Intermediate Microeconomics — Week 14

Professor Boyd **November 22, 2022**

Final Date and Time. According to the latest update from Space & Scheduling, our final will be on **Thursday, December 8** at **Noon** here in our regular classroom.

12.1.5 The Coordination Game II: Nash Equilibria **Repeated**

The resulting Nash equilibria are the same, but now the players care which equilibrium they get. In the original version, both players had a payoff of 20 in each equilibrium. Here player one gets 20 in the (L, L) equilibria, but only 10 in the (R, R) equilibria.

The situation is reversed for player two. This creates a conflict if they attempt to decide on an equilibrium ahead of time. Player one prefers (L, L), player two prefers (R, R).

How can they resolve this conflict?

Coordination Game II: Nash Equilibrium

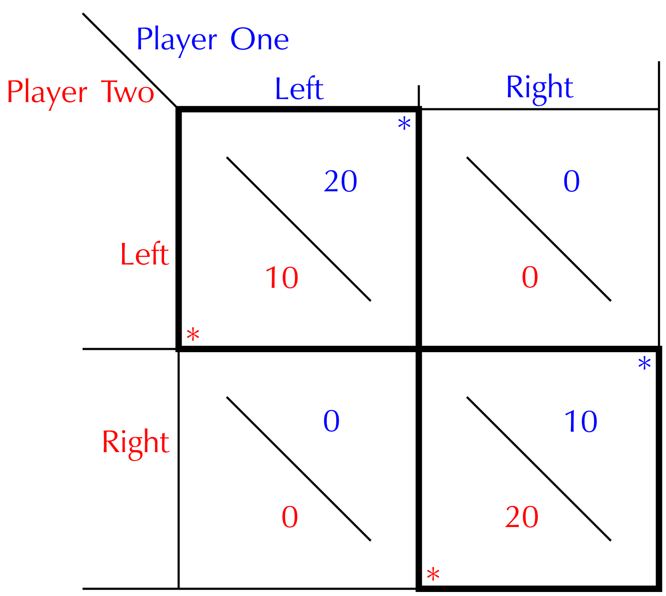

12.1.6 The Coordination Game II: Side Payments

New Slide

During discussion, one of you suggested that one of the players offer the other \$5 to agree to the first player's preferred outcome. In game theory, this is usually called a **side payment**.

This would work by having player one offer player two \$5 in exchange for agreeing to choose Left. We didn't specify when the payment is made.

Suppose the payment is made prior to the game being played. The payment terminates discussion and both players move immediately after. We include the side payment in the payoffs, which changes the game matrix. The payoffs now look like this.

Coordination Game II: Side Payments

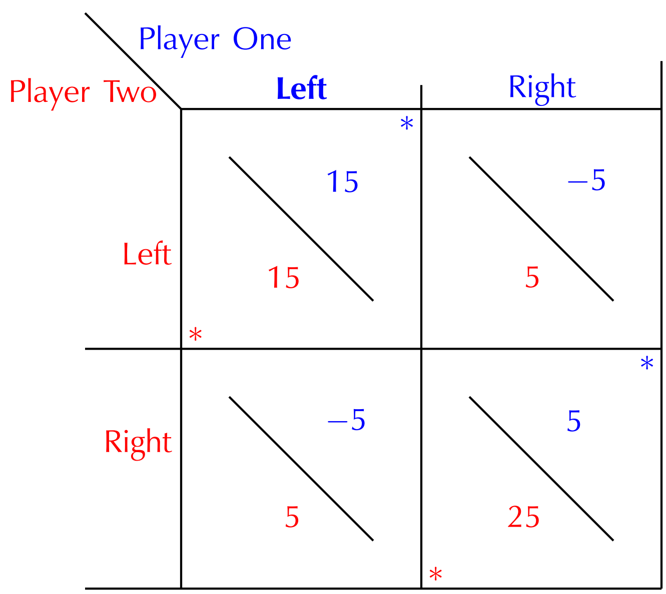

There are still two Nash equilibria, (L, L) and (R, R). However, player one expects player two to choose Left due to their agreement and immediately chooses the best response, Left, once they have an agreement. In turn, player two expects that player one will choose Left, and has best response Left as a result. The agreement is **self-enforcing**. The outcome of the new game following the agreement is that both players choose left and get \$15 in the new version of the game.

12.1.7 The Coordination Game II: Randomization **Expanded**

There's another way to get to a fair outcome without changing the game matrix.

Another way people actually deal with such conflicts is to randomize in a fair way. For example, player one could flip a fair coin (i.e., that flips to "heads" half the time, and "tails" half the time). Player two calls it.

If player two's call is correct, she chooses the equilibrium they coordinate on. She prefers (R, R), where her payoff is \$20 rather than \$10, and chooses that. If coin is tails, player one chooses. He prefers (L, L) where his payoff is \$20 rather than \$10, and chooses that.

In this augmented version of the game, there is a 50% chance the players will play (L, L), and a 50% chance they pick (R, R).

Then player one has a 50% chance of getting 20 and a 50% chance of getting payoff 10. Player two is in the exact same situation, with a 50% chance of getting 20 and a 50% chance of getting 10. The situations are exactly the same. The fact that they have the same chance of each payoff makes it fair.

In the literature on the economics of taxation, there are two basic concepts of fairness: horizontal equity and vertical equity. In the tax context, **horizontal equity** means that those in the same situation pay the same tax. **Vertical equity** means that those in different situations may pay different taxes. The concept is usually applied to the tax unit, which may be a family.

The randomization here is an example of horizontal equity.

12.1.8 Evaluating Random Payoffs: Expected Payoff

The fact that the random payoffs received by players one and two could be described exactly the same way conforms with our notions of fairness.

But what if no such description is possible? That would happen if player faced payoffs of 30 and 0 while player two continued to receive either 20 or 10.

One way to approach this is to consider the **expected payoff**. We compute the expected payoff by multiplying the probability of each outcome by the payoff. Then we add to obtain the expected payoff. In probability terms, we have calculated the **expectation** of the random payoff.

In our revised example, the expected payoff for player one is

$$0.50 \times 30 + 0.50 \times 0 = 15$$

(here $0.50 = 50\%$). The expected payoff for player two is

$$0.50 \times 20 + 0.50 \times 10 = 15.$$

Here the expected payoffs are the same, creating some degree of fairness.

12.1.9 Expected Utility

Another method is to use the **utility values** of the payoffs rather than monetary values. This yields expected utility. It is a useful way of discussing risks and how consumers and producers respond to them. We could just assume the payoffs are in utility terms.

This doesn't completely solve the problem. We usually think of utility as being unobservable and not comparable between consumers. If the payoffs were in utility terms, player one's payoffs would be in player one's utility, and player two's payoffs in player two's utility. Then a comparison of the expected utilities would be comparing player one's utility with player two's utility, which is generally invalid.

The problem is that the same payoffs that give player one utility levels of 10 and 20 may give player two utility payoffs of 5 and 20. So utility payoffs of 10 and 20 for both would represent different actual payoffs. Then expected utilities would be the same, but putting player one in player two's shoes would lead to a change in expected utility, not the same expected utility.

If the consumers agree that the expected payoff is what matters, we have no problem with the comparison. Otherwise, there may still be conflict even though the expected utility is the same in the two cases.

12.2 Rock, Paper, Scissors

One well-known game is "Rock, Paper, Scissors". The rules determining the winner are well-known, **Paper** covers **Rock**, **Scissors** cut **Paper**, **Rock** breaks **Scissors**. The players get a $+1$ payoff for winning, -1 for losing, and zero for ties. We express this as a game matrix.

Rock, Paper, Scissors

Two \ One	Rock	Paper	Scissors
Rock	0 \ 0	−1 \ +1	+1 \ −1
Paper	+1 \ −1	0 \ 0	−1 \ +1
Scissors	−1 \ +1	+1 \ −1	0 \ 0

12.2.1 Rock, Paper, Scissors: Best Responses

Now that we've marked the best responses, we can see that there is no Nash equilibrium. There are six boxes with a single star, indicating they are best responses for one player or the other. No box has two starts. There are no choices that are mutual best responses.

Rock, Paper, Scissors

Two \ One	Rock	Paper	Scissors
Rock	One: 0; Two: 0	One: +1 *; Two: −1	One: −1; Two: +1 *
Paper	One: −1; Two: +1 *	One: 0; Two: 0	One: +1 *; Two: −1
Scissors	One: +1 *; Two: −1	One: −1; Two: +1 *	One: 0; Two: 0

12.2.2 Nash's Solution: Randomization

Nash's solution to this problem is to change what we mean by a strategy to include random strategies. The strategies (= moves here) that are initially defined are called **pure strategies**. If you pick one, you use it with certainty—100% probability. However, he suggested assigning probabilities to the pure strategies to form **random** or **mixed strategies**.

12.2.3 Probability Distributions

The probabilities must obey two rules.

1. Each pure strategies is assigned a probability from zero (0%) to one (100%).
2. When added over all possible pure strategies, the probabilities must sum to one (100%).

Such probabilities describe a **probability distribution**.

Rule two means that some strategy will be chosen. By rule one, it is possible that some strategy may never be chosen, may have a probability of 0

12.2.4 Mixed Strategies as Probability Distributions

We will often write random strategies by listing the probabilities of each pure strategy in the order given. Here, the pure strategies are **(Rock, Paper, Scissors)**, so $(.5, .2, .3)$ means the random strategy that plays **Rock** 50% of the time, **Paper** 20% of the time and **Scissors** 30% of the time.

The pure strategies can also be written this way. The strategy $(1, 0, 0)$ is **Rock**, $(0, 1, 0)$ is **Paper**, and $(0, 0, 1)$ is **Scissors**.

12.2.5 The Language of Probability

When talking about probability, we use the term **event** to refer to a something that can possibly occur. Nash suggested we regard the possible pure strategies—**Rock**, **Paper**, and **Scissors**—as events. Our probability distribution describes the probability of each event, the probability that each pure strategy is chosen.

We compare strategies, both mixed and pure, by computing their expected payoffs. That introduces a complication. Both sides might be using mixed strategies. How do we compute the probabilities of strategy pairs such as $(\mathrm{Rock}, \mathrm{Paper})$ or $(\mathrm{Scissors}, \mathrm{Rock})$ if all we know is that player one is playing $(0.2, 0.3, 0.5)$ while player two plays $(0.6, 0.2, 0.2)$?

If the player make their moves simultaneously, what player one chooses doesn't affect what player two does, and vice-versa. In probability terms, the probability of the event that player one plays **Rock** does not depend on the probability that player two plays **Scissors** (or anything else).

Such events are called **independent**.

12.2.6 Conditional Probability

One important concept in probability is **conditional probability**. We write the probability of event A, if event B occurs as $P(A|B)$ this is the probability of A conditional on B. The formula for conditional probability is

$$P(A|B) = \frac{P(A \cap B)}{P(B)}.$$

This is illustrated in the diagram. Events A (red) and B (blue) can both occur together. That is, their intersection is non-empty. The conditional probability is the chance that A occurs if B occurs. The probability that both occur is $P(A \cap B)$, and the conditional probability is the fraction of B's probability $P(B)$ where both occur. I.e., $P(A \cap B)/P(B)$.

Conditional Probability

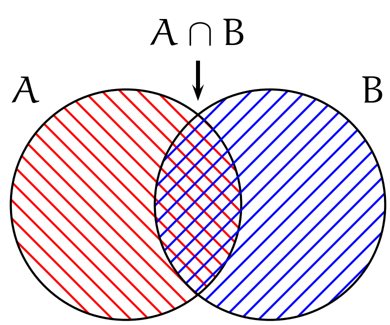

12.2.7 Independent Events

It is common to think of events as being sets. Then when A and B are events, then $A \cap B$ is the event where both A and B occur. When the probabilities $P(\cdot)$ obey

$$P(A \cap B) = P(A) \cdot P(B)$$

for all events A and B, the events are called **independent**.

Now if A and B are independent,

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{P(A) \cdot P(B)}{P(B)} = P(A).$$

In that case, as far as the probability of A is concerned, it irrelevant whether or not B occurs.

12.2.8 Mixed Strategy vs. Mixed Strategy

We're almost ready to compute probabilities when both sides use mixed strategies. We regard the randomization by player one as independent of player two's randomization and vice-versa. Then for each pair of pure strategies, independence means you multiply the probabilities of each to find the probability of the combination.

When player one is playing $(0.2, 0.3, 0.5)$ while player two plays $(0.6, 0.2, 0.2)$ we can now compute the probability of the outcome $(\text{Rock}, \text{Paper})$. It is $0.2 \times 0.2 = 0.04$. Similarly, the probability of $(\text{Scissors}, \text{Rock})$ is $0.5 \times 0.6 = 0.30$.

We can then calculate expected payoffs by multiplying the product of the probabilities together with the payoffs for any pair of strategies. Then we add up across all the strategy pairs.

12.2.9 Calculating Probabilities

For example, suppose player one picks $(.6, .3, .1)$ and player two uses $(.7, .1, .2)$.

Player one picks **Rock** with probability .6. The combination (R, R) has probability $.6 \times .7 = .42$, the combination (R, P) has probability $.6 \times .1 = .06$, the combination (R, S) has probability $.6 \times .2 = .12$. Note that the combined probability $.42 + .06 + .12 = .6$, the probability that player one chooses **Rock**.

Similarly, the probability that a choice of **Paper** by player one combines with player two's three pure strategies adds up to .3, and the probability that a choice of **Scissors** by player one combines with player two's three pure strategies adds up to .1.

12.2.10 The Probabilities Add Up

Suppose player one chooses the mixed strategy (p_1, p_2, p_3) and player two chooses (q_1, q_2, q_3). Let's add the probabilities of all possible events.

Rock, Paper, Scissors: Probabilities

Two \ One	Rock, p_1	Paper, p_2	Scissors, p_3
Rock, q_1	p_1q_1	p_2q_1	p_3q_1
Paper, q_2	p_1q_2	p_2q_2	p_3q_2
Scissors, q_3	p_1q_3	p_2q_3	p_3q_3

The sums are

$$\begin{aligned}\text{Total Prob} &= p_1q_1 + p_1q_2 + p_1q_3 \\ &\quad + p_2q_1 + p_2q_2 + p_2q_3 \\ &\quad + p_3q_1 + p_3q_2 + p_3q_3 \\ &= p_1(q_1 + q_2 + q_3) + p_2(q_1 + q_2 + q_3) \\ &\quad + p_3(q_1 + q_2 + q_3) \\ &= p_1 + p_2 + p_3 \\ &= 1.\end{aligned}$$

When each player's probabilities are independent, adding up all the probabilities gives us 1, as a probability distribution should.

12.2.11 Rock, Paper, Scissors: Nash Equilibrium I

Once we allow random strategies. The only Nash equilibrium is that both sides play each of the possible moves with probability 1/3.

Suppose player one does this. What is player two's best response?

If player two always chooses Rock, 1/3 of the time player one will also choose Rock, with payoff zero. One third of the time player one chooses scissors and loses, with payoff +1 to player two. The last third of the time player one chooses paper and wins, so player two get −1. Player two's expected payoff is

$$(1/3) \times 0 + (1/3) \times (-1) + (1/3) \times (+1) = 0.$$

A little thought reveals that if player two always chooses scissors, or always chooses paper, the expected payoff would still be zero. Indeed, any mixed strategy will always give expected payoff zero!

So it seems that everything is a best response—and it is. However, we will show that the only **mutual** best responses have probabilities $(1/3, 1/3, 1/3)$ vs. $(1/3, 1/3, 1/3)$ for (R, P, S).

12.2.12 Rock, Paper, Scissors: Nash Equilibrium II

If player one uses $(.4, .3, .3)$, i.e., **Rock** 40% of the time, and **Scissors** and **Paper** 30% each, the best response is not $(1/3, 1/3, 1/3)$.

Since player one picks **Rock** more than any other option, you can do better by picking to beat **Rock**, by always picking **Paper**. Then you win 40% of the time, are tied 30% of the time, and lose 30% of the time for an expected payoff of

$$0.4 \times +1 + 0.3 \times (-1) + 0.3 \times (0) = +0.1.$$

In contrast, with $(1/3, 1/3, 1/3)$ you get expected payoff zero the 40% of the time your opponent actually plays **Rock**, payoff zero the 30% of the time your opponent plays **Scissors**, and payoff zero the 30% of the time your opponent plays **Paper**. This adds to zero, and is worse than always choosing **Paper**.

It turns out that if your opponent chooses probabilities that are not equal, the best response is to always use the best response to the most likely option—**Paper** to **Rock**, **Scissors** to **Paper**, or **Rock** to **Scissors**. As a result, only the probabilities $(1/3, 1/3, 1/3)$ are **mutual** best responses. This is the only Nash equilibrium for Rock, Paper, Scissors.

12.3 Escape! A Game without an Equilibrium

There are simpler games without equilibria. One is the Escape! game.

In Escape!, there are two players, a guard and a prisoner. The prisoner breaks out of his cell and runs out on the street. He has a choice, run left or run right. Whichever way the prisoner goes, he turns a corner before the guard runs out on the street.

When the guard runs out after the prisoner, he has to decide whether to go left or right. If the guard follows the prisoner, he will eventually catch him. The guard will be rewarded and the prisoner punished.

If the prisoner escapes, he will enjoy freedom for a while, and get a positive payoff. The guard is not the person responsible for the prisoner's escape, and will not be penalized if he does not catch the prisoner.

Here's the game matrix.

Escape!

Guard \ Prisoner	Left	Right
Left	Prisoner: −100; Guard: +5	Prisoner: +50; Guard: 0
Right	Prisoner: +50; Guard: 0	Prisoner: −100; Guard: +5

12.3.1 No Pure Strategy Equilibrium in Escape!

The guard wants to catch the prisoner, and must go the same way to do this. That means the guard's best response to **Left** is **Left**, and his best response to **Right** is **Right**.

The prisoner wants to get away, which means going the opposite way as the guard. He must respond to the strategy **Left** with **Right**, and to the strategy **Right** with **Left**.

Each box contains a single star, so there is no pure strategy Nash equilibrium.

Escape! Best Responses

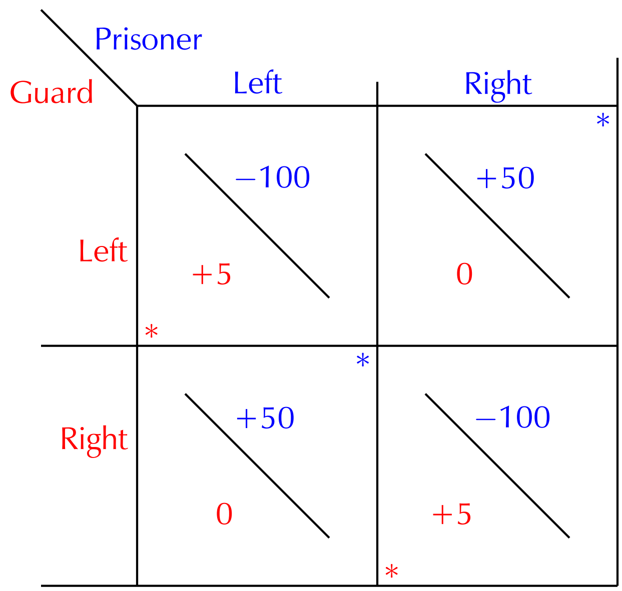

12.3.2 Nash Equilibrium for Escape!

Nash (1950, 1951) showed that any game with a finite number of players and a finite number of pure strategies must have an equilibrium, either in pure or mixed strategies.[1]

The Nash equilibrium is for both players to use the strategy $(.5, .5)$, go left with probability 0.5, and go right with probability 0.5. As before, we write the probabilities in the same order as the pure strategies in the game matrix.

[1] See John Nash (1950), "Equilibrium points in n-person games", *Proceedings of the National Academy of Sciences*, **36**:48–49, and John Nash (1951), "Non-Cooperative Games", *Annals of Mathematics*. **54**:286-295.

12.3.3 Mixed Strategies Can Make the other Player Indifferent

As is typical in mixed-strategy equilibria, the mixed strategies make the other player indifferent between at least two pure strategies.

If the guard uses $(.5, .5)$ and the prisoner chooses **Left**, the prisoner is caught half of the time and goes free half of the time. The prisoner's payoff is

$$.5 \times (-100) + .5 \times (+50) = -25.$$

If the guard uses $(.5, .5)$ and the prisoner chooses **Right**, the prisoner is still caught half of the time and still goes free half of the time. It's just the other half. The prisoner's payoff is

$$.5 \times (+50) + .5 \times (-100) = -25.$$

Either way, the prisoner's expected payoff is the same.

The guard is in a similar situation, but with an expected payoff of 2.5. Try calculating it.

12.3.4 Indifference Makes Randomization Possible

It is only when players are indifferent between using two strategies that they are willing to randomize. They will get the same payoffs by randomizing between them that they receive from either pure strategy.

Escape! has only two pure strategies for each player. By playing $(.5, .5)$, the guard has made the prisoner indifferent between his pure strategies, and any randomization involving them. In particular, $(.5, .5)$ is a best response by the prisoner.

When the prisoner chooses $(.5, .5)$, the guard is placed in a similar situation, and $(.5, .5)$ is a best response. We have mutual best responses—Nash equilibrium.

It is only the $(.5, .5)$ randomization that does this in Escape! If the prisoner chose $(.6, .4)$, the guard's best response would be **Left**, with a payoff of $.6(5) = 3$. Choosing **Right** only yields $.4(5) = 2$, and even $(.5, .5)$ yields only 2.5.

12.3.5 Combining Probabilities

Suppose the prisoner plays the mixed strategy $(p, 1 - p)$ and the guard responds with $(q, 1-q)$. Of course, $0 \leq p, q \leq 1$. What is the probability of each outcome?

To find that, we can use a probability tree. It shows that the prisoner goes left with probability p and right with probability $(1 - p)$. Depending on which the prisoner has chosen, the guard then goes left or right with probabilities q or $(1 - q)$.

When we get to the end, we multiply the probabilities along the branches leading there to find the probabilities of each outcome.

This procedure guarantees that the numbers at the ends are probabilities—that they are between 0 and 1 and sum to 1.

12.3.6 Probability Tree

The tree shows the various payoffs to the prisoner and guard (prisoner first) and the probabilities.

Escape! Probability Tree

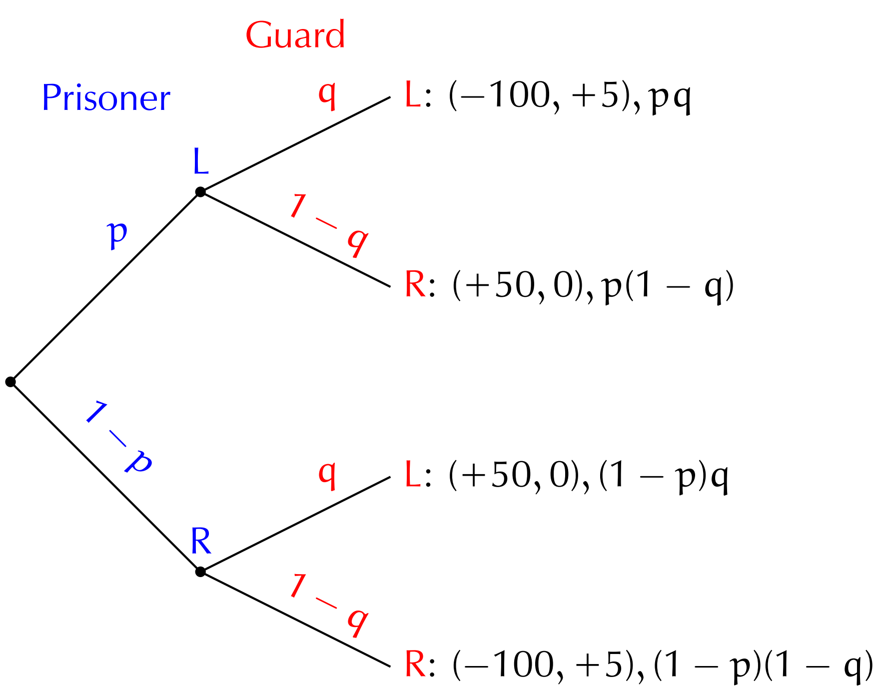

If the prisoner chooses $p = 0.5$, the expected payoff for the guard is

$$2.5q + 2.5(1 - q) = 2.5.$$

Any choice by the guard is a best response.

12.3.7 What if the Prisoner Favors Left?

Now suppose the prisoner has a preference for **Left**, and choose it with probability $p > 0.5$. Referring back the probability tree, we can calculate the guard's expected payoff:

$$\begin{aligned}
&5pq + 0p(1-q) + 0(1-p)q + 5(1-p)(1-q) \\
&\quad = 5pq + 5(1-p)(1-q) \\
&\quad = 5pq + 5 - 5p - 5q + 5pq \\
&\quad = 5 - 5p + 5(2p-1)q.
\end{aligned}$$

When $p > 0.5$, $2p-1 > 0$, so the guard's expected payoff is increasing in q. It's best for the guard to choose the largest possible q, which is 1. The guard's best response is to always go **Left**.

Similarly, if the prisoner has a bias toward the right, the guard should always pick **Right** to make the probability of catching the prisoner as high as possible.

The guard is in a similar situation. Unless the guard plays $(0.5, 0.5)$, the prisoner will always go the opposite way as the guard's bias.

12.3.8 The Other Form of Escape!

There are actually two forms of Escape! In the form we've seen, the guard is too slow getting out on the street to see where the prisoner went. The guard lacks information about what the prisoner did. If the guard got out a bit quicker, he would see the prisoner run around the corner and could follow, eventually catching the prisoner. Rather than show this using a game matrix, we will write the game in **extensive form**, using a **game tree**. It's like the probability tree, but showing only the moves, the payoffs, and the information. The **nodes** indicate points where a move must be made and/or a payoff paid.

Full Information Escape

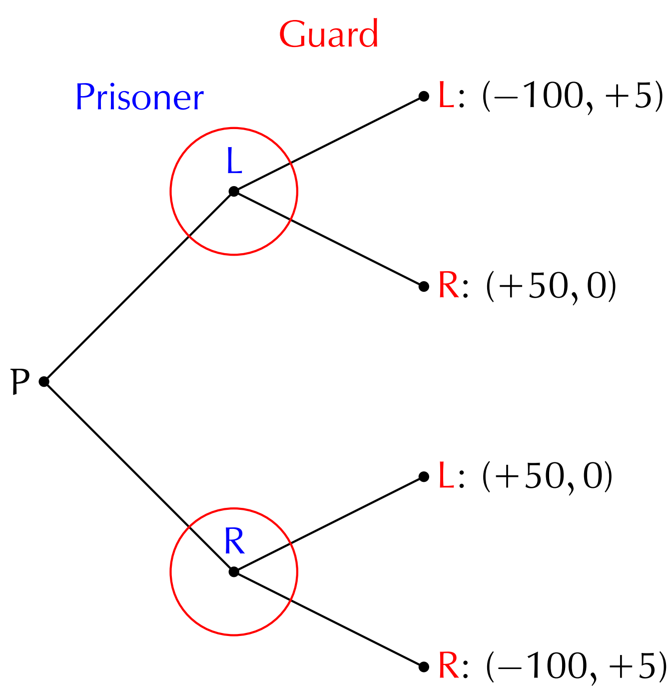

12.3.9 Extensive Form

The **root node**, marked P, belongs to the prisoner, who chooses either L or R. The two subsequent nodes, which are circled, belong to the guard, who also chooses either L or R. The circles are **information sets**. They indicate that the guard can tell which node has occurred. I.e., he saw the prisoner run around the corner. If the guard could not tell the difference, both nodes would be in the same information set, instead of in two different information sets. Payoffs are listed in the order the player move, (P, G), at the ends of each branch, the **terminal nodes**.

When each node is in its own information set, we call the game a **full information game** or **perfect information game**.

Full Information Escape

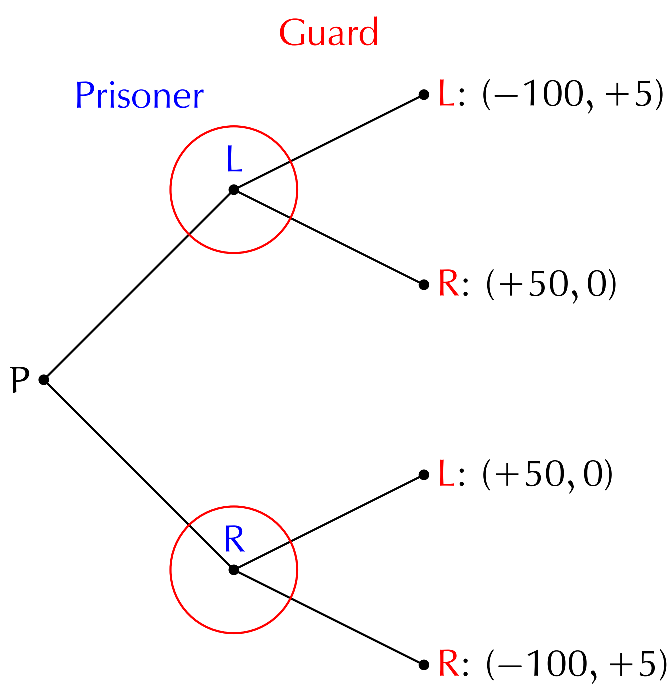

12.3.10 The Original Escape

In the original version of Escape, the guard couldn't tell whether the prisoner went left or right. We indicate that by putting them both in the same information set.

Because the guard has no information about which way the prisoner went, he must make the same move at both of his nodes. He cannot condition his move on what the prisoner did. If he chooses L in response to the prisoner's L, he must also choose L in response to the prisoner's R.

Here the prisoner also lacks information about what the guard will do. Neither side has any information about the other's moves. We don't need to draw an information set for the prisoner because he only has one node.

No Information Escape

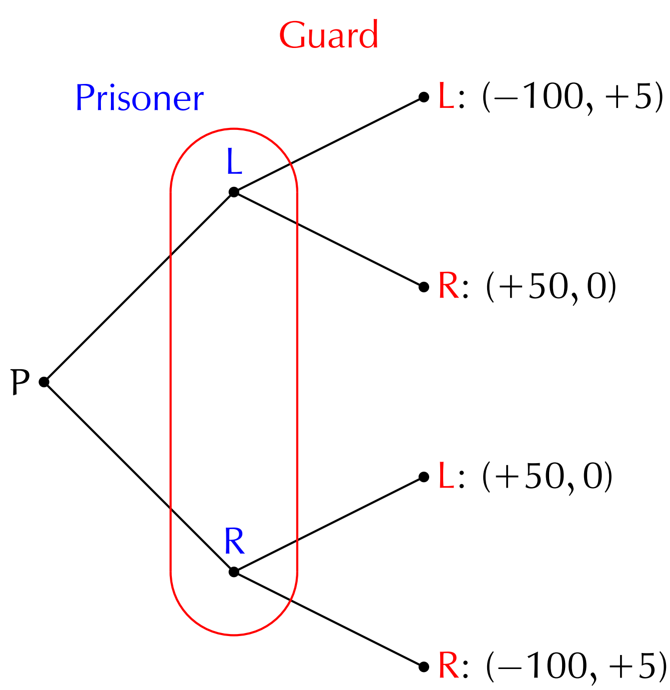

12.3.11 Nash Equilibrium in Full Information Escape

One way to think about the full information game is that the guard now has four possible pure strategies. He can always (1) go left or (2) go right. He can (3) follow the prisoner or (4) avoid the prisoner.

Following the prisoner is a dominant strategy, so the guard uses it. It then doesn't matter what the prisoner does. He can always go left, always go right, or randomize. No matter what the prisoner does, he will be caught.

Full Information Escape: Equilibrium

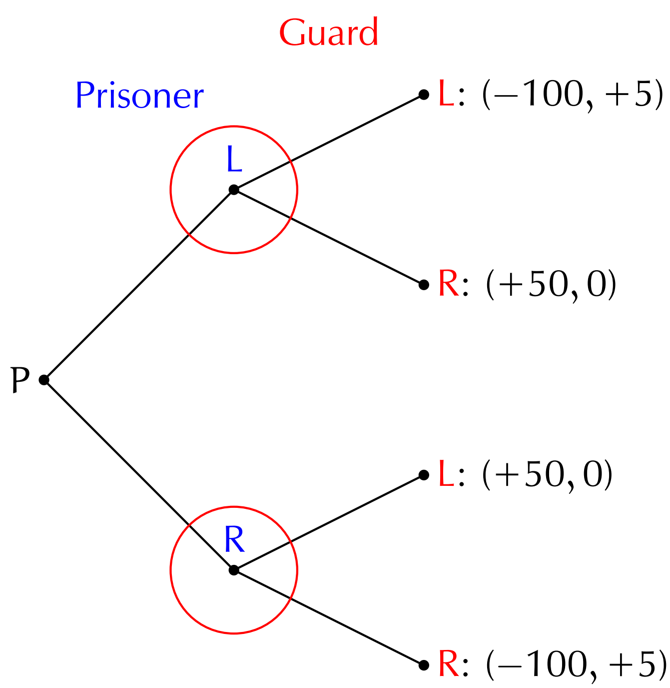

12.3.12 Backwards Induction

So how do we find Nash equilibria in a extensive form game? We listen to Sherlock Holmes! According to Holmes, "In solving a problem of this sort, the grand thing is to be able to reason backwards."[2]

For extensive form games, one method of finding Nash equilibria is to use **backwards induction**. We start by considering the payoffs at the terminal nodes.

On the next page, we note with a star (*) which moves at the subterminal nodes yield maximum payoff(s) for the player making the last move. We then replace the subterminal nodes by terminal nodes with those payoffs. Repeat this process until we reach the root node.

Let's try it with the full information version of Escape

Full Information Escape: Beginning Backwards Induction

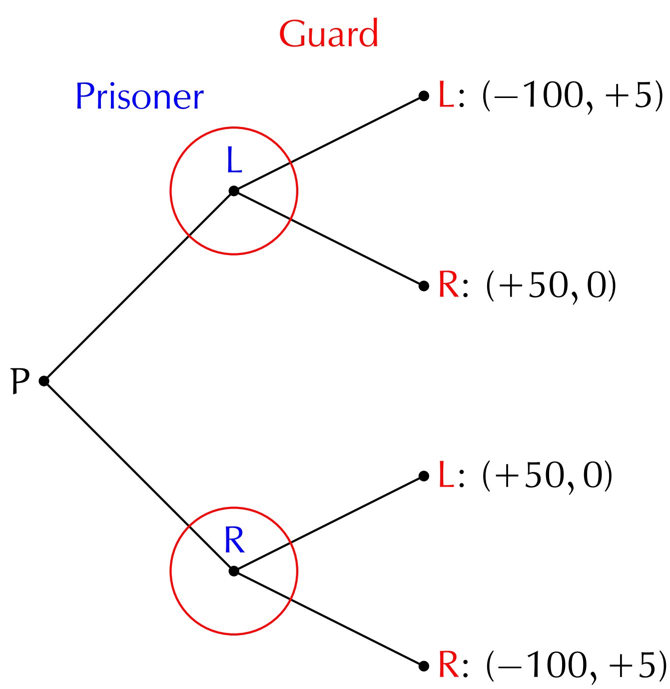

[2] Sherlock Holmes in *A Study in Scarlet* (Sir Arthur Conan Doyle).

12.3.13 The Best Subterminal Moves

We've marked the best responses at each subterminal node with a star (*).

We next replace the subterminal nodes by terminal nodes with those payoffs. We then repeat this process until we reach the root node.

Full Information Escape: Best Subterminal Moves

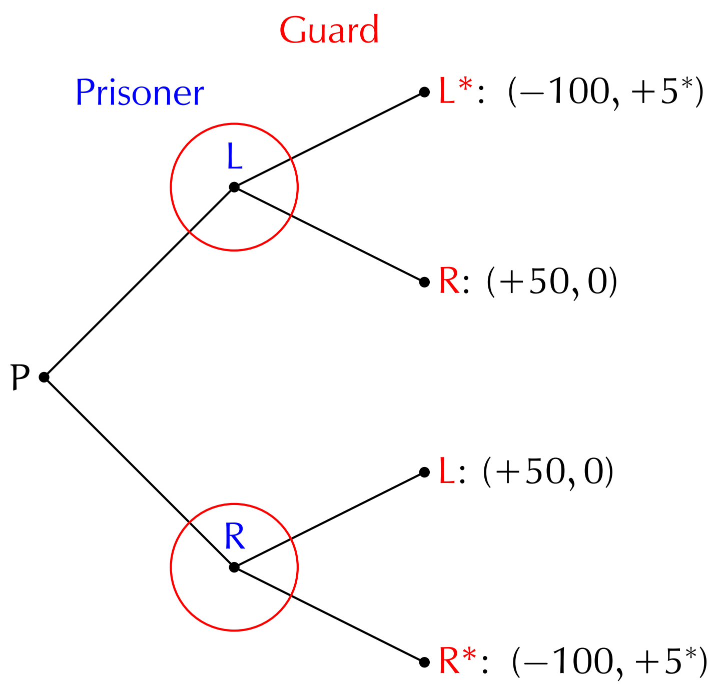

12.3.14 Reducing the Tree

We have reduced the tree. I've indicated the guard's moves that lead to the new terminal payoffs. At this point, the prisoner has a choice between a payoff of -100 and of -100.

It's not much of a choice. Any move by the prisoner, even a mixed strategy, will yield -100 as the maximum payoff.

The Nash equilibria that backwards induction gives us have the prisoner doing anything, and the guard following—replying to L with L and to R with R.

Reduced Escape Tree

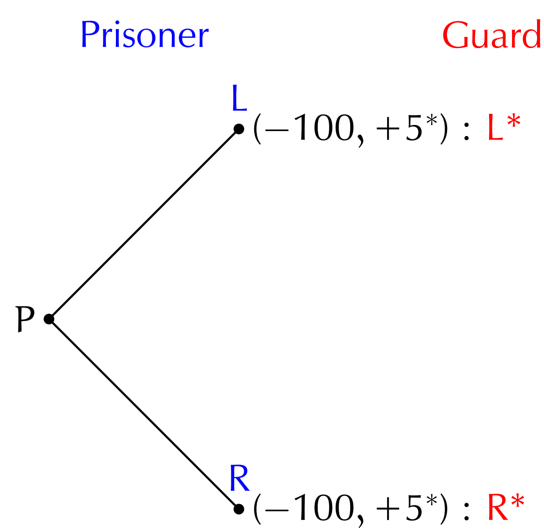

12.3.15 Limitations of Backwards Induction

Backwards induction has its strengths and weaknesses. As we have defined it, it only applies to perfect information games.

One strength is that if you take any subtree, and treat that as a game in its own right (a **subgame**), the resulting equilibrium is also an equilibrium in every subgame. Such an equilibrium is called **subgame perfect**.

A corresponding weakness is that backwards induction may not find all of the Nash equilibria.

12.3.16 Full Information Escape: Matrix Form

Let's look at the matrix form to see if we've found all the Nash equilibria in the full information version of Escape!

Full Information Escape!: Matrix Form

Guard \ Prisoner	Left	Right
Left	Guard: +5 *, Prisoner: −100	Guard: 0, Prisoner: +50 *
Right	Guard: 0, Prisoner: +50 *	Guard: +5 *, Prisoner: −100
Follow	**Guard: +5 *, Prisoner: −100 ***	**Guard: +5 *, Prisoner: −100 ***
Avoid	Guard: 0, Prisoner: +50 *	Guard: 0, Prisoner: +50 *

12.3.17 Nash Equilibria in Full Information Escape

There are two pure strategy equilibria: (L, F) and (R, F). As we've already observed, F in response to any mixed strategy by the prisoner is also an equilibrium. There are no other mixed strategy equilibrium because **Follow** is a strictly dominant strategy for the guard.

Later, we will examine a game with full information where one of the equilibria is not subgame perfect and cannot be found by backwards induction.

November 22, 2022