Sequential Games with Incomplete Information

Debraj Ray, November 2006

For the remaining lectures we return to extensive-form games, but this time we focus on imperfect information, reputation, and signalling games. Our first task 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.

1. Some Examples

Example 1. Consider the following extensive game:

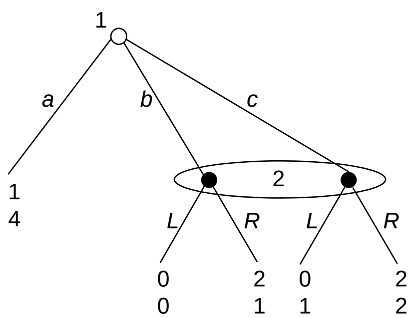

Notice that this game has no subgames. The information set where 2 moves might *look* like a subgame, but if you go back to the definition you will see that it isn't. A subgame starts at a singleton information set (and there are other restrictions as well).

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.

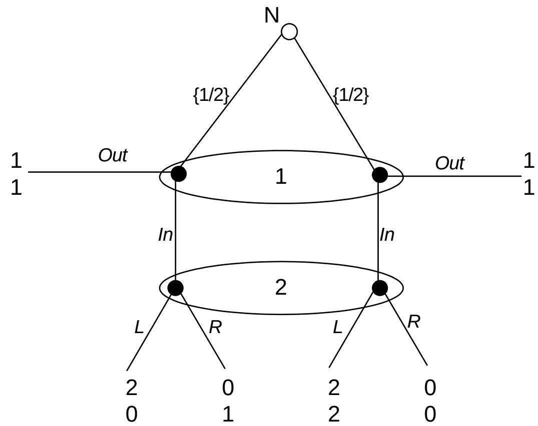

In this game observe that if 2 were to have to move she would not have an unambiguously "correct" course of action: it would all depend on her beliefs. Now consider the strategy profile in which 1 plays *Out*, and 2 plays R (if she finds herself having to move). This is Nash, it is subgame perfect. It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

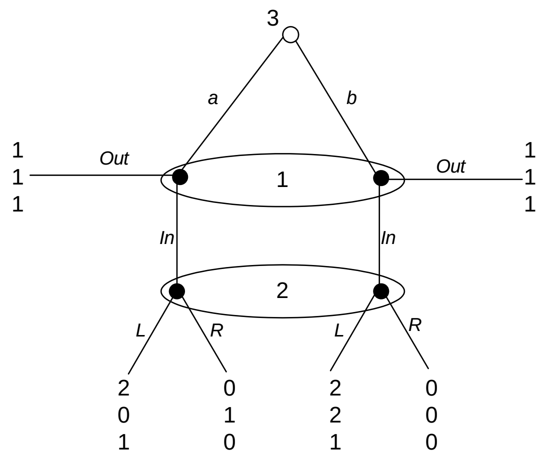

Consider the strategy profile in which 3 plays b, 1 plays *Out*, and 2 plays R (if she finds herself having to move). As before, this is Nash and subgame perfect, and satisfies both the criteria at the end of the last example. But what if this strategy profile is commonly believed by all: in that case 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. Sequential Equilibrium

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 question of who takes the action at any information set is easy; it's obviously the person who is supposed to move at that information set. In contrast, there is much more hidden in the notion of a belief system. Certainly beliefs at some information set are to be interpreted as the beliefs of the person who is scheduled to move there but more than that, the one belief system is commonly held by (and commonly known to) everybody. This is in the spirit of Nash equilibrium, and as we shall see it will have more implications than those captured in the examples so far.

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*. 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.

Second, beliefs μ must be "rational" given strategy profile σ; this is the notion of *consistency*. Consistency imposes 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) $(\sigma^m, \mu^m) \to (\sigma, \mu)$ as $m \to \infty$; (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 stratey 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 comes 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 strictly 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.

A little nore subtle are the "edge beliefs" $\alpha = 0$ and $\alpha = 1$; the former can be supported, for instance, using the completely mixed sequence $(1-\epsilon, \epsilon^2, \epsilon - \epsilon^2)$.

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 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):

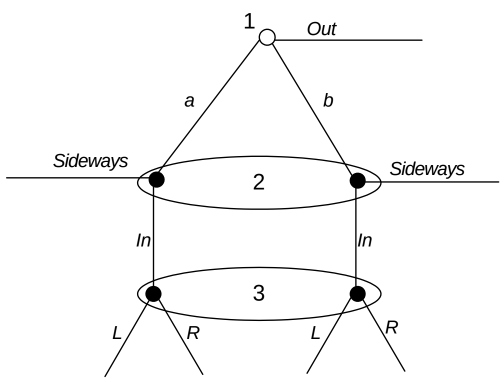

Consider the strategy profile in which player 1 chooses *Out* and player 2 chooses *Sideways.* It is fine to have the beliefs $(1, 0)$ at the first unreached information set and it is also fine to have the beliefs $(0, 1)$ at the second unreached information set, but sequential equilibrium will not allow 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.*

[1] Or 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.

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:

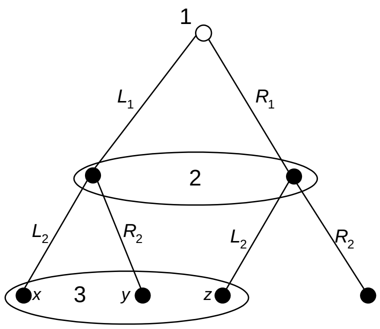

Consider the strategy profile in which player 1 plays R_1 for sure and player 2 plays R_2 for sure (what you see is only part of the game and I don't need to specify the rest of the profile). The question is: if player 3 finds herself having to move at the lowest information set in the diagram, what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

4. Application: 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.

4.1. **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.]

4.2. **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.

4.3. **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.[2]] Given these beliefs, the employer will "reward" any signal not equal to e^* with a payment of L. So — and this is the second condition — 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

[2] 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: $p : 1-p$.

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". 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. More on this later.

5. Application: Cheap Talk

The Spence model is an example of costly signaling. Now we do an example where all signaling is free: cheap talk, based on the classic by Crawford and Sobel (1983).

There are two agents, a sender (S) and a receiver (R), and a state of the world θ distributed on $[0,1]$. The sender knows the state, so θ may be identified with her type. The receiver only kows that θ is distributed on $[0,1]$ according to some cdf G.

The receiver's job will be to take some action; call it y. If he knew what the state was, he would try to maximize a vNM utility function $u_R(y,\theta)$ by choice of y. The sender, too, has preferences over y given θ though she cannot directly control y. Her preferences are summarized by a continuous function $u_S(y,\theta)$. Assume that

[A.1] Both u_S and u_R are continuous in both arguments and smooth, bounded, never flat and single-peaked in y for each θ.

[A.2] y and θ are complements in each utility function: for every $y > y'$ and $\theta > \theta'$,

$$u(y,\theta) - u(y',\theta) > u(y,\theta') - u(y',\theta'),$$

where u stands for both u_R and u_S.

Assumption [A.1] guarantees that u_R and u_S generate unique maxima $y_R(\theta)$ and $y_S(\theta)$ respectively for every θ. [A.2] guarantees that each of these functions is increasing in θ. [A.1] also assures us that the functions are continuous (by the maximum theorem).

[A.3] The sender wants a higher action than the receiver for every state. That is, for each θ, $y_S(\theta) > y_R(\theta)$.

The game is very simple. The sender announces the state (or perhaps a distortion of it), and the receiver takes an action after listening to the announcement. More formally, a strategy for the sender is an announcement m of the state as a function of her type θ, so it is a function $m(\theta)$. A strategy for the receiver is likewise a mapping from announcements into actions. Then payoffs are received.

It is obvious that if the receiver believes the sender, the sender wishes to exaggerate the state. Therefore the receiver has no reason to take the sender's statement at face value.

A nice example that fixes these ideas is the quadratic case:

$$u_S(y, \theta) = -(y - [b + \theta])^2, \quad u_R(y, \theta) = -(y - \theta)^2$$

for some $b > 0$. In this case it is obvious that $y_S(\theta) = \theta + b$, while $y_R(\theta) = \theta$.

As in the Spence model, a sequential equilibrium places no restrictions on the receiver's beliefs at off-equilibrium announcements: at values of m which are never announced under the strategy $m(\theta)$. Therefore the *flat* functions $m(\theta) = \bar{m}$ for some fixed announcement $\bar{m}$, and the choice of $y(m) = \bar{y}$ for some $\bar{y}$ that maximizes

$$\int u_R(y, \theta) dG(\theta)$$

is always an equilibrium. This is often called the *babbling equilibrium*, because the sender conveys nothing.

But there are other sequential equilibria. To describe them, the following result is central:

PROPOSITION **1.** *There exists $\epsilon > 0$ (uniform across all equilibria) such that if y and y' are two distinct actions taken in equilibrium, then $|y - y'| > \epsilon$.*

Proof. Let y and y' be two equilibrium actions. Suppose that type θ induces (via her announcement) the action y and type θ' the action y'. Then, because each type could have made the other type's announcements, it must be that

$$u_S(\theta, y) \geq u_S(\theta, y'),$$

while

$$u_S(\theta', y) \leq u_S(\theta', y').$$

By our assumptions on u_S, it is very easy to see that there exists $\hat{\theta} \in [\theta, \theta']$ such that if $\theta < \hat{\theta}$, then type θ prefers y to y' while if $t > \hat{\theta}'$ then type θ' prefers y' to y. Therefore when R chooses y, he is convinced that the true state is *below* $\hat{\theta}$, while if he chooses y', he is convinced that the true state is *above* $\hat{\theta}$. It follows from assumption [A.2] that

$$y \leq y_R(\hat{\theta}) \leq y',$$

but it is also true, given $\hat{\theta} \in [\theta, \theta']$ and [A.2], that

$$y \leq y_S(\hat{\theta}) \leq y'.$$

Apply [A.1] and [A.3] and use continuity of the y's and compactness of the θ's to conclude that y_R and y_S are separated by at least ϵ regardless of state. this completes the proof. $\square$

In the quadratic example, $y_S(\theta) - y_R(\theta) = b$ for all θ. Therefore any two actions must be separated by at least b. An equilibrium can at best involve a finite number of choices even though there is a continuum of types. This finiteness is a general result because even though the space of actions is in principle unbounded, every equilibrium action must be bounded above by $y_R(1)$.

We combine Proposition 1 with the following observation:

PROPOSITION **2.** *The set of values of θ that induces a particular equilibrium action must form an interval.*

Proof. Suppose that there is some action y and three ordered values of θ — call them θ_1, θ_2 and θ_3 — such that θ_1 and θ_3 induce the action y while θ_2 does not. Suppose that θ_2 induces y'; then θ_2 weakly prefers y' to y. If $y' > y$, then [A.2] tells us that θ_3 must *strictly* prefer y' to y, while if If $y' < y$, then θ_1 must *strictly* prefer y' to y. Either way we have a contradiction. $\square$

Now we can construct a typical equilibrium. Start by picking the lowest equilibrium action y_1. Then there exists a unique θ_1 such that the receiver picks y_1 as an optimal response to a message under the conjecture that the set of types in $[0, \theta_1]$ and no other type emits that message. But this uniquely pins down what y_2 must be (if there is a second equilibrium action). It must be chosen so that the threshold type θ_1 is indifferent between y_1 and y_2. Now keep proceeding recursively. There is an end-point condition to be met though. For some n along this recursion, θ_n must *exactly* coincide with 1. Any recursion that accomplishes this finds an equilibrium. Indeed, by the two propositions above, this (successful) recursion is a complete characterization of sequential equilibrium in the cheap talk model.

Let us illustrate the result in the quadratic example with uniform prior on θ in $[0, 1]$. Fix some $y_1 \geq 0$, and define $\theta_0 = 0$. Then at any stage $i \geq 0$ in which we are given θ_i and y_{i+1}, we can define θ_{i+1} by

$$y_{i+1} = \frac{\theta_i + \theta_{i+1}}{2} \tag{5}$$

(this captures the fact that y_{i+1} must be a best response for the receiver given his beliefs). Similarly, we can define y_{i+2} by the condition

$$\theta_{i+1} + b = \frac{y_{i+1} + y_{i+2}}{2} \tag{6}$$

(this makes θ_{i+1} exactly indifferent — as she must be — between inducing y_{i+1} and y_{i+2}).

Combining (5) and (6), we see that

$$\theta_{i+1} + b = \frac{\theta_i + 2\theta_{i+1} + \theta_{i+2}}{4},$$

or equivalently,

$$(\theta_{i+2} - \theta_{i+1}) - (\theta_{i+1} - \theta_i) = 4b, \tag{7}$$

which shows that the lengths of successive intervals climb by $4b$ at every step!

This proves right away that if $b \geq 1/4$, the *only* equilibrium involves babbling. Let us continue to see what we get when $b < 1/4$.

Recall the end-point condition that θ_n must equal 1 for some n. Using (7), an equivalent way of writing this is that

$$n\theta_1 + 4b[1 + 2 + \cdots + (n-1)] = 1,$$

or

$$n\theta_1 + 2b(n-1)n = 1$$

for some $\theta_1 \in (0, 1)$. So if b is small enough so that

$$2b(n-1)n < 1,$$

then there is a positive value of θ_1 that solves the required equation, and can accommodate n intervals.