Zero-Sum (and Constant Sum) Games

When the reduced payoff matrix has a single strategy for each player, this gives the optimal strategy for each player, assuming best play by the opponent. We saw however, that some games such as Chicken do not reduce to a single strategy for each player. It is clear in these games that being predictable is a clear disadvantage for wither player, however beyond hiding our intentions, it is unclear how to proceed. In this section, we will study zero-sum games. We will show how to arrive at an optimal strategy for a (small) two person zero-sum game. Sometimes these games will have an optimal fixed strategy for each player and some will require a mixed strategy as an optimal strategy assuming best play by the opponent.

Definition A two person **zero-sum game** is a game where the pair of payoffs for each entry of the payoff matrix sum to 0.
This means that one player's gain is equal to the other player's loss on any given play of the game.

Definition A two person **constant-sum game** is a game where the pair of payoffs for each entry of the payoff matrix sum to the same constant C.
The analysis of these games is the same as that of zero sum games, since subtracting the given constant from the column player's payoffs makes it a zero sum game. We will see that the analysis below depends entirely on the row player's payoffs.

Example: Rock Paper Scissors In the game of Rock-scissors-paper, the players face each other and simultaneously display their hands in one of the three following shapes: a fist denoting a rock, the forefinger and middle finger extended and spread so as to suggest scissors, or a downward facing palm denoting a sheet of paper. The rock wins over the scissors since it can shatter them, the scissors wins over the paper since they can cut it, and the paper wins over the rock since it can be wrapped around it. The winner collects a penny from the opponent and no money changes hands in the case of a tie.

Roger and Colleen play a game of Rock Paper Scissors. Fill in the pay off matrix for this game below. (remember that the first entry in each pair is Roger's payoff).

		Colleen		
		Rock	**Paper**	**Scissors**
Roger	**Rock**	(,)	(,)	(,)
	Paper	(,)	(,)	(,)
	Scissors	(,)	(,)	(,)

Note that since we know that this is a zero sum game, the column player's payoff can be calculated as the negative of the row player's payoff. Hence we can present all of the game information in an abbreviated form of the matrix listing only the row player's payoff's.

		Colleen		
		Rock	**Paper**	**Scissors**
Roger	**Rock**	0	-1	1
	Paper	1	0	-1
	Scissors	-1	1	0

By Convention the payoff matrix for a two player zero-sum game, shows the strategies for both players with the payoffs for the **row player** as entries. The payoffs for the column player for each situation can be calculated by taking the negative of the row player's payoff.

Recall that an **Equilibrium Point** of a game is a pair of strategies such that neither player has any incentive to change strategies if the other player stays with their current strategy. Note that in a zero sum game this corresponds to **an entry in the matrix which is simultaneously the minimum in its row and the maximum in its column.** To find such entries, we can calculate the minimum in each row and the maximum in each column and check if any entry simultaneously gives the minimum in its row and the maximum in its column.

An equilibrium point in a zero-sum game is sometimes called a **saddle point** because it is a minimum in one direction and a maximum in the other. As with the games in the previous section, a zero sum game can have one saddle point, more than one saddle point or no saddle points. If it has more than one saddle point, all of the equilibrium points must have the same payoff. If a matrix for a zero-sum game has a saddle point, then the optimum strategy for both players (assuming best play by the opponent) is at the saddle point. If not, the optimal strategy is a mixed strategy (see next section).

Rock Paper Scissors (a) Find any equilibrium points/saddle points for Rock Paper scissors that might exist.

(b) Find the reduced payoff matrix for Rock Paper Scissors.

Example: Football Run or Pass? [Winston] In football, the offense selects a play and the defense lines up in a defensive formation. We will consider a very simple model of play selection in which the offense and defense simultaneously select their play. The offense may choose to run or to pass and the defense may choose a run or a pass defense. One can use the average yardage gained or lost in this particular League as payoffs and construct a payoff matrix for this two player zero-sum game. Lets assume that if the offense runs and the defense makes the right call, yards gained average out at a loss of 5 yards for the offense. On the other hand if offense runs and defense makes the wrong call, the average gain is 5 yards. On a pass, the right defensive call usually results in an incomplete pass averaging out to a zero yard gain for offense and the wrong defensive call leads to a 10 yard gain for offense. Set up the payoff matrix for this zero-sum game.

		Defense	
		Run Defense	**Pass Defense**
	Run		
Offense			
	Pass		

(a) Does this matrix have a saddle point?

(b) Are there any dominated strategies?

Constant Sum Games

Example (Using Probability as Payoff) In our previous example of possible endgame strategies for basketball, we set up our matrix where the payoff for each team was the probability of a win for each team under the given circumstances. This is a **Constant sum game** since the probabilities add to 1. Note that if we subtract 1 from the payoff of the defense, we get a zero sum game here. Since we have all of the information is a matrix showing the payoff for the offense, we can use that as the payoff matrix for this game. The analysis we use to find the best strategy for a constant sum game (se next section) is the same as that for a zero-sum game.

		Defending	Team
		Defend 2	**Defend 3**
	Shoot 2	0.178	0.312
Offense			
	Shoot 3	0.5	0.228

Does this game have a saddle point or a dominated strategy for either player?

Example A baseball pitcher throws three pitches, a fastball, a slider and a change-up. As a measure of the payoff for this type of confrontation, we might use the expected number of runs the batter creates in each situation. (Note their are other possible measures that take into account the pitcher's abilities). We would expect that for any given pitch, the batter's performance is better if he anticipates the pitch. Lets assume that the batter has four possible strategies, To anticipate either a fastball, a slider or a change-up or not to anticipate any pitch.

		Pitcher		
		Fastball	**Change-up**	**Slider**
Batter	**Fastball**	0.3	0.3	0.35
	Change-up	0.25	0.4	0.4
	Slider	0.2	0.39	0.45
	None	0.3	0.39	0.4

(a) Is there a saddle point in the above matrix?

(b) Find the reduced payoff matrix for this game.

Example Squash is a game played on a court similar to a racquetball court. Suppose you have the option to hit the ball on your upcoming shot so that it will end up in the front of the court or at the back of the court. Your opponent is at the center of the court, but will start to move towards the front or back of the court as you take your shot before she figures out where the ball will end up. Based on previous play, we know that if you send the ball to the back of the court, there is an 80% chance that she will win the point if she anticipates correctly and there is a 10% chance that she will win the point if she anticipates incorrectly. If you send the ball to the front of the court, there is a 70% chance that she will win the point if she anticipates correctly and a 40% chance that she will win the point otherwise. Set up the payoff matrix with you (the person taking the shot) as the column player and your opponent as the row player.

Smart Football

Analysis and strategy by Chris

Smart Football has moved!

Please check out the new site, smartfootball.com. All future updates will be made there.

Monday, July 10, 2006

Run/Pass Balance and a Little Game Theory

Football is the most strategic of all sports. A big part of this is the unique feature that each game is 100-150 or so unique trials—the plays. This gives rise to the art and science of play-calling, “balance,” formations, sets, set-ups, counters, and whatever else that keeps us up at night thinking about this stuff. Stepping back for a second though, I wanted to simply look at the concept of balance and how we should best achieve it. Before that, however, I wanted to emphasize what I think are the most important offensive statistics.

Yards Per Carry and Average Yards Per Pass Attempt

I've long felt that the most important rushing stat, at least in terms of 1st and 2nd down performance, was yards per carry, and not total rushing yards or anything else. Quite simply, a running back who gets 120 yards on 15 carries plays in a better offense than a RB who gets 120 yards on 25 or 30 carries.

This same logic, however, applies to the passing game. More important to me than passing efficiency, or completion percentage (by itself), yards per completion, or any other statistic is Average Yards Per Pass play (including sacks). Bud Goode, legendary football statistician to the stars (Dick Vermeil, Bill Parcells, etc) has been harping on this stat for years.

The key is that it, in effect, combines completion percentage and yards per completion. The NFL QBs who have had the highest totals ever in a season are as diverse as Joe Montana (exceptionally high completion percentage) to more long-ball throwers. It penalizes the guy who inflates his completion percentage and the guy who points to his long-balls while ignoring how inefficient he is.

These two stats converge in the most important first and second down stats, which are average yards per play. The goal is to move forward to the other guy's goal-line, continually increasing your chances of scoring a TD. Further, you really want to do this on first and second down: Third down is a defense's down. The odds are in the defense's favor, and so are the strategies. Also, the teams with the best third down conversation rates are invariably the ones who have the shortest average distance to go on third down--further emphasizing that positive first and second down yards are the key.

So the goal is to find the mix of runs and passes that maximizes your teams' average gain per play.

[Note: This is not entirely true, as passes carry more risk. Turnovers make up the most important stats of all in terms of winning, and pass plays result in more turnovers than do run plays--both more fumbles and more interceptions. The answer, however, is not to ignore passing, but instead to require a "passing premium"--your passes should average more per play than your runs to counterbalance this risk.]

Offensive Identity and a Taste of Game Theory

How good you are in "absolute terms" at running or passing is a matter of talent, scheme, and reps. My argument here is that for the sake of "balance" it doesn't matter what you're better at, but, as Carolina offensive coordinator Dan Henning says, you pick a target mix and go for that, while adjusting to the defense.

This adjusting to the defense is where game theory comes in. The basic idea is that your offense and their defense have certain strengths and weaknesses, and, for the most part, everyone knows each other's weaknesses.

Imagine you are fortunate enough to have a future All-American guy at RB. He runs for a ton of yards as a junior, and now, a year later, you're ready to ride him to a state title. But everyone else knows about this guy now. They begin stacking the line. You've got this All-American at running back and you're averaging less per carry than you did three years earlier when you had three Academic All-Americans--and no football All-Americans--splitting time at RB. What's going on? What do you do?

You pass of course. You run bootlegs, you fake it to him, and you throw the ball. But how odd you say. You have the best running back you've had in 15 years, and you wind up running *less*? The answer is simply that everyone else knows you have this stud RB, so they commit so much effort and defensive scheme and structure to push your expected yards per rushing play down to a manageable number, your passing opportunities increase, even if you have less talent there than years past.

This same goes for great passing teams. (Think about all the spread offense teams that have used the defense's natural tendency to play pass against four wides to their running advantage.)

This little cat and mouse game is really an extension of the Nash Equilibrium from Game Theory (the subject of the movie A Beautiful Mind, about John Nash, the concept's namesake).

Application

The idea is if you are a very good passing team you pass most of the time, then you run when it is favorable and see positive results without having had to practice it too much. Same goes vice-versa--we all know how dangerous play-action passes are

from heavy run teams, especially say a veer option team.

Again, I don't think yards per rush and yards per passing attempt should be exactly equal--passes are riskier than running plays. Specifically, they more often result in lost yardage (sacks) and turn the ball over more often (both fumbles and interceptions). So you should expect your yards per pass attempt to be higher than yards per rushing attempt.

To reiterate the earlier points of how this can be counter-intuitive, look at Urban Meyer at the Florida Gators (stats below). Let's say next year, with a year in his system, the passing game stays the same but the running game improves by a full yard per play. Now, what happens? Well, first Meyer will run the ball more--less risky, same reward. But then the defense will see this and begin to step up to stop the run, and drive the average yards per run back DOWN. Yet, the defense will be weaker to the pass. The result?

Counterintuitively, the passing game yards per attempt could go up and Meyer should then actually pass **more**. Surprised? Just think about it: If the D had to do more to stop the run, the pass gets more attractive, so Florida starts getting maybe 6.7 or 7+ per play every time they throw it, so of course they are going to throw it, even if in absolute terms it was the run game, not the pass game, that improved. Regardless, the improvement in the run game should affect the entire offense's production, which is what is important.

The lesson? If your passing game is suddenly working better, it might not be because you are suddenly Bill Walsh. It might be because you've got a stud running back everyone wants to stop.

The point of this is that you can hang your hat on one thing, but you might be leaving production on the table by not running or passing enough.

Case-Studies:

I just pulled some basic stats off of espn.com for major college teams to give some examples. I don't mean this as a criticism of these teams since my stats include some downs like 3rd down that may inflate or skew the stats, and college football includes sacks as a running play. To counteract that I added the QB's rushing numbers to the passing stats (except for Vince Young). This may be problemmatic for Chris Leak at Florida, since Urban Meyer uses a system where the QB runs the ball, but Leak was not particularly good at this and did not run near as much as Meyer's previous quarterbacks.

I fairly randomly selected these teams, though I did want to highlight teams of interest and on different ends of the spectrum.

Texas Tech

Pass-happy Mike Leach at Texas Tech attempted 697 passes for 4857 yards, averaging 6.97 yards per pass attempt. (I also recognize how many of these are shovels and the like but I'm just being simplistic.)

They ran the ball 172 times for 1040 yards, or 6.05 per rushing attempt.

So we compare 6.97 per pass to 6.05 per rush. Putting the two together the average yards per play is 6.77. We can see you can make an argument that they should have passed MORE, since that would have raised their average yards per play, but a passing premium of about a yard seems about consistent with most other teams.

The result? Tech, for all its crazy stuff, is pretty **balanced.**

Florida Gators

Next I looked at the Florida Gators. They got 2801 yards on 490 passing attemts (5.72 average) and 1680 yards rushing on 350 attempts (4.71). Together, the total yards per play was 5.33. Again we see about a yard of "passing premium" indicating that Urban was pretty balanced but that his team was not as productive, on a per play and total basis, as Tech.

Minnesota Gophers

Let's look at the Gophers: They ran it 586 times for 3247 yards (5.54), and threw it 347 times for 2690 yards (7.75).

That's the biggest passing premium we've seen, over 2.25 yards. Unless Minnesota is extremely risk averse, it appears that the Gophers should have passed more than they did. This result makes sense with what we said above: Minnesota had one of the best backs in the country, Maroney, and another guy who got 1000 yards. Their QB, some guy named Cupito, I didn't even remember. But defenses and defensive coordinators know the same thing. They were all geared to stop Maroney and the Gophers zone run game.

Should they have gone pass happy? No, of course not. Yet, imagine if they had thrown 30-50 more passes instead of runs (only 2-4 more per game). With more passing, the yards per pass attempt would have gone down, but I don't imagine it would have gone down to less than 6 yards like the rushing average. Also, yards per rush would have probably gone up as well. Thus, Minnesota likely would have been more productive to the point of 3 or more points in several games. In the Big 10, that is the difference between winning and losing.

The fact is that Minnesota's strength was definitely running the ball, but everyone else knew it too: Minnesota could have seen some easy success in the passing game and helped out their offense in total by throwing a bit more.

[Note: My numbers are rough so I'm not really trying to criticize Minnesota per se, just use them as an example.]

Southern Cal

These numbers are less helpful for the truly dominant teams (and less important, being smart about things matters less when you've got all the best talent). Nevertheless, let's look at the teams in the National Title Game.

USC threw 523 passes for 4193 yards = 7.88 yards per pass attempt
and 474 rushes for 3344 yards = 7.06 per rush attempt

This indicates that USC, no surprise, was very balanced and efficient in its playcalling. Maybe they should have run a bit more since that "passing premium" was kind of low, but USC is also a very efficient passing team and they do not turn the ball over very much, so they can have a smaller passing premium and get away with it.

However, the stat that jumped at me was 1740/200 = 8.7. As in 8.7 yards per rushing attempt, as in Reggie Bush's yards per rushing attempt. As in, handing the ball off to Reggie Bush had a greater expected gain than did throwing the football, which is just unheard of. This implies that USC should have handed it to him more. Now there are other issues, like durability, and Reggie's receiving prowess, but that is such a substantial number you will not see anything like that.

Texas Longhorns

Texas's stats were interesting too.

336 passes for 3083 yards = 9.18 yards per pass attempt
605 rushes for 3574 yards = 5.9 per rush attempt

That's a huge discrepency--that dwarfs Minnesota's number earlier. Texas' numbers may be skewed because it was on the good-end of a lot of blowouts and probably ran the ball much more in the second half. Nevertheless, coupled with the fact that Vince Young was the nation's passing efficiency leader, this implies that Texas probably held Vince's hand to much and should have let him throw more (or he should have stayed in the pocket and thrown more). Especially since as a runner Young averaged nearly 7 yards a carry, better than all but one of Texas' running backs. This exceptionally high passing yards per attempt number is probably correlated with Vince's running ability--the D had to take men out of position to spy him on pass plays.

Conclusion

To reiterate, my stats here are a bit on the simple side but the point is not the stats, it's the thinking: Typically a fan or coach looks at numbers like 9 yards per pass and 6 yards per rush and says "well, you don't run it as well you throw it." I think the right response, though, is "you ran it too much" or "you didn't throw the ball enough." That's a very different approach. It makes perfect sense though. It's recognizing that you're coaching against a smart person on the other side who knows where your strengths are, and then exploiting that to your advantage.

I remember someone asking Hal Mumme when he was at Kentucky about how his teams' yards per carry had dropped around a yard or so from the season before. The reporter was incredulous and turned red faced at Mumme's response: Mumme told him that he saw the same thing, and that to fix it he would throw the ball more. The reporter cut him off and essentially called him an idiot, mentioning that everyone knows you run better by simply running more (wear them down!). I'm pretty sure Mumme's point was that he coached a passing team, and if his yards per carry was going down, at least one reason was that the defense was spending too much time on the run and that he, as playcaller, was not taking advantage of passing game weaknesses defenses were leaving open.

Posted by Chris at 7/10/2006 11:09:00 PM

References

Mathematics beyond The Numbers: Gilbert, G, Hatcher, R.
Game Theory and Strategy, Phillip D. Straffin, The Mathematical Association of America, New Mathematical Library.
Mathletics, Wayne L. Winston, Princeton University Press.
Thinking Strategically, Avinash K. Dixit, Barry J. Nalebuff, W.W.Norton and Company.
Strategies and Games, Prajit K. Dutta, The MIT Press.
Sports Economics, R. D. Fort, Pearson.
Blog: Mind Your decisions; Game Theory applied to Basketball, Shawn Ruminski