minimalism” of sequential equilibrium in its construction of theories about deviation. Here is another instance of this minimalism, related to the discussion we had before:



Consider the strategy profile in which player 1 plays R_1 for sure and player 2 plays R_2 for sure (what you see is only part of the game and I don't need to specify the rest of the profile). The question is: if player 3 finds herself having to move at the lowest information set in the

diagram, what beliefs can she have over the three nodes x , y and z ? It is easy to see that sequential equilibrium demands that node x must be given probability zero! Node x can be reached, but by *two* trembles. In contrast, nodes y and z are reachable with a single tremble. The completely mixed strategy profiles that converge to the profile under consideration must therefore have the weight on x going to zero “extra-fast” relative to the weights on the nodes x and y (try this out formally; good practice).