MAU34804

Lecture 20

2026-03-18

David Wilkins

Unfortunately David Wilkins, who developed this module and whose notes we've been using all semester, has died. For anyone interested his funeral will be at 14.15 on Tuesday 24 March in the Garden Chapel at Mount Jerome Cemetery.

Recovering the von Neumann theorem

Theorem 6.3 (Existence of Nash Equilibria) *for every collection of non-empty, compact convex* $S_i \subseteq \mathbf{R}^{m_i}$ *and continuous quasiconcave* $u_i \colon S \to \mathbf{R}$ *for* $1 \le i \le n$*, where* $S = S_1 \times \cdots \times S_n$*, there is a Nash equilibrium.*

Is this a true generalisation of the minimax theorem proved earlier?

Mathematically, yes. To get von Neumann's theorem it suffices to take $n = 2$, $S_1 = \Delta_P$, $S_2 = \Delta_Q$, $u_1 = f$ and $u_2 = -f$.

It's straightforward to check that the hypotheses on S_i and u_i are satisfied.

Relabeling $\mathbf{x}_1$, $\mathbf{x}_1^*$, $\mathbf{x}_2$ and $\mathbf{x}_2^*$ as $\mathbf{p}$, $\mathbf{p}^*$, $\mathbf{q}$ and $\mathbf{q}^*$ we can rewrite

$$u_i\left(\mathbf{x}_1^*, \ldots, \mathbf{x}_{i-1}^*, \mathbf{x}_i, \mathbf{x}_{i+1}^*, \ldots, \mathbf{x}_n^*\right) \le u_i\left(\mathbf{x}_1^*, \ldots, \mathbf{x}_{i-1}^*, \mathbf{x}_i^*, \mathbf{x}_{i+1}^*, \ldots, \mathbf{x}_n^*\right)$$

as

$$f(\mathbf{p}, \mathbf{q}^*) \le f(\mathbf{p}^*, \mathbf{q}^*) \qquad -f(\mathbf{p}^*, \mathbf{q}) \le -f(\mathbf{p}^*, \mathbf{q}^*)$$

or, equivalently,

$$f(\mathbf{p}, \mathbf{q}^*) \le f(\mathbf{p}^*, \mathbf{q}^*) \le f(\mathbf{p}^*, \mathbf{q}).$$

Difference of interpretation

Economically, it's not such a good generalisation.
Recall that

$$f(\mathbf{p}, \mathbf{q}^*) \leq f(\mathbf{p}^*, \mathbf{q}^*)$$

means Player A will be at least as well off choosing $\mathbf{p}^*$ as any other available strategy, assuming Player B chooses $\mathbf{q}^*$, while

$$f(\mathbf{p}^*, \mathbf{q}^*) \leq f(\mathbf{p}^*, \mathbf{q})$$

means they'll be no worse off if Player B chooses some other available strategy, assuming they've chosen $\mathbf{p}^*$.
The analogue of the two inequalities above in an n-person game would be

$$u_i\left(\mathbf{x}_1^*, \ldots, \mathbf{x}_{i-1}^*, \mathbf{x}_i, \mathbf{x}_{i+1}^*, \ldots, \mathbf{x}_n^*\right) \leq u_i\left(\mathbf{x}_1^*, \ldots, \mathbf{x}_{i-1}^*, \mathbf{x}_i^*, \mathbf{x}_{i+1}^*, \ldots, \mathbf{x}_n^*\right),$$

$$u_i\left(\mathbf{x}_1^*, \ldots, \mathbf{x}_{i-1}^*, \mathbf{x}_i^*, \mathbf{x}_{i+1}^*, \ldots, \mathbf{x}_n^*\right) \leq u_i\left(\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}, \mathbf{x}_i^*, \mathbf{x}_{i+1}, \ldots, \mathbf{x}_n\right)$$

We're going to get the first of these, but not the second.

Proof

Theorem 2.23 (Berge's Maximum Theorem) *Let X and Y be subsets of Euclidean spaces, let $f\colon X \times Y \to \mathbf{R}$ be a continuous function, and let $\Phi\colon X \rightrightarrows Y$ be a non-empty valued compact valued upper and lower hemicontinuous correspondence. Let*

$$m(\mathbf{x}) = \max\{f(\mathbf{x},\mathbf{y})\colon \mathbf{y} \in \Phi(\mathbf{x})\}, \quad M(\mathbf{x}) = \{\mathbf{y} \in \Phi(\mathbf{x})\colon f(\mathbf{x},\mathbf{y}) = m(\mathbf{x})\}$$

for all $\mathbf{x} \in X$. Then $m\colon X \to \mathbf{R}$ is a continuous function and $M\colon X \rightrightarrows Y$ is an upper hemicontinuous non-empty valued compact valued correspondence.

Apply this with $X = S = S_1 \times \cdots \times S_n$, $Y = S_i$, $\Phi(\mathbf{x}) = Y$ for all $\mathbf{x}$, and

$$f((\mathbf{x}_1,\ldots,\mathbf{x}_n),\mathbf{y}) = u_i(\mathbf{x}_1,\ldots,\mathbf{x}_{i-1},\mathbf{y},\mathbf{x}_{i+1},\ldots,\mathbf{x}_n).$$

These X, Y, Φ, and f satisfy the hypotheses so m_i and M_i satisfy the conclusion.
If $\mathbf{y} \in M_i(\mathbf{x})$ then

$$u_i(\mathbf{x}_1,\ldots,\mathbf{x}_{i-1},\mathbf{y},\mathbf{x}_{i+1},\ldots,\mathbf{x}_n) \geq u_i(\mathbf{x}_1,\ldots,\mathbf{x}_{i-1},\mathbf{z},\mathbf{x}_{i+1},\ldots,\mathbf{x}_n)$$

for all $\mathbf{z} \in S_i$.

Proof, continued

$M_i(\mathbf{x})$ is the set where u_i is equal to its maximum, or, equivalently, where it is greater than or equal to or greater than its maximum. $\mathbf{u}_i$ is quasiconcave, so by Lemma 6.2 the set $M_i(\mathbf{x})$ is convex.

In the statements above I'm always regarding u_i as a function on S_i by restricting the arguments other than the i'th one to be the corresponding entries of $\mathbf{x}$.

The Cartesian product of non-empty valued, compact valued, convex valued upper hemicontinuous correspondences is a non-empty valued, compact valued, convex valued upper hemicontinuous correspondence, so $M = M_1 \times \cdots \times M_n \colon S \rightrightarrows S$ is such a correspondence. By Proposition 2.11 its graph is closed.

All the hypotheses of the Kakutani fixed point theorem are satisfied so M has a fixed point, i.e. a point $\mathbf{x}^*$ such that $\mathbf{x}^* \in M(\mathbf{x}^*)$ or, equivalently, such that $\mathbf{x}_i^* \in M_i(\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)$ for each i.

In other words,

$$u_i(\mathbf{x}_1^*, \ldots, \mathbf{x}_{i-1}^*, \mathbf{x}_i^*, \mathbf{x}_{i+1}^*, \ldots, \mathbf{x}_n^*) \geq u_i(\mathbf{x}_1^*, \ldots, \mathbf{x}_{i-1}^*, \mathbf{x}_i, \mathbf{x}_{i+1}^*, \ldots, \mathbf{x}_n^*)$$

for all $\mathbf{x}_i \in S_i$.

Prisoner's Dilemma

The following is the original statement of the Prisoner's Dilemma "game" from Tucker (1950):
Two members of a criminal gang are arrested and imprisoned. Each prisoner is in solitary confinement with no means of communicating with the other. The prosecutors lack sufficient evidence to convict the pair on the principal charge, but they have enough to convict both on a lesser charge. Simultaneously, the prosecutors offer each prisoner a bargain. Each prisoner is given the opportunity either to betray the other by testifying that the other committed the crime, or to cooperate with the other by remaining silent. The possible outcomes are:

- If A and B each betray the other, each of them serves two years in prison.
- If A betrays B but B remains silent, A will be set free and B will serve three years in prison (and vice versa).
- If A and B both remain silent, both of them will serve only one year in prison (on the lesser charge).

Analysis

The strategy sets S_i are both the simplex $\{(p_1, p_2) \in \mathbf{R}^2_+ : p_1 + p_2 = 1\}$, where p_1 is the probability of staying silent and p_2 is the probability of betraying the other prisoner. We assume the utility of each prisoner is minus their expected prison sentence:

$$u_1((p_1, p_2), (q_1, q_2)) = -p_1 q_1 - 3p_1 q_2 - 2p_2 q_2 = -2q_2 - p_1,$$
$$u_2((p_1, p_2), (q_1, q_2)) = -p_1 q_1 - 3p_2 q_1 - 2p_2 q_2 = -2p_2 - q_1.$$

Each of these is quadratic, like the payoff in a von Neumann game, but $u_1 + u_2 = 2 - p_2 - q_2$ is not constant.

$$m_1((p_1, p_2), (q_1, q_2)) = -2q_2, \quad M_1((p_1, p_2), (q_1, q_2)) = P((q_1, q_2) = \{(0, 1)\},$$

$$m_2((p_1, p_2), (q_1, q_2)) = -2p_2, \quad M_2((p_1, p_2), (q_1, q_2)) = Q(p_1, p_2) = \{(0, 1)\},$$

$$M((p_1, p_2), (q_1, q_2)) = (P(q_1, q_2), Q(p_1, p_2)) = \{((0, 1), (0, 1))\}$$

so the only fixed point is $((0, 1), (0, 1))$, i.e. where each prisoner betrays the other.

The Ultimatum Game

- I discussed the ultimatum game in Lecture 6 as a variant of the cake cutting game.
- That discussion was informal but we can now treat it more formally.
- In the original version the first player was allowed to offer a slice of any size, but this would make the problem infinite dimensional and we wouldn't be able to apply our theorem on Nash equilibria to it, so I'll assume their is some fixed m such that the first player has to offer a fraction j/m of the cake to the second player.
- $j = 0$ and $j = m$ both have their own peculiarities, so I'll assume $m > 2$ and j is restricted to $0 < j < m$.
- The pure strategies for the first player will be these numbers j, while the pure strategies for the second player will consist of sets of such numbers, specifically those which the player is willing to accept.
- As usual, mixed strategies assign probabilities to pure strategies.
- The players choose their mixed strategies in isolation from each other and then the game proceeds by sampling these probability distributions and distributing the cake, or not, depending on the pure strategies selected.
- Utility is expected amount of cake.

Analysis

There are other ways to formulate the ultimatum game, often leading to different conclusions, but the version on the previous slide is the one we'll analyse.

As with other two person games we've considered $S_1 = \Delta_P$, $S_2 = \Delta_Q$, $P\colon \Delta_Q \rightrightarrows \Delta_P$ gives the set of utility maximising responses for the first player to a given strategy from the second player, and $Q\colon \Delta_P \rightrightarrows \Delta_Q$ gives the set of utility maximising responses for the second player to a given strategy from the first player.

The Nash equilibria are the fixed points of the correspondence $M\colon S \rightrightarrows S$ defined by

$$M(\mathbf{p}^*, \mathbf{q}^*) = (P(\mathbf{q}^*), Q(\mathbf{p}^*)).$$

We start by finding the correspondences P and Q. The first player's utility function is

$$u_1(\mathbf{p}, \mathbf{q}) = \sum_{j=1}^{m-1} p_j \left(\frac{m-j}{m} \sum_{K\colon j \in K} q_K \right).$$

The $\mathbf{p} \in P(\mathbf{q})$ are those where $p_j = 0$ unless j is a maximiser of $\frac{m-j}{m} \sum_{K\colon j \in K} q_K$.

Analysis, continued

The second player's utility function is

$$u_2(\mathbf{p}, \mathbf{q}) = \sum_K \left(q_K \sum_{j \in K} \frac{j}{m} p_j \right).$$

The quantity $\sum_{j \in K} \frac{j}{m} p_j$ has a maximum possible value of $\sum_{j=1}^{m-1} \frac{j}{m} p_j$, attained when $K = \{1, \ldots, m-1\}$. There may be other K's which attain this maximum though. The condition for K to be a maximiser is that if $p_j > 0$ then $j \in K$.

So the $\mathbf{q} \in Q(\mathbf{p})$ are those for which $q_K = 0$ unless K is a superset of $\{j \colon p_j > 0\}$.

So $(\mathbf{p}, \mathbf{q})$ is a fixed point of M if and only if the following two conditions are satisfied:

- $p_j = 0$ unless j is a maximiser of $\frac{m-j}{m} \sum_{K \colon j \in K} q_K$.
- $q_K = 0$ unless K is a superset of $\{j \colon p_j > 0\}$.

Suppose $(\mathbf{p}, \mathbf{q})$ is a fixed point with $p_i > 0$ and that $i < j$.

Then $q_K > 0$ only if $i \in K$ so $\sum_{K \colon i \in K} q_K = 1$ and

$$\frac{m-i}{m} \sum_{K \colon i \in K} q_K = \frac{m-i}{m}.$$

Analysis, conclusion

The conditions for a Nash equilibrium were that $p_j = 0$ unless j is a maximiser of $\frac{m-j}{m}\sum_{K\colon j\in K} q_K$, and that $q_K = 0$ unless K is a superset of $\{j\colon p_j > 0\}$.
If $p_i > 0$ and $i < j$ then

$$\frac{m-j}{m}\sum_{K\colon j\in K} q_K \le \frac{m-j}{m} < \frac{m-i}{m} = \frac{m-i}{m}\sum_{K\colon i\in K} q_K.$$

So j can't be a maximiser and so by the first condition $p_j = 0$. It follows that there is only one positive p, i.e. that the first player's optimal strategies are all pure.
If $(\mathbf{p}, \mathbf{q})$ is fixed point with $\mathbf{p}$ being the i'th pure strategy then the second condition says that $i \in K$ for all K with $q_K > 0$, while from the first condition we get

$$\frac{m-j}{m}\sum_{K\colon h\in K} q_K \le \frac{m-h}{m}, \qquad \sum_{K\colon h\in K} q_K \le \frac{m-i}{m-h}$$

for all h. This only imposes a restriction when $h < i$.

Comments

The Nash equilibria are the $(\mathbf{p}, \mathbf{q})$ such that $p_i = 1$ for some i and $q_K = 0$ whenever $i \notin K$ and $\sum_{K:\ h \in K} q_K \leq \frac{m-i}{m-h}$ for $h < i$.

Actually, I only proved that all Nash equilibria are of this form but it's easy to check that all points of this form are Nash equilibria.

The set of Nash equilibria is very large, with dimension $2^{m-1} - 1$, and there are lots of weird Nash equilibria.

For example, there's a Nash equilibrium where the first player only offers slices of size $\frac{1}{m}$ and the second player only accepts such slices.

Refusing better offers seems irrational, but it doesn't violate the conditions for a Nash equilibrium since such offers will never be made at this equilibrium.

There's another one where the first player only offers slices of size $\frac{m-1}{m}$ and the second player only accepts such slices.

The first player only ever makes the most generous possible offer, even though there's another equilibrium where the first player always makes the stingiest offer and it's always accepted.