Mixed Strategy Nash Equilibria (MSNE)

September 19, 2016

This is a long read - you can also refer to the textbook, or if you already kind of know what you're doing you can skip to the TL;DR section for the steps.

Intro

The basic idea with MSNE is that rather than settle on a single action as our equilibrium strategy, our strategy is a randomization device that tells us which action to play

- For example, think of a penalty kick in soccer/football. The player kicking the ball has to decide to kick the ball left or right, and simultaneously the goalkeeper has to decide to jump left or jump right. No single action can be an equilibrium (if the kicking player chooses to always kick left, then the goalkeeper will always jump left, which is bad for the kicking player). So the kicking player might decide to flip a coin secretly before he kicks, and kick left if the coin comes up heads, and kick right otherwise. Flipping a coin means that the kicking player is randomizing over kicking left or right with probability .5. How do we find strategies like this one? Is .5 the right probability to use? Read on.

Expected Value

First we need to know how to assign payoffs to random outcomes. In discussion group, we defined the payoff function as a function mapping outcomes (which are non-random) to a number. There's some subtlety about this, but I'll just skip to the punchline:

To get the payoff from a random event, we use a **probability-weighted sum of the payoffs from each possible outcome**.

What does this mean? Suppose tomorrow it will either be sunny or rainy. Suppose further that my payoff when it's raining is 1 (ie: $\pi(\text{rain}) = 1$), and my payoff when it's sunny is 2 (ie: $\pi(\text{sun}) = 2$). Finally, suppose that the probability that it will rain tomorrow is .75 (because we're in Vancouver), and so the probability it's sunny is .25. Then my expected payoff from tomorrow's weather is:

$$\begin{aligned} \mathrm{E}\left[\pi(\text{tomorrow})\right] &= .25\pi(\text{sun}) + .75\pi(\text{rain}) \\ &= .25(1) + .75(2) \\ &= 1.75 \end{aligned}$$

The E[] is the expectations operator - if you haven't already seen it, take a good look, you'll be seeing a lot of it.

Finding MSNE: Example

Consider the following game we talked about briefly in discussion:

PI\P2	L	R
U	1,-1	-1,1
D	-1,1	1,-1

Verify that there is no Nash Equilibrium in pure strategies. You can think of this game as exactly the soccer example above - if you replace U with L and D with R then P1 would be the goalkeeper and P2 would be the penalty kicker.

Now suppose that P2 thinks that P1 will play U with probability p (ie: some unknown number between 0 and 1), and D with probability $1 - p$. Given this strategy of P1, what is P2's payoff of playing L? Notice that P2 is treating P1's action as a random variable, and assigning a probability to each possible outcome - so we use expected value to figure out P2's payoff from playing L:

$$\mathrm{E}\pi_2(L) = p\pi_2(LU) + (1-p)\pi_2(LD)$$

Where $\pi_2(\cdot)$ is P2's payoff function. Remember that π is a function over outcomes (eg: LU) not actions (eg: L). So if we fill in the payoff numbers from the table of outcomes above we get

$$\begin{aligned} \mathrm{E}\pi_2(L) &= p(-1) + (1-p)(1) \\ &= -p + 1 - p \\ &= 1 - 2p \end{aligned}$$

Let's do the same thing to calculate the expected payoff for P2 playing R:

$$\begin{aligned} \mathrm{E}\pi_2(R) &= p(1) + (1-p)(-1) \\ &= 2p - 1 \end{aligned}$$

Now I'll go into more detail next week, but basically the trick here is that we find the Nash Equilibrium randomization probability by setting these two payoffs equal, ie: set

$$\mathrm{E}\pi_2(L) = \mathrm{E}\pi_2(R)$$

Now substitute in with the expressions we found above and solve the equation for p:

$$\begin{aligned}
\mathrm{E}\pi_2(L) &= \mathrm{E}\pi_2(R) \\
1 - 2p &= 2p - 1 \\
4p &= 2 \\
p &= \frac{1}{2}
\end{aligned}$$

What is this special p that we found? We imposed $\mathrm{E}\pi_2(L) = \mathrm{E}\pi_2(R)$, which means that we found the p that makes P2 have the same expected utility from playing left or right. Remember that p is the probability that **P1** plays U, so our result can be stated as:

> When P1 plays U with probability .5, P2 is indifferent between playing L or R

Now let's do this whole exercise again but for P1, supposing that P1 thinks that P2 will play L with probability q. Do the math and verify that you get $q = .5$. So we have the result:

> When P2 plays L with probability .5, P1 is indifferent between playing U or D.

OK now here's the tricky part conceptually: Suppose P1 is playing U with probability .5. Then P2 is indifferent between L or R, so *fix P2's strategy of playing L with probability .5*. Since P2 is indifferent, this strategy is a best response. Similarly, when P2 is playing L with probability .5, P1 is indifferent between U or D, so playing U with probability .5 is a best response. Therefore, the solution is:

> The MSNE is P1 plays U wp .5 and D wp .5, and P2 plays L wp .5 and R wp .5

where wp stands for 'with probability'.

TL;DR: Steps for finding MSNE

1. Consider P1. Assume that P2 plays one action with probability p and the other action with probability $1 - p$
2. Calculate the expected payoff for P1 playing their first action, given the randomization by P2 in (1).
3. Do (2) but for P1's other action
4. Set these two expected payoffs equal and solve for p
5. Consider P2. Assume P1 plays one action with probability q and the other action with probability $1 - q$
6. Same as (2), but now for P2's first action, given the randomization by P1 in (5)
7. Same as (3), but for P2's other action.
8. Same as (4), but solve for q
9. You're done! The MSNE is when P1 randomizes with probability q and P2 randomizes with probability p