Dynamic Games with Incomplete Information

Jonathan Levin

February 2002

Our final topic of the quarter is dynamic games with incomplete information. This class of games encompasses many interesting economic models — market signalling, cheap talk, and reputation, among others. To study these problems, we start by investigating a new set of solution concepts, then move on to applications.

1 Perfect Bayesian Equilibrium

1.1 Problems with Subgame Perfection

In extensive form games with incomplete information, the requirement of subgame perfection does not work well. A first issue is that subgame perfection may fail to rule out actions that are sub-optimal given any "beliefs" about uncertainty.

Example 1 Consider the following games:

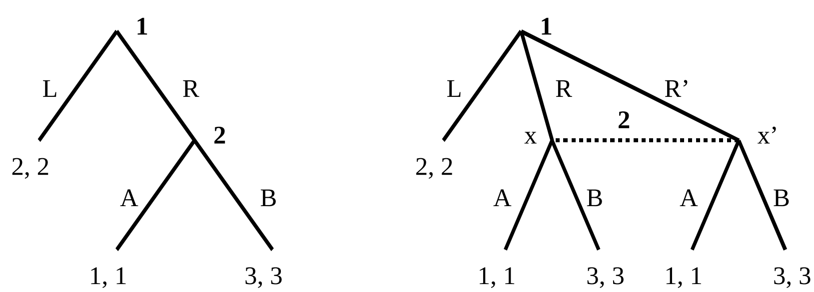

These two games are very similar. However, in the game on the left, (R, B) is the only SPE. In the game on the right, there are other SPE: $(pR + (1 - p) R', B)$ for any p, and $(L, qA + (1 - q) B)$ for any $q \geq 1/2$.

The problem here is that the game on the right has no subgames other than the game itself. So SPE has no bite. What can be done about this? One solution is to require players to choose optimally at all information sets. To make sense of this we need to introduce the idea of beliefs.

Example 1, cont. Suppose at the information set $h = \{x, x'\}$, we require player two to choose the action that maximizes his *expected* payoff given some belief assigning probability $\mu(x)$ to being at x and $\mu(x')$ to being at x', with $\mu(x) + \mu(x') = 1$. Then for *any belief* player two might have, choosing B is optimal.

A second issue is that subgame perfection may allow actions that are possible only with beliefs that are "unreasonable".

Example 2 Consider the following game:

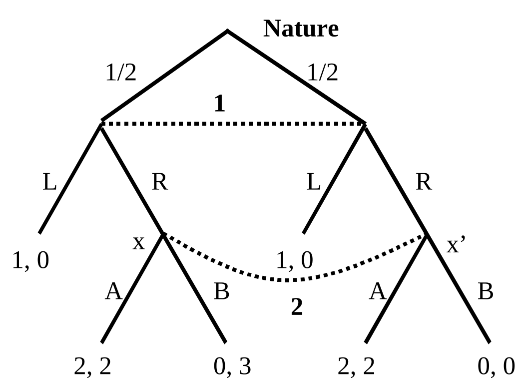

In this game, (L, B) is a subgame perfect equilibrium.

As in the previous example, (L, B) is an SPE in this example because there are no subgames. Note, however, that there are in fact beliefs for which B is an optimal choice for player 2. If player 2 places probability at least $2/3$ on being at x given that he is at $h = \{x, x'\}$, then B is an optimal choice. However, these beliefs seem quite unreasonable: if player 1 chooses R, then player 2 should place equal probability on being at either x or x'.

1.2 Perfect Bayesian Equilibrium

Let G be an extensive form game. Let H_i be the set of information sets at which player i moves. Recall that:

Definition 1 *A behavioral strategy for player i is a function $\sigma_i : H_i \to \Delta(A_i)$ such that for any $h_i \in H_i$, the support of $\sigma_i(h_i)$ is contained in the set of actions available at h_i.*

We now augment a player's strategy to explicitly account for his beliefs.

Definition 2 *An* ***assessment*** (σ_i, μ_i) *for player i is a strategy σ_i and belief function μ_i that assigns to each $h_i \in H_i$ a probability distribution over nodes in h_i. Write $\mu_i(x|h)$ as the probability assigned to node x given information set h.*

Example 1, cont. Player two's belief function μ_2 must satisfy $\mu_2(x) + \mu_2(x') = 1$. If player 2 uses Bayesian updating, and $\sigma_1(R) + \sigma_1(R') > 0$, then $\mu_2(x) = \sigma_1(R) / (\sigma_1(R) + \sigma_1(R'))$.

Example 2, cont. Player two's belief function μ_2 again sets $\mu_2(x) + \mu_2(x') = 1$. If player 2 uses Bayesian updating, and $\sigma_1(R) > 0$, then $\mu_2(x) = \mu_2(x') = 1/2$.

Definition 3 *A profile of assessments (σ, μ) is a* ***perfect bayesian equilibrium*** *if*

1. *For all i and all $h \in H_i$, σ_i is sequentially rational, i.e. it maximizes i's expected payoff conditional on having reached h given μ_i and σ_{-i}.*
2. *Beliefs μ_i are updated using Bayes' rule whenever it applies (i.e. at any information set on the equilibrium path).*

Example 1, cont. Sequential rationality implies that player 2 must play B, so the unique PBE is (R, B).

1.3 Refinements of PBE

While PBE is a bread and butter solution concept for dynamic games with incomplete information, there are many examples where PBE arguably allows for equilibria that seem quite unreasonable. Problems typically arise because PBE places no restrictions on beliefs in situations that occur with probability zero — i.e. "off-the-equilibrium-path".

Example 2, cont. Bayes' rule implies that if $\sigma_1(R) > 0$, then $\mu_2(x) = \mu_2(x') = 1/2$, so player 2 must play A. Therefore (R, A) is a PBE. However, (L, B) is also a PBE! If $\sigma_1(R) = 0$, then Bayes' Rule does not apply when player 2 forms beliefs. So μ_2 is arbitrary. If $\mu_2(x) > 2/3$, then it is sequentially rational for 2 to play B. Hence (L, B) is a PBE.

This example may seem pathological, but it turns out to be not all that uncommon. In response, game theorists have introduced a large number of "equilibrium refinements" to try to rule out unreasonable PBE of this sort. These refinements restrict the sorts of beliefs that players can hold in situations that occur with probability zero.

One example is Kreps and Wilson's (1982) notion of *consistent beliefs*.

Definition 4 *An assessment (σ, μ) is **consistent** if $(\sigma, \mu) = \lim_{n\to\infty} (\sigma^n, \mu^n)$ for some sequence of assessments (σ^n, μ^n) such that σ^n is totally mixed and μ^n is derived from σ^n by Bayes' Rule.*

Example 2, cont. In this game, the only consistent beliefs for player two are $\mu_2(x) = \mu_2(x') = 1/2$. To see why, note that for any player 1 strategy with $\sigma_1^n(L), \sigma_1^n(R) > 0$, it must be that $\mu_2^n(x) = \frac{1}{2}$. But then $\lim_{n\to\infty} \mu_2^n(x) = 1/2$, so $\mu_2(x) = 1/2 = \mu_2(x')$ is the only consistent belief for player 2.

Definition 5 *An assessment (σ, μ) is a **sequential equilibrium** if (σ, μ) is both consistent and a PBE.*

Sequential equilibrium is a bit harder to apply than PBE in practice, so we will typically work with PBE. It also turns out that there are games where sequential equilibrium still seems to allow unreasonable outcomes, and one might want a stronger refinement. We will mention one such refinement — the "intuitive criterion" for signalling games — below.

2 Signalling

We now consider an important class of dynamic models with incomplete information. These *signalling* models were introduced by Spence (1974) in his Ph.D. thesis. There are two players and two periods. The timing is as follows.

Stage 0 Nature chooses type $\theta \in \Theta$ of player 1 from a distribution p.

Stage 1 Player 1 observes θ and chooses $a_1 \in A_1$.

Stage 2 Player 2 observes m and chooses $a_2 \in A_2$.

Payoffs $u_1(a_1, a_2, \theta)$ and $u_2(a_1, a_2, \theta)$.

Many important economic models take this form.

Example: Job Market Signalling In Spence's original example, player 1 is a student or worker. Player 2 is the "competitive" labor market. The worker's "type" is his ability and his action is the level of education he chooses. The labor market observes the worker's education (but not his ability) and offers a competitive wage equal to his expected ability conditional on education. The worker would like to use his education choice to "signal" that he is of high ability.

Example: Initial Public Offerings Player 1 is the owner of a private firm, while Player 2 is the set of potential investors. The entrepreneur's "type" is the future profitability of his company. He has to decide what fraction of the company to sell to outside investors and the price at which to offer the shares (so a_1 is both a quantity and a price). The investors respond by choosing whether to accept or reject the entrepeneur's offer. Here, the entrepeneur would like to signal that the company is likely to be profitable.

Example: Monetary Policy Player 1 is the Federal Reserve. It's type is its preferences for inflation versus unemployment. In the first period, it chooses an inflation level $a_1 \in A_1$. Player 2 is the firms in the economy. They observe first period inflation and form expectations about second period inflation, denoted $a_2 \in A_2$. Here, the Fed wants to signal a distaste for inflation so that firms will expect prices not to rise too much.

Example: Pretrial Negotiation Player 1 is the Defendant in a civil lawsuit, while Player 2 is the Plaintiff. The Defendant has private information about his liability (this is θ). He makes a settlement offer $a_1 \in A_1$. The Plaintiff then accepts or rejects this offer, so $A_2 = \{A, R\}$. If the Plaintiff rejects, the parties go to trial. Here, the Defendant wants to signal that he has a strong case.

2.1 Equilibrium in Signalling Models

We start by considering Perfect Bayesian equilibrium in the signalling model.

Definition 6 *A perfect bayesian equilibrium in the signalling model is a strategy profile* $s_1(\theta)$, $s_2(a_1)$ *together with beliefs* $\mu_2(\theta|a_1)$ *for player two such that:*

1. *Player one's strategy is optimal given player two's strategy:*

$$s_1(\theta) \text{ solves } \max_{a_1 \in A_1} u_1(a_1, s_2(a_1), \theta) \text{ for all } \theta \in \Theta$$

2. *Player two's beliefs are compatible with Bayes' rule, i.e. if any type of player one plays* a_1 *with positive probability then*

$$\mu_2(\theta|a_1) = \frac{\Pr(s_1(\theta) = a_1)\, p(\theta)}{\sum_{\theta' \in \Theta} \Pr(s_1(\theta') = a_1)\, p(\theta')};$$

   if player one never uses a_1*, then* $\mu_2(\theta|a_1)$ *is arbitrary.*

3. *Player two's strategy is optimal given his beliefs and given player one's action:*

$$s_2(a_1) \text{ solves } \max_{a_2 \in A_2} \sum_{\theta \in \Theta} u_2(a_1, a_2, \theta)\, \mu_2(\theta|a_1) \text{ for all } a_1 \in A_1.$$

It is not hard to allow for mixed strategies, in which case player i's strategy is denoted σ_i.

As we will see in the examples to follow, it helps to think of PBE as falling into different categories:

1. Separating: Different types of player one use different actions, so player two perfectly learns player one's type in equilibrium.

2. Pooling: All types of player one use the same action, so no information is transmitted in player one's action.

3. Semi-Separating: Some actions of player one are chosen by several types of player one, other actions are chosen by a single type. Thus, there is some learning, but not perfect learning.

With this general framework established, we move on to applications.

2.2 Job Market Signalling

Consider a single worker, whose ability (productivity) is given by $\theta \in \{\theta_L, \theta_H\}$ with $\theta_H > \theta_L > 0$. The worker knows his own ability, and the labor market assigns prior probability λ to him having type θ_H. The worker first chooses his level of education e. Education is costly, and the cost $c(e, \theta)$ depends on the worker's ability.

Assumption Suppose that $c_e > 0$ (education is costly on the margin), and that $c_{e\theta} < 0$ (education is *less* costly on the margin for more able workers).

Once the worker chooses his education level, firms make wage offers. Suppose that workers all have a reservation wage 0 regardless of their ability. Assume also that the labor market is competitive.

The key idea we are headed for is that high ability workers may be able to communicate their ability by getting a lot of costly education. Indeed they may choose to become educated despite the fact that education has no direct effect on productivity!

To solve the model, we work backward. Let $\mu(e)$ denote the labor market's belief that a worker who has chosen education e is of high ability. Then the market assesses that the worker will produce θ_H with probability $\mu(e)$ and otherwise θ_L. Since the labor market is competitive, the wage for a worker with education e is given by:

$$w(e) = \mu(e)\,\theta_H + (1 - \mu(e))\,\theta_L.$$

Now consider the problem facing the worker, given beliefs $\mu(e)$ and the resulting wage $w(e)$. A worker with ability θ must solve:

$$\max_e w(e) - c(e, \theta).$$

To solve this problem, consider the worker's indifference curves in (e, w) space. Implicitly differentiating $u(w, e, \theta) = U$, we obtain:

$$\left.\frac{dw}{de}\right|_{u=U} = c_e(e, \theta) > 0$$

so indifference curves slope up. Also:

$$\frac{d}{d\theta}\left.\frac{dw}{de}\right|_{u=U} = c_{e\theta}(e, \theta) < 0$$

so indifference curves are flatter for high productivity workers. Graphically, this means that indifference curves exhibit the *Spence-Mirrlees Single Crossing Property.*

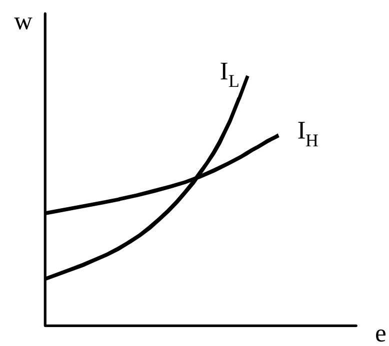

For any given $w(e)$, we can find the optimal choice of each worker type by selecting the point of tangency with an indifference curve.

Remark 1 *If we apply Topkis' Theorem to the worker's problem* $\max_e w(e) - c(e, \theta)$*, we see immediately that because* $c_{e\theta} < 0$*, then for any wage function* $w(e)$ *it must be the case that a worker of type* θ_H *selects a (weakly) higher education level than a worker of type* θ_L*. We could also derive this result in the indifference curve picture.*

The key question is where $w(e)$ (or equivalently $\mu(e)$) comes from. On the equilibrium path, it is implied by theworker's choice. However, for levels of education that are not chosen in equilibrium, it can be anything between θ_L and θ_H since PBE imposes no restriction on beliefs other than that $\mu(e) \in [0, 1]$.

This flexibility in setting $w(e)$ gives rise to many possible equilibria.

Separating Equilibria. We first look for separating equilibria where $e(\theta_H) \neq e(\theta_L)$.

Claim 1 In a separating equilibrium, $w(e(\theta)) = \theta$ for $\theta \in \{\theta_L, \theta_H\}$. (I.e. workers are paid their marginal product).

Proof. In a PBE, beliefs are derived from Bayes rule when possible. Type θ_L workers always choose $e(\theta_L)$, while type θ_H workers always choose θ_H. Thus, if $e(\theta)$ is observed, the market must believe the worker is type θ for sure. Thus, $w(e(\theta)) = \theta$. *Q.E.D.*

Claim 2 In a separating equilibrium, $e(\theta_L) = 0$.

Proof. Suppose not. By the previous claim, type θ_L workers are paid θ_L if they choose $e(\theta_L)$. Suppose instead they choose $e = 0$. Then $\mu(e = 0) \geq 0$, so $w(0) \geq \theta_L$. This deviation reduces education costs, but not wages, so it is beneficial. *Q.E.D.*

We are now ready to derive separating equilibria in the job market model. We need to find an education level $e(\theta_H)$ for high ability workers such that:

- High ability workers prefer to select $e(\theta_H)$ and get a wage θ_H rather than selecting $e(\theta_L) = 0$ and getting a wage $w(0) = \theta_L$. Thus:

$$\theta_H - c(e(\theta_H), \theta_H) \geq \theta_L - c(0, \theta_H) \tag{1}$$

- Low ability workers must prefer the opposite:

$$\theta_L - c(0, \theta_L) \geq \theta_H - c(e(\theta_H), \theta_L). \tag{2}$$

Claim 3 For any $e(\theta_H)$ that satisfies (1)–(2), there is a separating equilibrium in which high ability workers choose $e(\theta_H)$.

Proof. If (1),(2) hold then no worker will want to mimic a worker of the other ability. The only thing to worry about is that a worker might want to choose some $e \notin \{e(\theta_L), e(\theta_H)\}$. To ensure against this, we need to choose $\mu(e)$ sufficiently low so that for all $e \notin \{e(\theta_L), e(\theta_H)\}$

$$w(e) - c(e, \theta_H) \leq \theta_H - c(e(\theta_H), \theta_H)$$

and

$$w(e) - c(e, \theta_L) \leq \theta_L - c(0, \theta_L).$$

One choice that will definitely work is $\mu(e) = 0$, so that $w(e) = \theta_L$ for all $e \notin \{e(\theta_L), e(\theta_H)\}$. *Q.E.D.*

What is the range of values of $e(\theta_H)$ for which as separating equilibrium (i.e. the range of values of $e(\theta_H)$ that satisfy (1) and (2)? Note that by the single crossing property, if (2) is satisfied with equality, then (1) will be satisfied strictly. So let $\underline{e}$ be defined to solve:

$$\theta_L - c(0, \theta_L) = \theta_H - c(\underline{e}, \theta_L).$$

The separating equilibrium with $e(\theta_H) = \underline{e}$ is depicted below.

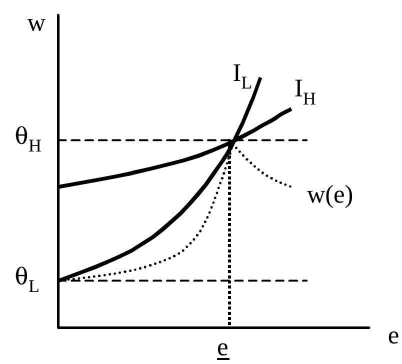

The single-crossing property also implies that if (1) is satisfied with equality, then (2) will be satisfied strictly. Define $\overline{e}$ to solve:

$$\theta_H - c(\overline{e}, \theta_H) = \theta_L - c(0, \theta_H).$$

The separating equilibrium with $e(\theta_H) = \overline{e}$ is depicted below.

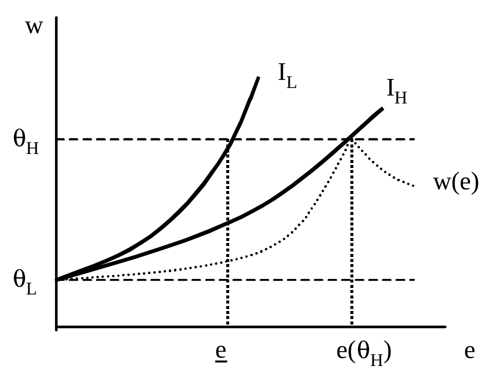

In fact, there is separating equilibrium for any $e \in [\underline{e}, \overline{e}]$, but not for any e outside of this interval.

Remark 2 *Note the important role played by the single crossing property. This is what allows high ability workers to choose a positive education level that is costly, but less costly for them than it would be for low ability workers. It is this* ***differential*** *cost that allows separation.*

Pooling Equilibria. In a pooling equilibrium, every worker chooses the same education level e^P with probability one. Therefore, the labor market must have beliefs $\mu(e^P) = \lambda$ in equilibrium, and thus:

$$w\left(e^P\right) = \lambda\theta_H + (1-\lambda)\,\theta_L = \theta^E.$$

Now, let $\hat{e}$ be defined so that:

$$\theta^E - c\left(\hat{e}, \theta_L\right) = \theta_L - c\left(0, \theta_L\right).$$

That is, $\hat{e}$ is chosen so that a low ability worker is just indifferent to acquiring education $\hat{e}$ and being paid θ^E, and acquiring no education and being paid θ_L.

Proposition 1 *For any $e^P \in [0, \hat{e}]$, there is a pooling equilibrium in which all workers choose e^P for certain.*

Proof. Let $e^P \in [0, \hat{e}]$ be given. Suppose that $\mu(e^P) = \lambda$ and that $\mu\left(e\right) = 0$ for all $e \neq e^P$, with wages given accordingly. Then $w\left(e\right) < w(e^P)$, so clearly no worker wants to deviate to $e > e^P$. Moreover, θ_L workers prefer e^P to any $e < e^P$ by definition of $\hat{e}$. And since θ_L workers prefer e^P to any lower e, so must θ_H workers (by the single crossing property).

There are a few general points to make about the signalling model. First, in the separating equilibrium, education does not increase productivity, but it does reveal information. Thus it is correlated with wages in equilibrium. In the pooling equilibrium, education neither increases productivity nor reveals information. Nevertheless, workers may incur education costs in equilibrium because there is a wage penalty for doing something unexpected (i.e. not becoming educated). Thus pooling pooling equilibria with positive education levels are very inefficient.

A special feature of the model is that education is not productive at all. We could, however, generalize the model to let the worker's productivity be a function of both ability and education: $\theta + e$, for example. Things would work out in a similar manner — with the key idea being that signalling would lead workers to invest more in education than is efficient.

2.3 The Intuitive Criterion

What can be said about the vast multiplicity of equilibria in the signalling model? Cho and Kreps (1987) argue that some equilibria are less appealing than others. We now consider their "intuitive criterion" for eliminating signalling equilibria.

Definition 7 *Let $BR(T, a_1)$ denote the set of player 2's best responses if player 1 has chosen a_1 and player 2's beliefs have support in $T \subset \Theta$:*

$$BR(T, a_1) = \cup_{\mu \in \Delta(T)} \arg\max_{a_2 \in A_2} \sum u_2(a_1, a_2, \theta)\, \mu(\theta)$$

Definition 8 *A PBE s^* of G* ***fails the intuitive criterion*** *if there exists $a_1 \in A_1$, $\theta' \in \Theta$ and $J \subset \Theta$ such that*

1. $u_1(s^*, \theta) > \max_{a_2 \in BR(\Theta, a_1)} u_1(a_1, a_2, \theta)$ *for all $\theta \in J$*
2. $u_1(s^*, \theta) < \min_{a_2 \in BR(\Theta/J, a_1)} u_2(a_1, a_2, \theta')$

Condition (1) implies that types in J would never try to play a_1 since even if they could convince player 2 that they were a particular type, they would do worse. Condition (2) implies that type θ' definitely does better by playing a_1 rather than the equilibrium so long as she can convince player 2 that her type is not in J.

Proposition 2 *In the job market signalling model, any pooling equilibrium fails the intuitive criterion.*

Proof. Consider a pooling equilibrium with some education level e^P. Then we look for a education level e that satisfies:

$$\theta^E - c\left(e^P, \theta_L\right) > \theta_H - c(e, \theta_L)$$

and

$$\theta^E - c\left(e^P, \theta_H\right) < \theta_H - c(e, \theta_H).$$

The Figure below show how to identify e.

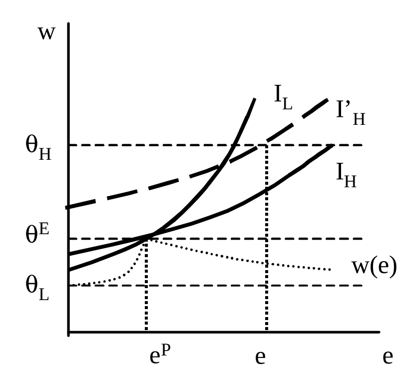

Tracing the indifference curves through $\left(w = \theta^E, e^P\right)$, we look for the education level that puts the low ability type on the same indifference curve when $w = \theta_H$. A slight increase to e will then suffice — at $(w = \theta_H, e)$, the low ability type will be on a lower indifference curve than at $\left(\theta^E, e^P\right)$, while the high ability type is on a higher indifference curve. *Q.E.D.*

It is also possible to show that of the separating equilibria derived above, only the most efficient — the separating equilibrium where the high ability cost type gets education $\underline{e}$ survives the intuitive criterion. This is typical of the intuitive criterion — in models with separating equilibria, it selects the most efficient. You are asked to work out another example on the problem set.

3 Cheap Talk

In signalling models, a player uses *costly action* to communicate information about his type. In some situations, however, effective costly signals may not be available. This raises the question of whether a player could credibly communicate information about his type simply through a costless process of communication. It is this possibility that we now explore.

3.1 The Model

The basic model of cheap talk is exactly like the signalling model, only player 1's action is a message that has no direct effect on payoffs.

Stage 0 Nature chooses type $\theta \in \Theta$ of player 1 from a distribution p.

Stage 1 Player 1 observes θ and chooses $m \in M$.

Stage 2 Player 2 observes m and chooses $a \in A$.

Payoffs $u_1(a, \theta)$ and $u_2(a, \theta)$.

We sometimes call player 1 the "sender" and player 2 the "receiver". We are interested in studying perfect bayesian equilibria of this game. The basic question is whether player 1's message will ever convey meaning in equilibrium.

3.2 Basic Observations

1. *Cheap talk doesn't work if different types of senders have the same preferences over actions.*

Example Consider a job market game where instead of going to school players just say "I'm high θ" or "I'm low θ". This game has payoffs:

	a_L	a_M	a_H
$\theta = L$	$1, 1$	$2, 0$	$3, -2$
$\theta = H$	$1, -2$	$2, 0$	$3, 1$

Claim. In this game there are no separating PBE.

Proof. Consider a candidate PBE where $s_1(\theta = L) = m$ and $s_1(\theta = H) = m'$ with $m \neq m'$. By Bayes Rule':

$$\mu_2(L|m) = 1 \qquad \Rightarrow \qquad a(m) = a_L$$

and so:

$$\mu_2(L|m') = 0 \qquad \Rightarrow \qquad a(m') = a_H$$

But then when $\theta = L$, player 1 would like to deviate since

$$u_1(a(m), L) = 1 < 3 = u_2(a(m'), L)$$

so this cannot be a PBE. *Q.E.D.*

2. *Cheap talk doesn't work if different types of senders have completely opposed preferences over actions.*

Example Suppose that Vince McMahon (player 2) is thinking about hiring Mike Tyson (player 1) as a WWF wrestler. In the Mike Tyson game, Mike can be either Normal or Crazy.

$$\Theta = \{\text{Normal, Crazy}\}$$

Mike wants the job if he is normal, but not is he is crazy. On the other hand, Vince wants to hire Mike if and only if Mike is crazy. We write the payoffs for Mike and Vince as follows:

	Hire	Don't Hire
$\theta = N$	$2, -2$	$0, 0$
$\theta = C$	$-2, 2$	$0, 0$

Prior to Vince's decision, Mike holds a press conference to announce his state of mind.

Claim. In this game, there are no separating PBE.

Proof. Suppose that $s_1(N) = m$, $s_2(C) = m'$ with $m \neq m'$. By Bayes' Rule:

$$\begin{aligned} \mu_2(N|m) &= 1 \qquad \Rightarrow \qquad a(m) = D \\ \mu_2(N|m') &= 0 \qquad \Rightarrow \qquad a(m') = H \end{aligned}$$

It follows that *both types* of Tyson want to switch messages:

$$\begin{aligned} u_1(a(m), N) &= 0 < 2 = u_1(a(m'), N) \\ u_1(a(m), C) &= 0 > -2 = u_1(a(m'), C). \end{aligned}$$

which means this cannot be an equilibrium. *Q.E.D.*

3. *Cheap talk can work great in coordination games.*

Example Consider a version of the "meeting place in New York" game where player 1 is already there and is not allowed to move. Suppose that player 1's "type" is his location in New York.

$$\Theta = \{\text{Grand Central Station, Empire State Building}\}.$$

Player 1 knows θ and calls player 2 on the phone. Player 2 listens and then chooses $\{G, E\}$. Payoffs are given by:

	G	E
$\theta = G$	$1, 1$	$0, 0$
$\theta = E$	$0, 0$	$1, 1$

This game has a separating equilibrium:

$$\begin{aligned} s_1(G) &= m & s_1(E) &= m' \\ \mu_2(\theta = G|m) &= 1 & \mu_2(\theta = G|m') &= 0 \\ a_2(m) &= G & a_2(m') &= E \end{aligned}$$

Remark 3 *Note that this last example also has a pooling or "babbling" equilibrium, where player 1 always says m (or m' or randomizes over all elements of M), and player 2 does not update her beliefs at all.*

3.3 A Richer Model

We now tackle a richer model of cheap talk where the possibilities for communication are more interesting. We assume that θ is uniformly distributed on the interval $[0, 1]$. Let player 2's action be denoted $a \in \mathbb{R}$. We assume that player 2 (the receiver) has payoffs $u_2(a, \theta) = -(a - \theta)^2$, while player 1 (the sender) has payoffs $u_1(a, \theta) = -(a - \theta - c)^2$.

Player 2's optimal action is to choose $a = \theta$, while player 1's is to choose $a = \theta + c$. Importantly, preferences are congruent in the sense that both like to take higher actions when the state is higher. However, player 1 is systematically biased by an amount c.

We investigate different types of perfect bayesian equilibria.

Babbling Equilibrium. Player 1 uses each message $m \in M$ with equal probability regardless of θ. Player 2 assigns a uniform belief to all values $\theta \in [0, 1]$ regardless of the message. She then chooses that action $a = 1/2$.

Clearly player 1 cannot deviate profitably given Player 2's beliefs. Similarly, Player 2's action is optimal given her beliefs. Finally, Player 2's beliefs are consistent with Bayes' Rule.

A Two Message Equilibrium. Let's try to construct an equilibrium where Player 1 uses two messages, m_1 and m_2. Let $a(m)$ denote player 2's action in response to message m. If the equilibrium conveys information, it must be the case that $a(m_1) \neq a(m_2)$. Assume without loss of generality that $a(m_1) < a(m_2)$.

Claim 1 In a two message equilibrium, player 1 uses a threshold strategy, choosing the message m_1 whenever $\theta \in [0, \theta^*)$ and m_2 whenever $\theta \in (\theta^*, 1]$, with θ^* indifferent (and choosing either).

Proof. To prove this note that player 1's payoff are as follows:

$$\begin{aligned} \text{Payoff to } m_1 &: \quad -(a(m_1) - \theta - c)^2 \\ \text{Payoff to } m_2 &: \quad -(a(m_2) - \theta - c)^2 \end{aligned}$$

So the incremental gain to choosing message m_2 versus m_1 is:

$$\Delta(\theta) = -(a(m_2) - \theta - c)^2 + (a(m_1) - \theta - c)^2$$

This is increasing in θ. Thus, it follows that if type θ prefers m_2 to m_1, so will every type $\theta' > \theta$. This gives a threshold characterization with θ^* being the type that is just indifferent. *Q.E.D.*

Claim 2 In equilibrium:

$$a(m_1) = \frac{\theta^*}{2} \qquad\qquad a(m_2) = \frac{1+\theta^*}{2}$$

Proof. This follows from Bayes' Rule and optimization by player 2. In equilibrium, player 2 must set $a(m) = \mathbb{E}[\theta|m]$. If all types between $[0, \theta^*)$ send message m_1, and all types between $(\theta^*, 1]$ send message m_2, then she must behave as stated. *Q.E.D.*

We are now ready to solve for the equilibrium. We use the fact that θ^* must be just indifferent between the two messages: $u_1(a(m_1), \theta^*) = u_1(a(m_2), \theta^*)$. This is equivalent to:

$$\theta^* + c - \frac{\theta^*}{2} = \frac{1+\theta^*}{2} - (\theta^* + c)$$

which simplifies to $\theta^* = \frac{1}{2} - 2c$. This equilibrium "works" so long as $\theta^* \in (0,1)$, i.e. so long as $c < 1/4$.

A few comments on the two message equilibrium.

1. Even if there are more than two messages available, this two message equilibrium can still be made into a PBE. For example, we can group all of the messages into two "classes" or message, with types in $[0, \theta^*)$ choosing from the first message class randomly and those in $(\theta^*, 1]$ choosing from the second.

2. Note also that the equilibrium is asymmetric — $\theta^* < 1/2$. This means that m_1 is a more informative message than m_2 in the sense that it narrows down the possibilities more.

A Three Message Equilibrium. What about an equilibrium with more communication? Let's try to construct an equilibrium with three messages m_1, m_2, m_3. Using precisely the same arguments, we can show that:

1. The equilibrium must divide types into three segments, with those in $[0, \theta_1)$ choosing m_1, those in (θ_1, θ_2) choosing m_2 and those in $(\theta_2, 1]$ choosing m_3, with θ_1, θ_2 indifferent between m_1, m_2 and m_2, m_3 respectively.

2. In equilibrium, player two must use:

$$a(m_1) = \frac{\theta_1}{2} \qquad a(m_2) = \frac{\theta_1 + \theta_2}{2} \qquad a(m_3) = \frac{1 + \theta_2}{2},$$

i.e. she must set $a(m) = \mathbb{E}[\theta|m]$ for $m = m_1, m_2, m_3$.

3. Finally, we have two indifference conditions. Type θ_1 must be indifferent between m_1 and m_2 :

$$(\theta_1 + c) - \frac{\theta_1}{2} = \frac{\theta_1 + \theta_2}{2} - (\theta_1 + c)$$

and type θ_2 must be indifferent between m_2 and m_3 :

$$(\theta_2 + c) - \frac{\theta_1 + \theta_2}{2} = \frac{1 + \theta_2}{2} - (\theta_2 + c).$$

Solving these two equations implies that $\theta_1 = \frac{1}{3} - 4c$ and $\theta_2 = \frac{2}{3} - 4c$.

This "works" as an equilibrium provided that $\theta_1 > 0$, or in other words if $c < \frac{1}{12}$. So if there is a three message equilibrium, there is also a two message equilibrium and a babbling equilibrium. Note also that the degree of bias must be smaller to have more communication (three messages in equilibrium as opposed to just two).

Arbitrary Message Equilibrium. It is possible to show that there are also 4,5,6,...,T message equilibria for some T which depends on c. In general, however, there is no fully separating (infinite message) equilibrium. The reason is that if messages were fully informative, then type θ would send the report corresponding to type $\theta + c$ and achieve a better outcome (i.e. the action $a = \theta + c$ rather than the action θ). Thus in any equilibrium, there is some loss of information. And moreover, the loss of information is directly related to the bias of the informed party. These results were originally worked out in a paper by Crawford and Sobel (1982).

4 Models of Reputation

Some of the most striking applications of incomplete information are to models of reputation. We will consider two kinds of reputation models.

- In the first model, we consider a long-run player who plays against a sequence of short-run playerss We find is that if short-run players entertain the possibility that the long-run player is irrational in a particular way, the long-run player may be able to build a reputation that works to his advantage.

- In the second model, we consider two long-run players. We find that again the slight possibility of irrationality changes the equilibrium greatly, for instance, allowing cooperation in the finitely repeated prisoners' dilemma.

4.1 Building a Reputation for Toughness

We start by considering the following entry game.

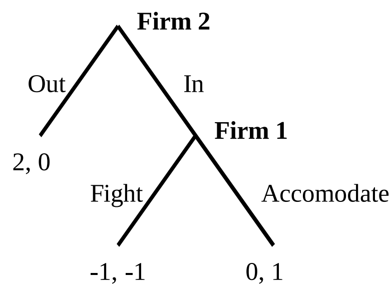

We consider a situation where this game is played T times. In each case, Player 1 is the same, but he faces a succession of Player 2s. Thus Player 1's objective is to maximize the undiscounted sum of his payoffs for all T periods, while each Player 2 wants to maximize her payoffs in the present period.

As an example of this type of situation, think of Player 1 as Microsoft, fending off a succession of lawsuits accusing it of anti-competitive behavior. There are a succession of small firms who feel they have a case against Microsoft and have to decide whether to launch a court battle or not. Microsoft then decides whether to settle (accomodate) or to fight tooth and nail (fight). The question is whether Microsoft might be able to fight some early lawsuits, and thus "build a reputation" for fighting so that later firms will not challenge it.

Proposition 3 *In the T-period complete information game, there is a unique SPE in which Player 1 accomodates in all periods and all entrants enter.*

Proof. Use backward induction. *Q.E.D.*

This result points to the strength of subgame perfection in complete information games. (Recall the centipede game or the Enron speculation game.)

Even if Microsoft is going to face a large number of challengers, it cannot build a reputation for being a fighter in this setting.

To incorporate reputation in this model, we relax the assumption that Player 1 is *necessarily* rational, and suppose instead that with small probability Player 1 actually enjoys fighting. To do this, suppose that Player 1 has two possible types:

$$\theta_1 \in \{\text{Crazy, Normal}\}$$

Player 1 is crazy with probability q. For a crazy type, we assume that Fight is a dominant strategy — a Crazy type of player 1 will fight in every period that an opponent enters.

Even in the $T = 1$ period game, this aspect of craziness may be enough to deter entry.

Proposition 4 *In the $T = 1$ period game with incomplete information:*

1. *If $q > 1/2$, the unique PBE has player 2 play Out.*
2. *If $q < 1/2$, the unique PBE has player 2 play In.*

Proof. This also can be solved by backward induction. *Q.E.D.*

However, the Normal player 1 behaves normally by accomodating the unique PBE. Things get more interesting when we go to the $T = 2$ period game.

Proposition 5 *In the $T = 2$ period game with incomplete information:*

1. *If $q > 1/2$, the unique PBE has the form:*

Period 1:	*Player 2 stays Out*
	Player 1 Fights
Period 2:	*Player 2 stays Out (unless he has seen Player 1 Accomodate)*
	Player 1 Accomodates

2. *If $q < 1/2$, the unique PBE has the form:*

Period 1:	*Player 2 stays Out if $q > 1/4$ and Enters if $q < 1/4$.*
	Player 1 Fights with Prob $q/(1-q)$ if 2 enters
Period 2:	*Player 2 enters (unless he has seen Player 1 Fight in which case he randomizes $\frac{1}{2}In + \frac{1}{2}Out$.*
	Player 1 Accomodates

Remark 4 *The strategies for Player 1 are for the case when he is Normal. If he is Crazy, he always Fights.*

We consider the two cases in turn.

Proof of (a): $q > 1/2$.

1. Because a Crazy Player 1 always plays Fight, Player 2 will not enter if she believes the probability that Player 1 is Crazy is greater than 1/2. Consequently, Player 2 certainly will not enter in the first period. Moreoever, if there is no first period entry, Player 2 will continue to believe in the second period that Player 1 is Crazy with probability $q > 1/2$. Consequently, she will not enter in the second period either.

We conclude that Player 2's equilibrium strategy is not to enter in either the first or the second period. However, we still need to describe what happens *off the equilibrium path* — i.e. when Player 2 *does enter* in the first period.

2. If Player 2 enters in the first period and Player 1 responds by Accomodating, this reveals that Player 1 is Normal. So in the second period, Player 2 will find it optimal to enter, and Player 1 will find it optimal to respond by Accomodating.

3. If Player 2 enters in the first period and Player 1 responds by Fighting, then regardless of the Normal Player 1's strategy, Player 2 must believe at the beginning of the second period that Player 1 is crazy with probability at least q. From Bayes' rule:

$$\Pr\left[\text{Crazy}|a_1 = \textit{Fight}\right] = \frac{q}{q + (1-q) \cdot \Pr\left[\text{Normal P1 } \textit{Fights}\right]} \geq q.$$

   It follows that in the second period, Player 2 will find it optimal not to enter, even though in the event of entry a Normal Player 1 would Accomodate.

We conclude that a Normal Player 1 would always Accomodate in the second period, but that following entry in the first period, Player 2 would enter in the second if and only if she saw Player 1 Accomodate in the first. Finally, we need to consider what a Normal Player 1 would do in the first period should Player 2 enter.

4. If Player 2 enters in the first period, then by playing Fight a Normal Player 1 can deter second-period entry. In particular, playing Fight

gives a payoff of $-1+2=1$. On the other hand, Accomodating gives a payoff of 0 today and encourages entry leading to 0 tomorrow. So Player 1 will optimally Fight should Player 2 enter in period 1.

This completely describes the perfect bayesian equilibrium if $q > 1/2$. *Q.E.D.*

Proof of (a): $q < 1/2$.

We start by considering what happens if Player 2 does not enter in the first period.

1. If Player 2 does not enter in the first period, then there is no learning about player 1's type. In the second period, we have a one-shot game with probability $q < 1/2$ that Player 1 is Crazy. Consequently, Player 2 will enter in the second period and a Normal Player 1 will accomodate.

Now consider what happens if Player 2 enters in the first period.

2. It *cannot* be a PBE for Player 1 to fight with probability 1 in the first period. If it was, then Player 2 would not update when she saw Player 1 fight and thus would enter in the second period regardless of whether or nor Player 1 fought or accomodated in the first period. Consequently, a Normal Player 1 would do better to accomodate in the first period.

3. It also *cannot* be a PBE for Player 1 to accomodate with probability 1 in the first period. If it was, then fighting would perfectly signal craziness. After seeing Fight, Player 2 would not enter the second period. But then the payoff to Fighting in the first period for Normal Player 1 would be $-1+2=1$, greater than the payoff 0 to accomodating.

We conclude that the PBE must have Player 1 randomize in the first period. For Player 1 to be willing to randomize, he must be indifferent to playing Fight and Accomodate in the first period. This means that:

$$\begin{aligned} 0 &= -1 + 2 \cdot \Pr\left[2 \text{ Out in period } 2 | a_1 = \textit{Fight}\right] \\ &\quad +0 \cdot \Pr\left[2 \text{ Enters in period } 2 | a_1 = \textit{Fight}\right]. \end{aligned}$$

For this to happen, it must be that:

$$\Pr\left[2 \text{ Out in period } 2 |\ a_1 = \textit{Fight}\right] = \frac{1}{2}.$$

We conclude from this that the PBE must have Player 2 randomize in the second period after seeing Fight in the first.

For Player 2 to be willing to randomize in the second period after seeing Fight in the first, it must be that:

$$\Pr\left[\text{Crazy}|a_1 = \textit{Fight}\right] = \frac{1}{2}.$$

In equilibrium, this belief will be obtained through Bayes' updating:

$$\Pr\left[\text{Crazy}|\ a_1 = \textit{Fight}\right] = \frac{q}{q + (1-q) \cdot \Pr\left[\text{Normal P1 } \textit{Fights} \text{ in Period 1}\right]}.$$

Combining these requirement, we find that the Normal Player 1 must fight entry in period 1 with probability $q/\left(1-q\right)$.

At this point, we have determined (a) what Normal Player 1 must do in period 1 if Player 2 enters, and (b) what Player 2 will do after (i) not entering in period 1, (ii) entering in period 1 and seeing Fight, and (iii) entering in Period 1 and seeing Accomodate.

The remaining question is whether Player 2 will enter in the first period. Her payoff to entering is:

$$-\left(q + (1-q) \cdot \frac{q}{1-q}\right) + (1-q) \cdot \frac{1-2q}{1-q} = 1 - 4q.$$

while staying out gives zero We conclude that Player 2 will optimally enter in the first period if $q < 1/4$, and will optimally stay out in the first period if $q > 1/4$. *Q.E.D.*

The general T-period version of this model has also been studied. It turns out that if T is large, then even if q is quite small, the Normal type of Player 1 will be able to mimic the Crazy type and obtain a large period period payoff. Remarkably, as $T \to \infty$, the Normal player's average payoff per-period approaches 2 (which is the highest payoff that player 1 could get by committing to a strategy in advance — i.e. his "Stackelberg" payoff). In this sense, the slight chance that a player is crazy may allow that player to do very well in equilibrium.

4.2 Cooperation in the Finitely Repeated Prisoners' Dilemma

In the above example, the long-run Player 1 was able to build a reputation to his own advantage. Reputation can also play a powerful role in situations where players may be able to *constructively* mislead others about their objectives. To see how this works, we consider Kreps, Milgrom, Roberts, and Wilson's (1982) treatment of the finitely repeated prisoners' dilemma.

The stage game is the standard prisoners' dilemma.

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

We assume the game is repeated T times with no discounting, with both player 1 and 2 being long-run players. Recall that if both players are maximizing their payoffs and there is complete information, then backward induction tells us that (D, D) will be played in every period of the unique subgame perfect equilibrium.

To add reputation, we imagine that player 2 is either "crazy" with probability q or "normal" with probability $1 - q$. A normal player maximizes his expected payoff in the standard way, but a crazy player mechanically plays the Grim Trigger strategy. That is a crazy Player 2 plays C so long as Player 1 has never played D, and plays D forever once Player 1 has played D for the first time.

If this this incomplete information game is played once, both Player 1 and the Normal Player 2 have a dominant strategy, which is to play D. However, if the game is played more than once, the possibility that Player 2 might be crazy gives Player 1 some incentive to play C in the early rounds. More subtle is the fact that it also gives Player 2 an incentive to play C — by doing so Player 2 can potentially convey the impression that he really is crazy.

Proposition 6 *In the $T = 2$ period model, the normal Player 2 always plays D, while Player 1 plays C in the first period if $q > 1/2$.*

Proof. In the second period, both the Normal Player 2 and Player 1 will clearly want to choose D regardless of what happened in the first period and regardless of Player 1's beliefs about Player 2's type. The Crazy Player 2 will play C if Player 1 played C in the first period and D if Player 1 played D.

In the first period, the Normal Player 2 will certainly play D since regardless of what she plays she expects Player 1 to play D in the last peiriod. The Crazy Player 2 will play C. Now consider Player 1. She has payoffs:

$$\begin{aligned} \text{Payoff to } C &: \quad q - (1 - q) + q \cdot 2 = 4q - 1 \\ \text{Payoff to } D &: \quad 2q \end{aligned}$$

Thus Player 1 optimally plays C in the first period if $q > 1/2$ and optimally plays D if $q < 1/2$. *Q.E.D.*

Proposition 7 *In the $T = 3$ period model, if $q > 1/2$, then both players cooperate in the first round.*

Proof. To analyze the three period game, we start by considering continuation play from date $t = 2$.

- If player 1 played D in the first period, then both players will play D thereafter.

- If player 1 played C in the first period, then let μ denote player 1's belief that player 2 is crazy. Equilibrium play in the last two periods will be just as in the 2-period model. In particular, player 1 will optimally play C in the second period if $\mu > 1/2$ and D in the second period if $\mu < 1/2$, and be indifferent if $\mu = 1/2$.

Now let's back up and try to construct an equilibrium in which both players cooperate in round one. Clearly if this happens, then no information will be revealed as to whether player 2 is truly crazy, so $\mu = q$.

- Supposing that $q > 1/2$, and that player 1 plays C in the first period, what should a normal player 2 do?

  $$\begin{aligned} \text{Payoff to } C &: \quad 1 + 2 + 0 \\ \text{Payoff to } D &: \quad 2 + 0 + 0 \end{aligned}$$

  Playing D results in D always being played in the last two periods, while playing C results in player 1 playing C in the next period. Thus a normal player 2 should play C in the first period so as not to reveal that he is normal.

- Supposing that $q > 1/2$, and that player 2 plays C in the first period, what should player 1 do?

  $$\begin{aligned} \text{Payoff to } C &: \quad 1 + q \cdot (1 + 2) + (1 - q) \cdot (-1) = 4q \\ \text{Payoff to } D &: \quad 2 + 0 + 0 \end{aligned}$$

So player 1 should also play C in the first period. *Q.E.D.*

Note that the equilibrium involves full cooperation in the first period, partial cooperation in the second period (only player 1 cooperates) and no cooperation in the third round. This is consistent with the way people play the finitely repeated prisoners' dilemma in experiments: some cooperation at the start, but hten tailing off.

A second point is that this is only one equilibrium, and we need parametric assumptions ($q > 1/2$) to get it. So cooperation is possible, but not necessary. The remarkable fact is that if the game is repeated enough, then cooperation is not only possible, but *necessary* in the sense that players cooperate in *almost every period.*

Proposition 8 *In any sequential equilibrium of the T-period game where T is "large," the number of periods where one player or the other plays D is bounded above by a constant that depends on q, but is independent of T.*

Proof (sketch)**.** The first point to note is that if it ever becomes known that player 2 is normal prior to some round t, then both players must select D in every following period. This follows from backward induction. Thus, if player 1 has never played D and player 2 plays D, D will be played in all following periods.

Similarly, if player 1 has never played D and player 2 plays D, then D will be played in all following periods. This happens because either player 2 is crazy, in which case he is triggered and will play D forever, or else he is normal. If he is normal and ever plays C, then he reveals normality and ends up playing (D, D) forever anyway. So player 2 will certainly play D forever following a D by player 1. And if player 2 is playing D forever, so should player 1.

We conclude that once D is played once, it will be played forever by player 1 and both types of player 2 in any sequential equilibrium.

Next, define $M = \frac{3-q}{q}$. We claim that if neither player has yet played D in any round up to and including t where $t < T - M$, then player 1 must select C in round $t + 1$. To see this note that once player 1 plays D, (D, D) will be played forever, so

$$\text{Player 1's payoff to } D: \text{ at most } 2 + 0 + ... + 0 = 2$$

On the other hand, Player 1 has the option of playing C and continuing to play C until either the last period, or until player 2 plays D for the first time. This strategy will give a payoff of *at least:*

$$q \cdot \left(\underbrace{1 + 1 + ... + 1}_{=T-t-2} + 2 \right) + (1-q) \cdot (-1) = (M+1)\, q - 1$$

So playing C is certainly optimal if:

$$(T - t - 1)\, q - 1 \geq 2 \qquad \Leftrightarrow \qquad T - t \geq \frac{3-q}{q},$$

which is the case since $T - t > M$.

Finally, we claim that if no player has yet played D in any round up to and including t where $t < T - M - 1$, then player 2 must select C in period $t + 1$. To see this, note that the crazy player 2 will certainly select C. If the normal player 2 selects D, then (D, D) will be played forever, so

$$\text{Player 1's payoff to } D: \text{ at most } 2 + 0 + ... + 0 = 2.$$

On the other hand, Player 2 has the option of playing C once and then D in all following rounds. This will give a payoff of *at least*

$$1 + 2 + 0 + ... + 0 = 3.$$

So player 2 certainly will play C in this case.

We conclude that in any period $t < T - M$, neither player will select D. Note that M is independent of T, so as $T \to \infty$, players cooperate for almost every period. *Q.E.D.*

We have not shown that an equilibrium actually exists, but there is a general existence result for sequential equilibrium, which tells us that there is an equilibrium in which (C, C) is played for a significant amount of time in equilibrium. More generally, equilibrium will have the following structure. Players start out by cooperating. As the end draws near, one or both players may start to mix between C and D with the probability that D is played in any given period increasing over time. Once a D has been played, (D, D) is played forever.

1. This pattern of cooperation matches up reasonably well with lab experiments: players start out cooperating and cooperation tails off as the end of the game approaches.

2. Interestingly, our last result did not require specific assumptions on q. Even if q is small, for large T there will be a lot of cooperation.

3. We relied heavily on a particular form of craziness to get this result. It turns out that if one allows for all conceivable forms of craziness, and if both players may be crazy, there is a folk theorem type result — anything can happen.

References

[1] Cho, I.K. and D. Kreps (1987) "Signalling Games and Stable Equilibria," *Quarterly Journal of Economics.*

[2] Crawford, V. and J. Sobel (1982) "Strategic Information Transmission," *Econometrica.*

[3] Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982) "Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory.*

[4] Kreps, D. and R. Wilson (1982) "Sequential Equilibrium," *Econometrica.*

[5] Spence, M. (1974) *Market Signalling*, Harvard University Press.