



Introducing Mixed-Strategy Equilibria: Payoffs vs. Probabilities

Students are usually taught to compute mixed-strategy equilibria in 2×2 games by determining what mixed strategy each player must play in order for the other player to be indifferent between their own pure strategies, evaluated using expected-utility reasoning. An alternative method, complementing the standard one, draws on each player's computations of certain disparities between the payoffs he or she may receive, and uses that information to compute the probabilities for the players' mixed strategies.

Ben Eggleston

University of Kansas

I would like to thank Clint Hurshman for valuable insights that contributed crucially to the ideas developed in this paper, and Donald Bruckner, Dale Miller, and two anonymous reviewers for helpful comments.

1. The Expected-Utility Method

In courses on game theory, the transition from pure-strategy equilibria to mixed-strategy equilibria is a crucial advance in students' understanding of the reach and importance of the equilibrium concept. The standard tool for making this advance is the 2×2 matrix in which no outcome is an equilibrium outcome, but there is nonetheless a mixed-strategy equilibrium. One such matrix is the following:

		q	$1 - q$
		C1	C2
p	R1	4, 6	8, 5
$1 - p$	R2	9, 3	2, 7

Perhaps the most common method for deriving the mixed strategies that are in equilibrium is to ascertain what mixed strategy each player must play in order for the other player to be indifferent between their own pure strategies, evaluated using expected-utility reasoning. For example, we can compute the row player's mixed strategy by setting the expected utilities of the *column* player's two pure strategies equal to each other, as follows:

$$\begin{aligned} EU(C1) &= EU(C2) \\ p \times 6 + (1 - p) \times 3 &= p \times 5 + (1 - p) \times 7 \\ p &= 4/5 \\ \text{strategy} &= (4/5 \text{ R1}, 1/5 \text{ R2}) \end{aligned}$$

Similarly, we can compute the column player's mixed strategy as follows:

$$\begin{aligned} EU(R1) &= EU(R2) \\ q \times 4 + (1 - q) \times 8 &= q \times 9 + (1 - q) \times 2 \\ q &= 6/11 \\ \text{strategy} &= (6/11 \text{ C1}, 5/11 \text{ C2}) \end{aligned}$$

So, the equilibrium strategy pair is [(4/5 R1, 1/5 R2), (6/11 C1, 5/11 C2)].

This is found in Mas-Colell, Whinston, and Green (1995), pp. 250–251; Heap and Varoufakis (2004), p. 71; and Spaniel (2012). Of these three sources, the first two are prominent textbooks and the third is a video in a series titled "Game Theory 101." Tellingly, this specific video is entitled "*The Mixed Strategy Algorithm*" (emphasis added).

2. The Payoffs vs. Probabilities Method

The above method is essential and I would not suggest abandoning it. I propose, however, that it can profitably be paired with another method. This method also imagines each player asking, "What do I have to do to make the other player indifferent between his or her options?" but explicates the psychology of the other player in a new way. (Below, I will use male pronouns for the row player and female pronouns for the column player.)

This new method takes its start from what I would characterize as a moderately sophisticated way that players might approach a game such as the one above. A couple of simple rules that players might employ are "choose the strategy that gives me a chance of getting my

highest possible payoff” and “choose the strategy that avoids any risk of my getting my lowest possible payoff.” In contrast to these simple rules, this approach is responsive to the relative magnitudes of the *disparities* between particular high and low payoffs – not just the maximum or minimum considered singly. For the row player, this approach involves ascertaining whether there is one column in which his choice of strategy (R1 or R2) will make a bigger difference to the payoff he will receive than it does in the other column. To articulate this in detail, we can imagine the row player in the above game reasoning as follows:

The payoff I will receive depends not only on the row I choose, but also on the column that the other player chooses. In column 1, my possible payoffs are 4 and 9; and in column 2, my possible payoffs are 8 and 2. The disparity between 8 and 2 is a little larger than the disparity between 4 and 9, so how I choose matters a little more if the column player chooses column 2. Accordingly, I should weight the possibility of her choosing column 2 a little more heavily than the possibility of her choosing column 1. In column 2, row 1 is better for me than row 2. Therefore, I should lean toward choosing row 1.

The row player is essentially asking, “Where does my choice of strategy matter more – column 1 or column 2?” and giving more weight to that column accordingly. Similarly, we can imagine the column player reasoning as follows:

The payoff I will receive depends not only on the column I choose, but also on the row that the other player chooses. In row 1, my possible payoffs are 6 and 5; and in row 2, my possible payoffs are 3 and 7. The disparity between 3 and 7 is four times as large as the disparity between 6 and 5, so how I choose matters a lot more if the row player chooses row 2. Accordingly, I should weight the possibility of his choosing row 2 four times as heavily as the possibility of his choosing row 1. Therefore, I should lean heavily toward choosing column 2.

This way of reasoning, although more sophisticated than the simple rules quoted above, is still fairly intuitive. However, there is one aspect of this reasoning that warrants additional clarification before we continue. One might wonder exactly why each player might weight the possibility of the other player’s choosing a certain option more heavily than the other possibility. For example, one might wonder why the column player might think that she should weight the possibility of the row player’s choosing row 2 more heavily than the possibility of his choosing row 1. After all, the fact that the row player’s choice of row 2 leads to possible payoffs for the column player with a greater disparity than if he chooses row 1 does not mean that he is more likely to choose row 2. (The row player is not, for example, trying to make the column player be in a row of the game in which there is a high disparity between her possible payoffs.) So, why might the column player be inclined to weight the possibility of the row player’s choosing row 2 more heavily than the possibility of his choosing row 1?

The inclination imagined for the column player does not depend on the idea that the greater disparity of her payoffs in row 2 makes it more likely that the row player will choose that row. To see this, suppose row 1 and row 2 correspond not to a choice made by another player, but to two different lottery tickets that the column player could be given, with the payoff of each ticket depending on which column she chooses, as in the original matrix. And suppose the column player does not have any beliefs about the probability of receiving ticket 1 versus ticket 2. In such a situation, it would be very natural for the column player to think, “If I receive ticket 1, it doesn’t matter very much whether I choose column 1 or column 2. However, if I receive ticket 2, it matters quite a lot whether I choose column 1 or column 2. Therefore, I should weight the possibility of receiving ticket 2 more heavily than the possibility of receiving ticket 1.”

This is the rationale for why each player might weight the possibility of the other player's choosing a certain option more heavily than the other possibility. It is intuitive to imagine each player reasoning in this way. More important for our purposes, it is plausible to imagine each player attributing such reasoning to the other player in order to figure out how to make the other player indifferent between their two options. I turn to this next.

Imagine the row player asking, "What do I have to do to make the column player indifferent between her options?" As we saw in section 1, one way the row player can cash out this question is to rephrase it as "What do I have to do to equalize the expected utilities of the column player's pure strategies?" Alternatively, however, the row player can reason as follows:

The column player knows that I will play R1 or R2 (whether as a pure strategy or as the result of playing a mixed strategy). She sees that if I play R1, then she will get either 6 or 5, and if I play R2, she will get either 3 or 7. The disparity between 3 and 7 is four times as large as the disparity between 6 and 5. Therefore, rather than weighting R1 and R2 equally in her choice between C1 and C2, she will weight R2 four times as heavily as R1. In order to counteract her weighting of R2 more heavily in her decision-making, I need to do something to make R1 weightier in her thinking. The only thing I can do to accomplish this is to increase my probability of playing R1, and decrease my probability of playing R2. How should I quantify this? Well, I just found that the column player is inclined to weight R2 four times as heavily as R1. To counteract this, I should make it four times as likely that I will play R1 as R2. The probabilities that accomplish this are $4/5$ and $1/5$. So, I will play the strategy ($4/5$ R1, $1/5$ R2).

Similarly (albeit somewhat more briskly), the column player can reason as follows:

The row player knows that I will play C1 or C2. He sees that the C1 disparity for him is 5 (4 vs. 9) and the C2 disparity for him is 6 (8 vs. 2). So, rather than weighting C1 and C2 equally in his choice between R1 and R2, he will be inclined to weight C2 a little more heavily, by a ratio of 6:5. To counteract this, I need to make C1 weightier in his thinking, by setting its probability a little higher than that of C2. Specifically, the probabilities need to be in the ratio of 6:5. So, I will play the strategy ($6/11$ C1, $5/11$ C2).

This form of reasoning can be called the "payoffs vs. probabilities" method, with the "vs." signifying how each player assigns probabilities to push *against* the payoff disparities that are present in the matrix.

The payoffs vs. probabilities method involves imagining that each player is trying to put the other player into a state of indifference between that other player's pure strategies. Considering this fact about this method, it might be objected that rational players seek to maximize their expected utility, not necessarily put other players into this state of indifference, or any other mental state. However, this method is only intended as a device to aid an *analyst* of a game (whether one of the two players, or a third party) in computing the equilibrium strategy pair. It simply capitalizes on the fact that imagining that the players engage in a certain form of reasoning turns out to be an effective way of computing the equilibrium strategy pair. Sometimes this information is of interest, and sometimes it is not. In any event, figuring out this information is what the payoffs vs. probabilities method is for – it is not intended to tell either player how they *should* play the game.

3. Discussion

It can of course be proved that the payoffs vs. probabilities method will always lead to the same probabilities as the expected-utility method summarized above. Perhaps more interesting, though, is how this form of reasoning is similar to and differs from the expected-utility method. An obvious similarity is that it involves each player putting the other player into a state of indifference: the row player makes the column player give equal weight to row 1 and row 2. For example, above, the column player gives more weight to row 2 to the extent that it has a greater payoff disparity – 4 vs. 1 – but gives more weight to row 1 to the extent that it has a greater probability of occurring (4/5 vs. 1/5). Because these ratios point in opposite directions but are equal in magnitude, they make the column player give equal weight to row 1 and row 2, as influences on her choice between playing C1 or C2. Likewise for the row player having been made to give equal weight to column 1 and column 2, as influences on his choice between playing R1 or R2.

Despite this similarity between the two methods in terms of each player putting the other into a state of indifference, there is an interesting contrast on this point. The expected-utility method involves imagining the row player directly calculating how to make the column player indifferent between her strategies C1 and C2, by equalizing their expected utilities. In contrast, the payoffs vs. probabilities method involves imagining the row player calculating how to make the column player find the two rows (not columns) equally worth considering in her reasoning about whether to play C1 and C2. The result is the same – indifference between playing C1 and C2 – but the path to get there is different.

This points to another contrast between the two methods: the second one does not involve imagining that either player is familiar with the concept of expected utility. Instead, it imagines the players understanding two other principles that rational players might employ when playing games like this.

The first principle is that when one player is examining the other player's two options (say, the row player examining C1 and C2), it is rational for the first player to differentially weight those options according to the disparities of the payoffs in them. For example, if the payoffs in column 1 were 100 and 0, and the payoffs in column 2 were 51 and 49, it would be rational for the row player to more heavily weight the prospect of ending up in column 1, and choose his row accordingly. Recall the analogy with lottery tickets from above – if ticket 1 had possible payoffs of 100 and 0 and ticket 2 had possible payoffs of 51 and 49, it would be rational for the row player to be more concerned with getting the 100 (vs. the 0) than with getting the 51 (vs. the 49), and choose his row accordingly.

The second principle is that when one player is examining the other player's two options, it is rational for the first player to differentially weight those options according to how likely they are to occur. For example, if C1 had a probability of 1/51 of occurring, and C2 had a probability of 50/51 of occurring, it would be rational for the row player to weight the prospect of the column player playing C2 more heavily than the prospect of her playing C1, and choose his row accordingly.

Neither of these two principles can stand alone, of course – each is only part of the story. Each correctly states one factor that is rational for players to take into account, and gives guidance as to how much. As implemented in the example above, they point in opposite directions, and with the same (mutually canceling) strength.

I mentioned above that this method does not involve the principle of maximizing expected utility. In response to this claim, it might be pointed out that just as this method always coincides with the expected-utility method, anyone who correctly employs the two principles I have just mentioned will always reach conclusions that coincide with the expected-utility method. I agree. My purpose here is not to defend the plausibility of any reasoning that conflicts with expected-utility reasoning, but simply to exhibit a different form of reasoning that is equally effective and illuminating. Paired with the expected-utility method, the payoffs vs. probabilities method can offer another perspective from which to give students insight into the subtle dynamics of indifference that lie at the heart of mixed-strategy equilibria.

4. Implementation

In this final section I discuss implementing the payoffs vs. probabilities method in classroom instruction. I have used this method successfully with advanced undergraduate students and graduate students, though as a more general rule I would suggest that it can be taught to any audience to which one would teach the expected-utility method, as the two methods seem to have comparable levels of difficulty.

Because the two methods are independent of each other, they can be taught in either order. I prefer to start with the expected-utility method, since students are more likely to have encountered it previously and are more likely to encounter it elsewhere in the future. Then I show them the payoffs vs. probabilities method as a complementary and (hopefully) intriguing alternative. I present it not as providing deeper insight into the expected-utility method, but as another manifestation of the way the concept of indifference is central to mixed-strategy equilibria. Indeed, I suspect that maintaining a distinction between the two methods, rather than looking for further links between them, is essential to avoiding confusion that could otherwise arise.

To elaborate further on the teaching of the payoffs vs. probabilities method, let me apply it to one additional example, without the discursive commentary that I offered in discussing the previous example. This example will also differ from the previous one in a specific way to include an additional feature these kinds of games might have. In the previous example, each player saw that the other player was inclined to weight one possibility that he or she might end up with (row 1 or row 2, or column 1 or column 2) more heavily than its alternative. However, this needn't be the case. I will continue with this point shortly; first, here is the example:

		q	$1 - q$
		C1	C2
p	R1	4, 3	7, 5
$1 - p$	R2	1, 8	9, 6

Let us apply the payoffs vs. probabilities method in a direct and step-by-step fashion. We will start with determining the row player's mixed strategy, though of course (as with the expected-utility method) one can start with either player.

First, the row player asks, "When the column player looks at this game, what disparity does she see in row 1?" Looking at the difference between 3 and 5, the row player finds that this disparity is 2.

Second, he asks, "When the column player looks at this game, what disparity does she see in row 2?" Looking at the difference between 8 and 6, he finds that this disparity is 2.

Because these two disparities are equal, the row player realizes that the column player is not inclined to weight the possibility of his playing row 1 more heavily than the possibility of his playing row 2, or vice versa. That is, the column player is already in a state of indifference regarding how seriously she takes the possibility that she will find herself in row 1, versus how seriously she takes the possibility that she will find herself in row 2. Unlike in the previous example, there is no imbalance that the row player needs to counteract; rather, he has the task of choosing probabilities for playing row 1 and playing row 2 that preserve this indifference. The way to do this is to pick the same probability for each row – namely, $1/2$. So, the row player's mixed strategy is $(1/2 R1, 1/2 R2)$.

Now let us determine the column player's mixed strategy. First, she asks, "When the row player looks at this game, what disparity does he see in column 1?" Looking at the difference between 4 and 1, she finds that this disparity is 3.

Second, she asks, "When the row player looks at this game, what disparity does he see in column 2?" Looking at the difference between 7 and 9, she finds that this disparity is 2.

So, she sees that the row player is inclined to weight the possibility that he will find himself in column 1 somewhat more heavily than the possibility that he will find himself in column 2. To counteract this inclination, she should be somewhat more likely to play column 2 than column 1. But "somewhat" is vague. How should she quantify that? Well, this is given by the ratio of the disparities just calculated – namely, 3:2. The probabilities that stand in this ratio are $3/5$ and $2/5$. Since the smaller probability needs to be attached to the column that the row player is already inclined to give more weight to – namely, column 1 – the column player's mixed strategy is $(2/5 C1, 3/5 C2)$.

As this review of the method illustrates, the mental effort it requires is minimal. Students might even be encouraged to employ both the expected-utility method and the payoffs vs. probabilities method in some circumstances (such as when answering test questions, if time permits) in order to double-check their work.

References

Heap, S. P. H. and Y. Varoufakis. 2004. *Game Theory: A Critical Text* (2nd ed.). London: Routledge, 2004.

Mas-Colell, A., M. D. Whinston, and J. R. Green. 1995. *Microeconomic Theory*. Oxford: Oxford University Press.

Spaniel, W. 2012. "Game Theory 101 (#8): The Mixed Strategy Algorithm." <<https://www.youtube.com/watch?v=aa8USttcDoE>>.