Weeks 4-5:
Mixed Strategies

Dr Daniel Sgroi

Reading: 1. Osborne chapter 4;
2. Snyder & Nicholson, pp. 247–252.

With thanks to Peter J. Hammond.

Mixed Strategies

So far we have considered only pure strategies,
and players' best responses to deterministic beliefs.

Now we will allow mixed or random strategies,
as well as best responses to probabilistic beliefs.

Many games have no pure strategy Nash equilibrium.
But we will discuss why every finite game
has at least one mixed strategy Nash equilibrium.

Matching Pennies

This is a classic "two-person zero sum" game.
Players 1 and 2 each put a penny on a table simultaneously.
If the two pennies match (heads or tails)
then player 1 gets both; otherwise player 2 does.

		P_2	
		H	T
P_1	H	1^* -1	-1 1_*
	T	-1 1_*	1^* -1

The method used earlier finds no pure strategy Nash equilibria.
Player 2 wants to deviate from any matching strategy profile;
player 1 wants to deviate from any non-matching strategy profile.
No Nash equilibrium exists.

Rock–Paper–Scissors (RPS)

In this child's game, recall that rock blunts scissors,
scissors cut paper, and paper wraps rock.
The following bimatrix assumes winning gives a payoff of 1,
losing a payoff of -1, and a draw is worth 0.

			P_2	
		R	P	S
	R	0 0	-1 1^*	1^* -1
P_1	P	1^* -1	0 0	-1 1^*
	S	-1 1^*	1^* -1	0 0

Again, there is no pure strategy equilibrium.
Starting with any pure strategy pair, at least one player is not responding best, and wants to change strategy.

Mixed Strategies

Mixed or randomized strategies offer
several important advances over what we have done so far.

1. Players can make choices like
   "I'll toss a coin and choose accordingly".
   Fatherly advice: "Never gamble."
   "If you have a hard decision to make,
   toss a coin and see if you're disappointed."
2. More importantly, players can have probabilistic beliefs.
   These are subjective or personal probabilities
   such that the player acts as if
   these were specified objective probabilities
   by maximizing subjective expected utility.

Example

Let E be the event that it is snowing at noon today.
Consider an indifference map describing preferences between the two commodities:

1. the probability p of winning £1000 if E occurs;
2. the probability q of winning £1000 if E does not occur.

Let $v(w)$ be the von Neumann–Morgenstern utility from winning an amount £w.
Let π be the subjective or personal probability that E occurs.
Then the obvious expected utility function is

$$\mathbb{E}v = Pv(1000)+(1-P)v(0) = v(0)+[v(1000)-v(0)]P = A+BP,$$

where $P = \pi p + (1-\pi)q$ is the compound probability of winning £1000, whether or not E occurs,
and A, B are constants with $B > 0$.

Probabilistic Beliefs as Constant MRSs

So the expected utility function

$$\mathbb{E}v = Pv(1000) + (1 - P)v(0) = A + BP$$

with $A = v(0)$, $B = v(1000) - v(0) > 0$ and $P = \pi p + (1 - \pi)q$ represents the same preferences as the normalized utility function

$$U(p, q) = P = \pi p + (1 - \pi)q$$

This is characterized by the parameter $\pi \in (0, 1)$.

The corresponding indifference curves in (p, q) space
are parallel straight lines whose slope is $-\pi/(1 - \pi)$,
the likelihood ratio of "rain" versus "no rain".

Generally, subjective probability ratios like this
are (constant) marginal rates of substitution (MRSs)
between two perfect substitutes — namely,
the objective probabilities of winning a prize in two different events.

Parallel Linear Indifference Curves

The utility function $U(p, q) = P = \pi p + (1 - \pi)q$ represents preferences described by parallel linear indifference curves on the unit square $[0, 1] \times [0, 1]$.

The probabilities p and q of winning £1000 in events E and not E respectively are perfect substitutes with constant MRS equal to the likelihood ratio $-\pi/(1 - \pi)$.

In the diagram $\pi > 1 - \pi$ and $-\pi/(1 - \pi) < -1$.

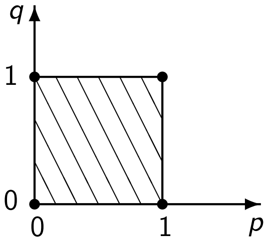

History of Ideas

This approach to subjective probability was pioneered
by Anscombe & Aumann in (1961).
The latter also won a Nobel Prize in economics
for game theory (though not for this specific work).
An excellent earlier book
is Savage's *Foundations of Statistics*.
This lays out a theory of subjective probability
without postulating objective probability.
But it is much more complicated, possibly unnecessarily so.

Finite Strategy Sets

Suppose players have finite strategy sets S_i $(i \in N)$.
A probability distribution $\sigma_i(\cdot)$ over a finite state space,
in our case player i's strategy set S_i,
is a mapping $\sigma_i : S_i \to [0,1]$ satisfying $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.

Definition

A mixed strategy for player i
is a probability distribution $\sigma_i \in \Delta S_i$.

Let $\sigma_i(s_i)$ denote the probability
that player i plays the specific strategy s_i.

We may identify each pure strategy s_i
with the degenerate probability distribution δ_{s_i}
that picks the specific pure strategy s_i with probability one.

Matching Pennies

Here each player i's strategy set is $S_i = \{H, T\}$.
Then $\sigma_i(H)$ and $\sigma_i(T)$ denote the probabilities that player i plays H and T, respectively.
And the set ΔS_i of mixed strategies is

$$\begin{aligned}\Delta S_i = \{(\sigma_i(H), \sigma_i(T)) \in \mathbb{R}^2 \\ \mid \sigma_i(H) \geq 0, \sigma_i(T) \geq 0, \\ \sigma_i(H) + \sigma_i(T) = 1\}\end{aligned}$$

i.e., the set of all pairs $(\sigma_i(H), \sigma_i(T)) \in \mathbb{R}^2$ that are non-negative, and sum to one.

We identify H and T with δ_H and δ_T respectively.
The set ΔS_i is a line interval joining δ_H and δ_T.

Rock–Paper–Scissors

In the game of rock–paper–scissors (RPS),
one has $S_i = \{R, P, S\}$, so

$$\begin{aligned}\Delta S_i = \{(\sigma_i(R), \sigma_i(P), \sigma_i(S)) \in \mathbb{R}^3 \\ \mid \sigma_i(R), \sigma_i(P), \sigma_i(S) \geq 0, \\ \sigma_i(R) + \sigma_i(P) + \sigma_i(S) = 1\}.\end{aligned}$$

Let ΔS_i denote the set of all such probability distributions.
When $\#S_i = 3$ with $S_i = \{R, P, S\}$,
then ΔS_i is a triangle with corners at $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$
in the space of triples $(\sigma(R), \sigma(P), \sigma(S)) \in \mathbb{R}^3$
representing probabilities of the three strategies in S_i.
More generally, the set ΔS_i is the unit simplex
in the $\#S_i$-dimensional Euclidean space $\mathbb{R}^{\#S_i}$.

Three Strategy Case

Suppose $S_i = \{a, b, c\}$.
Then ΔS_i is a triangle whose three corners are the degenerate mixed strategies as shown.

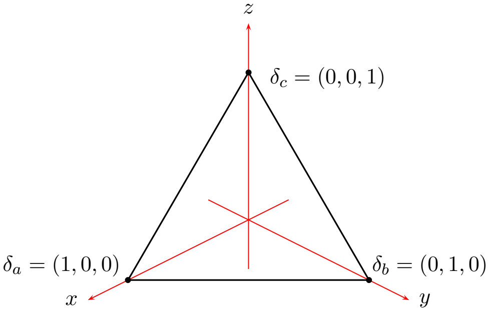

Supports

Given a mixed strategy $\sigma_i(\cdot)$, we distinguish between pure strategies that are chosen with positive probability, and those that are chosen with zero probability.

Definition

Say that a pure strategy $s_i \in S_i$ is in the support of $\sigma_i(\cdot)$ if and only if $\sigma_i(s_i) > 0$.

In the RPS game example,
suppose a player chooses R or P with equal probability,
and S with zero probability.
Then $\sigma_i(R) = \sigma_i(P) = \frac{1}{2}$ and $\sigma_i(S) = 0$.
What is the support?

Both R and P are in the support of $\sigma_i(\cdot)$, but S is not.

Infinite Strategy Sets

The Cournot and Bertrand duopoly examples
show that strategy sets need not be finite.
In case the strategy sets are subsets of the real line, a mixed strategy will be given by a cumulative distribution function:

Definition

Suppose $S_i \subseteq \mathbb{R}$.
A mixed strategy for player i
is a cumulative distribution function (or c.d.f.) $F_i : S_i \to [0, 1]$,
where $F_i(x)$ denotes the probability $\Pr\{s_i \leq x\}$ that $s_i \leq x$.

Say that $f_i(\cdot)$ is a (probability) density function for F_i
in case $F_i(x) = \int_{-\infty}^{x} f_i(s_i) ds_i$ for all $x \in \mathbb{R}$.
(Note that $f_i(x) = F_i'(x)$ wherever the latter exists.)
When F_i has such a density function,
its support is the union of all the closed intervals $[a, b] \subset \mathbb{R}$
such that $a < b$ and $f_i(x) > 0$ for all $x \in (a, b)$.

Cournot Duopoly Example

Consider the mixed strategy where firm *i* chooses a quantity between 30 and 50 with a uniform distribution. Then

$$F_i(s_i) = \begin{cases} 0 & \text{for } s_i < 30 \\ \frac{1}{20}(s_i - 30) & \text{for } s_i \in [30, 50] \\ 1 & \text{for } s_i > 50 \end{cases}$$

and we can take

$$f_i(s_i) = F_i'(s_i) = \begin{cases} 0 & \text{for } s_i < 30 \\ \text{undefined} & \text{for } s_i = 30 \\ \frac{1}{20} & \text{for } s_i \in (30, 50) \\ \text{undefined} & \text{for } s_i = 50 \\ 0 & \text{for } s_i > 50 \end{cases}$$

The support of this distribution is $[30, 50]$.

A Uniformly Mixed Strategy

These distribution and density functions are illustrated below.

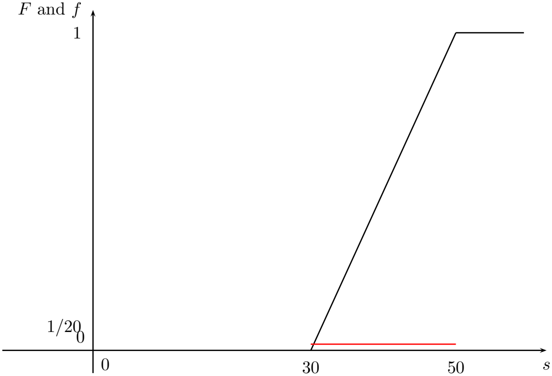

Probabilistic Beliefs

There are many reasons for player i to be uncertain
about the other players' strategies s_{-i}.
One possibility is that i believes the other players
are indeed choosing mixed strategies,
which immediately implies that s_{-i} is random.
Alternatively, player i may have incomplete information
about the other players' preferences, beliefs,
or other factors determining how they will play.

Definition

A (probabilistic) belief for player i
is a probability distribution $\pi_i \in \Delta S_{-i}$
over the joint strategies s_{-i} of i's "opponents"
— i.e., over $s_{-i} \in S_{-i} = \prod_{j \in N \setminus \{i\}} S_j$.
Thus, $\pi_i(s_{-i})$ denotes the probability
that player i assigns to the opponents playing $s_{-i} \in S_{-i}$.

Example

In the rock–paper–scissors game, player 1's beliefs
are modelled as a triple $(\pi_1(R), \pi_1(P), \pi_1(S))$
where, by definition, $\pi_1(R), \pi_1(P), \pi_1(S) \geq 0$,
and $\pi_1(R) + \pi_1(P) + \pi_1(S) = 1$.

The interpretation of $\pi_1(s_2)$ is the probability
that player 1 assigns to player 2 playing a particular $s_2 \in S_2$.

Recall that a mixed strategy for player 2
is a triple $\sigma_2(R), \sigma_2(P), \sigma_2(S) \geq 0$,
with $\sigma_2(R) + \sigma_2(P) + \sigma_2(S) = 1$.
This makes clear the analogy between beliefs π
and a mixed strategy σ;
both are members of the same space $\Delta(\{R, P, S\})$
of probability distributions over the strategy space $\{R, P, S\}$.

Expected Utility of Consequence Lotteries

In matching pennies, suppose player 2 chooses
the mixed strategy $\sigma_2(H) = \frac{1}{3}$ and $\sigma_2(T) = \frac{2}{3}$.
If player 1 plays H,
then the payoff is 1 with probability $\frac{1}{3}$,
and -1 with probability $\frac{2}{3}$.
If, however, player 1 plays T,
then the payoff is 1 with probability $\frac{2}{3}$,
and -1 with probability $\frac{1}{3}$.
So player 1's different actions
lead to different consequence lotteries.

It is usual to represent a player's preferences over such lotteries
by the expected value
of a von Neumann–Morgenstern utility function.

Expected Payoff

Definition

When the others play the mixed strategy $\sigma_{-i} \in \Delta S_{-i}$
player i's expected payoff from choosing $s_i \in S_i$
is $u_i(s_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_{-i})\, u_i(s_i, s_{-i})$.

Similarly, player i's expected payoff
from choosing the mixed strategy $\sigma_i \in \Delta S_i$
when the opponents play $\sigma_{-i} \in \Delta S_{-i}$
is $u_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i)\, u_i(s_i, \sigma_{-i})$.

This expectation of the expected payoff is the double sum

$$u_i(s_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) \left(\sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_{-i})\, u_i(s_i, s_{-i}) \right).$$

Thus, the lottery player i faces from choosing any $s_i \in S_i$
is created by the random selection of $s_{-i} \in S_{-i}$,
as specified by the probability distribution $\sigma_{-i}(\cdot)$.

Rock–Paper–Scissors Again

		P_2		
		R	P	S
P_1	R	0 0	-1 1^*	1^* -1
	P	1^* -1	0 0	-1 1^*
	S	-1 1^*	1^* -1	0 0

Suppose player 2 chooses $\sigma_2(R) = \sigma_2(P) = \frac{1}{2}$ and $\sigma_2(S) = 0$. Player 1's expected payoffs from the three different pure strategies are:

$$\begin{aligned} u_1(R, \sigma_2) &= \tfrac{1}{2} \cdot 0 + \tfrac{1}{2} \cdot (-1) + 0 \cdot 1 = -\tfrac{1}{2} \\ u_1(P, \sigma_2) &= \tfrac{1}{2} \cdot 1 + \tfrac{1}{2} \cdot 0 + 0 \cdot (-1) = \tfrac{1}{2} \\ u_1(S, \sigma_2) &= \tfrac{1}{2} \cdot (-1) + \tfrac{1}{2} \cdot 1 + 0 \cdot 0 = 0. \end{aligned}$$

Given these beliefs, player 1's unique best response is P.

Extension to Interval Strategy Sets

Suppose each player $i \in N$ has a strategy set
given by the interval $S_i = [\underline{s}_i, \bar{s}_i]$.

Suppose, say, player 1 is playing s_1,
and the other players $j = 2, 3, \ldots, n$ are each using
the mixed strategy given by the density function $f_j(\cdot)$.

Let $f_{-1}(s_{-1})$ denote the product $f_2(s_2)f_3(s_3)\cdots f_n(s_n)$.

Then player 1's expected utility is given by the multiple integral

$$\int_{\underline{s}_2}^{\bar{s}_2} \int_{\underline{s}_3}^{\bar{s}_3} \cdots \int_{\underline{s}_n}^{\bar{s}_n} u_1(s_1, s_{-1}) f_{-1}(s_{-1}) ds_2 ds_3 \cdots ds_n.$$

Example: Bidding for a Dollar

Suppose two players bid for a dollar.
Each can submit a bid that is any real number
(so we are not restricted to penny increments).
Thus $S_i = [0, \infty)$ for $i \in \{1, 2\}$.
As usual, the higher bidder gets the dollar,
but both bidders must pay their bids.
(This is called an all pay auction.)
In a tie both pay, and the dollar is awarded to each player
with a probability of $\frac{1}{2}$.
Thus, if player i bids s_i and the other player $j \neq i$ bids s_j,
then player i's expected payoff is

$$u_i(s_i, s_{-i}) = \begin{cases} -s_i & \text{if } s_i < s_j; \\ \frac{1}{2} - s_i & \text{if } s_i = s_j; \\ 1 - s_i & \text{if } s_i > s_j. \end{cases}$$

Bidding for a Dollar: Expected Utility

Suppose player 2 has a mixed strategy σ_2,
which is a uniform distribution over the interval $[0, 1]$.
So the cumulative distribution function is

$$F_2(s_2) = \begin{cases} s_2 & \text{for } s_2 \in [0, 1]; \\ 1 & \text{for } s_2 > 1. \end{cases}$$

Player 1's expected payoff from bidding $s_i > 1$ is $1 - s_i < 0$
since the bid will win for sure, but this would not be wise.

Because $\Pr\{s_1 = s_2\} = 0$,
the expected payoff from bidding $s_i < 1$ is

$$\mathbb{E}u_1(s_1, \sigma_2) = \Pr\{s_1 < s_2\}(-s_1) + \Pr\{s_1 > s_2\}(1 - s_1)$$

which reduces to $[1 - F(s_1)](-s_1) + F(s_1)(1 - s_1)$
or to $(1 - s_1)(-s_1) + s_1(1 - s_1) = 0$.

Mixed Strategy Nash Equilibrium

We restate the definition of Nash equilibrium for a game
in which each player i's pure strategy set S_i
is replaced with the mixed strategy set ΔS_i.

Definition

The mixed strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \ldots, \sigma_n^*)$ in
$\Delta S_1 \times \Delta S_2 \times \ldots \times \Delta S_n$ is a Nash equilibrium
if every player $i \in N$ is choosing
a best response σ_i^* in ΔS_i to $\sigma_{-i}^* \in \Delta S_{-i}$
— that is, $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$ for all $\sigma_i \in \Delta S_i$.

We can think of σ_{-i}^* as player i's belief π_i
about the opponents' pure strategies $s_{-i} \in S_{-i}$.
Rationality requires each player i
to respond best given the probabilistic beliefs π_i.
A Nash equilibrium requires all players' beliefs to be correct.

Which Pure Strategies are Possible?

Recall that pure strategy $s_i \in S_i$
is in the support of the mixed strategy σ_i if $\sigma_i(s_i) > 0$
— that is, if s_i is played with positive probability.

Suppose σ^* a Nash equilibrium profile.

Of course, σ_i^* must be a best response to σ_{-i}^*.

Assume more than one of player i's pure strategies,
say s_i' and s_i'', are both in the support of σ_i^*.
That is, suppose both $\sigma_i^*(s_i') > 0$ and $\sigma_i^*(s_i'') > 0$.

What can we then conclude about $BR_i(\sigma_{-i}^*)$,
player i's best response set to σ_{-i}^*?

Purified Best Responses

Proposition

If σ^* *is a Nash equilibrium,*
and s_i *is any pure strategy in the support of* σ_i^**,*
then $s_i \in BR_i(\sigma_{-i}^*)$*.*
Indeed, the same is true whenever σ_i^*
is a mixed strategy best response to σ_{-i}^**.*

This describes a procedure for purifying a mixed strategy,
to replace it with a pure strategy that is an equally good response.

Implications

This simple observation helps find mixed strategy Nash equilibria.

In particular, a player with a mixed strategy in equilibrium
must be indifferent between all the actions
being chosen with positive probability — that is, the actions
that are in the support of his mixed strategy.
Requiring one player to be indifferent
between several different pure strategies
imposes restrictions on other players' behaviour,
which helps find the mixed strategy Nash equilibrium.

For games with many players,
or with two players that have many strategies,
finding the set of all mixed strategy Nash equilibria is tedious,
often left to computer algorithms.
But we find all the mixed strategy Nash equilibria
for some simple games.

Matching Pennies

		P_2	
		H	T
P_1	H	1^* -1	-1 1_*
	T	-1 1_*	1^* -1

Recall that matching pennies
has no pure strategy Nash equilibrium.
Does it have any mixed strategy Nash equilibria?
Let p denote the probability that player 1 plays H,
so $1 - p$ is the probability that 1 plays T.
Similarly, let q be the probability that player 2 plays H,
so $1 - q$ is the probability that 2 plays T.

Player 1's Best Responses

With this notation, player 1's expected utilities
from the two pure actions H and T are:

$$\begin{aligned} u_1(H, q) &= q \cdot 1 + (1-q) \cdot (-1) = 2q - 1; \\ u_1(T, q) &= q \cdot (-1) + (1-q) \cdot 1 = 1 - 2q. \end{aligned}$$

Now H will be strictly better than T for player 1
if and only if $u_1(H, q) > u_1(T, q)$.
This is true iff $2q - 1 > 1 - 2q$, or iff $q > \frac{1}{2}$.
Similarly, T will be strictly better than H for player 1 iff $q < \frac{1}{2}$.
Finally, if $q = \frac{1}{2}$, player 1 is indifferent between H and T.
Hence player 1's best response correspondence is

$$BR_1(q) = \begin{cases} p = 0 & \text{if } q < \frac{1}{2}; \\ p \in [0, 1] & \text{if } q = \frac{1}{2}; \\ p = 1 & \text{if } q > \frac{1}{2}. \end{cases}$$

Player 2's Best Responses

Similarly, player 2's expected utilities
from the two pure actions H and T are

$$\begin{aligned} u_2(p, H) &= p \cdot (-1) + (1-p) \cdot 1 = 1 - 2p; \\ u_2(p, T) &= p \cdot 1 + (1-p) \cdot (-1) = 2p - 1. \end{aligned}$$

This implies that player 2's best responses are

$$BR_2(q) = \begin{cases} q = 1 & \text{if } p < \frac{1}{2}; \\ q \in [0, 1] & \text{if } p = \frac{1}{2}; \\ q = 0 & \text{if } p > \frac{1}{2}. \end{cases}$$

The Nash Equilibrium

Now, when player 1's best response mixes H and T,
both with positive probability,
then the expected payoffs from H and T must be equal.
This imposes a restriction on player 2's behaviour,
represented by q.
Namely, player 1 is willing to mix H and T
if and only if $u_1(H, q) = u_1(T, q)$, which holds if and only if $q = \frac{1}{2}$.

Similarly, player 2 is willing to mix H and T
if and only if $u_2(p, H) = u_2(p, T)$, which is true only when $p = \frac{1}{2}$.

So there is a Nash equilibrium pair of mixed strategies,
namely $(p, q) = (\frac{1}{2}, \frac{1}{2})$.

Finally, since there is no pure strategy Nash equilibrium,
and each player is only willing to mix strategies
when the other chooses H or T with equal probability,
this is the unique Nash equilibrium.

Graphs of Best Responses

Player 1's best response correspondence is in red.
Player 2's best response correspondence is in blue.
The only point of intersection is the unique Nash equilibrium, where $p = q = \frac{1}{2}$.

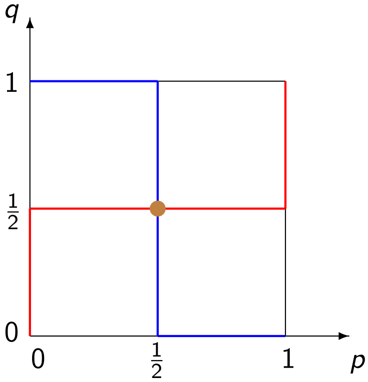

Rock–Paper–Scissors, I

Recall the rock–paper–scissors game
that was introduced in Lecture Notes 17.

Player 2

		R	P	S
Player 1	R	0 0	-1 1^*	1^* -1
	P	1^* -1	0 0	-1 1^*
	S	-1 1^*	1^* -1	0 0

Rock–Paper–Scissors, II

In two-person games like this
with more than 2 strategies for each player,
solving for mixed strategy equilibria
is much less straightforward than in 2×2 games.
There will usually be several equations in several unknowns.
To find the Nash equilibrium of the Rock–Paper–Scissors game,
we proceed in three steps:

1. first, we show that there is no Nash equilibrium
   in which at least one player plays a pure strategy;
2. then we show that there is no Nash equilibrium
   in which at least one player mixes only two pure strategies;
3. we find the solution,
   in which both players mix all three pure strategies,
   and show it is unique.

Step 1

Suppose *i* plays a pure strategy.
The payoff matrix implies that player *j* receives three different payoffs from *j*'s three pure strategies whenever *i* plays a pure strategy.

Therefore, player *j* has a unique best response, which must be a pure strategy.
Hence, *j* cannot be playing a mixed strategy if *i* plays a pure strategy.
Similarly, *i* cannot be playing a mixed strategy if *j* plays a pure strategy.

We conclude that there are no Nash equilibria where either player plays a pure strategy.

Step 2

Next we consider what happens if one player mixes
only two of the three strategies in $\{R, P, S\}$.
Suppose $\sigma_i(S) = 0$, for example.
Then P is better than R for player j,
so $\sigma_j(R) = 0$ when j responds best.

But then S is better than P for player i,
so $\sigma_i(P) = 0$ when i responds best.

Since $\sigma_i(R) = 1$ was ruled out in step 1,
there can be no Nash equilibrium with $\sigma_i(S) = 0$.

Similar reasoning shows that no Nash equilibrium
can have $\sigma_i(R) = 0$ or $\sigma_i(P) = 0$ either.
And a similar argument shows that there is no equilibrium
where the other player j mixes only two strategies.

Step 3: A Symmetric Equilibrium

By now you may have guessed
that the symmetric mixed strategies $\sigma_1^* = \sigma_2^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
form a Nash equilibrium.

Indeed, if i plays σ_i^*,
then j has an expected payoff of 0 from every pure strategy.

So $BR_j(\sigma_i^*)$ consists of the whole of ΔS_j,
whenever $i \neq j$ and $i, j \in \{1, 2\}$.

We conclude that (σ_1^*, σ_2^*) is a Nash equilibrium.

Unique Equilibrium

Note that player 1's payoffs from pure strategies are:

$$\begin{aligned} u_1(R, \sigma_2) &= -\sigma_2(P) + \sigma_2(S); \\ u_1(P, \sigma_2) &= \sigma_2(R) - \sigma_2(S); \\ u_1(S, \sigma_2) &= -\sigma_2(R) + \sigma_2(P). \end{aligned}$$

Since any Nash equilibrium
involves completely mixed strategies with full support,
all these three expected payoffs must be equal.
Let $\bar{u}$ denote their common value.
Adding the left-hand sides of these three equations gives $3\bar{u}$.
Adding their right-hand sides gives 0, as everything cancels.
Hence $\bar{u} = 0$, so the RHS of each equation equals 0.
It follows that the three probabilities $\sigma_2(R)$, $\sigma_2(S)$
and $\sigma_2(P)$ are all equal, so all are equal to $\frac{1}{3}$.
Moreover, there is no other Nash equilibrium.

Multiple Nash Equilibria, Both Pure and Mixed

Matching Pennies and Rock–Paper–Scissors
both have unique Nash equilibria in mixed strategies.
But mixed strategy equilibria need not be unique.
In fact, a game with multiple pure strategy Nash equilibria
will often have mixed strategy equilibria as well.

		P_2	
		L	R
P_1	U	0 0	3^* 5_*
	D	4^* 4_*	0 3

Here (U, R) and (D, L) are two pure strategy Nash equilibria.
It turns out that in cases like this,
with two distinct pure strategy Nash equilibria,
there will generally be a third equilibrium in mixed strategies.

Both Pure and Mixed Nash Equilibria

		P_2	
		L	R
P_1	U	0, 0	3^*, 5_*
	D	4^*, 4_*	0, 3

Let p, q denote $\sigma_1(U)$ and $\sigma_2(L)$ respectively.
Player 1 will mix when $u_1(U, \sigma_2) = u_1(D, \sigma_2)$,
or when $q \cdot 0 + (1 - q) \cdot 3 = q \cdot 4 + (1 - q) \cdot 0$,
which implies that $q = \frac{3}{7}$.
Player 2 will mix when $u_2(\sigma_1, L) = u_2(\sigma_1, R)$,
or when $p \cdot 0 + (1 - p) \cdot 4 = p \cdot 5 + (1 - p) \cdot 3$,
which implies that $p = \frac{1}{6}$.
This yields our third Nash equilibrium:
$(\sigma_1^*, \sigma_2^*) = \left((\frac{1}{6}, \frac{5}{6}), (\frac{3}{7}, \frac{4}{7})\right)$.

Finding All Nash Equilibria

All the equilibria emerge from carefully analysing best responses. Given the expected utility functions $u_1(U, \sigma_2) = 3(1 - q)$ and $u_1(D, \sigma_2) = 4q$, we have

$$BR_1(\sigma_2) = \begin{cases} p = 1 & \text{if } q < \frac{3}{7}; \\ p \in [0, 1] & \text{if } q = \frac{3}{7}; \\ p = 0 & \text{if } q > \frac{3}{7}. \end{cases}$$

Similarly, using the two expected utility functions $u_2(\sigma_1, L) = 4(1 - p)$ and $u_2(\sigma_1, R) = 5p + 3(1 - p)$,

$$BR_2(\sigma_1) = \begin{cases} q = 1 & \text{if } p < \frac{1}{6}; \\ q \in [0, 1] & \text{if } p = \frac{1}{6}; \\ q = 0 & \text{if } p > \frac{1}{6}. \end{cases}$$

Graphing these correspondences in the following diagram reveals all three Nash equilibria: $(p, q) \in \{(1, 0), (\frac{1}{6}, \frac{3}{7}), (0, 1)\}$.

Graphs of Best Responses

Player 1's best response correspondence is in red.

Player 2's best response correspondence is in blue.

The three points of intersection are the unique Nash equilibria at $(p,q) \in \{(1,0), (\frac{1}{6}, \frac{3}{7}), (0,1)\}$.

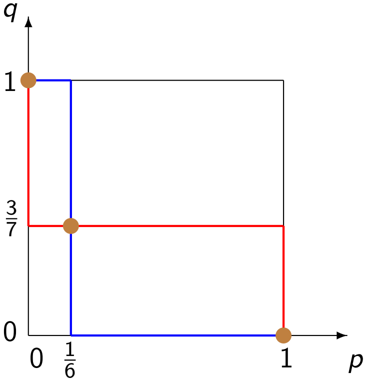

Nash's Existence Theorem

Theorem

Any finite game has a Nash equilibrium in mixed strategies.

COMMENT: In any finite game, each player's expected utility function is a continuous function of the probabilities used to describe a mixed strategy profile.

Proof.

Not trivial. After all, Nash's proof helped win a Nobel Prize! □

Implications

Nash's theorem revolutionized Game Theory.

It allowed us to say that we have some sort of predicted behaviour for *any* finite game.

The proof and the method of analysis are also very similar in spirit to the methods used in general equilibrium (fixed point theorems), ushering in a new era for microeconomics in which game theory became extensively used.

Note that as with GE theory this also hints at a nice feature: the number of equilibria will in general be odd (note that infinity can be classified as odd for this purpose!).

Uniqueness?

Despite the revolution ushered in by Nash's Theorem, the problem of multiple equilibria remained.

Nash's Theorem guarantees at least one equilibrium but not exactly one equilibrium.

The age of "refinement" which followed included a long search for a concept that could generate uniqueness, but to date nothing has yet been found that can produce uniqueness in a general setting.

Stability?

Also related to the quest for uniqueness is the issue of stability which also resulted in a wide variety of refinements.

Some Nash equilibria may be more stable in some sense than others (compare again with general equilibrium theory), and a huge literature sprang up around this idea in an attempt to identify what we can say about stability.

This included the birth of evolutionary game theory and links with biology and genetics.

But that really is beyond the scope of this course...