Correlated Equilibria, Repeated Games

Scribe: Devdatt Dubhashi

March 10, 2003

1 Correlated Equilibria

1.1 Elementary Probability Background

Given sets $[s] := \{1, \ldots, s\}$ and $[t] := \{1, \ldots, t\}$, a *joint distribution* on $[s] \times [t]$ are non-negative real numbers $p(i, j), i \in [s], j \in [t]$ such that $\sum_{i,j} p(i, j) = 1$. Given a joint distribution p, the *marginals* are $p(i) := \sum_j p(i, j)$ and similarly, $p(j) := \sum_i p(i, j)$.

A *special case*: if $\tilde{p}$ is a probability measure on $[s]$ and $\tilde{q}$ one on $[t]$ (think of them as strategies), then $p(i, j) := \tilde{p}(i) \cdot \tilde{q}(j)$ is a probability distribution on $[s] \times [t]$ called the *product measure.* We say that the choices from $[s]$ and $[t]$ are *independent.*

The *conditional probability* of "i given j" is $p(i \mid j) := p(i, j)/p(j)$, provided $p(j) \neq 0$.

1.2 Definition and Intuition

Assume we have a finite game where

- player I's moves are $[s]$, and
- player II's moves are $[t]$,

and we have a bi-matrix (A, B) of respective payoffs to I and II.

Definition 1 *A* **correlated equilibrium** *is a probability measure on* $[s] \times [t]$ *such that for all* $i, i' \in [s]$ *and* $j, j' \in [t]$,

$$\sum_j a_{i,j} p(j \mid i) \geq \sum_j a_{i',j} p(j \mid i) \quad \textit{provided } p(i) > 0, \tag{1.1}$$

$$\sum_i b_{i,j} p(i \mid j) \geq \sum_i b_{i,j'} p(i \mid j) \quad \textit{provided } p(j) > 0, \tag{1.2}$$

The idea is that (i, j) is picked according to a probability distribution known to both players by a *trusted third party* ("Monty"). Then Monty reveals i to player I and j to player II. But *each player is* **not** *told what the other is told.* Hence, if player I is told to play i, from his perspective, the distribution of player II's moves is the conditional distribution $p(j \mid i)$ i.e. play move j with this probability. Then $\sum_j a_{i,j} p(j \mid i)$ is the expected gain of player I if both play according to Monty's advice. On the other hand, $\sum_j a_{i',j} p(j \mid i)$ is the expected gain of player I if player II follows Monty, but player I switches instead to i'. Now equation 1.1 simply asserts that player I has no incentive to deviate unilaterally while equation 1.2 asserts that player II has no incentive to deviate unilaterally. That is, following Monty is an equilibrium. Note that *each player gets partial information which is private to that player.*

1.3 Example

Consider the bi-matrix:

$$\begin{pmatrix} (5,1) & (0,0) \\ (4,4) & (1,5) \end{pmatrix}$$

Our goal is to find a correlated equilibrium with as good a payoff to both players as possible. The expected payoff to player I of a correlated euilibrium p is $\sum_{i,j} a_{i,j} p(i, j)$.

1. Pick the Nash equlibria $(1, 1)$ and $(2, 2)$ equiprobably i.e. $p(1, 1) = 1/2 = p(2, 2)$ and $p(1, 2) = 0 = p(2, 1)$. Check that this is a correlated equlibrium. The payoffs are: $(\frac{1}{2}(5 + 1), \frac{1}{2}(5 + 1)) = (3, 3)$.

2. Consider instead the following joint distribution:

$$p(1,1) = p(2,1) = p(2,2) = \frac{1}{3}, p(1,2) = 0.$$

   verify that this is a correlated equilibrium. The payoffs are:

$$(\frac{1}{3}(5+4+1), \frac{1}{3}(5+4+1)) = (\frac{10}{3}, \frac{10}{3}).$$

1.4 Example: A three person game

$$\begin{pmatrix} (0,1,3) & (0,0,0) \\ (1,1,1) & (1,0,0) \end{pmatrix}, \quad \begin{pmatrix} (2,2,2) & (0,0,0) \\ (2,2,0) & (2,2,2) \end{pmatrix}, \quad \begin{pmatrix} (0,1,0) & (0,0,0) \\ (1,1,0) & (1,0,3) \end{pmatrix}$$

Player I chooses a row, player II a column and player III the matrix.
$(1,1,1)$ is a pure Nash equlibrium (unique!).
A correlated equlibrium can do better:

$$p(1,1,2) = 1/2 = p(2,2,2).$$

Player III gains by limiting his own information (!) since knowing this the other players are induced into playing in a way desirable to player III.

The equations (1.1) and (1.2) can be rewritten as *linear* inequalities for the $p(i,j)$. Hence the set of correlated equilibria is a compact convex set and contains the convex hull of the Nash equlibria.

Remark 2 (Correspondence between Correlated and Nash Equlibria) *A joint distribution which is a product measure is a correlated equilibrium iff it is a Nash equlibrium. In particular, each Nash equlibrium is a correlated equilibrium.*

Exercise 3 *It follows that any convex combination of Nash equlibria a correlated equilibrium? Is this obvious directly?*

Questions about max payoffs over all correlated equlibria could be interpreted as maximising the sum of the two payoffs (or some linear combination of them). This is then a *linear programming* problem.

2 Repeated Games

The same game is played many times. The payoff is the sum of the payoffs of each play. Suppose we have a game with move spaces X and Y. The space of moves of the n-move game is not simply X^n and Y^n: the moves might depend on the *history*.

To take this into account, player I's moves are given by a set of functions $\{f_i, i \in [n]\}$ where

$$f_k : X^{k-1} \times Y^{k-1} \to X.$$

(Note f_1 is just an element of X.) The interpretation is that $f_k((x_i, y_i), i < k)$ is player I's kth move if the history upto this point is

- player I's moves: $x_1, x_2, \ldots, x_{k-1}$
- player II's moves: $y_1, y_2, \ldots, y_{k-1}$.

Our goal is to understand the n-move game G^n in terms of the original game G.

2.1 Example:The Prisoners' Dilemma

The Prisoner's Dilemma (PD) is the game with bi-matrix:

$$\begin{pmatrix} (5,5) & (0,6) \\ (6,0) & (1,1) \end{pmatrix}$$

Note that $(2,2)$ is the unique Nash equlibrium.

An easy induction shows:

Theorem 4 *If the Prisoner's Dilemma is played n times, there is a unique Nash equlibrium, namely, always do* $(2,2)$*.*

More generally,

Observation 5 *If* (p,q) *is a Nash equlibrium for* G*, then* (p^n, q^n) *is a Nash equlibrium for* G^n*.*

But there can be many others! (Peter's claim: this unique if G has a unique Nash equlibrium.)

For $0 < \delta < 1$, the δ-*discounted* version $G(\delta)$ of a game G is the infinite game where the payoff on the nth game is $(1-\delta)\delta^{n-1}$ of the original game. That is, suppose player I's moves are $(f_i)_{i \geq 1}$ and player II's moves are $(g_j)_{j \geq 1}$. Then the payoff for player I is:

$$A^{\infty}((f_i),(g_j)) = \sum_{t \geq 1} (1-\delta)\delta^{t-1} A(x_t, y_t),$$

(where x_t, y_t are the moves played on move t according to (f_i) and (g_j). Another interpretation is that this is the original game G played a random number of times, whose distribution is a geometric distribution with parameter δ.

Theorem 6 *If* δ *is small enough, then there is a unique Nash equlibrium for* $PD(\delta)$*, namely, do* $(2,2)$ *always.*

The intuition is that if δ is small, then everything essentially is determined by the first game.

If δ is large, on the other hand, there are other strategies. Consider the *unforgiving strategy,* s*: start with move 1 and stick with it as your opponent also does 1. But as soon as your opponent does 2, switch to 2 and stick with it forever.*

Theorem 7 *If δ is large enough, then (s, s) is a Nash equlibrium.*

The intuition is that one establishes *trust* for later games by playing 1 and retaining it unless there is betrayal. Since δ is large, it pays to establish this trust for later games.

2.2 Stupefying Example

Consider the game:

$$\begin{pmatrix} (8,8) & (1,2) \\ (2,1) & (0,0) \end{pmatrix}$$

This looks innocuous with $(1,1)$ clearly the only Nash equlibrium. *However for δ large enough, the game $G(\delta)$ has other Nash equlibria!.*