Operations Research on NCAA Football Re-injury Prevention

Nelson Chung[1]
[1]U. S. Census Bureau, 4600 Silver Hill Rd., Suitland, MD 20746

Abstract

Beane (Lemire 2015) identified injury prevention as the next frontier in sports analytics. In this study, I demonstrate that statistics is essential by examining re-injury prevention, through testing whether NCAA coaches and players, facing the need to limit repetitions as key players recover from injury, allocate those repetitions in a manner that maximizes win probability. I calculate in-game win probabilities using the methodology of Stern (1991) and Winston (2009), NCAA historical point spread data from Covers.com, and NCAA expected points of each down-distance-position state from Knowlton (2015). As a test case, I consider Brigham Young University's use of dual-threat quarterback Taysom Hill in his 2016 return from a prior-year Lisfranc injury. I extract play-by-play data from BYUCougars.com using R's rvest package and regular expression functions. I determine the situations in which Hill's rushing, based on his yards per carry prior to the game on standard and passing downs, maximizes changes in win probability.

Key Words: sports analytics, injury prevention, football, data science, probability

1. Introduction

The icon of the analytics revolution in sports, Billy Beane, has said that injury prevention was the next frontier in analytics (Lemire 2015). Using Brigham Young University’s (BYU) dual-threat quarterback Taysom Hill as a case study, I examine re-injury prevention by determining whether BYU allocated his runs in a way that maximized win probability for the 2016 football season. Hill suffered three season-ending injuries prior to 2016: a knee injury in 2012 against Utah state, a leg fracture in 2013, also against Utah State, and in 2015, a Lisfranc injury against Nebraska. In 2016, he played most of the season while having his rushes limited, until hyper-extending his elbow in the last regular-season game against Utah State.

In this study, I assume that quarterbacks are more likely to injure themselves if they rush. I find that in 2016, off all possible situations in which BYU could have used Hill, they used him in a manner that was modestly better than using him at random. There was a high degree of variation in the amount of leverage in Hill’s rushes. In most games, Hill was used in situations that ranked in the 90th percentile in leverage. In two games, against Utah and Cincinnati, he was used in the highest-leverage situation of each game. When Hill was used in low-leverage situations, he gained more yards than expected, suggesting that coaches and Hill had more information about the defense in those situations than just down, distance, field position, time, and score.

2. Literature Review

2.1 Similar Methodologies in Popular Analytics

Analysts have employed similar methodologies in popular analytics works. Woolner (2006) attempted to answer the question of whether baseball teams can use a starting rotation with only four pitchers instead of the usual five, without overextending and consequently injuring them. His approach is opposite mine: he determined how much strain pitchers can withstand given the fixed workload of three days' rest, while mine is to estimate the maximum contribution to win probability given repetitions fixed to prevent re-injury. Tom Tango et al. (2014) also studied the allocation of talent by determining the batting lineup that would maximize the number of runs scored. In my study, the objective is to maximize win probability, of which expected points is a component. Dawkins (2012) examined why Washington Nationals phenom Steven Strasburg kept re-injuring himself.

2.1 Re-injury Prevention in Scholarly Literature

Holme et al. (1999) assessed the impact of supervised rehabilitation programs on re-injuries after acute ligament strains. Woods et al. (2004) found that the re-injury rate for soccer players who underwent hamstring injuries was higher than for other injuries. In a longitudinal study of twenty-three UEFA soccer teams over seven consecutive seasons, researchers in Sweden found that 12 percent of all injuries were re-injuries, and re-injuries caused longer absences than non-reinjuries (Ekstrand, et al. 2010).

3. Methodology

3.1 Determining Optimal Use in Face Possible Re-injury

Players recovering from injury need to limit their workload during games in order prevent re-injury. Determining whether BYU used Hill optimally requires assessing the impact his use has on win probability. Ideally, this entails calculating how having Hill run in every possible situation changes win probability, and inspecting whether the instances in which BYU *actually* had him run created the largest positive changes in win probability among all possible situations. But doing what is optimal in every situation, based only on down, distance, and field position, and information on your own team, is not possible under game-theoretic conditions because the opponent will prepare himself for the play and negate your advantage. Winston (2009, 158-64), for instance, has demonstrated in a two-person zero-sum game, it may be optimal to pass less than before *even if* the team averaged more yards per pass than before, and vice versa. Thus, to determine optimal use, I examine how win probability gains (or losses) in which BYU used Hill to run lie along the entire distribution of win-probability gains. To summarize, I (1) calculate the changes in win probability for all possible plays in which BYU would run Hill, and (2) determine where the plays in which Hill was used lies along the distribution of (1).

3.1.1 Win Probability

Historically, the point spread has outperformed all regression models in predicting the margin of victory in football games (Fair and Oster 2007). Stern (1991) built a win probability model based on the point spread. The margin-of-victory of an NFL game is distributed normally with a mean equal to the spread and a standard deviation of 13.861. Using this, Winston (2009) constructed an equation for in-game win probabilities with a pdf over the game margin. The win probability would equal the entire area under the curve from $(0.5 + x)$ to ∞, plus the half the area from $(-0.5 + x)$ to $(0.5 + x)$, all divided by the entire area under the curve. The x equals the game margin plus expected points. The reasoning behind this is since game outcomes are discrete, so the [-0.5 + x, 0.5 + x] interval represents a tie game, which would send the game into overtime, and overtime games are won with a fifty percent chance. The following equation illustrates this relationship:

$$(1)\ W(x,\mu,\sigma) = \int_{0.5+x}^{\infty} \frac{e^{-\left(\frac{x-\mu}{2\sigma}\right)^2}}{\sqrt{2\pi}}\,dx + \frac{1}{2}\int_{-0.5-x}^{0.5+x} \frac{e^{-\left(\frac{x-\mu}{2\sigma}\right)^2}}{\sqrt{2\pi}}\,dx,$$

where $x = m + E(pts)$, with x as the expected game margin, m the game margin, and $E(pts)$ denoting expected points. The μ represents the point spread for the game, and should be adjusted down proportional how much of the game there is to play. The intuition is this: since the point spread equals how much the team would score above the other team in a sixty-minute period, if there is only one quarter remaining, one would expect the team to score a quarter of the point spread more points than its opponent. The standard deviation should also be adjusted down by dividing by the square root of the inverse of time proportion of the game remaining. For instance, BYU was a 7.5-point underdog against West Virginia. With about 6.141 minutes remaining in the game, $\mu = (7.5)(6.141/60) = 0.7676$, with $\sigma = 15.82/\mathrm{sqrt}(60/6.141) = 5.061$. Figure 1 provides the graphical illustration.

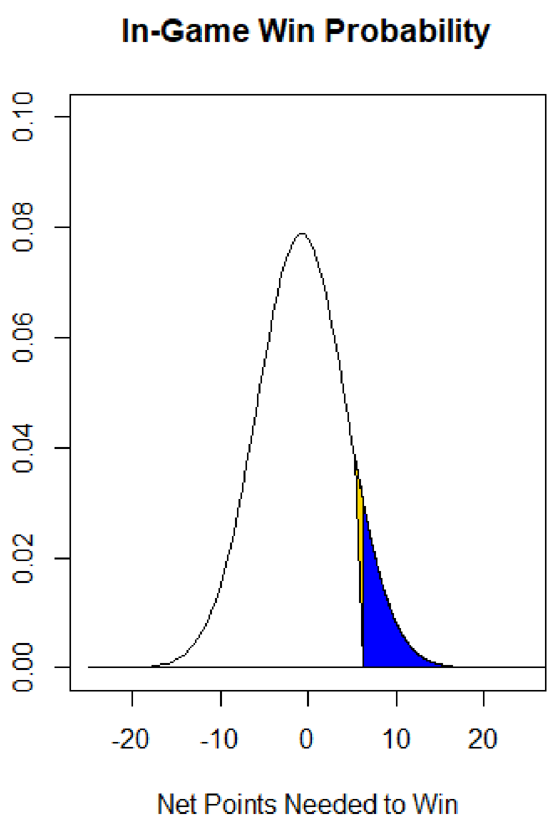

Figure 1: Sample Win Probability Situation. BYU was playing 7.5-point favorite West Virginia at Fedex Field, trailing by 10 points with about 6 minutes remaining in the game, with the ball at the West Virginia 19 yard line facing a 1st and 10. Having the ball in that predicament equals 4.224 expected points. Since football scores are discrete, a tie (represented by the gold region) means losing by 0.5 to winning by 0.5. A win is represented by the blue region, wining by 0.5 points or more. To win outright during

regulation, BYU would have to score 10 – 4.224 + 0.5 = 6.276 points. Total win probability equals half the area of the gold region added to the entire area of the blue region, all divided by the total area under the curve.

3.1.2 Expected Points

What matters to win probability is not only the current score, but also the expected score. Carter and Machol (1971) first introduced the concept of expected points. To determine the impact the player has on expected points, we assume that each time he rushes, he expects a certain number of yards. I assume the number of expected yards should be Hill's cumulative average before each game, on rushing downs and passing downs. This distinction is important for game-theoretic reasons—defenses gear for rushes on down-distance situations when offenses usually pass and for passes when offenses usually pass. Connelly (2015) has determined passing downs to be second-and-5 or more, third-and-8 or more, and fourth-and-8 or more.

To determine expected changes in expected points, suppose that Hill averages 8 yards per carry on passing downs, and 7 yards per carry on standard downs. In BYU's game against Michigan State, he was facing second-and-8 from the Michigan State 22-yard line. NCAA football teams average 3.793404 points from that position. Since this is a passing down, we expect Hill to gain 8 yards by rushing. This means the new expected state would be first-and-10 from the 14-yard line, with expected points of 4.384309. Hill is expected to gain 4.384309 – 3.793404 = 0.5909049 expected points were Hill to rush. For this study, what I am doing differs slightly in that I will be calculating is changes in win probability, of which expected points is a component, instead of expected points per se.

3.2 Ranking Hill's Runs among All Plays in Terms of Changes in Win Probability

After I calculate changes in win probability for every down-distance-field position and time left remaining in game, if Hill were to run, I rank them from highest to lowest. Then I take the plays in which Hill actually did run, and compute the percentile they fall on. Afterwards, I examine boxplots of the data and compute the mean percentile of Hill's actual runs.

To determine expected changes in expected points, assume that Hill averages 8 yards per carry on passing downs, and 7 yards per carry on standard downs. To illustrate: in BYU's game against Michigan State, Hill was facing second-and-8 from the Michigan State 22 yard line. NCAA football teams score 3.793404 points from that position. Since this is a passing down, we expect Hill to gain 8 yards by rushing. This means the new state would be first-and-10 from the 14 yard line, with an expected points of 4.384309. Hill is expected to gain 3.793404 – 4.384309 = 0.5909049 expected points if he were to rush. For this study, what I am doing differs slightly in that what I will be calculating is changes in win probability, of which expected points is a component, instead of expected per se.

4. Data

4.1 Priors

Before analyzing BYU's use of Hill during the 2016 season, two pieces of information are required to assess his impact on win probability: first, information on the distribution of margins-of-victory for NCAA games; second, information on the expected outcome if Hill were to be used.

4.1.1 Betting Spread for Prior Normal
Covers.com has betting spreads and game results from the 2010-11 NCAA football season on. From the 2010-11 to the 2015-16 seasons, the spread served as an unbiased predictor of games, with deviations distributed normally with a mean of 0 and a standard deviation of 15.82.

4.1.2 Hill's Yards Per Carry
To compute the expected outcome were Hill to run, I scrape play-by-play data from BYUCougars.com, the official athletic site of the BYU Cougars, for all games in which he played prior to the 2016 season, using R's `rvest` library. Simple yards per carry, available without any calculation, is insufficient for my purposes for the following reasons: (1) the need to compute yard per carry for standard and passing downs separately, and (2) the need to distinguish between sacks and quarterback keepers (planned or improvised) that resulted in negative yardage, the latter of which should be excluded. Fortunately, the data from BYU athletics makes this distinction.

To maintain time consistency, I assume that BYU coaches and Hill both know Hill's career yards per carry leading up to the game. The following table displays Hill's yards per carry for standard and passing downs leading up to each game:

Table 1. Taysom Hill Yards per Carry Prior to Each Game

	Yard per Carry	
	Standard Downs	Passing Downs
Arizona	7.356828	8.107843
Utah	7.245763	8.048544
UCLA	7.307054	8.233645
West Virginia	7.173387	7.875000
Toledo	7.084942	8.043103
Michigan State	7.084942	8.170940
Mississippi State	7.033835	8.058333
Boise State	7.025830	8.190083
Cincinnati	7.007246	8.047244
Southern Utah	7.000000	8.193798
UMass	6.955326	8.193798
Utah State	6.915825	8.298507

4.2 Game Information for 2016
With prior information in place, I now extract the information required to determine whether BYU used Hill optimally as he recovered from injury in 2016. This entails betting spreads for each game and in-game information.

4.2.1 Betting Spread for Each Game
Like historical spreads, spreads for each 2016 game are available from Covers.com. Note that the number, which represents the number of points BYU is favored by, is the negative of the actual point spread.

Table 2. Points Spread for Each Game

	BYU Favored by
Arizona	1.5
Utah	-3
UCLA	-3.5
West Virginia	-7.5
Toledo	3
Michigan State	-3.5
Mississippi State	7.5
Boise State	-7.5
Cincinnati	8
Southern Utah	32
UMass	28
Utah State	17

4.2.2 In-Game Information
I scraped all required in-game information from BYUCougars.com, extensively using regular expression capabilities in R. This information entails play-by-play data, which includes down, distance, field position, quarter, time remaining in quarter, yards gained on play, and whether Hill ran. The time remaining in the quarter is only available for beginning and end of drives and when timeouts are called. For plays in between those available, I estimate by assuming a linear relationship between the each play and the time. For instance, suppose that for a four-play drive, it is known that the first play occurred with 10:00 left in the second quarter, and the last play occurred with 8:30 remaining in the same quarter. I would estimate the other two plays. I would estimate the two plays occurring in between to be at 9:30 and 9:00 in the quarter.

4.3 Win Probability Calculations
I calculate changes in expected win probabilities for before and after each play using formula (1) in section 3.1.1 and R's `pnorm()` function.

5. Results and Discussion

5.1 Results on BYU's Use of Hill
After calculating the changes in expected win probability for all of BYU's offensive plays from scrimmage, I examine where along the distribution do plays in which BYU used Hill lie, in terms of percentile. The 99th percentile the play in which theoretically running Hill would produce the largest change in win probability in the entire game. The first percentile represents the lowest.

5.1.1 Game Averages
The following table contains the average rank, in percentile, for plays in which Hill rushed, among all plays in each game. The higher the percentile, the more leverage the use of Hill. The mean percentile of all the game means is the 53rd percentile. This indicates in

determining when to run Hill, on average, BYU used him with only modestly better leverage than using him at random.

Table 3. Rank of Taysom Hill Rushing Plays (in Percentile), Changes in Win Probability

	Average Percentile
Arizona	0.458333333
Utah	0.611244019
UCLA	0.436507937
West Virginia	0.392307692
Toledo	0.748858447
Michigan State	0.533950617
Mississippi State	0.468312757
Boise State	0.630864198
Cincinnati	0.477272727
Southern Utah	0.725
UMass	0.559925094
Utah State	0.593181818

5.1.2 Boxplot

To gain a better picture of the distribution of plays in which Hill ran, I generate a boxplot using the `ggplot2` library, found on Figure 1. Although BYU, on average, did not appear to be using Hill's limited runs in an optimal manner, they *did* use him when the highest-leverage situations presented themselves. In games against Utah and Cincinnati, Hill took the highest-leverage situation of the entire game. Against Utah, with 00:18 left in the game, BYU was facing a first and goal at the Utah 7-yard line, and was down by 7 points. Since this is a standard down, if Hill were to run in that situation, he would gain an expected 7.25 yards (see table 1), which would result in a touchdown, evening the score and brining BYU's win probability to slightly under 50 percent. Hill *did* run and scored the touchdown, although the team unsuccessfully tried for a two-point conversion instead. Against Cincinnati, BYU was facing a third down and 7 at the Cincinnati 8-yard line, with the score tied with 00:48 seconds left in the first half. This is a passing down, so we expect Hill to gain 8.05 yards (see table 1) if he were to run, which would result in a touchdown. He did, and scored a touchdown, enabling BYU to take a 10-3 lead. In these two games, Hill ran when running accrued the greatest benefit the entire game.

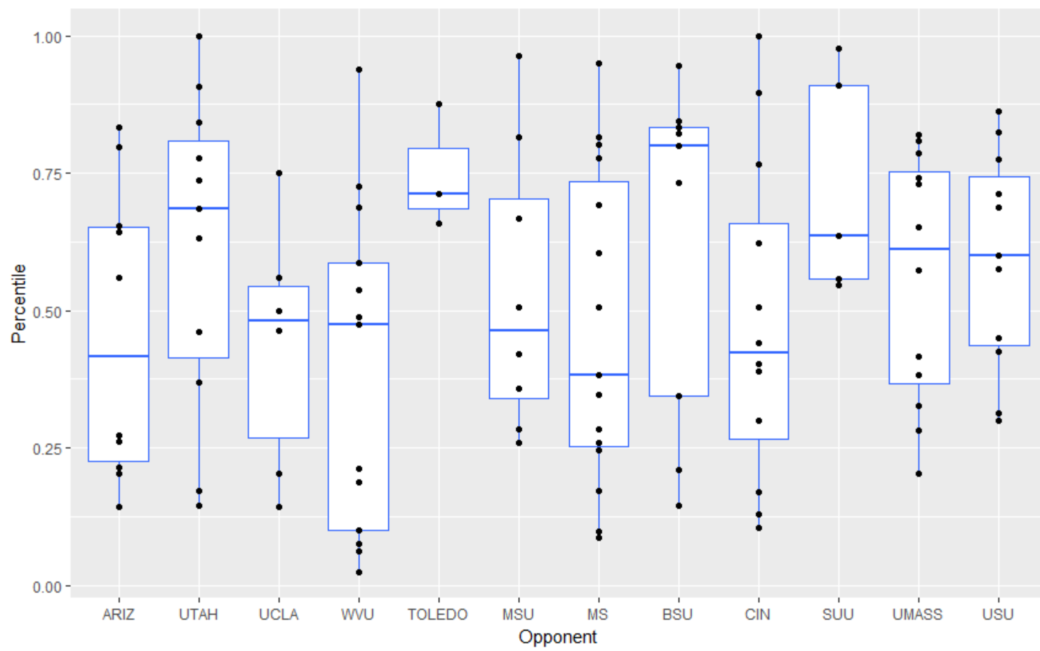

Figure 1: Changes in Win Probability, Instances in which Hill Ran for Each Game (ranked in terms of percentile among all possible plays)

5.1.3 Did Coaches or Hill Have More Information?

The offensive coordinator, Ty Detmer in this case, and quarterback Hill should have more information in making their decisions than the score, down, distance, field position, and time remaining. They may see something in the defensive alignment to base their decision whether the quarterback should run. To determine whether this is true, we explore the low-leverage situations in which Hill ran. The data suggests that there is truth to this. For Hill's runs that rank under the 25th percentile for all potential run situations in games, Hill averaged 6.82 yards per run, compared to 6.09 yards per run for all his other runs, indicating reasons for running in low-leverage situations.

5.2 Discussion

The broadest results from this study show a high degree of variation in the leverage of Hill's rushes. BYU used Hill's rushes in both high- and low-leverage situations, with the average being modestly above using him at random. This also shows that BYU *did* take much of the highest-level situations. Furthermore, there appears to be justification for using Hill in low-leverage situations.

5. Conclusion and Further Steps

As Hill returned from injury for the 2016 season, coaches sought to limit his repetitions, in this case the number of times he rushed. This study showed that he was used with modestly higher leverage than being used at random; that BYU did use him in many of the highest-leverage situations; and, when BYU used him in low-leverage situations, they had more information that justified doing so. This examination should be extended to other dual-threat quarterbacks and running backs for more external validity. Collaboration with subject-matter experts, namely sports doctors, should be done on Lisfranc injuries and hyper-extended elbows.

Acknowledgements

Special thanks to Phil LeClerc, Aref Dajani, Steve Clark, William Wisniewski, and Phyllis Singer for their feedback on both the form and the substance of the presentation. Much gratitude to Rolando Rodriguez for help with data manipulation in R. Physicist Chad Junkermeier and Kyle Hart deserve credit for changes to the key slide. Most importantly, thanks to Zachary Knowlton for sharing his expected points.

References

Carter, V. & Machol, R. (1971). Technical Notes—Operations research on football. *Operations Research* 19, 541-544.

Connelly, B. (2015). The 2015 advanced college football stats glossary. *Football Study Hall*. Retrieved from https://www.footballstudyhall.com/2015/2/9/8001137/college-football-advanced-stats-glossary

Dawkins, C. (2014). Could Stephen Strasburg's injury have been prevented? In *Extra innings: More baseball between the numbers* (pp. 283-295). New York: Basic.

Ekstrand, J. M., Hagglund, M., & Walden, M. (2009). Injury incidence and injury patterns in professional football: The UEFA injury study. *British Journal of Sports Medicine* 45, 553-558.

Fair, R. C. and Oster, J. F. (2007). College football rankings and market efficiency. *Journal of Sports Economics* 8, 3-18.

Holme, E., Magnusson, S. P., Becher, K., Bieler T., Aagaard, P., & Kajer, M. (1999). The effect of supervised rehabilitation on strength, postural sway, position sense, and re-injury risk after acute ankle ligament sprain. *Scandinavian Journal of Medicine and Science* 9, 104-109.

Knowlton, Z. & Fellingham, G. (2015). Ranking NCAA teams through expected points. In *JSM Proceedings*. Alexandria, VA: American Statistical Association.

Lemire, J. (2015, December 16). MLB's next frontier: Injury prevention. *USA Today*. Retrieved from https://www.usatoday.com/story/sports/mlb/2015/12/16/mlbs-next-frontier-injury-prevention/77412998/

Stern, H. (1991). On the probability of winning a football game. *The American Statistician* 45, 179-183.

Tango, T., Lichtman, M., & Dolphin, A. (2014). *The Book: Playing the percentages in baseball*. Lincoln, NE: University of Nebraska Press.

Winston, W. L. (2009). *Mathletics: How gamblers, managers, and sports enthusiasts use mathematics in baseball, basketball, and football.* Princeton, NJ: Princeton University Press.

Woods, C., Hawkins, R. D., Maltby, S., Hulse, M., Hodson, T. A. (2004). The football association medical research programme: An audit of injuries in professional football—analysis of hamstring injuries. *British Journal of Sports Medicine* 38, 36-41.

Woolner, K. (2006). Five starters or four? On pitching and stamina. In *Baseball between the numbers: Why everything you know about the game is wrong* (pp. 74-86). New York: Basic.