Decision Theory in Football

By Brian Burke

In Decision Theory, there are generally two kinds of analysis. Descriptive analysis is what people actually do, and prescriptive analysis is what people should do. Rarely are the two things the same. For example, when I use the win probability model to evaluate 4th down decisions, I'm doing prescriptive analysis. Trying to explain whatever the heck coaches are actually doing would be descriptive analysis.

To be fair, coaches are not computers. They are subject to all the imperfections of human decision making. In this post, I'll examine some of the ways that coaches may be making decisions, including minimax, minimax-regret, prospect theory, and expected utility. I'll also discuss the potential for how much of a difference a pure prescriptive analysis can make when applied in real games.

NFL Orthodoxy

NFL football has evolved as extremely conservative game. By that I mean that coaches adhere to the wisdom passed down from previous generations and are reluctant to deviate from the established orthodoxy. In the real world, away from sports, this approach usually makes sense. Unlike sports, the world is not bounded by sidelines, end zones, and 15-minute quarters. It is highly uncertain and far less predictable than we'd like to think. It makes sense to adhere to what is known to work rather than try to engineer an optimized outcome in a highly uncertain environment.

But in football, we have the stats. We know the probabilities. And we know the possible consequences. 'Conservative,' as I defined it, is therefore often not the best approach. I think the reason that so many coaches adhere to the same orthodoxy, whether in terms of playbooks or 4th down doctrine, is because they aren't conscious of the level of certainty available to them.

Minimax

One of the more conservative approaches is the minimax criterion. Minimax says pick the option that assures you the highest minimum utility. Let's say you have the choice between going on a picnic and going bowling. You'd really rather go on the picnic, but it might rain. Your payoff matrix would look like this:

Payoff Matrix

	No Rain	Rain
Picnic	4	0
Bowling	1	1

SPECIAL NOTE

Welcome to the Advanced Football Analytics archive. This is where you can find over 1,500 ground-breaking articles full of the most innovative research and analysis of modern NFL football.

New content can be found at our main site, found here.

SEARCH ADVANCED FOOTBALL ANALYTICS

REQUIRED READING

SI Article on Football Analytics
When to Intentionally Allow a TD
Don't Pay CJ
The Fourth Down Study
4th Downs in the New OT
How Much Should a Win Cost?
QB Aging Curve
Hawks, Doves, and Home Field Advantage
Earthquakes, Kevin Bacon, and Pro Bowls
Unicorns, the Tooth Fairy, and Field Goal Range
First Down Probability
Bricklayers vs. Gladiators
The Concorde of the NFL

If it doesn't rain, the picnic pays off, but if it rains you've lost the afternoon. Bowling is not as much fun as the picnic, but it wouldn't matter if it rains. Minimax says go bowling because 1 is its minimum payoff while 0 is the minimum payoff for the picnic.

Minimax-Regret

Another decision method is known as the minimax-regret criterion. This method seeks to minimize potential regrets. Imagine coming out of the bowling alley and being greeted by a sunny blue sky. 'Darn. Should have gone on the picnic.' In this case, if you go bowling and it doesn't rain, you've gained 1 unit of utility but lost out on 4 units, for a net regret of 3. If you go on the picnic and it does rain, you've gained 0 utility but lost out on 1 unit, for a net regret of 1. If you want to minimize your regret, you'd choose the picnic.

Notice that I haven't mentioned the weather forecast yet. These methods are best relied upon when there is a very high level of uncertainty in the "states of nature" that will determine the payoffs.

Now consider a football example. Say a coach has three plays that make sense for a given situation, and the opposing defense can call one of three kinds of defenses. An example payoff matrix might look something like this:

Hypothetical Football Payoff Matrix

	Def X	Def Y	Def Z
Play A	-4	4	12
Play B	-2	3	8
Play C	3	2	1

Note that this is not game theory. We're not looking for a Nash equilibrium. The offensive coordinator is thinking of the defense as a "state of nature." It's something he has no control over and is difficult to predict.

In this case, both Plays A and B have the possibility of negative payoffs. Play C guarantees at least a payoff of 1, and therefore would be the minimax decision.

The regret method says something different. Assume the defense had called Def X. The best payoff possible given Def X would be 3 with Play C, so had we called Play C there would be no regret. But had we called Play B, we would have earned a -2 payoff, which equates to a regret of -5. In other words, we could have had 3, but we got -2. And had we called Play A, we would have earned a -4, which is a regret of -7.

If we repeat the regret calculation for each possible defense, we get a whole new regret matrix:

Regret Matrix

	Def X	Def Y	Def Z
Play A	-7	0	0
Play B	-5	-2	-4
Play C	0	-2	-11

Given this regret matrix, the minimax-regret criterion would look for the choice that assures us of the best worst-case scenario. For Play A, the worst regret is -7. For Play B, it is -5. And for Play C, it's -11. Therefore, we'd pick Play B because it is the least

ARCHIVE

costly in terms of maximum possible regret.

Of course, coaches or anyone else would never actually draw up a matrix and do the math to make a decision. But just like in the picnic-bowling example, our brains are attempting poor analog versions of these kinds of decision criteria, and emotions play a large role.

Expected Utility

What if we reduce the uncertainty in the defense? We can't predict exactly which one we'll see, but we can estimate the probabilities that we can expect each defense. The expected utility of a choice is the weighted average of the possible payoffs. For simplicity, say each defense is equally likely with a 1 in 3 chance. Now we can estimate the expected utility for each play choice. In the example above, the expected utility for Play A is (1/3)(-4) + (1/3)(4) + (1/3)(12) = 4. The expected utility for Play B is 3, and for Play C it's 2. The expected utility method therefore says Play A is the best choice.

The three methods each call for a different decision. Each method is logical and consistent in its own way, but there is only one truly correct method in football, only one prescriptive analysis. Remember, in football we can know the probabilities and the payoffs, or at least have a solid league-wide baseline for them. The expected utility method is the only correct method.

The math behind expect utility analysis couldn't be any easier. It's 5th grade arithmetic. The challenge is knowing the utility function. Yards, and even points, don't equate to utility. A 7-yard gain is usually good, but it's relatively useless on 3rd and 8. And a 3-point field goal doesn't help late in the 4th quarter when down by 7.

Fortunately, there is win probability (WP). WP is the one and only correct utility function for any game, including football. Winning is all that matters, whether by 1 point or 100 points. WP is also perfectly linear, which is essential to valid expected utility analysis. A 0.40 WP is exactly twice as good as a 0.20 WP, and 0.80 WP is twice as good as 0.40 WP.

Prospect Theory

But even if coaches were to somehow use expected WP analysis when making decisions (say by using 'quick reference' cards like they sometimes do for 2-point conversion decisions), it's likely they still wouldn't be very rational.

Prospect theory says that people fear losses more than they value equivalent gains. Humans evolved with a tendency to try to avoid loss. We're usually more upset with ourselves when we misplace a $20 bill than we are happy when one falls out of the laundry. This tendency has been borne out time and time again in clinical experiments and other studies.

In football, this means that decisions are warped because coaches would fear a loss in WP more than an equivalent gain in WP. The chart below illustrates this concept. According to prospect theory, the "joy" from a 0.05 gain in WP is less than the "pain" from a 0.05 loss in WP.

► March (8)
► February (9)
► January (26)
► 2008 (171)
► 2007 (206)

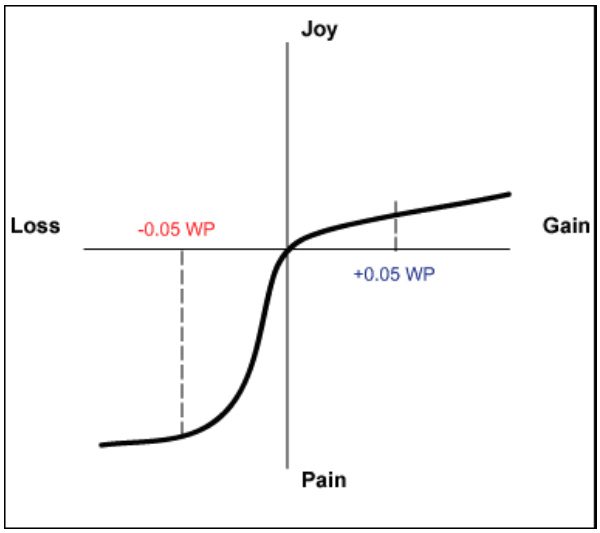

This asymmetry would affect tactical decisions in many ways, but the most obvious may be 4th down doctrine. Say a team finds itself in a situation where punting would result in a 0.50 WP, but the expected utility analysis says going for the conversion would result in a net 0.55 WP. If the goal is to win the game, the correct decision in this case is to go for it. Period.

The analysis isn't so straightforward for the coach (even if he could do all the math on the spot). Say the failed conversion results in a 0.45 WP and the successful conversion results in a 0.65 WP. A 50% chance at successful 4th down conversion therefore results in a net 0.55 WP.

But the coach sees the 0.45 WP as a possible *loss* of 0.05 WP, and he sees the 0.65 as a *gain* of 0.15 WP. Because he fears the loss far more than he values the potential gain, even one 3 times as large, he'll prefer the sure-thing option and punt.

Further, it's possible to actually measure the risk aversion of coaches by comparing the WP advantages in situations where they went for the conversion to the WP advatanges in situations where they forego the conversion attempt.

An Advantage

The coach who can resist this human tendency and make decisions based purely on expected utility will have an advantage. Just how big an advantage, no one can ever know. Actually, that's not true--I'll tell you right now. Just by following a pure expected utility analysis on 4th down, a coach would win an average of an extra 1.4 games per year.

I calculated this based on a play-by-play database from the past 9 seaons. For each 4th down in which a team kicked either a FG attempt or punt, I calculated the difference between going for it and kicking. Wherever the difference was positive, I summed the increase in WP for going for it. The grand total for nearly 2400 games was +203.1 WP, which equates to an increase of 0.17 WP for every game. But since there are always two teams competing in every game, this means that we need to halve that, which is 0.086. The bottom line is that a pure expected utility approach to 4th down decisions would increase a team's chances of winning a game from 0.50 WP to about 0.59 WP. This is equivalent to an extra 1.4 wins per season (0.086*16).

That's a bold claim, I realize. But if you trust my WP model, which is really nothing more than a smoothed empirical observation of how often teams actually won in given game situations in real NFL games, then the claim is not so bold. It's not a