Fantasy Football Draft Analysis

Prepared for: Client
Date:

Table of Contents

Introduction

In this report, we present the methodology we used to develop a Fantasy Football drafting tool that determines the optimal player selection at each decision point in a draft. The tool is based on a two-pronged approach that combines (a) mathematical models for point projections, and (b) qualitative adjustments, based on content analysis, to account for factors such as players changing teams or the projected performance of rookies.

The report is organized as follows. In the first section, we develop a methodology to create point projections for each player based on their characteristics, such as age, past performance, etc. The second section presents the results of a Content Analysis used to create qualitative adjustments to the point projections. Section III shows how to use the inputs of our projections, combined with Modern Portfolio Theory, to compute a drafting strategy that maximizes risk-adjusted points earned. Next, the fourth section presents a selection of players based on point projections.

Finally, in Section V we develop a drafting engine that uses an artificial intelligence algorithm to simulate the draft given the available information about the drafting styles and propensities of other players in the pool.

SECTION I: Player Projections and Point Calculations

Creating a Fantasy Football projection for a given player is incredibly complex. There are more variables in football than in most other sports because a player's production is very dependent on his teammates and situation, and the various positions progress in a dramatically different fashion over the years. For example, running backs are often very productive as young players, but their production falls off steeply as they approach 30, while quarterbacks are often not that effective until their mid-to-late 20s, though they can be productive for many years after that.

Contained within this section are 2008 projections for NFL quarterbacks, running backs, wide receivers, tight ends, kickers and team defenses, along with information and templates that one could continue to use into the future. These projections focus heavily on the quantitative approach, which will very accurately project the holistic evolution of a group of players. We also consider a qualitative approach, however, which will take into account specific football knowledge for individual players. For example, anyone that watches football *knows* that Randy Moss' statistics two seasons ago (in his last season with the Raiders, where he essentially quit playing) have little to no bearing on his upcoming season with the Patriots. Each positional quantitative projection system varies slightly because of positional nuance. A discussion of each position follows, in order of ascending complexity. In addition, we estimate a number of risk factors that allow for the assessment of the volatility that holding a specific player may bring. These factors are discussed prior to the estimation of players by position.

Risk Factors

Our projections take into account three separate risk factors that are more completely addressed in the Portfolio Theory section of this report. This section will provide a brief introduction into how we have thought about and operationalized risk.

Standard Deviation of Production as a Measure of Risk:

The first factor that we have chosen to incorporate into our study that no other such analysis available on the market uses is the standard deviation of production. Every other projection system on the market uses an entire season's worth of games as a data point. We, on the other hand, have actually studied the game-by-game statistics of every player, averaged them, and then taken the standard deviation of each statistic per game. Thus, while our season averages are the same, we now have data on how variable a player's performance is likely to be.

In certain situations, you may require stable players like Marshawn Lynch, who has a low standard deviation because he gained between 70 and 107 yards in all but two games (156 and 66 were the high and low) and scored a touchdown in 7 out of 13 games but never scored more than one touchdown per game. In other situations, you may want to choose a player with a higher standard deviation, such as Larry Johnson. In the last 24 games, Johnson has been held five times under 40 yards rushing and failed to score touchdowns in 10 of the games. Yet he also exploded for a four touchdown and a three touchdown game and scored twice several times and also gained over 150 yards four times. Marshawn's youth and Johnson's age and injuries also play into their safety and riskiness and are factors that one should consider when making a pick.

Projection Differences Model of Risk:

Our model also considers another risk factor for all players. Instead of simply relying on historic risk factors such as standard deviation, our model also takes into account large gaps in per game point production from last year to this year. These large gaps, which are operationalized through the subtraction of the previous year's point per game totals from the projections of points earned from this season divided by the number of games, yield vital information regarding one or more implicit risks for the model. This second risk factor is called Projection Differences (PD). This risk factor takes into account increases or decreases in point per game production between seasons.

Latent Risk:

The last risk factor that we consider for all players is also the most complicated. This risk factor takes into account the facets of risk that may be missed by the standard deviation factor and the projection differences factor. Overall, we estimate that the combined risk attribution to standard deviation and projected differences does not take into account a number of other potential sources for risk. The latent factor, operationalized through intensive study, surveying, and statistical analyses is designed to take a number of implicit sources of risk into account.

Confidence Intervals:

The three metrics above were combined to create a final three-factor measure of risk. This methodology is appropriately discussed in the Portfolio Theory section;

however, the use of the three-factor risk figure is prominently shown in each of the following player projection sections. The use of the risk measure for the purposes of this section is the construction of confidence intervals around the expected points figures. These confidence intervals relate to the concept of having a "breakout" season. Specifically, we took the 80% universe of top seasons that could happen based upon the standard deviation-scaled, three-factor risk associated with each player. These tables are displayed at the end of each player projection discussion.

Running Backs

Running backs are the simplest players to model because including more than one season of data does not increase the accuracy of the projections.[1] While that does not mean that such data is not relevant in certain specific cases, our goal here is to produce a template to generally predict a running back's performance versus age.

Aging:

To accomplish this estimation, we analyzed data from the top 50 running backs in NFL history (who, therefore, were starters for the duration of their career) and created the following matrix:

Age	Carries	RushYds	YPC	RushTD	RecYds	ScrimYds	TotTD	1,000s*	10s**
23-	242.3	1064.4	4.39	8.28	273.2	1337.6	9.44	48.5	40.6
24	267.3	1187.8	4.44	8.42	355.5	1543.3	9.6	69.4	49
25	297.4	1306.9	4.39	9.48	352.6	1659.5	10.92	71.4	42.9
26-27	298.3	1282.7	4.3	9.16	348.3	1631	10.48	73.5	41.8
28-29	279.2	1188.1	4.26	8.59	306.5	1494.6	9.86	53.9	30.3
30	241.1	1000.2	4.15	6.56	269.2	1269.4	7.64	33.3	26.2
31	206.4	858.8	4.16	5.72	242.6	1101.4	6.59	26.5	14.7
32	176.6	682	3.86	5.08	229.2	911.2	5.81	10.3	17.2

[1] http://sports.espn.go.com/fantasy/football/ffl/story?page=nfldk2k830rbs

33+	162.4	601.9	3.71	5.29	151.3	753.2	5.78	7.7	12.8

*1,000s is percentage of starters at this age who were 1,000 yard rushers.
**10s is the percentage of starters at this age who managed a 10 touchdown season.

We transformed this data into a more usable set, which demonstrated the average rise or decline in production versus age:

Age	Carries	RushYds	YPC	RushTD	RecYds	ScrimYds	TotTD	1,000s	10s
24	0.10	0.12	0.01	0.02	0.30	0.15	0.02	0.43	0.21
25	0.11	0.10	-0.01	0.13	-0.01	0.08	0.14	0.03	-0.12
26	0.00	-0.02	-0.02	-0.03	-0.01	-0.02	-0.04	0.03	-0.03
27-28	-0.06	-0.07	-0.01	-0.06	-0.12	-0.08	-0.06	-0.27	-0.28
29-30	-0.14	-0.16	-0.03	-0.24	-0.12	-0.15	-0.23	-0.38	-0.14
31	-0.14	-0.14	0.00	-0.13	-0.10	-0.13	-0.14	-0.20	-0.44
32	-0.14	-0.21	-0.07	-0.11	-0.06	-0.17	-0.12	-0.61	0.17
33	-0.08	-0.12	-0.04	0.04	-0.34	-0.17	-0.01	-0.25	-0.26

From this simple table, we should be able to project how the average running back of a given age will do when he progresses one year. This may initially sound oversimplified, but it is a far more accurate model in the long run than one might think. For example, this data predicts massive fall offs for 29 and 30-year-old backs — while there are some notable exceptions throughout history, historical data demonstrates that some of the most frequently overrated running backs have been older backs coming off of surprisingly good seasons. For example, Shaun Alexander followed his 27-touchdown performance in 2005 at age 28 with 7 touchdowns in 10 games in 2006 at age 29, and a 4-touchdown (6.4/16) performance in 2007 at age 30. Bad picks such as these can often be avoided by following the projections given by this matrix.

In fact, there have been several other studies that suggest running backs hit a wall as they approach age 30, and our data corroborates this.[2] Thus, despite the occasional aberration, one should have fairly negative expectations of older running backs. On the other hand, running backs aged 24 to 25 experience tremendous gains. This has been reflected in our projections, as we have calculated that on average, each 24 year old will add 10.3% more carries, 11.6% more yards and 1.7% more touchdowns. Twenty five year olds are projected to experience 13.75% touchdown growth as well as similar gains in yardage and carries.

Of course, not every member of the group will experience such terrific gains. However, the *averages* across the board should hold constant, meaning that these young players are worthy of their inflated points projections. While it is true that some will fail to live up to these expectations, it is also necessarily true that others will *exceed* these expectations. Further, some averages show counterintuitive results. It is unlikely that 33-year-old running backs improve their rushing touchdowns. Instead, the 0.04 growth indicates an outlier. We tested that outlier, finding that it is not significantly different than zero. As a result, for the purposes of our calculation, it was treated as such. Overall, our projections provide a balance between proven talent and potentially explosive growth.

Confidence Intervals:

As mentioned in the introduction, we used the three-factor risk metric to create confidence intervals around the 80% universe of breakout seasons. Both the low end

[2] http://sports.espn.go.com/fantasy/football/ffl/story?page=nfldk2k830rbs

(20% universe) and the high end (80% universe) are displayed to give a sense of the size of the overall confidence interval. They are displayed below:

Name	Mean	20% Universe	80% Universe
Adrian L. Peterson	262.60	211.00	314.00
LaDainian Tomlinson	262.56	213.00	312.00
Brian Westbrook	262.46	212.00	313.00
Marshawn Lynch	235.34	186.00	285.00
Joseph Addai	230.41	180.00	281.00
Larry Johnson	220.79	169.00	273.00
Brandon Jacobs	207.75	144.00	271.00
Jamal Lewis	204.02	138.00	270.00
Steven Jackson	200.71	160.00	242.00
Maurice Jones-Drew	198.73	149.00	249.00
Reggie Bush	194.99	148.00	242.00
Clinton Portis	191.54	147.00	236.00
Frank Gore	185.93	147.00	225.00
Ronnie Brown	182.75	116.00	250.00
Marion Barber III	182.69	140.00	225.00
Laurence Maroney	177.63	126.00	229.00
Willis McGahee	169.69	131.00	208.00
Earnest Graham	160.58	101.00	220.00
LenDale White	160.33	112.00	209.00
Ryan Grant	156.43	105.00	208.00
Edgerrin James	153.37	93.00	214.00
Chester Taylor	153.18	112.00	195.00
Fred Taylor	146.83	97.00	196.00
Willie Parker	143.45	97.00	190.00
Kevin Jones	143.16	87.00	199.00
Ahman Green	141.99	83.00	201.00
Darren McFadden	141.00	94.00	188.00
Jonathon Stewart	138.00	92.00	184.00
Rudi Johnson	137.49	82.00	193.00
Justin Fargas	137.19	79.00	195.00
Kevin Smith	136.00	90.00	182.00
Matt Forte	134.00	88.00	180.00
Kenny Watson	132.84	73.00	193.00
Ron Dayne	127.63	67.00	188.00
Deuce McAllister	125.84	75.00	177.00

Name	Mean	20% Universe	80% Universe
Shaun Alexander	120.64	74.00	168.00
Thomas Jones	119.92	72.00	168.00
Rashard Mendenhall	116.00	71.00	161.00
Cadillac Williams	112.96	51.00	175.00
Tatum Bell	110.70	69.00	153.00
DeAngelo Williams	110.21	62.00	158.00
DeShaun Foster	107.53	58.00	157.00
Felix Jones	105.00	63.00	147.00
Warrick Dunn	103.97	51.00	157.00
Selvin Young	99.73	51.00	148.00
Adrian Peterson	97.16	60.00	134.00
Jerious Norwood	92.78	44.00	142.00
Julius Jones	76.28	35.00	118.00
Ladell Betts	54.80	6.00	103.00
Michael Turner	36.32	0.00	89.00

*NOTE: Confidence Interval would indicate total points scored could be less than 0. Therefore, size of the confidence interval was artificially constrained to zero on the low end.

The above table demonstrates the upper and lower confidence intervals of performance for each player. Players with high risk metrics relative to their points are worthy of consideration in risky circumstances because they may have a far better than projected season. Those players have the largest spread between the 80% and 20% universes of performance. That said, their risk must be balanced against their expected point totals. The remainder of the confidence interval tables in this chapter will be presented without this note.

Wide Receivers

Wide receivers are somewhat more difficult to predict than running backs. Data suggests that there is a much steeper learning curve to playing wide receiver in the NFL than there is to playing running back, but receivers also age much more gracefully.

Analysis of receivers' progression throughout the first 10 years of their careers is as follows:

Yr	Rec	Yards	Y/R	TD	1000s	10s
1	38.90	618.50	15.89	3.94	6.00	2.00
2	53.00	865.50	16.32	6.00	26.00	18.00
3	57.90	921.90	15.93	7.55	36.00	24.00
4	68.00	1039.80	15.30	7.01	44.00	14.00
5	76.70	1138.40	14.84	8.28	58.00	24.00
6	76.60	1135.00	14.82	8.22	58.00	22.00
7	74.30	1096.60	14.75	7.56	56.00	26.00
8	71.90	1046.70	14.57	6.65	50.00	16.00
9	69.9	1008.9	14.42	6.9	42	24
10	68.50	980.60	14.31	6.62	40.80	12.20

This data falls in line with another study, which demonstrated that wide receivers typically hit their peak around age 27.[3] Note that this does not mean that the best WRs are 27 years old but rather that when compared with themselves, wide outs are often at their best when they are 27. Of course, there are many examples of players who peak earlier or much later, but this data is representative of wide receivers *as a whole* over the last 10 years. A chart showing average improvements with age is as follows:

Yr	Rec	Yards	Y/R	TD	1000s	10s
1 to 2	0.37	0.40	0.09	0.52	3.33	8.00
2 to 3	0.09	0.07	-0.02	0.26	0.38	0.33
3 to 4	0.17	0.13	-0.04	-0.07	0.22	-0.42
4 to 5	0.13	0.09	-0.03	0.18	0.32	0.71
5 to 6	0.00	0.00	0.00	-0.01	0.00	-0.08
6 to 7	-0.03	-0.03	0.00	-0.08	-0.03	0.18
7 to 8	-0.03	-0.05	-0.01	-0.12	-0.11	-0.38
8 to 9	-0.03	-0.04	-0.01	0.04	-0.16	0.50
9 to 10	-0.02	-0.0281	-0.0076	-0.0406	-0.0286	-0.4917

[3] http://www.middlebury.edu/services/econ/repec/mdl/ancoec/0718.pdf

Old age does not typically become a factor for wide receivers until their sixth or seventh season. An age regression chart, similar to that of running backs, more accurately analyzes production of wide receivers over 28. This chart is as follows:

Age	Rec	Yards	Y/R	TD	1,000s	10s
27	75.20	1145.10	15.23	8.42	64.00	30.00
28-31	72.10	1048.90	14.55	7.00	49.50	18.20
32-33	65.00	919.20	14.15	6.13	38.80	12.90
34	59.80	855.20	14.31	4.86	25.80	6.50
35	52.80	734.10	13.90	4.23	28.00	0.00
36+	45.90	602.50	13.12	3.33	8.60	0.00

It is further distilled into a progress prediction chart:

Age	Rec	Yards	Y/R	TD	1,000s	10s
27 to 28	-4.12	-8.40	-4.46	-16.86	-22.66	-39.33
28-32	-9.85	-12.37	-2.75	-12.43	-21.62	-29.12
32-34	-8.00	-6.96	1.13	-20.72	-33.51	-49.61
34 to 35	-11.71	-14.16	-2.87	-12.96	8.53	-100.00
35 to 36+	-13.07	-17.93	-5.61	-21.28	-69.29	0.00

Two-Season Basis for Projections:

Furthermore, wide receivers' production shows a greater correlation to the previous two seasons combined (weighted about 2:1 to the more recent season) than to just the previous season alone as for running backs. Thus, our projections modeled the 2007 data through one appropriate year on the table and modeled the 2006 data through two years. So, a third-year wide receiver will have his (likely somewhat low) first-year yardage increased by 37% and then 9%, and then his (likely much higher) second-year yardage increased by 9%.

Because wide receivers naturally have much more volatility than running backs, using 32 data points rather than 16 and then optimizing the weighting system towards more recent data helps create more accurate projections. In addition, as with the running backs, we have included data on wide receivers' standard deviation of production. For example, we project Chad Johnson and Santonio Holmes will both score 173 points, but Holmes is much more consistent with a standard deviation of 28, while Chad Johnson has many more multi-touchdown, tremendous games and also many more games where he fails to make any impact whatsoever, yielding him a standard deviation of 41.

Strength of Schedule:

In an effort to achieve even better projections, we have also taken care to factor in the strength of the opposing pass defenses that a player has faced. We ranked the 2007 and 2006 NFL pass defenses from top to bottom by number of fantasy points allowed to receivers, and then calculated the ratio of each team's score to the mean NFL score. For example, in 2007, Detroit had the worst passing defense, allowing 1.31 times as many fantasy points to WRs than the NFL average, while Indianapolis had the best pass defense, allowing just 0.8 times as many points. Thus, a receiver who scored 22 points against Indianapolis accomplished a task equally as difficult as a receiver who scored 36 points against Detroit.

After we multiplied each receiver's stats against the strength of the pass defense he faced in that game, we arrived at the receiver's 'true' statistics for the year. These statistics let us know more precisely how good a given receiver is, which is important in projecting how well he will do in 2008. With this data in hand, we then calculated the strength of schedule for the upcoming season (by using last season's data) and projected future performance. Thus, a receiver who comparatively faced poor pass defenses over the last two years and who will face a comparatively good set of pass defenses in 2008 will likely see his production drop compared to the strict years-in-league-based projection. Conversely, a receiver who faced tough pass defenses and will now face more porous secondaries will likely see his production rise more than otherwise expected. For example, in 2008 the Bears' and Packers' wide receivers face pass defenses nearly 8% softer than the norm, while the Bills face a strength of schedule nearly 7% easier than the average. Note that our need to project how ably a receiver will be able to deal with pass defenses in 2008 is the only reason we care about how 'good' he truly is. A 'good' receiver who earned 200 points facing extremely tough pass defenses and will again face extremely tough pass defenses is just as attractive as an 'average' receiver who earned 200 points by facing extremely easy pass defenses but will again face extremely easy pass defenses in 2008.

The relative strengths of the pass defenses in 2007 and 2006 are displayed below:

	2007 Points	Multiplier	Ratio
Detroit	605.1	0.77	1.31
New Orleans	584.4	0.79	1.26
Minnesota	554.5	0.84	1.2
Cincinnati	542.7	0.85	1.17
Cleveland	542.1	0.85	1.17
Arizona	527.6	0.88	1.14
Atlanta	527.5	0.88	1.14
Houston	518.2	0.89	1.12
Baltimore	517.7	0.9	1.12
St. Louis	511.3	0.91	1.1
San Francisco	508.3	0.91	1.1
Buffalo	495.4	0.94	1.07
Chicago	484.8	0.96	1.05
NY Giants	475.7	0.97	1.03
Green Bay	474.6	0.98	1.02
Carolina	474.6	0.98	1.02
Miami	469.9	0.99	1.01
Washington	462.4	1	1
Jacksonville	461.6	1	1
San Diego	461.2	1	1
Denver	459.4	1.01	0.99
Dallas	454.9	1.02	0.98
Tennessee	444.7	1.04	0.96
New England	442.1	1.05	0.95
Philadelphia	440.9	1.05	0.95
Seattle	440.5	1.05	0.95
NY Jets	423.4	1.09	0.91
Oakland	415.2	1.12	0.9
Pittsburgh	414.4	1.12	0.89
Kansas City	404.2	1.15	0.87
Tampa Bay	380.8	1.22	0.82
Indianapolis	372.4	1.24	0.8

	2006 Points	Multiplier	Ratio
Detroit	484	0.9	1.11
New Orleans	441.4	0.99	1.01
Minnesota	471.8	0.92	1.08
Cincinnati	525.8	0.83	1.21
Cleveland	444.2	0.98	1.02
Arizona	495.4	0.88	1.14
Atlanta	486.8	0.89	1.12
Houston	476.4	0.91	1.09
Baltimore	397.1	1.1	0.91
St. Louis	429.5	1.01	0.99
San Francisco	507.1	0.86	1.16
Buffalo	409.9	1.06	0.94
Chicago	419.6	1.04	0.96
NY Giants	490.9	0.89	1.13
Green Bay	480.9	0.91	1.1
Carolina	432	1.01	0.99
Miami	432.7	1.01	0.99
Washington	529.1	0.82	1.21
Jacksonville	379.8	1.15	0.87
San Diego	435.2	1	1
Denver	419	1.04	0.96
Dallas	500.6	0.87	1.15
Tennessee	504.2	0.86	1.16
New England	380.3	1.15	0.87
Philadelphia	408.7	1.07	0.94
Seattle	463.6	0.94	1.06
NY Jets	436.2	1	1
Oakland	343.3	1.27	0.79
Pittsburgh	465.3	0.94	1.07
Kansas City	441.4	0.99	1.01
Tampa Bay	491.4	0.89	1.13
Indianapolis	350.8	1.24	0.81

Quarterback Changes:

The last question that we strove to answer in our quest for the most accurate receiver projections possible was whether we should factor in quarterback changes and improvements into wide receiver progression. It would seem logical that younger quarterbacks will improve, older quarterbacks might regress, and new free agent or rookie starters might be better or worse than the QB the previous season, affecting wide receiver output. We determined that the best way to study this question would be to analyze how receivers' production was affected during a season in which the starting quarterback missed 3 to 8 games. This is an ideal situation because it comes as close to holding all other variables constant as one can get in the real world. In theory, everything other than the quarterback ability should be the same.

We examined 15 seasons where a Top-10 quarterback (in fantasy points produced the previous season) was injured and his (presumably much worse) backup had to play between 3 and 8 games. During those 15 seasons, the top receiver scored 1716 points in 144 games while the starting QB was playing and 776 points in the 67 games that the backup played. This computes to 11.9 points per game with the starter and 11.57 points per game with the backup. Given this tiny difference compared with a huge change in quarterback talent (top 10 to second stringer), it seems reasonable that comparatively medium and small changes in talent (quarterbacks improving or starters getting traded) will affect the top receiver very little, if at all. Thus, we decided against trying to include a change in quarterback proficiency in our progressions, as counter-intuitive as that may seem.

Confidence Intervals:

Again, we used the three-factor risk metric to create confidence intervals around the 80% universe of breakout seasons. They are displayed below:

Name	Mean	20% Universe	80% Universe
Randy Moss	225.70	192.00	259.00
Andre Johnson	220.19	182.00	258.00
Larry Fitzgerald	212.91	176.00	249.00
Greg Jennings	207.92	159.00	257.00
Marques Colston	203.92	165.00	243.00
Terrell Owens	200.75	169.00	232.00
Braylon Edwards	197.83	166.00	229.00
Reggie Wayne	196.49	164.00	229.00
Dwayne Bowe	183.84	134.00	234.00
Roy Williams	181.24	138.00	225.00
T.J. Houshmandzadeh	175.18	139.00	211.00
Plaxico Burress	174.59	135.00	214.00
Santonio Holmes	173.10	129.00	218.00
Chad Johnson	173.01	136.00	210.00
Marvin Harrison	171.75	126.00	218.00
Anquan Boldin	166.50	122.00	211.00
Calvin Johnson	155.76	115.00	196.00
Steve Smith	155.59	125.00	186.00
Wes Welker	154.29	114.00	194.00
Torry Holt	154.07	121.00	187.00
Jerricho Cotchery	150.13	109.00	191.00
Brandon Marshall	149.50	109.00	190.00
Bernard Berrian	144.29	100.00	188.00
Lee Evans	141.30	90.00	193.00
Donald Driver	141.23	98.00	184.00
Laveranues Coles	140.10	98.00	183.00
Anthony Gonzalez	139.99	90.00	190.00
Roddy White	131.86	90.00	173.00
Javon Walker	126.61	76.00	177.00
Hines Ward	125.43	85.00	166.00

Name	Mean	20% Universe	80% Universe
Joey Galloway	124.40	82.00	167.00
Reggie Brown	122.60	83.00	163.00
Kevin Curtis	119.95	81.00	159.00
Sidney Rice	114.34	66.00	162.00
Darrell Jackson	107.99	67.00	149.00
Bobby Engram	107.28	61.00	154.00
Jerry Porter	105.68	62.00	149.00
Chris Chambers	105.20	58.00	153.00
Derrick Mason	104.22	58.00	150.00
Mike Furrey	101.92	56.00	148.00
Shaun McDonald	96.97	50.00	144.00
Donte Stallworth	93.25	58.00	129.00
Isaac Bruce	89.18	41.00	137.00
Nate Burleson	86.36	38.00	135.00
Mark Clayton	85.58	46.00	125.00
Muhsin Muhammad	81.22	43.00	119.00
Michael Clayton	47.18	15.00	79.00

Tight Ends

To gather progression data on tight ends, we analyzed the year-to-year progression of active (starting or seeing significant action) tight ends between 1992 and 2007. We analyzed 60 players and created a progression matrix in the same manner as with running backs and receivers, which is as follows:

Year	Catches	Yards	Yards per Carry	TD
1 to 2	0.130	0.127	-0.003	0.297
2 to 3	0.078	0.089	0.005	-0.021
3 to 4	0.039	0.062	0.009	0.062
4 to 5	-0.052	-0.076	-0.015	0.000
5 to 6	-0.044	-0.050	-0.028	-0.104
6 to 7	-0.045	-0.052	-0.062	-0.035
7 to 8	-0.076	-0.061	-0.059	-0.047
8 to 9	-0.038	-0.011	0.024	-0.146
9+	-0.015	-0.067	-0.033	-0.104

As one can observe from this chart, there is a significant increase across the board after a TE's rookie season and further improvement in years three and four. Yards and catches fall slightly in year five, and year six marks the beginning of a slow, steady decline.

Tight ends are unique because there tend to be several (from 3 to 10, depending on the year) players who significantly outperform the rest of the group. These players obviously skew the trend a little bit, and some players (like Tony Gonzalez) inexplicably continue to get better even as they age through their late 20s. Nevertheless, this model will accurately describe tight ends as a group.

Unlike wide receivers, data showed that opposing teams' pass defense had very little to do with tight end production. This is possibly because 'pass defense' is comprised mainly of the skills of cornerbacks and the rush caused by defensive ends, while linebackers typically cover tight ends, who are thrown quick passes that a strong rush cannot prevent.

Confidence Intervals:

Again, we used the three-factor risk metric to create confidence intervals around the 80% universe of breakout seasons. They are displayed below:

Name	Mean	20% Universe	80% Universe
Antonio Gates	139.59	116.00	163.00
Jason Witten	127.78	105.00	150.00
Kellen Winslow Jr	125.06	99.00	151.00
Tony Gonzalez	122.73	103.00	143.00
Chris Cooley	117.40	101.00	134.00
Dallas Clark	112.51	85.00	140.00

Name	Mean	20% Universe	80% Universe
Heath Miller	97.50	83.00	112.00
Owen Daniels	96.89	70.00	124.00
Benjamin Watson	96.84	63.00	131.00
Todd Heap	93.70	69.00	118.00
Jeremy Shockey	90.35	65.00	115.00
Tony Scheffler	89.54	65.00	114.00
Alge Crumpler	88.21	59.00	118.00
Vernon Davis	87.89	64.00	112.00
Desmond Clark	77.60	63.00	93.00
Zach Miller	73.39	58.00	89.00
Donald Lee	73.01	56.00	90.00
Greg Olsen	68.15	45.00	92.00
Randy McMichael	62.04	49.00	75.00
Chris Baker	57.29	43.00	72.00
Bo Scaife	55.75	41.00	70.00
L.J. Smith	53.95	36.00	71.00
Eric Johnson	51.49	37.00	66.00
Ben Utecht	51.05	40.00	62.00
Marcedes Lewis	46.89	33.00	61.00
Jeff King	44.02	32.00	56.00
David Martin	42.18	30.00	55.00
Daniel Graham	38.62	27.00	50.00
Marcus Pollard	31.56	20.00	43.00
George Wrighster	30.67	15.00	47.00
Visanthe Shiancoe	28.82	16.00	42.00
Reggie Kelly	23.08	15.00	31.00

Quarterbacks

Aging:

Projecting quarterback progression is a five-step process. First, as with the players at other positions, we created a matrix to calculate the effect of aging. The matrix was produced by analyzing the year-to-year development of the top 40 seasons of players at each year of play between 1987 and 2007, limited to players with at least 100 pass attempts. A hypothesis test demonstrated that like those of receivers, quarterbacks' projections were much more highly correlated with years in the league than with age, which should make intuitive sense.

Initially, we thought that the best method would be to take a large sample of 30 to 35-year-old quarterbacks and study their year-to-year progressions. However, this would have created a biased sample for two reasons. First, quarterbacks are unlikely to still be playing into their 30s unless they are extremely talented. Thus, our sample would have been biased towards good quarterbacks and would therefore not have been applicable to the population of general quarterbacks starting in 2008. Second, this would have created the problem that the only 'rookies' in our formula would have been players who were rookies 10 to 15 years ago. Many things have changed in that time — the game has become more complex, teams are passing more, and quarterbacks are completing a higher percentage and becoming more dangerous runners, to name a few. It would be wrong to project 2008's young quarterbacks using the progressions of young quarterbacks as long as two decades ago. Taking the top-40 players each year (out of 50 or 60) with at least 100 pass attempts solves many of these problems and gives us a sample with as little bias and as much applicability to modern football as possible, while still factoring in long-term historical trends.

Quarterbacks as a group were modeled based upon their historical aggregate performance, which overall indicates a general improvement during their first few years, followed by a relatively consistent decline. The progression matrix is shown below:

Yr	Attempts	Yards	TD	INT
1 to 2	0.071	0.073	0.149	-0.202

2 to 3	0.008	0.004	0.044	0.010
3 to 4	0.028	0.015	0.112	-0.089
4 to 5	-0.007	0.007	-0.075	-0.004
5 to 6	-0.003	0.013	0.055	-0.011
6 to 7	0.015	0.004	0.009	-0.015
7 to 8	-0.003	-0.007	-0.048	0.007
8 to 9	0.000	-0.012	0.029	-0.061
9 to 10	-0.010	-0.005	-0.004	0.149
10 to 15	-0.023	-0.031	-0.080	0.072

We used two seasons of data for quarterbacks, optimally weighted at 2:1, except in special cases such as rookies, players who missed 2007 because of injury, or those who started in 2007 but were on the bench in 2006. We put the previous season through one year of regression on the matrix, the season two years prior through two years of regression, and then combined them, giving twice as much weight to the more recent statistics. We find this to be an acceptable tradeoff in order to obtain an unbiased sample.

Generally, it appears that barring two aberrations, the quarterback model of progression follows a logical evolution. The two aberrations are a 7.5% decline in TDs in year 4 to 5 and a 2.9% gain in TDs in year 8 to 9. These numbers, while they fail to fit the general progression, do not tell the the entire story. Specifically, consistency in the numbers should be taken into account. We can measure consistency generally by the standard deviation of a given vector of numbers. The standard deviation of TDs comes in at 7.5%, whereas the standard deviation of Yards comes in at 2.7%, making those numbers less than half as volatile. The general progression of quarterbacks is thus most clearly seen in the Yards data.

Strength of Schedule:

Then, just as we did with wide receivers, we adjusted the quarterback's game-by-game stats for strength of schedule. We ranked the 2007 and 2006 NFL pass defenses from top to bottom by number of fantasy points allowed to opposing passing offenses and then calculated the ratio of each team's score to the mean NFL score. For example, in 2007, Detroit had the worst pass defense, allowing 1.31 times as many fantasy points to opposing offenses than the NFL average, while Indianapolis had the best pass defense, allowing just 0.8 times as many points. Thus, a quarterback who scored 22 points against Indianapolis accomplished a task equally as difficult as a quarterback who scored 36 points against Detroit.

After we multiplied each quarterback's stats against the strength of the pass defense he faced in that game, we arrived at the player's 'true' statistics for the year. This lets us know more precisely how good a given quarterback is, which is important in projecting how well he will do in 2008. With this data in hand, we then calculated the strength of schedule for the upcoming season (by using last season's data) and projected future performance. Thus, a player who comparatively faced poor pass defenses over the last two years and who will face a comparatively good set of pass defenses in 2008 will likely see his production drop compared to the strict years-in-league-based projection, as with the Buc's Jeff Garcia. Conversely, a player who faced tough pass defenses and will now face more porous secondaries will likely see his production rise more than otherwise expected. Note that our need to project how ably a quarterback will be able to deal with pass defenses in 2008 is the only reason we care about how 'good' he truly is. A 'good' player who earned 200 points facing extremely tough pass defenses and who will again face extremely tough pass defenses is just as attractive as an 'average' player who earned 200 points by facing extremely easy pass defenses but will again face extremely easy pass defenses in 2008.

Changes in Wide Receiver production:

In order to further fine-tune the projections, we also modeled the change in talent level of wide receivers. Unlike the study that showed quarterback skill has very little effect on receiver production, receiver skill does in fact have an extremely strong correlation with quarterback production. In other words, while receivers produce about the same statistics regardless of who is throwing to them, a quarterback's statistics will fluctuate with the talent of his wide receivers. We modeled each team's first through third receivers' change in talent using independent projections in order to reduce the chance of increased related errors and then adjusted the quarterback's projection.

We found that on average, #1 receivers are responsible for 48% of the fantasy points generated by the trio, #2 receivers are responsible for 32%, and #3 receivers are responsible for 20%, and so we used those weights to determine the overall change. For example, the Rams' receiving corps takes a big hit — #1 receiver Torry Holt is aging, #2 Issac Bruce was replaced with aging #3 Drew Bennett, and #3 Drew Bennett was replaced by a rookie, Donnie Avery. Thus, the receiver production level dropped by slightly more than 5%, which is reflected in the Rams' quarterback projections. On the other hand, the Browns' quarterback will benefit with a 6.5% bonus because of #1 Braylon Edwards' continued improvement as well as the signing of Dante Stallworth to replace Joe

Jurevicius in the #2 spot, who will then in turn replace the mediocre Tim Carter in the #3 spot. Some teams experience both gains and losses, like the Steelers, who will see Hines Ward's aging mitigated by the growth of budding star Santonio Holmes as well as an upgrade at #3 because Limas Sweed was drafted.

Adjusting for Pass Attempts:

The final adjustment that we have included is an optional alteration for pass attempts. We took the player's average attempts per game, multiplied that by a 16-game season to get the projected average attempts for the season, and then replaced that figure with the average number of passes thrown by the team for the last two seasons. For a player like Peyton Manning, this does not change his statistics much because he throws 97% of his team's passes anyway. This does, however, make a huge difference to a player who did not finish games and therefore had a 'game played' on record but was not able to accumulate full statistics.

For example, Tony Romo has four games on record where he did not throw a single pass, and one game where he only threw two. Instead of registering four games with only two passes, this adjustment will project Romo's efficiency statistics (yards per attempt, TDs per attempt, INTs per attempt) onto the number of attempts that were made in those games. Other examples include Kurt Warner and Matt Leinart, who in 2007 had 19 'games played' between them during the 16-game season. Thus, having full games on record during the three games where they split time unfairly punished both players. The 'Adjusted for Pass Attempts' measurement projects what will happen in 2008 if a given player is the full-time starter. As one can see by looking at the ranking chart, the largest differences are with Tony Romo, David Garrard, Jeff Garcia, Matt Schaub, Sage Rosenfels, Jake Plummer, Tarvaris Jackson, Matt Leinart, Kyle Boller, and Brad Johnson. Rosenfels, Plummer, Boller, and Johnson are unlikely to start, so that huge difference is meaningless, but with Romo, Garrard, Garcia, Jackson and Leinart it suggests that there could be some hidden value.

Confidence Intervals:

Again, we used the three-factor risk metric to create confidence intervals around the 80% universe of breakout seasons. They are displayed below:

Name	Mean	20% Universe	80% Universe
Tom Brady	307.50	278.00	337.00
Peyton Manning	287.94	257.00	319.00
Donovan McNabb	268.40	231.00	305.00
Ben Roethlisberger	255.97	223.00	289.00
Tony Romo	255.46	223.00	287.00
Drew Brees	253.98	221.00	287.00
David Garrard	253.45	214.00	292.00
Derek Anderson	249.46	217.00	282.00
Carson Palmer	246.60	214.00	279.00
Jake Delhomme	231.22	188.00	274.00
Brett Favre	229.17	188.00	270.00
Jay Cutler	222.74	189.00	257.00
Jon Kitna	218.96	183.00	255.00
Matt Hasselbeck	217.40	188.00	247.00
Jason Campbell	204.53	168.00	242.00
Eli Manning	203.02	170.00	236.00
Vince Young	201.88	164.00	240.00
Kurt Warner	201.70	155.00	248.00
Jeff Garcia	193.81	161.00	227.00
Philip Rivers	193.56	164.00	224.00
Matt Schaub	182.91	147.00	219.00
Chad Pennington	178.87	135.00	223.00
Mark Brunell	172.07	128.00	217.00
Marc Bulger	171.37	148.00	195.00
Sage Rosenfels	165.81	126.00	205.00

Name	Mean	20% Universe	80% Universe
Damon Huard	158.64	120.00	198.00
Alex Smith	158.44	124.00	192.00
Tarvaris Jackson	155.98	121.00	191.00
Joey Harrington	153.48	113.00	194.00
Matt Leinart	150.36	113.00	188.00
Rex Grossman	146.55	105.00	188.00
Cleo Lemon	142.97	102.00	184.00
Trent Edwards	140.02	104.00	176.00
Charlie Frye	136.59	92.00	181.00
J.P. Losman	134.75	95.00	175.00
Kyle Boller	122.88	83.00	162.00
David Carr	119.18	79.00	159.00
Aaron Rogers	119.00	81.00	157.00
Brian Griese	116.45	74.00	159.00
Matt Ryan	102.00	60.00	144.00
John Beck	70.00	29.00	111.00

Team Defense

Estimation:

Projecting team defenses is in some ways less volatile than anticipating player performance, but it is also much less accurate. Enormous swings in the production of team defenses are lower because unlike individual players, who can get benched, arrested, injured, or can catch fire and have an abnormally good season, a defense plays every game no matter what and also incorporates as many as 20 different players. Therefore, it is not as susceptible to an individually very good or bad season. Unlike an individual player, a team defense cannot score zero points in a game or season but is also equally unlikely to triple its total from the previous year. However, the standard deviation of any one team is actually relatively high — about 18% across the board. It makes sense here to compute a league-wide average for defensive variance, as going back into history to calculate variances for individual teams would not yield any clearer results, as those past teams often bear little resemblance to the current lineup.

Projections of team defenses are historically much less accurate than player projections. This has to do with the often unquantifiable qualities of defensive players (unless you have statistics to evaluate defensive linemen who draw double teams, corners who hold the best receiver to a lower production total than he otherwise would have had, safeties that work well with their teammates and improve overall awareness, etc.), as well as the sheer number of factors involved. Players are traded, they age, and coaches change schemes and adapt, while opposing offensive players improve or regress. It is almost impossible to project defenses from a micro-level and create accurate projections like the ones we have made for offensive skill position players, which we believe to be the best on the market. Instead, the most reliable method is to take a macro approach and aggregate the season projections of several different top experts, coming to a general consensus on draft order.

We weighted three sets of projections in order to calculate sacks, interceptions, TDs, fumbles recovered, and yards allowed. Then, we computed the number of points generated from opponents' scoring, depending on the scoring system used. For example, Yahoo's system gives 10 points for a shutout, 7 for holding the opponent to 1-6, 4 for holding it to 7-13, 1 for 14-20, and 0 for 21-27. Defenses are actually penalized 1 point for allowing 28 to 34 points, and 4 points for allowing more than 35. We then took the number of fantasy points generated from the scores of the games and multiplied by the ratio of the projected total team points allowed in 2008 to the actual total team points allowed in 2007. This gives us a strong mean projection of total fantasy points scored by a team in 2008, although as stated before, the actual outcomes are highly volatile.

The logical conclusion of this data is that while the highest scoring defense would often be a very useful asset to a fantasy football team relative to players at other positions, and therefore would merit a high draft spot, it is very difficult to determine exactly which team it will be. When you group the highest performing teams of years past and compare them with the rest of the teams in the league, there is a significant difference. However, when you group the highest *projected* teams, the results are much weaker. We have done our best to mitigate this situation by aggregating several different picks instead of relying on our formulas alone, hopefully canceling out any biases. We believe that the top teams may be worth a mid-level pick in 2008.

It should also be noted that different scoring systems project defenses differently, as some like the Yahoo system place an emphasis generating turnovers, while the CBS System places a larger emphasis and a greater proportion of fantasy points on holding the opposing offense to a low point total. A player should be careful to select the defense that best matches his or her scoring system.

Confidence Intervals:

Again, we used the three-factor risk metric to create confidence intervals around the 80% universe of breakout seasons. They are displayed below:

Name	Mean	20% Universe	80% Universe
San Diego Defense	169.33	142.00	197.00
Minnesota Defense	161.07	137.00	185.00
Pittsburgh Defense	153.59	134.00	173.00
Seattle Defense	147.08	126.00	168.00
New England Defense	143.28	120.00	167.00
Indianapolis Defense	139.35	116.00	162.00
Jacksonville Defense	138.62	117.00	160.00
Dallas Defense	134.95	115.00	155.00
Tennessee Defense	133.69	115.00	153.00
Tampa Bay Defense	133.20	115.00	151.00
Chicago Defense	131.13	115.00	147.00
Green Bay Defense	129.91	109.00	151.00
Washington Defense	124.22	108.00	141.00
NY Giants Defense	120.38	105.00	136.00
NY Jets Defense	118.97	99.00	139.00
Philadelphia Defense	118.77	102.00	135.00
Baltimore Defense	117.66	103.00	132.00
Cleveland Defense	112.20	97.00	127.00
Buffalo Defense	109.23	93.00	125.00
Atlanta Defense	107.91	89.00	127.00
Carolina Defense	101.61	89.00	115.00
Kansas City Defense	99.99	83.00	117.00
Arizona Defense	99.07	86.00	112.00
Oakland Defense	95.01	82.00	108.00
Houston Defense	94.30	82.00	107.00
San Francisco Defense	91.55	75.00	109.00
New Orleans Defense	91.05	79.00	103.00
Cincinnati Defense	89.76	78.00	101.00
Denver Defense	86.15	75.00	98.00
St. Louis Defense	77.90	63.00	93.00
Detroit Defense	77.61	63.00	92.00
Miami Defense	71.99	62.00	82.00

Kickers

Estimation:

Projecting kickers is a similar task to that of modeling team defenses because a kicker's production is a function of how good his offense is at getting within range of field goals and scoring touchdowns, which require extra points. Thus, Kickers' scores are a reflection of the production of the entire offensive unit and are therefore extremely variable. This, combined with the fact that the effects of aging are almost negligible on kickers, allows us to best model kicker production with a weighted average of several different projections rather than using our numbers alone, just as we did when we projected team defenses.

Like team defenses, kickers experience a lot of volatility and inverse relationships — more touchdowns means more extra points but generally fewer field goals. Therefore, even if we know that an offense is going to be much better or worse than the previous year, we cannot project whether the kickers' points will go up or down. For example, Sebastian Janikowski saw the number of extra points he made last year practically double from 16 to 32, while his field goal attempts also increased from 18 to 23. However, from 2002 to 2003, extra points for the Raiders were almost halved, dropping from 50 to 29, while field goals remained nearly constant from 26 to 22. Every season in his eight year career, Janikowski's field goals moved in the same direction as his extra points — more FG's and XP's or less of both each year. On the other hand, during 8 out of the 10 seasons of Ryan Longwell's career, extra points and field goals moved in the *opposite* direction of one another. This demonstrates the independent movement of and therefore the difficulty in predicting kickers' fantasy scores.

Furthermore, even the bonus points attributed to long kicks are difficult to model. With few exceptions, it seems that the notion of a specific kicker being particularly accurate or having a particularly long leg is more of a backwards-looking induction rather than an actual, proven correlation. Almost all kickers in the NFL can make a huge percentage of their kicks if given the correct circumstances, and their results fluctuate according to the laws of randomness, with almost no correlation from one season to the next. For example, the aforementioned Sebastian Janikowski was supposed to have the biggest leg of any kicker in the modern era, but he has converted less than half (16 out of 34) of his 50+ yard attempts, while the small in stature Martin Gramatica, initially regarded as a short range, high accuracy kicker, has converted almost two thirds (16 out of 25) of his career 50+ yard attempts.

We aggregated the forecasts of four fantasy experts and ranked the kickers from best to worst. The difference between the highest projected kickers in the league (Nick Folk and Nate Kaeding at 132.5) and the 10th rated kicker in the league (Mason Crosby at 120) is only 12.5 fantasy points over the course of the season. As a group, no kicker in the NFL missed a single kick from less than 19 yards last year. Further, all kickers made more than 95% of their kicks from less than 29 yards. This rendered the unique modifier for Fox's system, which subtracted a point for a short miss, as relatively meaningless. Bonuses for long kicks in the other systems are already modeled into our picks, as the fantasy experts whose projections we aggregated have taken into account which kickers are very slightly (and the difference is quite small) more likely to be able to take and make longer kicks.

In order to calculate volatility, we did something different from all of our other positions, which was to examine the spread of the projections that we aggregated. If all four of the projections were relatively close together, then we can make a reasonable assumption that the kicker's production will be somewhere near that forecast. If they are all over the board, then we can still take the mean prediction but will acknowledge the high variance due to the inconsistencies.

Because there is typically a far greater point differential in the projections of players at other positions, and because kickers are relatively consistent and unlikely to be injured, we recommend taking a kicker extremely late in the draft process. The rationale for this is that even though kickers may have a reasonably high mean score compared with mediocre receivers or running backs, the fact is that kickers are difficult to project accurately, and even then, there is not a huge amount of difference between them. For example, the highest rated kicker in 2008 is Nick Folk, a second-year player, while several proven veterans cannot even find a job. Even if you knew who the highest scoring kicker was going to be, you would likely still gain greater expected value by adding to your pool of backup running backs and wide receivers to ensure against starter injury or to potentially pick an unexpected high performer.

Confidence Intervals:

Again, we used the three-factor risk metric to create confidence intervals around the 80% universe of breakout seasons. They are displayed below:

Name	Mean	20% Universe	80% Universe
Nick Folk	143.40	138.00	149.00
Nate Kaeding	143.06	135.00	151.00
Adam Vinatieri	141.84	134.00	149.00
Shayne Graham	139.87	132.00	147.00
Stephen Gostkowski	138.28	130.00	147.00
Phil Dawson	131.75	126.00	137.00
Josh Scobee	131.03	119.00	143.00
Neil Rackers	130.31	122.00	139.00
Mason Crosby	130.15	122.00	138.00
Rob Bironas	124.21	119.00	129.00
Kris Brown	124.18	120.00	128.00
Robbie Gould	121.87	114.00	129.00
Josh Brown	121.68	114.00	129.00
Jeff Reed	121.68	116.00	127.00
Ryan Longwell	117.65	111.00	124.00
David Akers	116.68	109.00	125.00
Lawrence Tynes	115.96	109.00	123.00
Martin Gramatica	112.77	102.00	124.00
Matt Bryant	112.43	103.00	122.00
Jason Hanson	108.46	94.00	123.00
John Kasay	108.21	104.00	113.00
Rian Lindell	107.05	102.00	112.00
Matt Stover	106.37	96.00	117.00
Matt Prater	103.86	83.00	124.00
Mike Nugent	101.99	95.00	109.00
Jason Elam	101.52	94.00	109.00
Joe Nedney	101.40	94.00	109.00
Shaun Suisham	101.11	93.00	109.00
Jay Feely	96.18	91.00	102.00
Sebastian Janikowski	94.14	90.00	99.00

Rookie and Youth Estimation

In order to have a complete data set for every position, a number of rookies and young players need to be taken into account. Some of these players will be picked despite the risk that can surround untested individuals. This section is devoted to the estimation of those players. The estimates are presented by position.

Running Backs:

There are a number of running backs this year that are likely to appear in Fantasy Football lineups. These rookies and youth players do not have enough historical data to allow us to estimate their performance numbers through statistical techniques, and, as a result, we applied logical relative criteria to their estimates. First, there is Darren McFadden. He is the overall #1 running back pick, a contender for the Heisman every single year at Arkansas, and basically carried the team single handedly. The Raiders already have a powerful rushing attack, but they also have two other backs to compete with McFadden for time. Those players are Justin Fargas, a proven, middle-of-the-road veteran, and Michael Bush, a former projected first-round pick who was injured in college, slipped to the third round, and has not yet proven himself in the NFL. As a result, McFadden has an extremely high standard deviation as a rookie uncertain of his role. His points projection puts him at 24th among running backs at 141, with a standard deviation of 48.

Next, McFadden's backfield mate at Arkansas, Felix Jones, should get a lot of burn in Dallas. Marion Barber will get the majority of the snaps, but Jones could carry a decent portion of the load as well. He is projected as the backup, with 105 points and a standard deviation of 37.

Rashard Mendenhall, who is a rookie, should see the field plenty in his first year in Pittsburgh, with their run-oriented offense. He is projected as the backup to Willie

Parker but will doubtless get a fair share of carries. Mendenhall was widely considered as a 'steal' when the Steelers grabbed him with their first pick. If he is as explosive as some pundits predict, it is possible that he could be splitting carries or even starting by the end of the season. He is projected at 116 points with a standard deviation of 43.

Kevin Smith was drafted in the third round and has a great shot at starting for Detroit, as his only competition is the somewhat mediocre Tatum Bell, who underwhelmed last year. Smith ran for the second-most yards in a season of any running back in NCAA history (trailing only the legendary Barry Sanders), and is a workhorse who does everything well but isn't outstanding at any one aspect. We estimate a 70% chance of him starting in the fall with a projected 136 points and a standard deviation of 36.

Matt Forte is a huge question mark. He put up very impressive statistics but did so at Tulane, an afterthought in college football. However, he continued to produce even when they played top-level teams such as LSU, so there is reason to suspect that he may continue to be successful in the NFL. Chicago released Cedric Benson, so it seems as if the coast is clear for Forte to start and have a successful year. He is projected at 134 points with a standard deviation of 32.

Jonathan Stewart is one of the more ridiculous physical specimens that we have seen in a long time, but during his time at Oregon he was plagued by criticism that he lacked vision. He is a big, powerful player and he should fit well within Carolina's smash mouth system. He will likely split time with scat back DeAngelo Williams, but it remains to be seen who will get the bulk of the carries. He is projected at 138 points with a standard deviation of 42.

Wide Receiver:

There was only one rookie or youth addition to our WR lineup. That was Calvin Johnson. He is a 2nd-year player who is projected to do extremely well this season. We rank him 17th with a projected points total of 155.76 and a standard deviation of 19.5.

Quarterbacks:

As previously noted, we estimated projections for players who are either young or rookies without any data if we believe that they are likely to contribute to fantasy leagues. First, there is Aaron Rogers.

After several years of patiently waiting in the wings, Aaron Rodgers is finally the starter in Green Bay, and despite the recent drama surrounding Brett Favre's return, that is unlikely to change. We have him projected for 19 touchdowns, 14 interceptions, 3350 passing yards, 110 rushing yards and a rushing touchdown. He is projected at 199 points for Yahoo and ESPN and 237 for CBS Sportsline, NFL, and Fox Sports. His standard deviation is expected to be 45.

Widely regarded as the top quarterback in the nation last year at Boston College, Matt Ryan has been handed an enormous contract and the keys to the Falcons' offense. He might sit out a few games to get acclimated but will likely at some point during the season take over as the Falcons' starter. Things will be rough for him, though, as the Falcons' are not exactly bristling with offensive weapons, much like Ryan's team during his last year in college. Ryan will probably have a pretty rough rookie season. We have him projected for 13 touchdowns, 14 interceptions, 1950 yards, and no significant rushing stats. He is projected at 102 points for Yahoo and ESPN and 128 for CBS Sportsline, NFL, and Fox Sports. His standard deviation is expected to be 37.

John Beck is facing competition from veteran Josh McCown and rookie Chad Henne but is still projected as the starter, at least for now. He is improving steadily, and if given free reign over the offense, should put up some reasonable numbers. That said, we find it unlikely that he finishes the season as the Dolphins #1 man at QB. We project him at 8 TDs, 7 INT, 1300 yards and no significant rushing stats. That puts him at 70 points for Yahoo and ESPN and 199 for CBS Sportsline, NFL, and Fox Sports.

SECTION II: Qualitative Adjustments to Players

Statistics can only go so far in projecting player performances. Players retire, move up or down on the depth chart, get hurt in the offseason, or have recognizably anomalous seasons that historically are almost impossible to replicate. Three such potentially anomalous seasons are Adrian Peterson's 4^{th} best (adjusted for 16 games) season *ever* by a rookie running back, Tom Brady's record-breaking 50 passing touchdown season, and the highly related, also record-breaking 23 receiving touchdown season of Randy Moss.

Adrian L. Peterson

One such season was that of Minnesota running back Adrian L. Peterson. In just 14 games, Peterson ran for 1341 yards, added an additional 268 receiving yards, and scored 13 touchdowns. Adjusted for a full season, that would have been 1838 total yards and 15 touchdowns, which would have been good for 274 fantasy points, earning the fourth highest total. Peterson trailed only the 278 points of the more experienced Joseph Addai, the 301 of veteran Brian Westbrook, and the 303 of veteran Ladanian Tomlinson.

The standard projection model calls for an approximate increase of 10% on rookie statistics during a running back's second year. However, it seems that Peterson is hardly the prototype of a rookie running back and that he would better fit a model of past players in a similar position. To generate a projection for Adrian Peterson, we instead conducted an examination of the rookie to sophomore progressions of the top 20 rookie running backs of the last 25 years, removing those who missed significant time their second season. The data is as follows:

PLAYER	YEAR	AGE	SEASON	GAMES	FAN POINTS	FP/16 GAMES	PERCENT CHANGE
Eric Dickerson	1983	23	1	16	341.2	341	
Eric Dickerson	1984	24	2	16	306.4	306	-11.36%
Edgerrin James	1999	21	1	16	315.9	316	
Edgerrin James	2000	22	2	16	338.3	338	6.62%
Clinton Portis	2002	21	1	16	289.2	289	
Clinton Portis	2003	22	2	13	274.5	338	14.40%
Fred Taylor	1998	22	1	15	266.4	284	
NO TOP100 2ND YEAR							NO DATA
Curtis Martin	1995	22	1	16	264.8	265	
Curtis Martin	1996	23	2	16	250.5	251	-5.71%
Kurt Warner	1983	22	1	16	261.4	261	
NO TOP100 2ND YEAR							NO DATA
Barry Sanders	1989	21	1	15	259.2	276	
Barry Sanders	1990	22	2	16	274.4	274	-0.76%
Mike Anderson	2000	27	1	14	255.6	292	
NO TOP100 2ND YEAR							NO DATA
Marshall Faulk	1994	21	1	16	252.4	252	
Marshall Faulk	1995	22	2	16	239.3	239	-5.47%
Herschel Walker	1986	24	1	16	241.4	241	
Herschel Walker	1987	25	2	12	208.6	278	13.21%
Adrian Peterson	2007	22	1	14	238.9	273	
NOT INCLUDED							NO DATA
Maurice Jones-Drew	2006	21	1	16	227.7	228	
Maurice Jones-Drew	2007	22	2	15	171.5	183	-24.47%
LaDainian Tomlinson	2001	22	1	16	220.3	220	
LaDainian Tomlinson	2002	23	2	16	307.2	307	28.29%
Robert Edwards	1998	24	1	16	216.6	217	
NOT TOP 100							NO

PLAYER	YEAR	AGE	SEASON	GAMES	FAN POINTS	FP/16 GAMES	PERCENT CHANGE
2ND YEAR							DATA
Ickey Woods	1988	22	1	16	216.5	217	
NO TOP100 2ND YEAR							NO DATA
Jerome Bettis	1993	21	1	16	209.3	209	
Jerome Bettis	1994	22	2	16	155.8	156	-34.34%
Ricky Watters	1992	23	1	14	207.8	237	
Ricky Watters	1993	24	2	13	193.6	238	0.33%
Willis McGahee	2004	23	1	16	207.7	208	
Willis McGahee	2005	24	2	16	172.5	173	-20.41%
Eddie George	1996	23	1	16	203	203	
Eddie George	1997	24	2	16	186.3	186	-8.96%
Jamal Lewis	2000	21	1	16	202	202	

Thus, for the 13 players who (as we project AP to) finished among the top 100 second-year backs, we saw slight improvements by Edgerrin James, Clinton Portis, Herschel Walker, LaDanian Tomlinson, and Ricky Watters. We saw slight declines from Erick Dickerson, Curtis Martin, Barry Sanders, Marshall Faulk, Maurice Jones-Drew, Jerome Bettis, Willis McGahee and Eddie George. Overall, the average change from year 1 to year 2 for the greatest rookie running backs of the last 25 years was actually a 3.471% *decrease* in production. This is reasonable because while most of these players went on to be great, these outstanding seasons are likely to experience some regression towards the overall mean.

Based on what data we have, it seems that these players are all reasonable comparisons to Adrian Peterson, and the probability distribution of the chances of becoming a star, a solid career player, or a one or two-year flame-out are fairly reasonable. Thus, we actually predict Adrian Peterson's second year stats to be:

PLAYER	RUSHES	YARDS	REC YARDS	TDs	FAN PTS
Adrian L. Peterson	262.6	1479.4	295.7	14.3	263.3

Tom Brady and Randy Moss

Tom Brady and Randy Moss are coming off record-setting seasons, where they set records for touchdowns thrown (50) and most receiving touchdowns in a season (23). This is an interesting scenario, because unlike other QB/WR tandems who got together early, and aged into greatness (Peyton Manning/Marvin Harrison, Joe Montana/Jerry Rice) they were both stars before coming together later in their careers. Thus, unlike a particularly good Manning/Harrison or Montana/Rice season, which stands out from the rest in terms of production, it is difficult to determine whether the outrageously successful 2007 Brady/Moss combination was an abnormally good year, or simply an average season for two of the best players in NFL history combined.

When you further examine their 2007 seasons, you find that while they are extraordinary, the number of completions, catches, and yards that Brady and Moss experienced were very good but not record-setting. In 2008, Moss had 'only' the 57th most catches in a season in the last 25 years and the 24th most yards. (In fact, he had more catches *and* yards in 2003). Brady had the 5th most completions and the 3rd most yards. What *was* incredible about their 2007 year was the number of touchdowns that were converted. Thus, we initially sought to examine the 20 most prolific QB to WR touchdown seasons, and see how those tandems fared the next year, and then project this onto the Brady/Moss combination. The 20 seasons are as follows:

NAME	POS	YR	AGE	EXP	G	REC	RECYD	YD/REC	RECTD	PERCENT CHANGE
Randy Moss	wr	2007	30	10	16	98	1493	15.23	23	
NO 08										N/A
Jerry Rice	wr	1987	25	3	12	65	1078	16.58	22	
Jerry Rice	wr	1986	24	2	16	86	1570	18.26	15	-48.86%
Mark Clayton	wr	1984	23	2	15	73	1389	19.03	18	
Mark Clayton	wr	1985	24	3	16	70	996	14.2	4	NOT USED
Sterling Sharpe	wr	1994	29	7	16	94	1119	11.9	18	
RETRIED IN 1995										N/A
Cris Carter	wr	1995	30	9	16	122	1371	11.24	17	
Cris Carter	wr	1996	31	10	16	96	1163	12.1	10	-41.18%
Randy Moss	wr	2003	26	6	16	111	1632	14.7	17	
Randy Moss	wr	2004	27	7	13	49	767	15.65	13	-5.88%
Randy Moss	wr	1998	21	1	16	69	1313	19.03	17	
Randy Moss	wr	1999	22	2	16	80	1413	17.66	11	-35.29%
Carl Pickens	wr	1995	25	4	16	99	1234	12.46	17	
Carl Pickens	wr	1996	26	5	16	100	1180	11.8	12	-29.41%
Jerry Rice	wr	1989	27	5	16	82	1483	18.09	17	
Jerry Rice	wr	1990	28	6	16	100	1502	15.02	13	-23.53%
Braylon Edwards	wr	2007	24	3	16	80	1289	16.11	16	
NO 08										N/A
Muhsin Muhammad	wr	2004	31	9	16	93	1405	15.11	16	
Muhsin Muhammad	wr	2005	32	10	15	64	750	11.7	4	NOT USED
Terrell Owens	wr	2001	28	6	16	93	1412	15.18	16	
Terrell Owens	wr	2002	29	7	14	100	1300	13	13	-7.14%
Marvin Harrison	wr	2001	29	6	16	109	1524	13.98	15	
Marvin Harrison	wr	2002	30	7	16	143	1722	12.04	11	-26.67%
Marvin Harrison	wr	2004	32	9	16	86	1113	12.94	15	
Marvin Harrison	wr	2005	33	10	15	82	1146	13.98	12	-14.67%
Randy Moss	wr	2000	23	3	16	77	1437	18.66	15	

NAME	POS	YR	AGE	EXP	G	REC	RECYD	YD/REC	RECTD	PERCENT CHANGE
Randy Moss	wr	2001	24	4	16	82	1233	15	10	-33.33%
Terrell Owens	wr	2007	34	12	15	81	1355	16.73	15	
NO 08										N/A
Jerry Rice	wr	1995	33	11	16	122	1848	15.15	15	
Jerry Rice	wr	1996	34	12	16	108	1254	11.6	8	-46.67%
Jerry Rice	wr	1993	31	9	16	98	1503	15.34	15	
Jerry Rice	wr	1994	32	10	16	112	1499	13.4	13	-13.33%
Jerry Rice	wr	1986	24	2	16	86	1570	18.26	15	
Jerry Rice	wr	1987	25	3	16	64	1306	20.4	9	-40.00%
Andre Rison	wr	1993	26	5	16	86	1242	14.44	15	
Andre Rison	wr	1994	27	6	14	81	1088	13.4	8	-39.05%

As is apparent from looking at the seasons, Randy Moss, Jerry Rice, and Marvin Harrison have a heavy presence. We dropped several data points to get a more accurate picture — Muhsin Muhammad's 2004 season was a statistical fluke (he scored 16 touchdowns in '04 and never scored more than eight in the rest of his 12-year career) and three of the 20 best seasons (from Owens, Moss, and Edwards) actually occurred in 2007 and therefore cannot be compared. Lastly, Shannon Sharpe retired after his career year in 1994 where he scored 18 touchdowns. Also, note that the percentage change factors into games played and uses touchdowns per 16 games.

Using the rest of the data, which includes many receivers who, like Moss, have been dominant for many years, we found that in those 14 seasons the receivers scored 228 touchdowns, and in the following year they as a group netted only 158 touchdowns, or 69.3% of their previous year's total. In 2007, Brady to Moss scored 22.7 strength-adjusted scores, and after our initial ageing regressions we projected 19.5 Brady to Moss touchdowns in 2008. In light of this new analysis, we are revising this projection to 15.7 touchdowns, while keeping our yardage and catches predictions constant. Thus, the new projections for Brady and Moss are as follows:

PLAYER	CATCHES	YARDS	TDs	FAN PTS
Randy Moss	88	1314.9	15.7	225.7

PLAYER	YARDS	TDs	INTs	FAN PTS
Tom Brady	4422.5	37	8.4	307.5

Rejected Adjustments

Sports analysts often draw spurious conclusions — attempting to assign reason to things that happened due to random variance, or asserting seemingly logical statements that for whatever reason do not hold up to statistical analysis. We studied many other qualitative adjustments but did not find enough conclusive evidence to include them. For example:

- We hypothesized that teams from warm weather locations would do poorly when they played in a cold weather, outdoor stadium after October. We intended to adjust the strength of schedule for these more difficult games, but when we analyzed the data we could find no evidence to support this hypothesis.
- Similarly, we found no correlation to the common wisdom that teams traveling from West to East play worse than teams traveling from East to West.
- We considered manually raising Vince Young's projections given the return of Mike Heimerdinger, who had great success with Steve McNair as well as the

general feeling among analysts that he was 'due' for a breakout year. However, when we examined his body of work and compared with other similar players, we found that no player that has been as inefficient as Young has over the last two years has ever become a top-15 quarterback the following year.

- We considered lowering Greg Jennings' projections given the retirement (or trade) of Brett Favre, one of the greatest quarterbacks of all time. Yet, citing the study discussed in the quarterback section earlier, where we discovered that #1 wide receivers tend to be just as effective with a backup quarterback as they are with a top 10 (in fantasy points) starter, we decided against it. Furthermore, it seems probable that Aaron Rodgers will play much better than an average backup quarterback (which is why he is now the starter) and, thus, any notions of a falloff due to quarterback talent discrepancies proved to be ill-conceived.
- We also hypothesized that home players performed better on Monday and Sunday night games, and again, found no evidence.
- Along the 'he will do better once he gets a few games under his belt' line of thought, we hypothesized that rookie players would perform better towards the end of their first season than the beginning. Again, we found no evidence of this, and also amusingly revisited the Ryan Leaf saga, where he won his first two starts with style, then proceeded to go 2-21 in his next 23 games.

SECTION III: Portfolio Theoretic Approach to Drafting

Introduction

Every investment process consists of two broad tasks. First, you must assess the risk and the expected return attributes of the entire set of possible investment vehicles. After characterizing the individual assets, you must form the optimal portfolio of the individual assets. This second task involves the determination of the best risk-return opportunities available from the feasible investment portfolios. Further, you must choose the best portfolio from this feasible set. If you follow these steps, you will create the best possible portfolio that yields the maximum return with the least risk. In this case, a three-factor model is utilized in order to generate the best approximation of risk.

While this investment process is commonly used for constructing securities portfolios, it can be applied equally well to the construction of fantasy football teams. Fantasy football teams are composed of a set of players (investment vehicles) whose performance determines the amount of points the fantasy team receives (return). The risk involved is that the player's performance varies from week to week and from season to season. You must consider all of these variables when selecting your team.

To create the best fantasy football team, you must first sort through the individual football players and characterize their performance based on the number of points and the volatility of the points they will contribute to the fantasy team. Then, you must devise different fantasy teams comprised of these individual players, accounting for the interactions among the players (e.g. you would not want a fantasy football team including only players from the San Diego Chargers). Finally, you would compare the different teams, examining the expected number of points they will produce and their variability. Depending on the prize structure in your league, you either want to choose the team that maximizes the expected earned points while minimizing the variability of the expected earned points, or you want to choose the team that simply maximizes expected points.

How do you value individual players? More importantly, how do you value a fantasy football team? How would the addition of Reggie Bush impact the performance of your team? Clearly, you want the most points possible, but how do you construct a team to optimize the amount of points you receive? To answer these questions, portfolio theory declares you use a wide variety of tools: expected return, risk, correlation, and the Sharpe ratio.

Expected Return

The most important statistical concept in optimizing any portfolio is that of the expected return, which is the probability-weighted average of all possible outcomes. It is very similar to a mean or average. The difference is that the latter two names are used if you are working with past outcomes, while the expected return is the correct terminology if you are working with future values. In a world of uncertainty, you often encounter scenarios where the return will be greater than expected and scenarios where the return will be less than expected. For example, here is a simplified table hypothetically describing how Randy Moss might perform in game 1:

Performance	Points Earned	Probability
Better than Average	90	0.2
Average	50	0.6
Worse than Average	30	0.2

Expected Points Earned = 90(.2) + 50(.6) + 30(.2) = 54

Theoretically, if Randy Moss were to play game 1 infinitely many times, the average points he would earn for a fantasy football team would be 54. Of course, in this scenario distribution, Randy Moss could never earn 54 points — the expected value does not need to be a possible realization (actual outcome) of the event.

To make it easier to work with uncertainty, statisticians have invented the concept of the random variable. It is a variable whose outcome has not yet been determined. In the Randy Moss example above, the random variable would be "points earned" that would acquire the value 50 with a 60% probability, 90 with a 20% probability and 30 with a 20% probability. The "expected points earned" would be 54. After game 1 is complete, "points earned" is no longer a random variable; "points earned" becomes a realization because you now know their true value. Also, if you are certain about an outcome before it happens, then the actual realization and the expected value are the same. If the Bears have a bye week, you know with 100% certainty that the Bear's defense will not directly contribute to your earned points.

Overall, future values or expectations are uncertain. As a result, they must be estimated. While there are a variety of methodologies available for the estimation of future values, only a few are regularly utilized on a large scale. For the purposes of this report, historical data projections are used. These data provide the basis for reliable estimates in the case where large amounts of volatility are present. On the margin, an additional pass, touchdown, or interception has a small effect on the final estimate. There are, of course, seasons where exceptional performance can occur; however, they are not as far from the average season as the 1987 financial crash was from the average daily return. Further, for football players, you only need a couple of years of statistical data to make reliable projections about the players' future performance.

Keep in mind that you are not interested in what happened in past years but rather what will happen this year. Historical data is the best data to use, but you should not trust it blindly. There are many factors that are not incorporated in historical data that may be important. Events that did not occur in the past may occur in the future. Players age, are injured, have personal problems, or perform badly for no foreseeable reason. These variables are not incorporated in the historical data but could still affect players' future performance. However, for football players, past recent data provides a good representation of their expected future performance.

In the drafting model, the expected return is represented as the expected points that a player will contribute to the fantasy football team. The drafting model uses statistics from past seasons to predict future performance. It uses statistical methods to assess historic data and applies time series analysis and logical rigor to predict the number of fantasy points players will contribute based on the fantasy football league's scoring rules.

All of the variables included in the scoring rules are taken into account. In addition, other predictor variables not related to the scoring rules are used to predict future performance. The health of the player is one such variable. It is based upon the number of injuries, their number of missed games, and their historical health record. The age of the player is another variable that affects their performance. It affects performance differently for different positions. For instance, wide receivers generally improve over the first couple of playing years but then experience a decrease in performance in their later years. Therefore, being young would improve the receivers expected performance, but being old may actually decrease the receivers expected performance relative to what the historic data would suggest. Using historical predictors based on past playing statistics and incorporating predictor variables based on tested statistical models increases the probability that the expected number of points players contribute to the fantasy football team is as accurate as possible.

The next concept that is related to expected return is the concept of the requisite weighting of a given factor in a portfolio. Typically, in finance, the common way to think about weighting is the percentage of capital allocated to a given investment divided by the total capital allocated to the portfolio. It is tempting to apply the same methodology to the point system described in the first section of this report. However, a simple logical example will provide an alternative that is better suited to the application of a portfolio theoretic model to football.

Imagine a two-player portfolio where Player A is expected to yield 600 points this season, while Player B is expected to yield 400 points. According to the typical financial methodology described above, Player A would receive a 60% weight (600/(600+400)) while Player B would receive a 40% weight (400/(600+400)). This portfolio would then have an expected return of 60%*600 + 40%*400 = 520. However, logic dictates that this is incorrect. Specifically, we know that a portfolio of two players (such as Player A and Player B) would have an expected return of 1,000 points (rather than 520). Therefore, we can see that for a portfolio theoretic model of fantasy football players, the appropriate weight for every player is 100%.

Risk

In portfolio theory, you need to measure the (average) reward that you would expect to receive from an investment. Usually, you use the expected return of the investment as your measure. You also need to measure a second characteristic of investments: risk. Thus, you need summary measures of how the possible outcomes are distributed around the expected value, or some other methodology or combination of methodologies, to assess risk.

Risk is the measure of the variability surrounding your expected number of points. For the purposes of building the most accurate and comprehensive risk model available, we took three factors into account: (1) standard deviation; (2) projection differences; and (3) latent risk factors. These determinants of the three-factor model will be discussed in order below.

Standard Deviation:

Of the three measurements of risk utilized for this model, the most common assessment of risk is the standard deviation, which takes the square root of the sum of squared deviations from the mean. The standard deviation is the square root of the variance, which is the sum of the squared deviations from the mean. For example, here is how you would calculate risk, denoted σ, for the simplified Randy Moss example shown above:

$$Var = E((X - E(X))^2)$$
$$Var = .2(90-54)^2 + .6(50-54)^2 + .2(30-54)^2 = 384$$

$$\sigma = \sqrt{Var} = E((X - E(X))^2)$$
$$Risk = \sqrt{.2(90-54)^2 + .6(50-54)^2 + .2(30-54)^2} = 19.6$$

Together, the expected return and the standard deviation provide one way for you to characterize your investment. You would say that you expect Randy Moss to score 54 points (expected return), plus or minus 19.6 (standard deviation). More specifically, the standard deviation suggests that Randy Moss will yield 54 points $\pm$ 19.6 points with 68.2% certainty, 54 points $\pm$ 39.2 points (twice the standard deviation) with 95.4% certainty, and 54 points $\pm$ 58.8 points (three times the standard deviation) with 99.6% certainty.

Historical risk is similar to historical return in which you use past values to obtain meaningful measurements of future values. While historical returns tend to be mediocre predictors of future performance, historical risks tend to be good predictors of future volatility. Thus, if Tom Brady's performance was consistent throughout the past two years, chances are that his performance will remain consistent for years to come. This does not mean that Tom Brady will continue to play well. If he plays poorly for a couple games and his risk remains low, he will most likely continue playing poorly.

Projection Differences:

Another way to measure the variability attributed to a given player earned is the calculation of projection differences in the development of the three-factor model to generate comprehensive assessments of risk for each player. Projection differences are calculated as the difference between the projection of expected points per game for the current season and the total number of points per game earned the previous year for a given player. In addition to standard deviation, which takes into account within-season variability, Projection Differences (PDs) are useful in assessing our best guess of between-season variability in expected points per game earned by a given player. This methodology is also unique to our model and is explained thoroughly in this section. We applied this methodology, as with the latent factor analysis, to the four main positions that have the largest impact on the value of a Fantasy Football team (QB, RB, WR, and TE).

The calculation of the PDs is a multi-step process. First, earned points from the previous season are calculated. Then, the projection of points for this season is used as a base from which the previous season's projections are subtracted:

Projection Differences = Expected points per game (t) – Actual Points per game Earned (t-1)

This calculation yields a definition of risk that is scaled to the range of points earned per game by each player in each position. Players with large PDs could have experienced a recent year, or years, that results in our model regressing that player to his mean, thus boosting his PD. Take the case of Marvin Harrison, for instance. He is projected to experience a 75% improvement in points per game from the previous season. Overall, Harrison has been an exceptional player his whole career but played only five games in '07. He put up reasonable stats through three games, then was injured early in game 4 after catching only one ball for 8 yards. He then tried to play in game 7, put up a mediocre 3 catches for 16 yards, and didn't play the rest of the season. Thus, his '07 stats are highly misleading. Harrison did not finish two out of the five games he played, and he only had one touchdown, which creates a skew in the projections. Given that he is healthy for '08, our models see no reason to project that factors other than aging will affect him. Therefore, we put his '06 season through two regressions and also still factored in the '07 sample with a 20% weight. We project him at 77/1117/10.0, which is probably a conservative estimate given that he went for 95/1366/12 in '06 and still has a healthy Peyton Manning throwing to him.

There are also a number of other "outliers" (here defined as > two standard deviations from the mean) that warranted our investigation. For instance, Mark Clayton, who is a 26-year-old receiver who should be consistently producing, experienced a steep decline in production in his 3rd season in 2007. He caught 48 passes and did not score a single touchdown. As a result, this is likely a bout of negative variance, so we project him to bounce back and experience a 61% increase in points per game over the coming year.

While both Harrison and Clayton are outliers, we chose not to address their numerical contribution to the model because the qualitative assessments of their numbers matched what was being projected. More serious deviations from the average were addressed statistically through normalization. The normal, or central tendency, of a given series of data is known as the average. Thus, for serious outliers whose explanations cannot be qualified, their numerical results were replaced with the average for their position. For example, Charlie Frye played one game as QB last season and threw an interception. As a result, he recorded only 0.46 points. We have him projected to score a more normal 100+ points this season, which leaves his percent increase in points per game at 1,756%. As a result of the decimal point and normal progression, we assigned his PD to be a more reasonable 1.66 (indicating a 13% improvement from what "should" have been the case had he not only played one game and threw for a loss). Other instances similar to this, where the previous season's points were extremely low, thus lending themselves to overstated improvement, were corrected through the process of normalization. Finally, we need to deal with the rare instances of rookies and those individuals who have played very few games. These players do not have previous data from which projection differences could be calculated and so some kind of estimate would need to be ventured. We decided to simply give them projection differences of twice the average because, given their short or non-existent NFL playing history, we know they are riskier than the average player. An estimate of twice as risky is likely in the ballpark and therefore acceptable for these purposes.

The calculation of the PDs resulted in a few interesting cases that also warrant discussion. One such case deals with Ronnie Brown, who had an incredibly productive 2007 season that was cut short due to injury. As a result, his stellar performance may have been disproportionally good, which would temper our projections and result in a 38% decrease in points per game going forward. A more interesting example comes in the form of Kurt Warner and Matt Leinart. These QBs have been engaged in a battle for the past two years, which has affected their projections. Two years ago, Warner played poorly and was replaced by Leinart. Last year, the reverse happened and Warner took the leading role. Warner's points per game are thus affected by a short, great season last year, which is being regressed to his mean performance this year. Further, Warner's age negatively impacted the points per game measure. On the other hand, Leinart is due for a bump in his stats in his 3rd year and a likely larger than normal reduction in interceptions (which have been unusually high for him in the past). This is also a regression to the mean, which our models use prominently to address issues like these.

The most obvious regression to the mean that directly impacts points per game is the regression of Adrian Peterson. Peterson, as mentioned previously, experienced one of the best rookie seasons of all time. Our models initially projected a modest decline, however a qualitative adjustment, discussed earlier, further downwardly biased his performance. These factors lead to the 24% decline in points per game expected of Peterson.

The overarching purpose of the PD measure is to capture a specific type of uncertainty in statistical analyses called “modeling risk.” Modeling risk generally refers to the rigidness of most mathematical formulae with respect to the reality of the phenomenon that is being modeled. The examples discussed above indicate one area where the modeling risk that we take into account is not a result of the phenomenon we seek to study. Instead, it is a small sample size bias for which we developed a solution. Another source of deviation that may hamper our overall risk model is the treatment of injuries that do not result in a stoppage of play. Our three-factor model relies, in large part, on the mutual exclusivity of the factors. Injury risk is taken into account exclusively by the latent factor discussed later in this section. One aspect of injury risk that our PD model might also take into account would be if a player was injured and *continued to play*. In this case, their overall points per game would be lower during that time, and thus, their risk carried into the next season would be indicated as being higher. While ideally, we wouldn’t be taking injury into account at all (one of the main reasons for selecting a points per game unit of analysis), this small burring of the boundaries between the latent factor and the PD factor is extremely small, affecting no more than 2-3% of the sample.

As we move towards the end goal of building a three-factor risk model, we must discuss a simplified interpretation of PDs with respect to standard deviation. We introduce this discussion by framing the case in terms of instances where a player has a combination of high or low standard deviations and PDs. The set of mutually exclusive definitions of the interaction between PD and standard deviation includes two categories for each type of risk. One category is “High” risk and the other is “Low” risk. Thus, the general case of how to classify PDs can be seen through a mutually exclusive 2x2 matrix of states.

Standard Deviation \ Projection Differences	High	Low
High	Very Risky	Moderately Risky
Low	Moderately Risky	Not Risky

When both measures of risk are high, the player is deemed very risky. This result indicates that there is a great deal of variability in that player’s within-season points earned in addition to the points per game difference between seasons. The opposite case would be where the player had a low standard deviation in addition to a low projected difference rating. In this case, the player has shown himself to be a safe bet with respect to performance, both within a given season and between seasons.

The mixed cases are examples where the risk factors do not coincide in their assessment of a given player. In both examples, either the within-season variability measure is high while the between-season variability measure is low, or the within-season variability measure is low while the between-season variability measure is high. These are the most important cases because it allows our model to add depth to the assessments of risk we utilize to generate the optimal portfolio of players.

Another aspect of PDs that we want to explore is their relative variability and bias. First, one bias that is apparent is that of the wide receiver positive difference to be skewed to the high side, averaging over 20% compared with that of 13% for QBs, 12% for RBs, and 12% for TEs. This skew can be traced to the fact that many WRs in the sample are young, and young WRs are projected to improve disproportionately more with respect to the other positions. This is likely because WRs tend not to be as solid at the beginning of their careers as other positions due to the sheer number of WRs in the league.

In terms of variability, WRs also exhibit the largest dispersion of observations. The following charts display the progression of variability from least to most with respect to PDs in percentage terms on the same scale (for comparison).

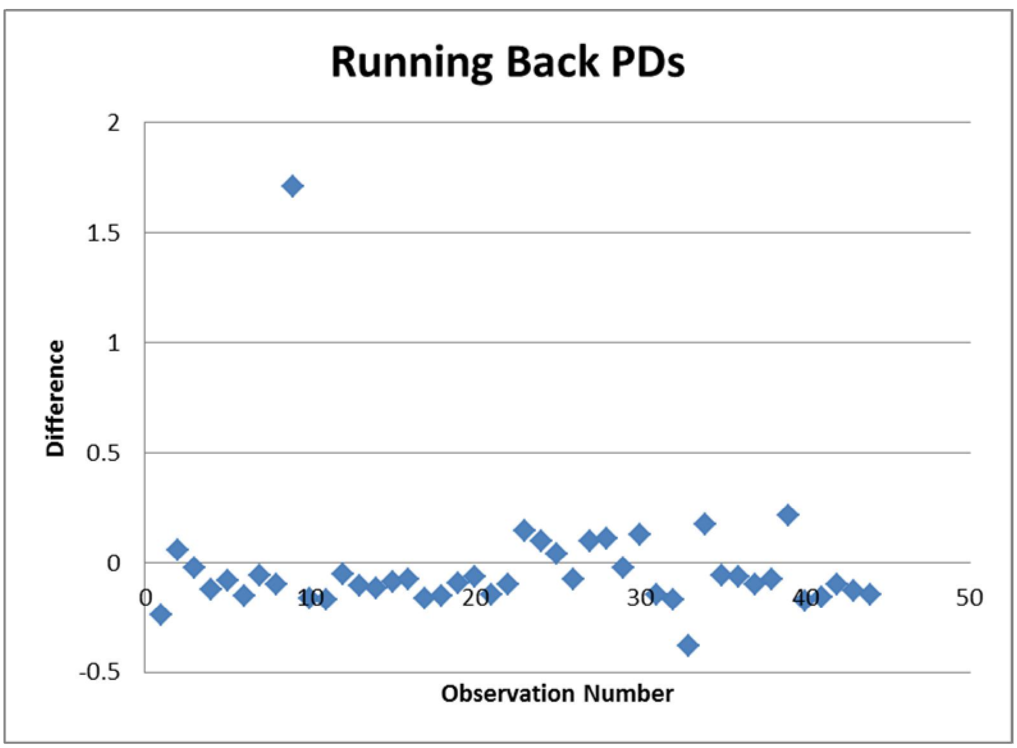

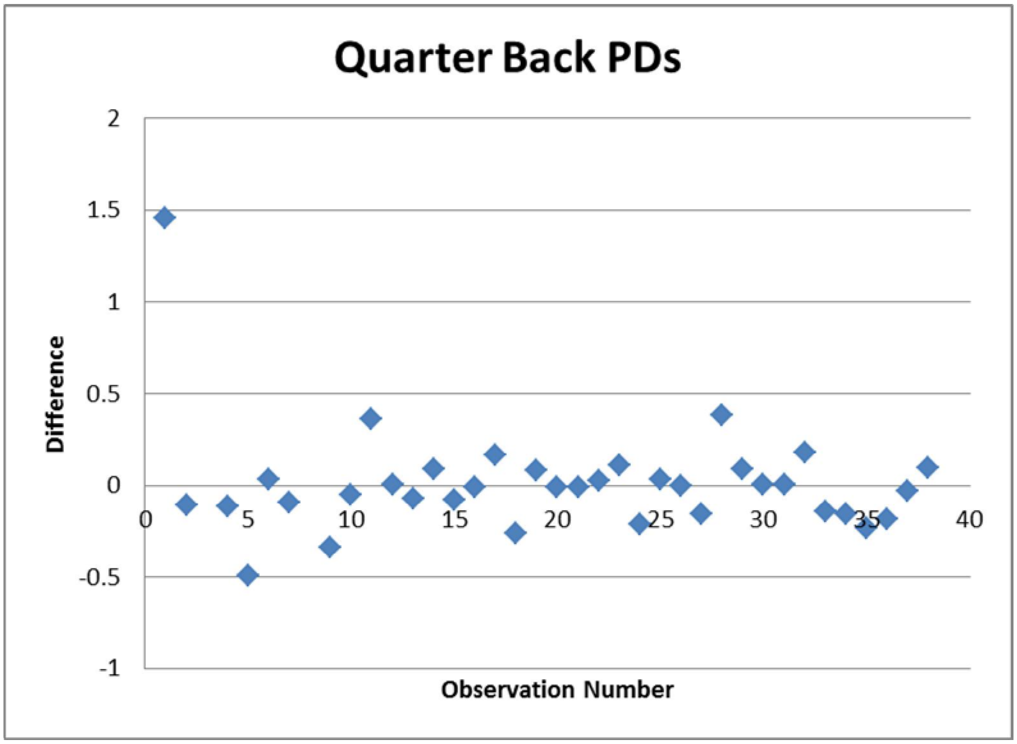
Quarter Back PDs
Difference
Observation Number
2
1.5
1
0.5
0
-0.5
-1
0
5
10
15
20
25
30
35
40

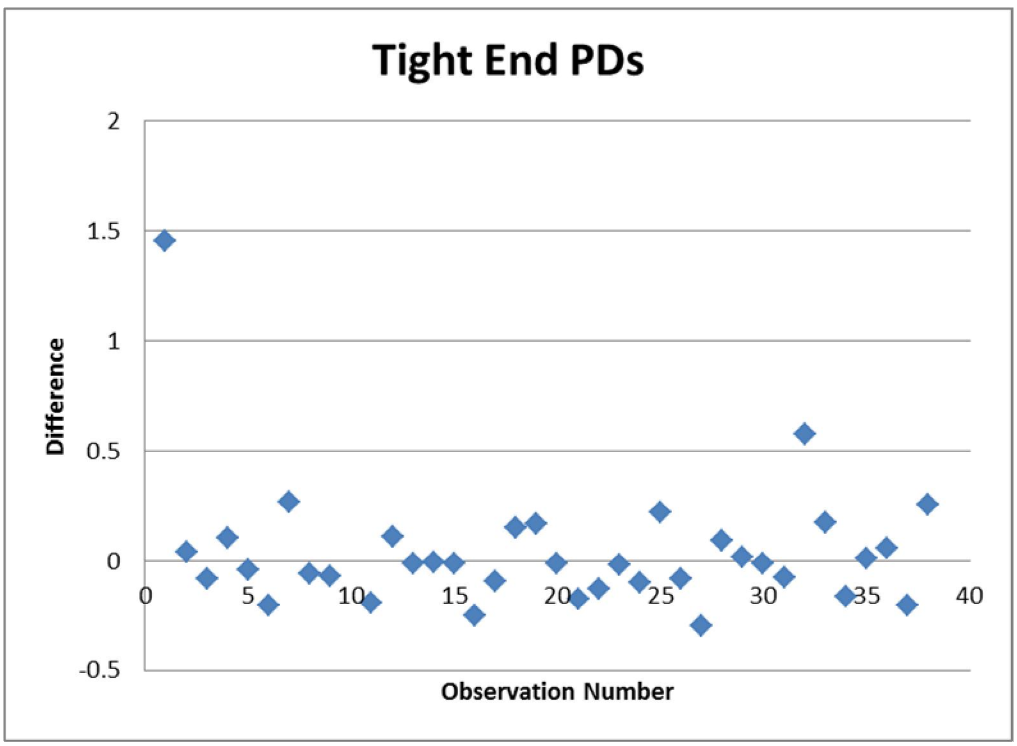
Tight End PDs
Difference
Observation Number
2
1.5
1
0.5
0
-0.5
0
5
10
15
20
25
30
35
40

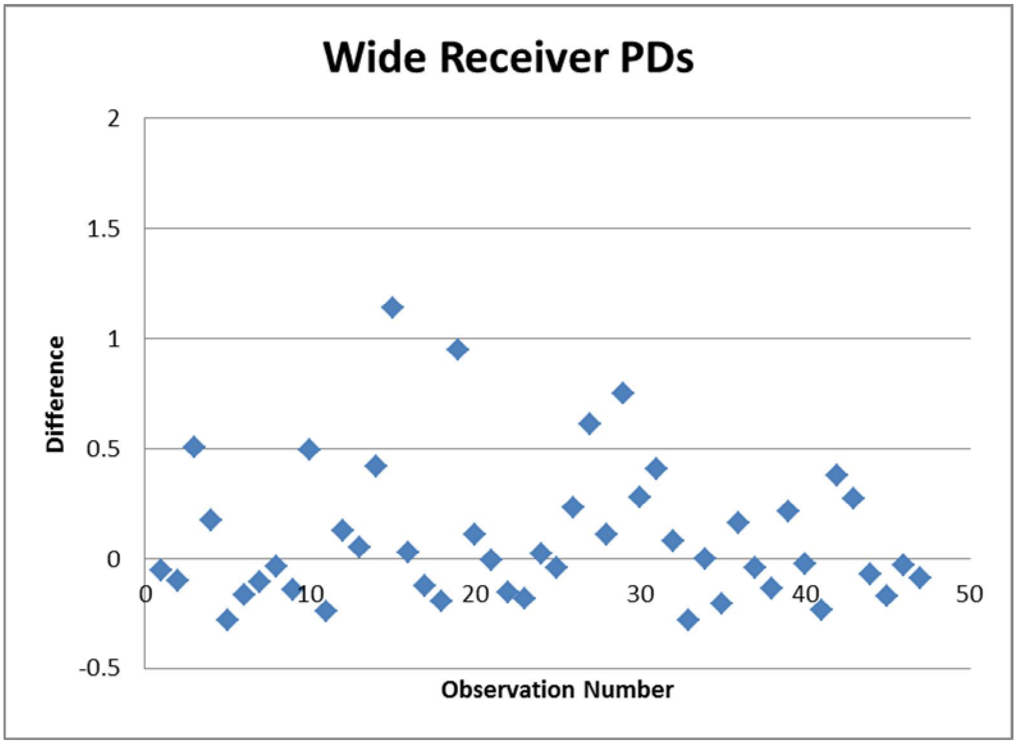

In order to verify that the Projection Differences model works correctly (or approximately correctly), we assessed the logical assertion that the overall net differences among all selected players in a given position should be relatively low after allowing for some mitigating factors, and that some should be negative. The results of the calculation of the net differences (as opposed to the absolute differences which is what we use as the risk metric) are as follows:

Position	Cumulative Change in PPG / Cumulative PPG by Position
Running Back	-7.24%
Quarter Back	-2.78%
Wide Receiver	1.53%
Tight End	-0.53%

Overall, each position has a relatively low cumulative PD difference. In percentage terms, running backs overall appear to be set for a decline, as is the case with quarterbacks and tight ends. Wide receivers, however, should see a modest increase based on the PD calculation and the phenomenon therein as described above.

Latent Risk:

The final, and most complex and comprehensive factor in the three-factor risk model is the one that takes into account the latent measure of risk. The philosophy behind this risk measure states that overall, the standard deviation and the projection differences, on their own or combined, do not do a good enough job of capturing the true risk of a player. The reason has to do with a common problem in statistical analysis: the omitted variable problem. Specifically, statistical models can do an excellent job describing or explaining the factors that they explicitly take into account. However, when applied to the object that they are designed to assess, those same models can sometimes fail. One possible reason for that failure could be the non-inclusion of a variable, or set of variables, that contribute to the explanation of the dependent variable.

In this case, our dependent variable is "risk." We have two solid metrics that provide insight into the within-season and between-season risks and that can describe those two measures well. However, we know that there are a number of other factors that contribute to the explanation of risk that we cannot directly observe or quantify. Assessing those factors, however, can be an extremely difficult endeavor, especially when they fall into a "miscellaneous" type category as they do here.

The reason for this is that, unlike other player variables (such as a average number of points scored, age, proportion of time the player spends injured, etc.), player miscellaneous risk cannot be directly measured. “Miscellaneous risk” is not a variable that can be observed from the player statistics, nor can it be inferred through some transformation of those statistics. Variables or constructs that cannot directly be measured from the relevant unit of analysis (in this case, the players) are defined as latent variables.

It may be useful to start with an example from social sciences, as the concept of latent variables is most commonly used in that field. This example should clarify how this is related to the assessment of player miscellaneous risk. In social sciences, variables such as “attitudes” towards something, “perceptions” about an issue, personality traits, and so forth, are considered to be latent variables, as it is not possible to directly measure and quantify the “attitudes” of an individual. However, it is necessary to operationalize and potentially quantify (although by more “indirect” methods) this kind of variable in some way in order to study them or include them as explanatory variables in a model.

The usual way in which latent variables are quantified is by way of indicator variables. These indicator variables are those that can be directly observed and quantified and that are supposed to be strongly related to the latent variable of interest, so that they can be used as an approximation of that latent variable. Continuing with examples from social sciences, the level or type of “attitudes” towards an issue is usually operationalized by the responses to a questionnaire in which individuals are asked to indicate their level of agreement with certain statements. For example, a questionnaire designed to measure the level of "attitudes toward computers" could include several statements such as "Computers make me feel uncomfortable," or "Learning about computers is worthwhile," to which the respondent would indicate his or her level of agreement. In this example, the levels of agreement with these statements would be the indicator variables (because they are directly observable — respondents would specifically state that they agree, disagree, or are neutral about each specific statement in the questionnaire). Then, these responses would be aggregated through some mathematical transformation in order to arrive at a final quantification of the latent variable, "attitudes toward computers."

It is very important to note that the determination of the indicator variables that theoretically are related to the latent variable being measured is a very complex process. In particular, the latent variable scale, as measured by a given set of items, must be evaluated carefully to determine whether it both valid and reliable. Validity refers to the extent to which the chosen indicator variables are indeed measuring the latent variable they are supposed to measure. Reliability refers to the extent to which these indicator variables are inter-correlated (that is, if the indicator variables are all measuring the same latent variable, then one would expect to see a relatively high inter-relationship among them). Assessing the validity and reliability of a scale can be a difficult process, and it usually involves, in addition to relatively complex statistical procedures, qualitative input from a panel of experts to determine whether the proposed indicator variables appear to be measuring the latent variable they are supposed to measure (this is called "face validity"). In fact, the panel of experts may also be used at the very beginning of the design of a procedure to measure a latent variable, by having them suggest possible indicators, discussing them with the rest of the members, confirming them with a second group of experts.

This distinction between latent variables and observed or indicator variables can be easily applied to the measurement of risk. As mentioned previously, player miscellaneous risk can be considered as a latent variable because it is not directly observable. Therefore, the measurement of player miscellaneous risk is much more complicated than the measurement of other player variables (such as a player's average number of points per season). However, miscellaneous risk can still be quantified by way of a number of indicator variables that should be strongly correlated to this type of risk.

Once again, it is very important to stress that this type of risk (i.e., miscellaneous) is conceptually different from the two other risk factors that have been incorporated into the model (which are standard deviation and projection differences). Those other two risk factors can be measured directly from the player statistics — these factors are thus observed variables and can thus be incorporated in a straightforward way in any computation of the overall risk of a player (which would involve some mathematical transformation applied directly to those values). By contrast, this third risk factor (miscellaneous risk) is not directly measurable, and thus, it is not immediately clear how it is combined with the other two (measured) factors. Although the easiest solution to such a problem is to simply exclude from the model whatever cannot be directly measured (i.e., simplify the model to include only the observed variables), this would constitute a clear loss of potential information, which could cause the model to be unrealistic. As discussed above, this could result in model misspecification.

It is thus important to include this unobserved risk factor into the model. Therefore, in order to be able to incorporate this third type of risk into the final mathematical model for drafts, we have to indirectly quantify it in some way. Hence, there is a need for using indicator variables that may help get an estimate of the level of the player's miscellaneous risk and produce a more accurate final model for overall player risk.

Latent Risk Factor Creation and Experimental Design

In order to determine the indicators of miscellaneous risk, we took an extremely in-depth approach that involved multiple phases of investigation, survey construction, analysis, testing, and verification. The first step in this process was the construction of a focus group to help assess the risk factors that the other historical variables may not be accurately taking into account. The focus group was comprised of a convenience sample of 15 "experts" in Fantasy Football. This group of individuals was instructed to relay their thoughts, perceptions, and experiences regarding the player specific risks involved in the selection of a Fantasy Football team.

The methodology utilized to extract those thoughts, perceptions, and experiences was a semi-structured interview that is a standard type of instrumentation in psychological research. This instrumentation fits within the theoretical framework of phenomenology. Phenomenology is perfect for the identification of the latent risk factors in Fantasy Football players because it is based upon the belief that unobservable matters should not be simply categorized as such and then dismissed as un-analyzable. Instead, practitioners of phenomenology assert that evaluation of the unobservable is possible due to a given participant's subconscious or conscious awareness of the underlying factors attempting to be evaluated. As a result, phenomenology provides an effective means for exploratory research based upon the rich textured narrative descriptions of the lived experiences of Fantasy Football players.

The experimental design of the study utilized to generate the latent factor is as follows. As noted earlier, the first step was an aggregation of a convenience sample of 15 experts in order to record their thoughts and perceptions regarding miscellaneous player risk through a semi-structured interview. Following that interview, the recording was transcribed and the resulting document uploaded into NVivo© qualitative analysis software for review and coding (discussed later). The coding of the interviews for analysis allowed for a comprehensive review of the themes brought up by the participants. Three risk factors were derived from the focus group. The sample size of the focus group was chosen to be 15 because, for qualitative analyses, fewer responses are necessary to achieve a given depth of understanding. Indeed, standard academic research interview sizes range from 5 to 20 people, depending on the phenomenon.

Once those themes were structured, the original 15-member expert focus group was given a list of players to rank from 0-100, according to the risk factors identified in their report. The results from this survey were used to model the third factor discussed below. In order to validate the risk assessments generated by the model, we then conducted a larger survey, described later, to ensure that the risk scoring utilized in the latent risk model mapped correctly to a general assessment of overall player risk. Prior to conducting the larger study, a power analysis was done in order to calculate the optimal number of respondents necessary for detection of statistically significant results through the tests. The purpose and calculation of a power analysis are both detailed in the following section. Those participants in the larger survey were asked to assign a single risk score to each player from 0-100. These results were compared to those of the focus group and a factor analysis was conducted to ensure that the assessment of latent risk in our risk model resulted in a good measure of overall risk.

Latent Risk Factor Analysis

The focus group was approached with a series of questions in a semi-structured interview that allowed them to share their thoughts and experiences on Fantasy Football player risk. The interview was meant to direct the experts' thoughts to the main factors that could determine risk that we cannot quantitatively observe. Because the latent risk model must be mutually exclusive of the other risk models, it must, at least in large part, deal with the future or with other unobservable factors. The result of the focus group was a recording of the thoughts and perceptions of the experts regarding miscellaneous player risk. This recording was then transcribed into a document that could be utilized for analysis.

In order to prepare for the analysis, the interview transcript was evaluated for content analysis using the NVivo© qualitative analysis software program in order to identify salient elements relating to importance and unobtrusive themes. Data from the transcribed interview was coded to reduce attributions to the component elements of cause, outcome, and links between cause and outcome. In other words, the basic meaning of the lived experiences and thoughts of the expert participants was derived. Coding was guided by the Leeds Attribution Coding Systems. The six stages of attributional coding include:

1. Source identification
2. Extract attributions from transcripts
3. Separate cause and outcome elements of the attributions
4. Code speaker, agent or cause of the attribution, and target of the outcome
5. Coding attributions on causal dimensions
6. Analysis

The objective was to identify content for the elements that were present and countable from the interviews. The intent was to analyze the data and establish common themes, patterns, terms, or ideas that can inform a deeper understanding of the issue surrounding the research problem while articulating a description of Fantasy Football player risk.

A further objective of the analytic method was to identify the manifest content for the elements that were physically present and countable from the interviews. The combined sources of research data are appropriate to the research design and strategy to obtain valid and reliable empirical information.

We utilized a version of qualitative data analysis standardized by Moustakas (1994), which was a modification of the van Kaam (1959) method of analysis. The steps for analyzing the data from the transcribed interview include:

1. Listing and preliminary grouping of every relevant experience.
2. Reduction and elimination of extraneous data to capture essential constituents of the phenomenon
3. Clustering and Thematizing the Invariant Constituents to identify core themes of the experience.
4. Final identification and verification against the complete record of the research participant to ensure explicit relevancy and compatibility.
5. Construct for each co-researcher an individualized textural description of the experience based upon the verbatim transcripts using relevant and valid invariant constituents and themes.
6. Construct for each co-researcher and an individual structural description of the experience based upon individual textural description and imaginative variation.
7. Construct for each participant a textural-structural description of the meaning and essence of the experiences.

The result of this rigorous analysis was a short list of themes that were further reduced into the most basic components that determine latent risk. These three factors, which make up the final factor of the three-factor model, are as follows:

1. Attitude risk
2. Projected season injury risk
3. Benching risk

The phenomenological assessments of the above risk measures provided insight into the forward-looking risk factors that contribute to a more complete assessment of player risk. First, Attitude risk was mentioned consistently by the panel of experts as a non-previously quantified risk that needs to be taken into account. A prime example of this would have been Randy Moss a number of years ago, prior to joining the Patriots. Now, with a solid team and an exceptional quarterback, it is expected that Randy Moss's Attitude risk is lower than it was previously. In this way, historical assessments like Projected Differences may factor in a risk measure in a way that does not accurately reflect the future. The qualitative assessment of attitude risk, however, captures that nuance and allows for a more complete understanding of player risk.

Next, projected season injury risk was taken into account as a further means of assessing the likelihood of individual players missing games due to injuries. This measure does more than just take into account previous injuries. It factors in expert opinions on aspects that may not have occurred during season play (and would thus not be taken into account by the historic Projected Differences measure). Benching risk is similar to projected season injury risk. This metric takes into account the experts' assessments of the likelihood of a player getting benched during the coming season. Again, the inclusion of this qualitative based measure of risk adds an important level of rigor to the model.

These contributors to the final latent risk measure of the three-factor model still had to be operationalized. In other words, a quantitative measure of their attribution to each player was required to implement the three-factor model. The first step to accomplish this goal is the use of the focus group to give each player a value from 0-100, where 0 is least risky and 100 is most risky, on each of the three latent variables that could affect risk for each of the players in the four main positions in the model. Kickers and team defense were not included in this process because they play a very small role overall in the draft order and would be overly complicated to model in this way, respectively.

Once the 0-100 ranking score was recorded for each of the players described above and for each of the three latent risk measures, the scores were averaged on a per player basis to yield a raw latent risk measure. It is important to note that, a priori, it is not obvious that three scores can be averaged in order to arrive at an overall final latent risk score. This is because there may be interactions among the different risk factors (i.e., the effect of each factor on the final latent risk measurements may be dependent on the value for some of the other factors). In such cases, a linearly weighted average of the factors may not be an appropriate measure of the overall risk scale. Moreover, the internal consistency reliability of these three factors of latent risk must be assessed. In other words, it is important to ensure that these three factors are indeed measuring the same underlying construct, which is miscellaneous risk.

The most common way in which internal consistency reliability for a scale is assessed is through Cronbach's alpha. Cronbach's alpha is a statistical coefficient that ranges from 0 to 1 and measures the extent to which a number of items or factors are inter-correlated, with values closer to 1 indicating a higher correlation among those factors. The formula for Cronbach's alpha is given by:

$$\alpha = \frac{rN}{v + r(N-1)},$$

where *r* is the average inter-correlation among all possible pairs of factors (measured through Pearson's correlation coefficient), *v* is the average variance across all factors, and *N* is the number of factors used in the analysis.

Nunnally and Bernstein (1994) suggest that Cronbach's alpha values of 0.7 or higher are enough to conclude that a scale exhibits adequate internal consistency reliability (that is, that the inter-correlation among the components of that scale is high enough to conclude that the components are measuring the same underlying construct). We computed Cronbach's alpha based on the scores for each of the three factors as given by the 15 respondents in order to verify whether these factors exhibited internal consistency reliability. Cronbach's alpha was 0.89, suggesting excellent reliability for these factors, and providing evidence that they are measuring the same underlying construct (miscellaneous risk).

After finding evidence that the three identified factors were indeed measuring the same underlying construct, we sought to determine whether they could be combined through a linearly weighted average. As explained above, the factors should satisfy *additivity* in order to do this (i.e., there should be no interactions among the three factors). In order to determine whether the factors satisfied the additivity condition, Tukey's test of additivity was conducted on the data. The outcome of this test is a test statistic that follows an F distribution under the null hypothesis that the model is indeed additive. Results of this analysis showed that the null hypothesis that the model was additive could *not* be rejected at the .05 level. Therefore, we have shown that the three factors that we assumed would make up miscellaneous risk (a) indeed measured the same construct and (b) could be combined in an additive fashion (such as an average, as we used) in order to arrive at a final score for miscellaneous risk.

Given that the miscellaneous risk measure was then to be combined with two other observable measures of risk (PD and standard deviation), miscellaneous risk was then normalized to ensure that neither an artificially high nor an artificially low weight was assigned to miscellaneous risk when computing the overall risk measure (the normalization technique discussed in the next section). The final step in the process is the testing of validity to ensure that the score for latent risk factors actually provides a good measurement of overall player risk. This step was taken after the generation of the final three-factor model.

Three-Factor Model of Risk:

In order to effectively combine the three measures of risk into one solid measure, the three measures needed to be scaled to an identical level. First, to avoid confusion, it is necessary to understand that the three factors discussed in this section are not the same as those discussed in the previous one. In the previous section, we discussed three factors that, when combined, produced a measure of miscellaneous (latent) risk. In this section, we discuss how we combine miscellaneous risk, standard deviation, and PD in order to arrive at measure of overall risk (and thus the "factors" referred to in this section are miscellaneous risk, standard deviation, and PD).

Specifically, a three-step statistical process is used to generate a scaled three-factor risk measure. The first step is to normalize each player's standard deviation score by the variability in all players' standard deviation scores for a given position. This calculation results in scores that represent the multiple of a given player's standard deviation over the standard deviation of all players in a given position. Simply put, this methodology normalizes each player's standard deviation with respect to the volatility of all players' standard deviations within a given position. This step is useful because it allows for the comparison of this measure of risk with other measures of risk normalized in a similar manner. Normalization in this context means reducing a quantity by an appropriate base, or denominator, in this case the standard deviation of the standard deviations of all players.

The same methodology is then utilized for the PDs and for the latent risk measure. This calculation results in a normalized quantity of each of the three risk measures that can then be combined and scaled to be meaningful both intuitively, and in the portfolio theory sense. The next steps accomplish this by first averaging the three normalized risk measures into one number. This number represents the base quantity of the three-factor model. This base quantity is then multiplied by the average standard deviation of all players in a given position in order to make standard deviation comparable to the scaled three-factor risk measure. The same scaling operation is also applied to the projection differences and latent risk measure in order to make all risk measures comparable within the context of portfolio theory. The table for these risk measures for the top-20 running backs is displayed below:

	Points	Stdev	PD	Normalized Stdev	Normalized PD	Normalized Latent	3 Factor	Scaled 3 Factor
Adrian L. Peterson	262.60	58.27	5.10	6.21	4.90	2.00	4.37	40.99
LaDainian Tomlinson	262.56	54.76	2.81	5.84	3.60	3.25	4.23	39.66
Brian Westbrook	262.46	49.26	2.29	5.25	3.30	4.40	4.32	40.49
Marshawn Lynch	235.34	27.62	1.31	2.95	2.74	7.00	4.23	39.66
Joseph Addai	230.41	48.93	1.17	5.22	2.67	5.00	4.29	40.27
Larry Johnson	220.79	55.11	1.74	5.88	2.99	4.40	4.42	41.46
Brandon Jacobs	207.75	37.90	0.33	4.04	2.19	10.00	5.41	50.73
Jamal Lewis	204.02	56.29	1.73	6.00	2.98	7.80	5.60	52.46
Steven Jackson	200.71	31.12	1.07	3.32	2.61	4.55	3.49	32.75
Maurice Jones-Drew	198.73	36.96	1.25	3.94	2.71	6.10	4.25	39.86
Reggie Bush	194.99	36.57	1.37	3.90	2.78	5.25	3.98	37.29
Clinton Portis	191.54	44.01	2.14	4.69	3.21	3.50	3.80	35.65
Frank Gore	185.93	39.49	0.63	4.21	2.36	3.40	3.32	31.16
Ronnie Brown	182.75	48.00	7.02	5.12	5.99	6.10	5.74	53.79
Marion Barber III	182.69	43.21	0.94	4.61	2.53	3.80	3.65	34.20
Laurence Maroney	177.63	41.13	1.02	4.39	2.58	6.25	4.40	41.30
Willis McGahee	169.69	31.26	1.78	3.33	3.01	3.50	3.28	30.77
Earnest Graham	160.58	42.92	1.98	4.58	3.12	7.50	5.07	47.51
LenDale White	160.33	29.63	0.38	3.16	2.22	7.00	4.13	38.69
Ryan Grant	156.43	40.10	0.64	4.28	2.37	6.50	4.38	41.08

Once again, as explained in the previous section, it is important to conduct statistical tests to assess whether three factors can indeed be combined in this way (i.e., by way of an average). Therefore, we conducted the same tests as those described earlier (Cronbach's alpha and Tukey's test on additivity on the measures of standard deviation, PD and latent risk). Results of these analyses suggested that (a) these three factors exhibited excellent internal consistency reliability (Cronbach's alpha = 0.84), and (b) the null hypothesis of additivity was not rejected as per Tukey's test. Therefore, we have found evidence that these three factors measure the same underlying construct (overall risk), and that they can be combined through an additive operation (such as an average) in order to arrive at an encompassing measure of risk.

Once the scaled three-factor risk measure was operationalized, a convergent validity test was required. Convergent validity is the extent to which two or more different measures of the same theoretical construct are related. Going back to the examples based on the social sciences, a new questionnaire designed to measure "attitude toward computers" should yield outcomes that are related to those of other questionnaires designed to measure the same thing. If an individual has very positive attitudes toward computers, then he or she should score high in all questionnaires designed to measure this construct, assuming that these questionnaires are correctly designed.

In this particular case, we are interested in testing whether our own measure, which we have hypothesized measures "overall risk," is indeed measuring that construct. Therefore, if we found that the measure of overall risk produced by our model is similar to the measure of risk produced by a different model or thought process, we would have found evidence that our model indeed measures overall player risk by proving that it exhibits convergent validity.

In order to execute the test, a live survey was conducted at a minor league baseball gathering on July 8, 2008, in Bridgeport, Connecticut. The purpose of this test was to determine the degree to which our operationalized latent risk measure was a good indicator of the overall risk associated with a given Fantasy Football player or set of players. The functional survey asked the participants to give a number from 0 to 100, where 0 is least risky and 100 is most risky, that describes the general level of risk of each player in the sample.

The first step in conducting this convergent validity test was a determination of the number of respondents necessary to yield statistically significant results. The proper way to determine what the requisite sample size would be is to conduct a statistical analysis is to perform what is known as a power analysis. Performing a power analysis for sample size estimation is an important aspect of experimental design. Without these calculations, the resultant set of participants may be too high or too low with respect to the specific needs of the analysis at hand. If the chosen sample size is too low, the experiment will lack the precision to provide reliable answers to the questions it is intends to investigate. If sample size is too large, then time and resources will be wasted, often for minimal gain. When performing a statistical power analysis, one needs to consider the following important information:

1. *Significance Level or Confidence Interval*: A common, yet arbitrary, choice is 95%.
2. *Power to detect an effect*: This is expressed as power=1 - ß, where ß is the probability of a false negative. Power=0.80 is also a common, yet arbitrary, choice.

3. *Effect size (in actual units of the response) the researcher wants to detect*: Effect size and the ability to detect it are indirectly related; the smaller the effect, the more difficult it will be to find it.
4. *Variation in the response variable*: The standard deviation, which usually comes from previous research or pilot studies, is often used for the response variable of interest.
5. *Sample size*: A larger sample size generally leads to parameter estimates with smaller variances, giving you a greater ability to detect a significant difference and results in shorter confidence intervals.

These five components of a power analysis are not independent: in fact, any four of them can be utilized to calculate the fifth. The usual objectives of a power analysis are to calculate the sample size for given values of the other four items. This was the methodology undertaken for the face validity survey. Specifically, we collected 25 completed surveys from convenience sampled participants at the event and then analyzed those samples to determine the total number we would have to collect.

Based upon the desired effect size, the desired significance level, and the variability in the sample, we determined that there would have to be about 240 observations in order for us to be comfortable with the results of the Face Validity test. We conducted the survey live at the Bridgeport Bluefish Minor League Baseball gathering and were able to acquire the requisite 240 person sample. Once we collected the observations, we assessed the sample's risk assessment compared with our own three-factor measure.

Two statistical procedures were conducted in order to compare our measures with the outcome from this survey: (1) averaging the responses from the survey and (2) calculating the Peason's correlation coefficient. First, for each player, we computed the average risk level from the 240 responses. Thus, for each player, we had two measurements: the risk level as produced by our three factor model, and the average risk level as estimated by the 240 survey participants. We computed Pearson's correlation coefficient to assess the extent to which these two measures were correlated. Pearson's correlation coefficient varies from -1 to 1, with values close to -1 indicating strong negative correlation (i.e., high scores in one measure are associated with low scores in the other one); values close to 1 indicating strong positive correlation (i.e., high scores in one measures are associated to low scores in the other one); and values close to 0 indicating no relationship. Clearly, high values of Pearson's correlation coefficient would provide evidence of convergent validity, as it would suggest that our model's measure of risk is similar to that of average of the 240 participants. We obtained a coefficient of 0.76, which indicates a very strong correlation between both measures. This provided evidence toward the hypothesis that the outcome of our model indeed measures overall player risk.

The results of the comparison indicated that our survey methodology met the criteria for convergent validity because the responses were, for the most part, consistent with the relative measures of risk contained within our variable.

Implementation of Portfolio Theory

There are many themes within portfolio theory that center on risk. The first is the basic tenet that investors avoid risk and demand a reward for engaging in risky investments. The reward is taken as a risk premium, the difference between the expected rate of return and that available on alternative risk-free investments. A prospect that has zero risk premium is called a fair game. Investors who are risk averse reject investment portfolios that are fair game or worse. Risk-averse investors are willing to consider only risk-free or speculative prospects with positive risk premia. Loosely speaking, a risk-averse investor penalizes the expected rate of return of a risky portfolio by a certain percentage to account for the risk involved — the greater the risk, the larger the penalty.

The second theme allows you to quantify investors' personal trade-offs between portfolio risk and expected return. To do this you use a utility function, which assumes that investors can assign a welfare or "utility" score to any investment portfolio depending on its risk and return. The utility score may be viewed as a means of ranking portfolios. Higher utility values are assigned to portfolios with more attractive risk-return profiles. Portfolios receive higher utility scores for higher expected returns and lower scores for higher volatility. An investor's risk aversion will have a major impact on the investor's appropriate risk-return trade-off.

Finally, the third fundamental principal is that you cannot evaluate the risk of an asset without considering the portfolio of which it is a part. This suggests that the proper way to measure the risk of an individual asset is to assess its impact on the volatility of the entire portfolio of investments. Two common approaches to reducing risk in a portfolio are hedging and diversification. Hedging occurs when you invest in an asset with a payoff pattern that offsets exposure to a particular source of risk (e.g. flood insurance). Diversification occurs when you mix different investments within a portfolio that reduces the impact of each on the overall portfolio performance. If one investment component goes down, the other investment component should sometimes go up. This imperfect correlation ("non-synchronicity") reduces the overall portfolio risk. Taking this approach, you find that seemingly risky securities may be portfolio stabilizers and actually low-risk assets.

For example, if the Pittsburgh Steelers are playing the Baltimore Ravens, you may want both Ben Roethlisberger and the Ravens' Defense on your fantasy football team. The chances are slim that both components will perform well, but the chances are high that at least one of the two components will perform well. In this scenario you would reduce the expected points you would earn, but you would also reduce the volatility of the points you would earn. This is the principle behind diversification: A decrease in the expected return is associated with a decrease in the volatility of the return.

In the drafting model, risk is represented as the volatility present in the underlying factors that determine the expected points players will contribute to the fantasy football team. In other words, risk is not solely determined based upon the points the players would produce each game, but by the individual components that are used to establish the points produced. A running back's individual statistics for every past game are analyzed to determine the volatility in the individual statistics. For example, the rushing yards have a different volatility than the number of touchdowns scored. It is possible that Reggie Bush always scores two touchdowns, but that his rushing yards can vary from 25 to 200 based upon the strength of the run defense of the opposing team. Incorporating every component that is used in the analysis of the expected points earned per game provides an optimal analysis of total player risk.

Correlation:

To obtain a measure of association that conveys the degree of intensity of the co-movement between two variables, you relate the covariance to the standard deviations of the two variables. Each standard deviation is the square root of the variance. Thus, the product of the standard deviation has the dimensions of the variance that are also shared by the covariance. Therefore, we can define the covariance and the correlation coefficient, denoted ρ, as:

Covariance:

$$Cov(r_1, r_2) = E((r_1 - E(r_1))(r_2 - E(r_2)))$$

Correlation Coefficient:

$$\rho_{1,2} = \frac{Cov(r_1, r_2)}{\sigma_1 \sigma_2}$$

The subscripts on *r* identify the two variables involved. Because the order of the variables in the expression of the covariance is of no consequence, the above equation shows that the order does not affect the correlation coefficient either. The correlation coefficient can vary from -1.0, perfect negative correlation, to +1.0, perfect positive correlation. Perfect negative correlation suggests that when one variable increases, the other variable always decreases. Perfect positive correlation suggests that when one variable increases, the other variable will also increase. A correlation coefficient of 0 suggests that the variables are independent of one another.

The drawback to correlation coefficients is that they do not represent scale or causality. If the correlation coefficient is 1, we know the two variables always perfectly move in the direction, but not by how much. If you have three variables (A, B, and C) that are all perfectly positively correlated, they all have correlation coefficients of 1 with respect to each other. However, if A increases by 1 and B may increase by 1, C may increase by 1,000,000. The correlation coefficient ignores this scale effect. Also, because the correlation coefficient measures the degree of association, it tells us nothing about causality. The direction of causality has to come from theory and be supported by specialized tests.

The correlation matrix is a visual tool to help represent the relationship between different variables. For example, below is a simplified correlation matrix describing the relationships between Tom Brady, Randy Moss, and the Packers defense:

	Tom Brady	Randy Moss	Packers Defense
Tom Brady	1	0.15	0.02
Randy Moss	0.15	1	0.03
Packers Defense	0.02	0.03	1

In this hypothetical model, there is a positive relationship between Randy Moss and Tom Brady. This would suggest that when Tom Brady plays well, Randy Moss has a tendency to play well, too. Intuitively, this makes sense. Tom Brady and Randy Moss' performance should be linked; they are one of the most formidable quarterback-wide receiver duos in football history. However, there is not a strong relationship between the Packers defense and Randy Moss or Tom Brady. Tom Brady and Randy Moss's performance should not significantly affect the Packers defense's performance (unless, of course, they are playing each other). The slightly positive correlation between them is shown to demonstrate that there are many indirect factors that may arise from the enormous amount of interactions that occur in the NFL. This example is just a simplified explanation of the many variables (both observed and unobserved) that may cause relationships between fantasy football draft prospects.

In order to generate the highest quality correlation matrix, this model considered a number of rules that could be applied to fill in each relationship. These rules were initially approached from a player-specific perspective whereby positions and player examples were analyzed. An example of this logic can be seen from the perspective of the degree to which a quarterback should correlate to his wide receiver. We know that logically, there should be a correlation there. But should a better quarterback be better correlated to his wide receivers than an average quarterback would be? A better quarterback will throw for more yards, catches, and touchdowns, and therefore improve his wide receiver's performance. It would appear that the better quarterback and wide receiver combination should thus be more correlated than an average combination of quarterbacks and wide receivers. However, the better production numbers put up by a high performing quarterback and wide receiver combination are simply the result of more completions at the same correlation coefficient. That given correlation coefficient might appear to be higher because the combination of the two star players is producing more than it otherwise would. Instead, it is just a factor of higher production numbers at the same level of correlation.

Another example can be shown through the presumably negative correlation between running backs and wide receivers. While this overall figure is likely true (more passes result in fewer rushes and vice versa), a better running back will open up the passing game, allowing the wide receivers to score more points with fewer passes. He would also acquire more first downs that could allow the passing game to continue. On the other hand, logically, a better running back would get more carries and take up some of the points that might have been given to the wide receivers. On the whole, however, these examples demonstrate that while positional correlations may be predictably utilized, rules-based methodologies that hinge on specific player and position combinations do not contribute additional information to the correlation matrix.

As a result, we utilized historical studies and qualitative assessments to approximate the relationships between the positions, and thus, the individual players. We first generated a number of rules for the pairwise correlations between positions to create the overall correlation matrix. It was found through historical studies that the strongest positive correlation was between the quarterback and his wide receivers. Specifically, those two positions share a correlation coefficient of 13.20%. Next highest is the quarterback to the tight end at 8.76%. The overall correlation coefficients utilized throughout this analysis are presented in the table below:

Position	Correlation
QB to WR	13.20%
QB to TE	8.76%
QB to RB	4.54%
WR to TE	0.60%
RB to TE	-1.45%
RB to WR	-1.68%
PK to QB	4.00%
PK to RB	4.00%
PK to WR	2.00%
PK to TE	2.00%

As can be noted from the table, those positions that compete for points through yardage have a negative correlation. In this case, running backs and tight ends, and running backs and wide receivers, share a negative correlation coefficient because those positions cannot share in yardage gained in a single play.

In addition to creating correlations generated from very general historical data between positions, we have also incorporated several manual adjustments. Intuitively, it may seem that powerful tandems such as Brady and Moss would have a higher correlation than the typical QB/WR combo. However, when we examined the data it turned out that because both players were stars in their own right, they still produced outstanding numbers whether or not they played on the same team. The highest correlation that we found was that mediocre quarterbacks were heavily reliant on a star receiver if they had one. We made the following analytical adjustments to the correlation matrix:

- The Falcons' starting quarterback (be it Joey Harrington last year, or Matt Ryan or Chris Redman this year) will have a higher correlation average to Roddy White.
- Though he is a fairly decent player, we also bumped up Derek Anderson's correlation to Braylon Edwards given that Edwards had one of the top 20 receiving years of all time last year. Our decision was also based on the fact that we have little evidence that Anderson would continue to produce as much without Edwards as he does with him (unlike Brady, Manning, or Roethlisberger, etc. — all of whom have produced both with and without star WRs)
- Brodie Croyle, as a first year starter will likely be heavily dependent on consistent producer Tony Gonzalez, as Gonzalez is a perennial all star and young quarterbacks tend to rely slightly more on star tight ends if they have them.

- We also bumped Leinart/Warner's reliance on (and therefore correlation to) Larry Fitzgerald.
- We also bumped Carolina's starting quarterback's reliance on Steve Smith, be it Matt Moore, David Carr or Jake Delhomme (likely the latter).

Also, in a very few instances, some quarterbacks are also slightly dependent on big-time running backs to open up the passing game. However, in the few instances that we studied when the star running back missed games, the quarterback's efficiency dropped as expected, but the extra passes he threw compensated for the resulting de-emphasis on the running game.

The drafting model incorporates these correlations among the players to optimize the fantasy football team. If your team has high correlation coefficients amongst all of the players, it would be riskier than an alternative team that has low correlation coefficients amongst all of the players. On a highly correlated team, if one player performs really well and produces a lot of fantasy points, your entire team will most likely produce a lot of fantasy points, which would be good. On the flip side, if one player performs badly, the chances are high that your entire team will perform badly, which is not so good. Conversely, if the players are not highly correlated, one player who performs badly does not necessarily mean that your entire team will perform badly. Generally, a highly correlated team will have a higher degree of volatility in the fantasy points earned than a team that is less highly correlated. A test of the correlation assumptions for this model will follow the presentation of the prize structure and the Sharpe Ratio.

Prize Structures

Our Portfolio Theory model can be used in two ways, based upon the prize structure of a given league. There are two main prize structures that are used by Fantasy Football leagues. The prize structure used would then dictate the strategy that you would apply to your draft. The two main prize structures are the “Winner Take All” model and the “Top 3 Paid” model. As their names suggest, they indicate the distribution of prizes among the winners. In the “Winner Take All” model, only the first place finisher earns any money. If your league uses this prize structure, then your goal should be the unconstrained maximization of expected points. Risk is not necessarily important here because stable performers would not be preferred to those who have the greatest chance of producing large point totals. Second place wins nothing so you should always shoot for first.

On the other hand, if your league pays the top three spots, then risk becomes a crucial element in your drafting strategy. In this case, drafting only the players who are expected to perform the best may not maximize your expected value. Instead, a portfolio of players that has stable performers, even if it does not have the highest expected points, could end up maximizing your chance to receive a portion of the league prize.

In order to accommodate these two prize structures, our model can be set to maximize expected points, or to take risk and correlation into account to build a more “stable” team. The metric we use for that purpose is called the Sharpe Ratio.

Sharpe Ratio

Since William Sharpe derived the Sharpe ratio in 1966, it has been one of the most utilized metrics in finance, and much of this popularity can be attributed to its simplicity. The ratio's credibility was further established when Sharpe won a Nobel Memorial Prize in Economic Sciences in 1990 for his work on the capital asset pricing model (CAPM). The ratio describes how much excess return you are receiving for the extra volatility that you endure for holding a riskier asset. Simply put, Sharpe's ratio is a measure of the excess return per unit of risk. It is used to characterize how well the return of an asset compensates the investor for the risk taken. The formula is as follows:

$$S_X = \frac{E(r - r_f)}{\sigma_X}$$

In the above formula, S_x is the Sharpe's ratio for investment x, $E(r-r_f)$ is the return investment x provides in excess of the risk-free return, and σ_x is the standard deviation of investment x's returns.

From an investment perspective, when comparing two assets with the same excess return, the asset with the higher Sharpe ratio gives more return for the same risk. You can increase the Sharpe ratio by choosing an asset with higher excess returns (and the same risk), or with lower risk (and the same excess return). Ideally, you would choose an asset with greater excess return and lower risk.

So, if Plaxico Burress was expected to produce 30 points every week for your fantasy football team and Marvin Harrison was expected to only produce 28 points, it would appear that Burress would perform better. However, if Burress took much larger risks (maybe his acrobatic abilities make him more prone to injury but allow for more receptions), it may actually be the case that Harrison has a better risk-adjusted return. You may pick Harrison as opposed to Burress even though Harrison's expected earned points are lower because Harrison would be a more consistent contributor to your team.

In the drafting model, the Sharpe ratio is adapted to fit fantasy football parameters. Excess returns are represented as the number of points your fantasy football team is expected to produce. The volatility of the portfolio is represented as the variability in the expected points produced by the entire team. We created a proxy for volatility with the three-factor risk measure detailed above. Thus, Sharpe's ratio for your fantasy football team represents the number of points your team would earn per unit of volatility. In a league where the winner is not awarded the entire prize, the Sharpe ratio approach to maximizing expected points per unit of volatility is an excellent way to structure thinking about a draft.

The Sharpe Ratio model is a unique tool that nobody else is using in the world of fantasy football. Considered one of the most fundamental formulas within the finance world, the Sharpe Ratio maximizes risk adjusted return for investments. In fantasy football, we have used the Sharpe Ratio to choose draft picks by treating a group of players as a portfolio, each player having a mean expected return along with several measures of volatility.

The simulation's calculation of the Sharpe Ratio factors in correlations between players, and calculates the degree to which players on the same team may have related fantasy scores — in some cases, this can be positive, while in others it generates too much risk, and the formula will distinguish between those situations mathematically.

The Sharpe Ratio simulator is not without its flaws — it ignores positions and the strategies of other players, and it maximizes risk adjusted return rather than win percentage. The Sharpe Ratio, however, is not meant to be a standalone drafting tool but rather is intended as a unique and powerful supplement to a draft. For example, it can help players think 'outside the box' about draft picks and point out how unlikely players can actually balance out a 'portfolio' extremely well and mitigate risk. It also points out a strategy that seems very strange to most fantasy players (who typically draft for need), which is to stockpile good players at a particular position.

Further, the Sharpe Ratio allows players to examine the attractiveness of trades — simply plug in your current squad, and compare the players that you already have to those that are being offered to you. The Sharpe Ratio formula will then compute a risk-adjusted return and can advise you as to how the deal will affect your team's 'portfolio.' A modified variation of this could potentially be extremely useful midseason when you are well ahead of the pack and want to reduce risk, or when you need to catch up and want to intentionally increase the volatility of your team. Thus, when used in conjunction with the game theory draft model as well as a solid fundamental basis of football knowledge upon which allows one to recognize unusual outputs, the Sharpe Ratio simulator can be an innovative and potentially invaluable tool in winning your league.

Summary:

Portfolio theory is an extremely valuable tool that can be easily tailored for the selection of the best fantasy football team possible. Fantasy football teams and investment portfolios have many similar aspects, which is the reason applying portfolio theory works so well in the selection of the ultimate team. The players are your investments and their performance (which determines the number of points you receive) is the return on your investments. The variability in football players' performance from week to week is similar to the ups and downs of the stock market; it is the risk involved in the investment. Thus, you should pick your fantasy football team like you pick your investment portfolio. You should strive for the maximum return with the least risk.

The drafting tool incorporates statistics from past years and projections for the future year to optimize your fantasy football team. The expected points earned and the variability in the points earned is projected for every football player. The co-movement among the points earned by the players is analyzed to judge the overall risk of your team. Highly correlated teams are riskier than lowly correlated teams; if a player performs well, your team performs well, but if a player performs poorly, your team will also perform poorly.

The trade-off between performance and risk makes you walk a thin line. You want the most points possible, but you must take on a lot of risk to obtain the most points possible. How do you judge how much risk you should undertake in order to make sure you gain a sufficient amount of points? Luckily, Sharpe's ratio provides the means to make this decision. The drafting tool optimizes the Sharpe ratio for your fantasy football team. When the Sharpe ratio is optimized, you know that your team will produce the most points while exposing you to minimal risk.

Correlation Testing

In order to demonstrate how we arrived at the correlation model we utilized, we conducted a statistical test in order to determine to what degree correlation was actually important and whether it made sense to fill in every single statistical correlation. As can be imagined, filling in a 44,000 entry correlation matrix should be worth the effort with respect to gains in accuracy of the model. So, we decided to conduct a simulation in MATLAB to see how a randomized set of correlations affects the target of our maximization, the Sharpe Ratio.

The methodology for this test was to assign each of the 44,000 correlations a randomly determined value from 0% to 100%. The distribution of those random values was set as uniform so the mean correlation of the set of 44,000 correlations would be 50%. Once the correlations were assigned, a 12-player team was simulated and the Sharpe Ratio was calculated. This process was repeated 50,000 times for each risk metric in order to see the degree to which the correlation would affect the resulting Sharpe Ratio. After each completed test (one test was conducted per risk metric), the mean Sharpe Ratio, its standard deviation, and a 95% confidence interval around the mean were calculated. A histogram of the results was also created.

For the Standard Deviation risk metric, the test found a mean SR of 7.1 with a standard deviation of 0.29. The 95% confidence interval was [6.58, 7.72]. The resulting spread of the 50,000 SRs is shown in the histogram below:

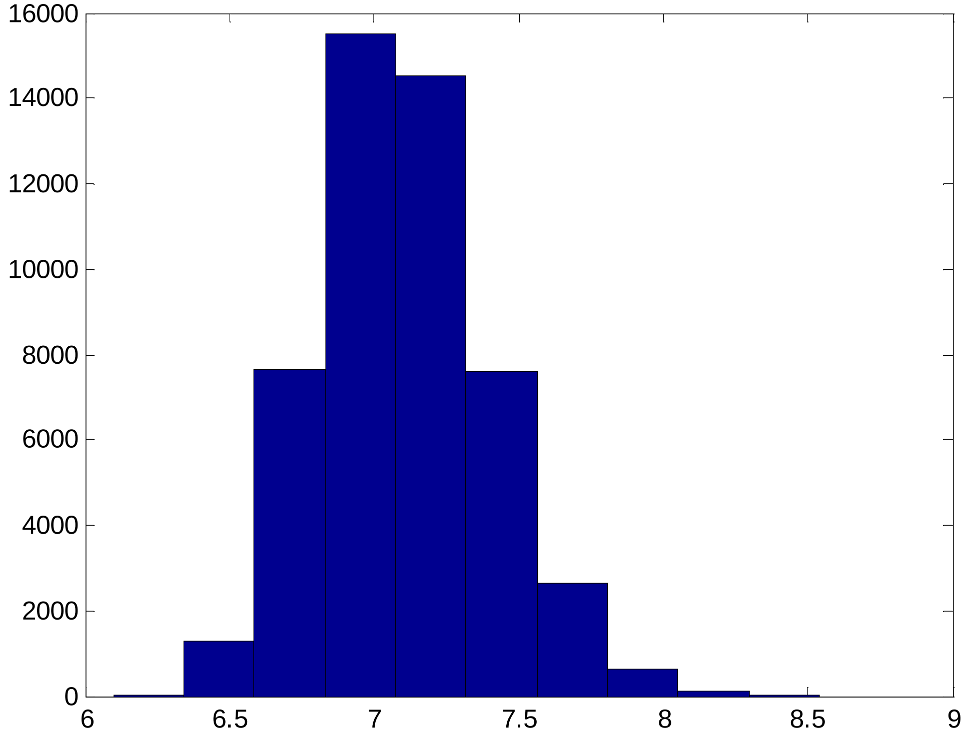

For the Projection Differences risk metric, the test found a mean SR of 10.92 with a standard deviation of 0.51. The 95% confidence interval was [10.02, 12.02]. The resulting spread of the 50,000 SRs is shown in the histogram below:

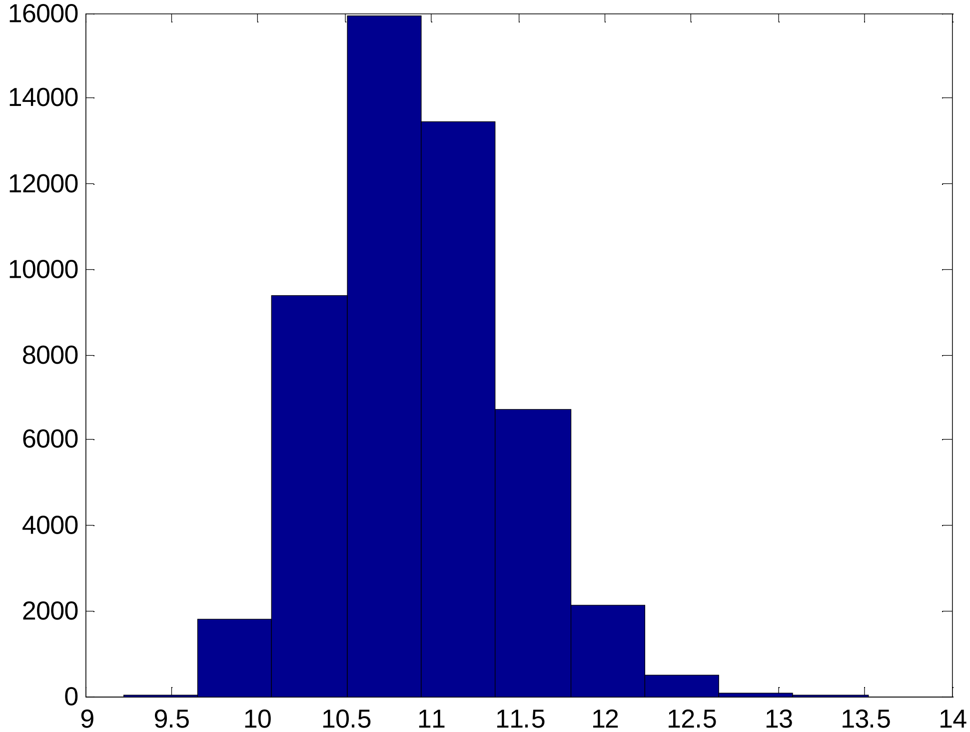

For the two-factor risk metric, the test found a mean SR of 12.67 with an identical standard deviation of 0.51. The 95% confidence interval was [11.75, 13.77]. The resulting spread of the 50,000 SRs is shown in the histogram below:

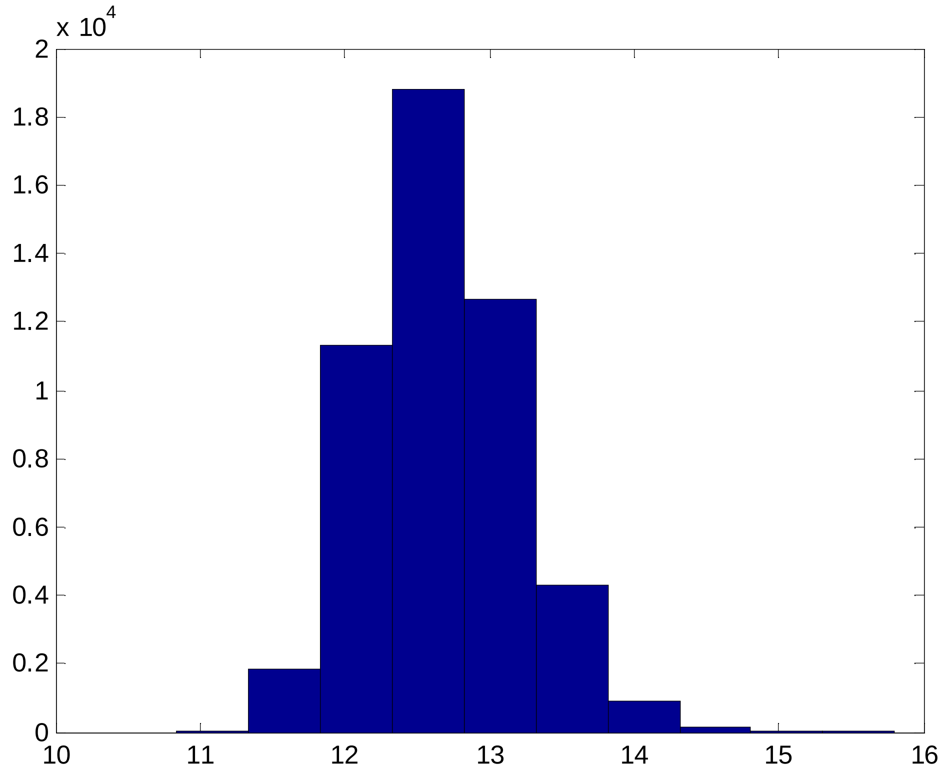

These tests demonstrate that given random correlations with a mean of 50%, the overall effect of those correlations on the Sharpe Ratio is minimal yet important. As a result of these tests, we decided to utilize a combination of the position-based correlation methodology, in addition to a number of qualitative adjustments made manually to the model to reflect a bit more of the reality of the situation.

Simulations

This section of the report is dedicated to the extensive analysis that we put towards developing the ideal simulation to test our model. This thinking also carries over to the Game Theoretic model. First, we will discuss the most popular approaches to drafting players. Then, we will outline the variations to those basic and most commonly used strategies. Finally, we will set up the simulation and present and discuss the results.

Value Based Drafting:

One commonly utilized methodology for drafting individual players is the value-based approach. In this methodology, points projections are created for each player and a "cheat sheet" is presented whereby the players are ranked by those estimates of points. Proponents of the strategy insist the unique cheat sheets demonstrate the relative depth at each position and ensure that you will pick the most valuable player available when it is your turn to draft. This basic approach to drafting can result in poor team choices because it often indicates that you select players in positions that are already filled, or which it would not make sense to fill at a given point in the draft.

Stud Running Back Theory:

Another commonly practiced drafting strategy is based upon the precept that running backs are the core of any fantasy football team because teams would, on the margin, rather run than pass. This predisposition for passing could allow running backs to be a stable baseline for a fantasy team to be built around. This strategy entails the selection of running backs, and their backups very early in the draft. The issue surrounding this methodology is that it may leave a great deal of points "on the table." This is especially true in leagues like CBS, where six points are awarded for each touchdown pass thrown by a quarterback.

Overall, we feel that there are a number of issues surrounding this highly utilized methodology. First, the point differentials between running backs and other positions can be large at specific times in the draft, and it may be hard to make up those point differentials easily throughout the season. Further, there is a specific pool of players from which individuals are drafted and the same ones are consistently picked. This is to say that there is a zero-sum game that occurs between teams for expected points and consistency, whereby one team's gain is another team's loss. Forgoing a solid, second-round pick after selecting a star running back in favor of another running back may not maximize a team's potential.

Another significant issue with this drafting strategy is that it does not take into account what the other players will do at any given decision point. Our final model, the Game Theoretic model, does indeed consider those issues. The Portfolio Theory model does not consider them, though it is instructive to think about how drafting running backs early may leave excellent picks available for your opponents.

Average Draft Position:

The last common strategy that we will discuss is based upon the concept of the Efficient Market Hypothesis in finance. Basically, this theory states that all information relating to the prices of securities is already discounted by the market. Similarly, the average drafting position of a given player among a selection of drafts is likely to contain all relevant information regarding that player. Therefore, adherence to a sample of average draft positions could serve as an excellent methodology to develop a solid team. The information contained here may be a mingling of the previous two strategies, depending on whether players follow one or the other as well as the sample of players from whom the average draft positions were calculated.

Our Approach:

Based upon the above three methodologies, we constructed rules for our opponents and ourselves. We placed those rules in a drafting strategy matrix and assigned Experts and Amateurs specific ones to follow. The matrix is included here:

	Drafting Strategy Matrix	US	Amateurs	Experts	Ranking
1	Do not draft a kicker until the last round	x		x	10
2	If a Team Defense is among top 5 or TE is among top 6, they may be picked as early as round # N/2. If not, in last 3 rounds.	x	x	x	3
3	Do not draft a second TE unless you have a top 6 TE. Do not draft second TE until last 3 rounds.	x	x	x	1
4	Always draft a third RB prior to drafting a second QB	x	x	x	7
5	If you get Brady or Manning, don't draft a second QB until you already have 2 WRs, 2 RBs	x		x	9
6	If you get Peterson or Tomlinson or Westbrook, don't get a third RB until you have 2 WRs, 2 QBs	x		x	10
7	If you get Tomlinson, Peterson, or Westbrook, disregard rule #5	x	x	x	
8	Do not draft a second defense if you have a top 5 defense	x	x	x	5
9	Do not draft a second WR, QB, or RB until the 4th round. If a RB, QB, or WR is picked prior to the fourth round, do not pick another RB, QB, or WR until at least the 4th round	x			7
10	Always draft your (X+1)th WR prior to drafting an Xth QB or RB provided x>=3	x		x	1
11	Do not draft a backup QB, WR, or RB that has same bye week as your starter(s)	x	x	x	10
12	For 2nd backup WRs and RBs, take highest potential performance figure (80% Breakout Universe)	x			8
13	Pick lowest ADP #		x	x	10

These rules are hardcoded into the statements and based off a logical model that approximates opponent strategies and develops appropriate ones for ourselves. For example, if you already have two quarterbacks and only one wide receiver, it does not make any sense to draft a third quarterback (no matter how good he is) instead of a second WR. The second WR will start for you and score points, and he will likely have many more expected points than the player who would replace him if you waited another round. The third string QB would likely never score a single point for you the entire season unless there is an injury to one of your first two players.

Rules 5 and 6 are player-specific implementations of this concept. Manning and Brady are so good that it is unlikely that you would ever want to start your #2 quarterback over them, even if Manning/Brady appeared to be in a slump and/or the #2 guy appeared to be doing well. Thus, your backup will not see playing time unless there is an injury or a bye week. On the other hand, a backup WR or a backup RB might be extremely valuable because you are unlikely to have the top guy. After all, you already have Manning/Brady and are more likely to have a guy who might be outperformed by someone below him on your depth chart. For example, if you have Brady as your QB, Larry Johnson as your RB, and Andre Johnson as your first WR, it makes more sense to grab a guy like Clinton Portis or Dwayne Bowe, both of whom have a shot at actually doing better than the Johnsons and could score you more points with a mid-season switch. In other words, you must consider not only the number of points that a player will generate in a given season, but also how likely he is to start for you given the other players that you have.

The same logic holds true for Peterson, Tomlinson, and Westbrook for running backs. However, it should be noted that even though Randy Moss is a dominant WR, you should not hold back drafting another WR once you get him because one typically starts two WRs, and thus your second and third WRs are still very valuable. Rule 8 was implemented in order to take into account the highly unlikely scenario in which a team ends up with two of the five best players. In that instance, the rules would be overly restrictive and thus would have to be relaxed.

Rule 3 involves the application of a concept known as "handcuffing." This occurs when you have spent a pick on a person that you want to protect. For instance, picking Antonio Gates in round 8 instead of another WR or RB means that you have given up the opportunity to earn points with the RB or WR. As a result, you have a vested interest in Gates and want to make sure that you can earn points with that position should he become injured. In this case, you would pick another TE later in the draft. Because TEs have about the same value apart from the top 5, this strategy would allow you to gain an edge in TEs and protect that edge later in the draft.

The rest of the rules are modeled around the fact that a second tight end, a 4th quarterback, a 4th RB or a 6th WR are unlikely to ever see the field and so should be prioritized accordingly. Of course, it is still worthwhile to pick the best players even late in the draft, as you never know when some of your top ranked players' prognostications will be false, or when one of your starters will get injured, and you will want the ability to decide who is doing the best at that point in the season between several different backups. This model does not take into account the probability that a given player will start for your team and thus these rules help restrict the model in order to maximize expected value.

Furthermore, this maximizes expected value by taking into account the lessons we learned when examining kickers and team defenses earlier and limits the temptation to take kickers before the last round (taking kickers in the last round maximizes expected value, as the players you get in the second or third to last rounds instead of them are more likely to make an impact on your points total). It allows for taking a top-five defense early, but advises against the drafting of mediocre, uncertain defenses at any level before the last or second-to-last round. Based upon the above rules, the two simulations are now presented with commentary.

There were two draft simulations that were conducted in order to assess the drafting ability of the two versions of the portfolio theoretic model. The two simulations were executed with respect to the two most commonly used prize structures. First, a "Winner Take All" prize structure was applied whereby risk was not taken into account. The reason for this is that in some Fantasy Football pools, only the winner receives any compensation. In these cases, volatility should not be considered as an objective function that requires consideration in the maximization process. Instead, only projected points should be maximized. Second, the "70%/20%/10%" prize structure was set forth for the draft simulation of the two-factor model. In this case, volatility becomes very important because finishing in the top three spots provides compensation. Therefore, excessive risk is not rewarded. In the first case, finishing second is the same as finishing last, whereas with the "70%/20%/10%" prize structure, second place is still rewarded.

Winner Take All Simulation:

As noted above, the objective function to be maximized for this simulation was simply the projected points. Only the best player in terms of expected points was taken at each decision point. For this simulation, we play against eight Amateurs (1, 2, 4, 5, 6, 8, 11, 12) and three Experts (3, 9, 10). In order to remove any advantages from going early or late, we have positioned ourselves at #7. The results of this draft were as follows:

Round	Team 1	Team 2	Team 3	Team 4
1	LaDainian Tomlinson	Adrian L. Peterson	Brian Westbrook	Steven Jackson
2	Larry Fitzgerald	Brandon Jacobs	T.J. Houshmandzadeh	Tony Romo
3	Jamal Lewis	Steve Smith	Marques Colston	Maurice Jones-Drew
4	Santonio Holmes	Julius Jones	Ben Roethlisberger	Brandon Marshall
5	Darren McFadden	Thomas Jones	Derek Anderson	Rudi Johnson
6	Rashard Mendenhall	Laveranues Coles	Kevin Curtis	Kevin Smith
7	Antonio Gates	Jason Witten	Kellen Winslow Jr	Tony Gonzalez
8	Philip Rivers	David Garrard	Bobby Engram	Derrick Mason
9	Minnesota Defense	Aaron Rogers	Jerious Norwood	Donte Stallworth
10	Nick Folk	Stephen Gostkowski	Shaun McDonald	Mason Crosby
11	Adam Vinatieri	Indianapolis Defense	Ladell Betts	Isaac Bruce
12	L.J. Smith	Owen Daniels	Tony Scheffler	Green Bay Defense
13	Jeff Garcia	Phil Dawson	Baltimore Defense	Greg Olsen
14	Ernest Wilford	Ronald Curry	Jeff Reed	Kurt Warner
Exp Points:	2304.82	2255.59	2213.26	2160.49
3 Factor:	0.00	0.00	0.00	0.00
Sharpe	0.00	0.00	0.00	0.00

Round	Team 1	Team 2	Team 3	Team 4
Ratio:				

Round	Team 5	Team 6	Team 7	Team 8
1	Joseph Addai	Tom Brady	Peyton Manning	Randy Moss
2	Laurence Maroney	Braylon Edwards	Andre Johnson	Willis McGahee
3	Chad Johnson	Reggie Bush	Ronnie Brown	Drew Brees
4	Carson Palmer	Roy Williams	Dwayne Bowe	Willie Parker
5	LenDale White	Calvin Johnson	Marvin Harrison	Hines Ward
6	Chris Chambers	Roddy White	Chester Taylor	Selvin Young
7	Dallas Clark	Chris Cooley	San Diego Defense	Jerricho Cotchery
8	Felix Jones	Ahman Green	Bernard Berrian	Justin Fargas
9	Matt Schaub	Darrell Jackson	Kevin Jones	New England Defense
10	Nate Kaeding	Tatum Bell	Brett Favre	Matt Leinart
11	Josh Brown	Jason Campbell	Reggie Williams	Shayne Graham
12	Alge Crumpler	Chicago Defense	Ron Dayne	Benjamin Watson
13	NY Giants Defense	Donald Lee	Desmond Clark	Rob Bironas
14	Santana Moss	Jason Hanson	Josh Scobee	Patrick Crayton
Exp Points:	2184.26	2291.87	2494.86	2185.36
3 Factor Risk:	0.00	0.00	0.00	0.00
Sharpe Ratio:	0.00	0.00	0.00	0.00

Round	Team 9	Team 10	Team 11	Team 12
1	Marion Barber III	Frank Gore	Larry Johnson	Clinton Portis
2	Terrell Owens	Ryan Grant	Marshawn Lynch	Reggie Wayne
3	Torry Holt	Plaxico Burress	Earnest Graham	Wes Welker
4	Edgerrin James	Greg Jennings	Anquan Boldin	Michael Turner
5	Matt Hasselbeck	Donovan McNabb	Fred Taylor	Jonathon Stewart
6	Joey Galloway	Donald Driver	Lee Evans	Matt Forte
7	Anthony	Javon Walker	Eli Manning	Jay Cutler

Round	Team 9	Team 10	Team 11	Team 12
	Gonzalez			
8	Deuce McAllister	DeAngelo Williams	Reggie Brown	Marc Bulger
9	Sidney Rice	Jake Delhomme	Jerry Porter	Nate Burleson
10	Vince Young	Pittsburgh Defense	Jon Kitna	Mark Clayton
11	Kenny Watson	Mike Furrey	Robbie Gould	Jason Elam
12	Heath Miller	Todd Heap	Vernon Davis	Jeremy Shockey
13	Dallas Defense	Shaun Alexander	Jacksonville Defense	Tennessee Defense
14	David Akers	Matt Stover	Muhsin Muhammad	Neil Rackers
Exp Points:	2202.45	2300.05	2269.03	1977.86
3 Factor Risk:	0.00	0.00	0.00	0.00
Sharpe Ratio:	0.00	0.00	0.00	0.00

In the winner-take-all model, we seek to maximize the expected value of our points. In order to do this, we attempt to draft the best starters at running back, quarterback, and wide receiver, as they represent the majority of points scored by a fantasy football team. We hold off on tight ends and kickers because as a whole, they tend to score far fewer points than the 'big three' positions. In addition, the production of the middle tier does not lag that far behind the production of the top tier, so holding off until late and taking an average player allows you to upgrade your points potential at the more important positions.

Defenses are interesting because while they score a lot of points, they are very difficult to predict. If you could anticipate with certainty which defenses would be the best, then they would get drafted relatively early. Unfortunately, there is so much change from year to year by defensive units that only the top few merit a selection higher than being taken in the very last few rounds.

Another aspect to maximizing expected value is the concept of drafting not only the best player available but also the one that fits your needs. For example, if you have Drew Brees and Carson Palmer on your team, you don't know which one you will start — you hope that at least one of them has a great season, but you might switch them off from game to game depending on how they've been playing, how their team is doing, and what opponents they face. On the other hand, if you have Tom Brady, you are going to start him every single game no matter what (unless he is hurt or on a bye), and thus, your backup is much less valuable than he would be if your starter was Drew Brees. Thus, you must factor in not only how many points a player is likely to score, but also *how many games that player is likely to start for your team.* Therefore, if you draft Brady or Peyton Manning, it would be wise to hold off on a backup quarterback until much later, allowing you to draft better wide receivers and running backs. Likewise, if you have Adrian Peterson, Ladanian Tomlinson, or Brian Westbrook, you are likely to hold off on a second running back. Below is a pick-by-pick analysis of our draft:

1. With Tom Brady off the board, and the top five running backs all drafted, we take Peyton Manning and secure one of the top two quarterbacks, who are both far ahead of any other QBs in the league at this point.
2. With most of the stud running backs gone by this point, we know that we can get an average RB at any time, so we instead draft Andre Johnson, a high quality WR.

3. Ronnie Brown is a solid RB this late. We held off until the third round, and got Brown, who we might have considered taking in the second round.
4. Dwayne Bowe is a young receiver who we project to have enormous development. He will likely be our second starter at WR and gives us great depth here.
5. Given the question marks with Bowe, because he is so young, and the general volatility of receivers in general, we take Marvin Harrison here. He is an incredibly solid veteran, and is correlated with Peyton Manning. Given that we probably need Manning to have a good year in order to win our league, we have to then assume that Harrison will likely be very solid as well.
6. Chester Taylor gives us depth at RB behind Brown.
7. This is pretty early to take a defense, but given that we have San Diego projected at #1 (although like we mentioned earlier, those rankings are highly volatile) this is a pretty solid pick that could pay off in the long run. Since this is a 'winner-take-all' draft, we are more willing to take risky picks such as this one.
8. Bernard Berrian adds depth at WR.
9. Kevin Jones adds some needed depth at RB.
10. Brett Favre gives us depth at QB. He will be playing somewhere, and he will likely be productive. Given that Manning is probably going to start for us no matter what, and has no past injury history, Favre gives us a shot if Manning goes down to a freak accident.
11. Reggie Williams adds further depth at WR.
12. Ron Dayne rounds out our RB corps

13. Desmond Clark gives us a functional starting TE. If he gets hurt or we need someone during his bye week, we'll just release our worst (or non) performing backup and pick up a second TE midway through the season. Given that we have such a great QB and a solid RB/WR corps and the top defense, we are just trying to lose as few points as possible as TE. It's a small sacrifice to maximize value in more important areas.
14. Josh Scobee gives us a functional, starting kicker.

As opposed to our model, the Amateurs make several obvious mistakes that are common for bad players. Generally, the Amateurs fail to maximize expected value because they too frequently pick good players despite the fact that those players are unlikely to ever start for them. For example, player 1 drafts two kickers, and despite the fact that he drafts stud LaDainian Tomlinson in the first round, he also takes Jamal Lewis in the third round, even though he has only one WR and no QBs. He also takes Rashard Mendenhall and Darren McFadden before he even takes his first QB. These mistakes may be a function of the Amateur erring more towards the side of the Stud Running Back Theory.

You can see similar mistakes made by Player 12 (who takes 2 K's) and player 2 (who takes 4 RB's with his first 5 picks.) While it is true that two running backs start, and that they score a lot of points, at some point one must consider the fact that one player is not very likely to get many games for your team, and in that case you could really maximize value by taking a player who will start for you every week.

The Experts perform much better. Expert #3 gets a top-three RB, then two stud WR's, then two top QBs who will contend for the starting job, followed by a top-five TE and then a slew of backups, leaving the defense and kicker until the very end. Expert #9 similarly stacks up with two WR's, two RBs, and then a solid QB (Hasselbeck) and some young, speculative talent (Anthony Gonzalez, Sidney Rice, Vince Young). Expert #10 also takes a pair of solid running backs followed by a pair of good receivers and then Donnovan McNabb. He closes with solid backups spaced out so as to maximize value from players who are most likely to get playing time.

Despite the fact that we draft seventh, by pursing an optimal strategy we are able to 'win' the simulation by maximizing the expected point value of our team. Of course, winning the simulation does not equate to winning Fantasy Football. Overall, the place finish in terms of expected points in the simulation should be considered, though it should not be weighted too heavily. Instead, we should concentrate on the composition of our team and the timing of our picks.

Top 3 Paid Simulation:

Rather than award a prize to the top player only, other leagues give prizes to the top three players. In this case, we need to factor the riskiness of players into account, rather than going full out to maximize expected point value. Our model treats each player as one component of an overall portfolio, factoring in a players' riskiness as well as his correlation to other players in the portfolio, and then uses the Sharpe Ratio to select the players who will give the whole package the greatest risk adjusted return. The results of this draft were as follows:

Round	Team 1	Team 2	Team 3	Team 4
1	LaDainian Tomlinson	Adrian L. Peterson	Brian Westbrook	Steven Jackson
2	Andre Johnson	Brandon Jacobs	Larry Fitzgerald	T.J. Houshmandzadeh
3	Jamal Lewis	Steve Smith	Marques Colston	Maurice Jones-Drew
4	Santonio Holmes	Julius Jones	Ben Roethlisberger	Carson Palmer
5	Darren McFadden	Thomas Jones	Derek Anderson	Rudi Johnson
6	Kevin Smith	Kevin Curtis	Chris Chambers	Roddy White
7	Antonio Gates	Jason Witten	Kellen Winslow Jr	Tony Gonzalez
8	Philip Rivers	David Garrard	Bernard Berrian	Felix Jones
9	Bobby Engram	Minnesota Defense	Jerious Norwood	Aaron Rogers
10	Nick Folk	Stephen Gostkowski	Shaun McDonald	Mason Crosby
11	Adam Vinatieri	Isaac Bruce	Indianapolis Defense	Josh Brown
12	Baltimore Defense	L.J. Smith	Owen Daniels	Green Bay Defense
13	Greg Olsen	Jeff Garcia	Kevin Jones	Donald Lee
14	Kurt Warner	Mike Nugent	Jeff Reed	Ronald Curry
Exp Points:	2355.66	2159.59	2407.48	2180.05
3 Factor:	116.38	113.24	117.40	107.90
Sharpe Ratio:	20.24	19.07	20.51	20.20

Round	Team 5	Team 6	Team 7	Team 8
1	Joseph Addai	Tom Brady	Peyton Manning	Randy Moss
2	Tony Romo	Laurence Maroney	Braylon Edwards	Willis McGahee
3	Chad Johnson	Reggie Bush	Chester Taylor	Drew Brees
4	Brandon Marshall	Willie Parker	Roy Williams	Greg Jennings
5	LenDale White	Calvin Johnson	Jerricho Cotchery	Hines Ward
6	Joey Galloway	Selvin Young	Laveranues Coles	Matt Forte
7	Dallas Clark	Chris Cooley	San Diego Defense	Rashard Mendenhall
8	Ahman Green	Derrick Mason	Reggie Williams	Justin Fargas
9	Donte Stallworth	Matt Schaub	Adrian Peterson	Darrell Jackson
10	Matt Leinart	Nate Kaeding	Tatum Bell	Vince Young
11	Shayne Graham	Robbie Gould	Brett Favre	Mike Furrey
12	Alge Crumpler	Benjamin Watson	Tony Scheffler	Chicago Defense
13	NY Giants Defense	Dallas Defense	Shaun Alexander	Rob Bironas
14	Jason Hanson	Santana Moss	Kris Brown	Randy McMichael
Exp Points:	2158.5	2209.12	2324.21	2198.37
3 Factor Risk:	110.11	110.69	111.08	116.85
Sharpe Ratio:	19.60	19.96	20.92	18.81

Round	Team 9	Team 10	Team 11	Team 12
1	Marion Barber III	Frank Gore	Larry Johnson	Clinton Portis
2	Terrell Owens	Ryan Grant	Marshawn Lynch	Reggie Wayne
3	Torry Holt	Plaxico Burress	Earnest Graham	Wes Welker
4	Edgerrin James	Greg Jennings	Anquan Boldin	Michael Turner
5	Matt Hasselbeck	Donovan McNabb	Fred Taylor	Jonathon Stewart
6	Joey Galloway	Donald Driver	Lee Evans	Matt Forte
7	Anthony Gonzalez	Javon Walker	Eli Manning	Jay Cutler
8	Deuce McAllister	DeAngelo Williams	Reggie Brown	Marc Bulger
9	Sidney Rice	Jake Delhomme	Jerry Porter	Nate Burleson
10	Vince Young	Pittsburgh Defense	Jon Kitna	Mark Clayton
11	Kenny Watson	Mike Furrey	Robbie Gould	Jason Elam
12	Heath Miller	Todd Heap	Vernon Davis	Jeremy Shockey
13	Dallas Defense	Shaun Alexander	Jacksonville Defense	Tennessee Defense
14	David Akers	Matt Stover	Muhsin Muhammad	Neil Rackers
Exp Points:	2202.45	2300.05	2269.03	1977.86
3 Factor Risk:	0.00	0.00	0.00	0.00
Sharpe Ratio:	0.00	0.00	0.00	0.00

Below is a pick-by-pick analysis of the overall draft:

1. Again, we start off with Peyton Manning with the same rationale as before.
2. Here, we grab Braylon Edwards rather than Andre Johnson, because his three-factor risk is lower.
3. We take Chester Taylor again, but this time ahead of Ronnie Brown, as his three factor risk is again lower.

4. Roy Williams gives us another very solid, reliable starter at WR
5. Jerricho Cotchery adds some depth. Notice how long we are holding off on a backup QB here, once again.
6. Laveranues Coles adds more WR depth, and adds another potential starter.
7. Again, we take the San Diego defense, the #1 projected squad.
8. Reggie Williams is another very safe, reliable receiver to round out our receiving corps.
9. Adrian Peterson (the not-as-great one) gives us a very stable backup to Taylor.
10. Tatum Bell adds depth at RB.
11. Once again, we find ourselves with Brett Favre for the exact same reasons as before.
12. Tony Scheffler is an extremely consistent tight end.
13. Shaun Alexander finalizes our RB depth.
14. Kris Brown, a kicker in the last round.

Notice that our team is not only the third overall best team from a points perspective but that it is also the fourth least volatile team from a risk perspective. No other team managed to get in the top six for both. As a result, we have maximized our Sharpe Ratio and have actually achieved the highest Sharpe Ratio of any team, meaning that we anticipate the greatest long term, risk-adjusted projections. In other words, we have assembled a team that in the long run is both likely to do well and do so consistently. As a result, we will reach the top-three paid spots a disproportionately large percentage of the time, such that our investment has a positive expected return.

Testing of the Portfolio Theory Model

Once the portfolio theory model was finalized, it required a great deal of testing. The rules were reviewed and drafts were simulated to make sure no logical loop would cause a player to stop picking or not be able to continue picking. The rigorous process involved reviewing each draft for logical errors as well as failure to adhere to the rules of the simulation. This process was undertaken hundreds of times to ensure consistency and applicability of the drafting strategies.

SECTION IV: Gambles and Sure Things Player Categories

Unlike other Fantasy Football projections, which simply rank players from best to worst and perhaps also project points, we present in this section a number of lists that rank Fantasy Football players based upon our measure of their riskiness. Players from each main position are ranked below according to their three factor percentage of points and their risk adjusted points. The interpretation and analysis of the tables follow their presentation.

Running Backs

Three-Factor % of Points

Rank	Gambles	% 3 Factor	Rank	Sure Things	% 3 Factor
1	Michael Turner	116%	1	LaDainian Tomlinson	15%
2	Ladell Betts	71%	2	Brian Westbrook	15%
3	Cadillac Williams	44%	3	Adrian L. Peterson	16%
4	Julius Jones	44%	4	Steven Jackson	16%
5	Jerious Norwood	42%	5	Frank Gore	17%
6	Warrick Dunn	41%	6	Marshawn Lynch	17%
7	Selvin Young	39%	7	Joseph Addai	17%
8	Ron Dayne	38%	8	Willis McGahee	18%
9	DeShaun Foster	37%	9	Clinton Portis	19%
10	Kenny Watson	36%	10	Marion Barber III	19%

Top 30 Risk Adjusted Running Backs

Rank	Player	Risk Adjust
1	LaDainian Tomlinson	6.62
2	Brian Westbrook	6.48
3	Adrian L. Peterson	6.41
4	Steven Jackson	6.13
5	Frank Gore	5.97
6	Marshawn Lynch	5.93
7	Joseph Addai	5.72
8	Willis McGahee	5.51
9	Clinton Portis	5.37
10	Marion Barber III	5.34
11	Larry Johnson	5.33
12	Reggie Bush	5.23
13	Maurice Jones-Drew	4.99
14	Chester Taylor	4.61
15	Laurence Maroney	4.30
16	LenDale White	4.14
17	Brandon Jacobs	4.09
18	Jamal Lewis	3.89
19	Willie Parker	3.89
20	Ryan Grant	3.81
21	Darren McFadden	3.75
22	Jonathon Stewart	3.71
23	Kevin Smith	3.70
24	Fred Taylor	3.70
25	Matt Forte	3.62
26	Ronnie Brown	3.40
27	Earnest Graham	3.38
28	Adrian Peterson	3.30
29	Tatum Bell	3.28
30	Rashard Mendenhall	3.23

The top of this ranking chart should look familiar, as LT has been the top running back for several years running now. As this is a risk-adjusted chart, Brian Westbrook is ranked ahead of Peterson. This is because while we project Peterson to have slightly more expected points than Westbrook, we also acknowledge that he is somewhat more risky because we only have one year of data on him and cannot draw any hard conclusions yet. Thus, when adjusted for risk he slips barely below Westbrook by 0.05 risk-adjusted points.

Like Peterson, Marshawn Lynch also slides a bit due to his aforementioned projection differences variability rating, which he has because of his high projected improvement as a second-year player. The big gainers on the risk adjusted model are players in their prime who on one hand are old enough to have proved themselves, but are also not so old as to be injury prone or hitting the wall. Steven Jackson, Westbrook, and Frank Gore are all great examples of this.

It should also be noted that conventional wisdom surrounding a few of these players may be incorrect. LenDale White has a reputation for being out of shape and lazy since he performed so poorly at the combine and earned a bad reputation at USC. Still, since he has joined the NFL he has shown slow but extremely steady improvement. Similarly, the first thing that comes to the mind of many football fans when they hear the name 'Willis McGahee' is a visual of his graphic and traumatizing knee injury in college. However, McGahee has not experienced a significant injury since then, and has given every indication that he will continue with his clean bill of health with a probability equal to that of any other running back.

Quarterbacks

Quarterbacks projected points and volatilities vary based upon the scoring system utilized. Effectively, there are two scoring systems since the main determinant of change from the projections is points per passing touchdown. For Yahoo and ESPN, there are 4 points awarded per passing touchdown, whereas there are 6 points awarded for passing touchdowns for CBS Sportsline, Fox Sports, and NFL.com. The appropriate tables are presented separately with the ones for Yahoo and ESPN coming first.

Three-Factor % of Points

Rank	Gambles	% 3 Factor	Rank	Sure Things	% 3 Factor
1	John Beck	46%	1	Tom Brady	8%
2	Matt Ryan	33%	2	Peyton Manning	9%
3	Brian Griese	29%	3	Tony Romo	10%
4	David Carr	27%	4	Drew Brees	10%
5	Charlie Frye	26%	5	Ben Roethlisberger	10%
6	Kyle Boller	26%	6	Derek Anderson	10%
7	Aaron Rogers	26%	7	Carson Palmer	11%
8	J.P. Losman	24%	8	Matt Hasselbeck	11%
9	Cleo Lemon	23%	9	Marc Bulger	11%
10	Rex Grossman	23%	10	Donovan McNabb	11%

Top 30 Risk Adjusted Quarterbacks

Rank	Player	Risk Adjust
1	Tom Brady	13.00
2	Peyton Manning	11.52
3	Tony Romo	9.97
4	Drew Brees	9.68
5	Ben Roethlisberger	9.62
6	Derek Anderson	9.56
7	Carson Palmer	9.38
8	Matt Hasselbeck	9.17
9	Marc Bulger	9.16
10	Donovan McNabb	9.05
11	Jay Cutler	8.21
12	David Garrard	8.12
13	Philip Rivers	8.06
14	Jon Kitna	7.68
15	Eli Manning	7.60
16	Jeff Garcia	7.31
17	Brett Favre	6.94
18	Jason Campbell	6.91
19	Jake Delhomme	6.73
20	Vince Young	6.63
21	Matt Schaub	6.32
22	Alex Smith	5.83
23	Tarvaris Jackson	5.50
24	Kurt Warner	5.45
25	Sage Rosenfels	5.25
26	Chad Pennington	5.09
27	Damon Huard	5.07
28	Matt Leinart	4.98
29	Trent Edwards	4.86
30	Mark Brunell	4.83

The above model demonstrates the value of quarterbacks when you take into account their risk as well as their mean projection (reward). Of course, any quarterback chart is topped with Tom Brady and Peyton Manning, and this one is no exception.

Interestingly, they are followed by Romo and Brees, who have only the 5th and 6th highest projections, respectively. These two players get bumped up the chart due to their season-to-season consistency and replace the oft-injured Donnovan McNabb (who is brilliant when healthy). This risk model factors into account McNabb's physical fragility as well as the fact that if he goes down there is a slight chance that touted backup Kevin Kolb takes the position and doesn't relinquish it.

Similarly, Kurt Warner and Matt Leinart are both hurt by the quarterback race in Arizona, as nobody knows who will start and how short the starters' leashes will be. If Leinart starts and does poorly in the first few games, he might not get a chance to dig himself out, and he could get benched in favor of Warner. On the other hand, if Leinart starts and does well, Warner might never see the field. Thus, there is a huge risk associated with drafting either quarterback, as there is a significant possibility that they do not even play. The same 'unlikley to start' tag is applied to Huard, Pennington, and Brunell, as well as several others, such as Alex Smith or Jason Campbell.

Quarterbacks (6 points per Touchdown)

Three-Factor % of Points

Rank	Gambles	% 3 Factor	Rank	Sure Things	% 3 Factor
1	John Beck	38%	1	Tom Brady	6%
2	Matt Ryan	26%	2	Peyton Manning	7%
3	David Carr	23%	3	Tony Romo	8%
4	Brian Griese	23%	4	Derek Anderson	8%
5	Charlie Frye	22%	5	Ben Roethlisberger	8%
6	Kyle Boller	21%	6	Drew Brees	9%
7	J.P. Losman	21%	7	Carson Palmer	9%
8	Cleo Lemon	20%	8	Matt Hasselbeck	9%
9	Rex Grossman	19%	9	Donovan McNabb	9%
10	Joey Harrington	18%	10	Jay Cutler	10%

Top 30 Risk Adjusted Quarterbacks

Rank	Player	Risk Adjust
1	Tom Brady	15.61
2	Peyton Manning	13.97
3	Tony Romo	12.47
4	Derek Anderson	11.87
5	Ben Roethlisberger	11.84
6	Drew Brees	11.72
7	Carson Palmer	11.41
8	Matt Hasselbeck	11.35
9	Donovan McNabb	10.65
10	Jay Cutler	9.93
11	Philip Rivers	9.81
12	David Garrard	9.66
13	Marc Bulger	9.51
14	Eli Manning	9.47
15	Jon Kitna	9.28
16	Jeff Garcia	8.66
17	Brett Favre	8.38
18	Jake Delhomme	8.36
19	Jason Campbell	8.09
20	Aaron Rogers	7.78
21	Vince Young	7.39
22	Matt Schaub	7.35
23	Kurt Warner	6.79
24	Sage Rosenfels	6.69
25	Tarvaris Jackson	6.40
26	Alex Smith	6.38
27	Damon Huard	6.10
28	Chad Pennington	6.09
29	Matt Leinart	5.76
30	Mark Brunell	5.72

Switching to the 6 points per touchdown system obviously benefits quarterbacks who throw more touchdowns. It also increases the volatility of the highest-rated quarterbacks (and therefore those that throw a lot of TDs). Due to the law of large numbers, the small numbers of touchdowns that a quarterback throws are more volatile than the larger numbers of yards and completions.

With the six-point-per-passing-TD system, Roethlisberger, Anderson, and Rivers all get a moderate bump. Other players like Drew Brees drop a bit, as they play a more conservative style. This system places less emphasis on INTs because now one touchdown cancels out three INTs rather than just two, which further benefits these 'gunsligners,' who throw deep more often and throw riskier passes and more touchdowns.

Wide Receivers

Three-Factor % of Points

Rank	Gambles	% 3 Factor	Rank	Sure Things	% 3 Factor
1	Michael Clayton	55%	1	Randy Moss	12%
2	Nate Burleson	45%	2	Terrell Owens	13%
3	Isaac Bruce	43%	3	Braylon Edwards	13%
4	Shaun McDonald	39%	4	Reggie Wayne	13%
5	Muhsin Muhammad	37%	5	Larry Fitzgerald	14%
6	Mark Clayton	37%	6	Andre Johnson	14%
7	Mike Furrey	36%	7	Marques Colston	15%
8	Chris Chambers	36%	8	Steve Smith	16%
9	Derrick Mason	35%	9	T.J. Houshmandzadeh	16%
10	Bobby Engram	35%	10	Chad Johnson	17%

Top 30 Risk Adjusted Wide Receivers

Rank	Player	Risk Adjust
1	Randy Moss	8.41
2	Terrell Owens	7.99
3	Braylon Edwards	7.81
4	Reggie Wayne	7.56
5	Larry Fitzgerald	7.28
6	Andre Johnson	7.23
7	Marques Colston	6.57
8	Steve Smith	6.31
9	T.J. Houshmandzadeh	6.11
10	Chad Johnson	5.90
11	Torry Holt	5.78
12	Plaxico Burress	5.47
13	Greg Jennings	5.33
14	Roy Williams	5.22
15	Santonio Holmes	4.85
16	Wes Welker	4.84
17	Calvin Johnson	4.82
18	Marvin Harrison	4.68
19	Brandon Marshall	4.66
20	Anquan Boldin	4.65
21	Dwayne Bowe	4.60
22	Jerricho Cotchery	4.59
23	Laveranues Coles	4.12
24	Bernard Berrian	4.09
25	Donald Driver	4.09
26	Roddy White	3.96
27	Hines Ward	3.87
28	Kevin Curtis	3.86
29	Reggie Brown	3.84
30	Joey Galloway	3.68

Randy Moss obviously tops the WR chart, but it is somewhat less obvious that Owens would come in second. While Owens is likely the second most talented WR in the league, he has had nagging injury and character issues in the past. However, Owens' team relies on him heavily, so in terms of football risk, he has very little. The Cowboys will find ways to get him the ball and keep his production high. Braylon Edwards and Reggie Wayne also get big bumps for their year-in-and-year-out consistency, while younger players who we have projected to improve a lot this year, such as Marques Colston (5th on our projections, 7th risk adjusted) and Dwayne Bowe (10th on our projections, 20th risk adjusted), slide a bit due to the inherent risk of the increasing inaccuracy of larger projected differences.

Tight Ends

Three-Factor % of Points

Rank	Gambles	% 3 Factor	Rank	Sure Things	% 3 Factor
1	George Wrighster	41%	1	Chris Cooley	11%
2	Visanthe Shiancoe	35%	2	Heath Miller	12%
3	Marcus Pollard	30%	3	Tony Gonzalez	13%
4	Reggie Kelly	29%	4	Antonio Gates	14%
5	Benjamin Watson	28%	5	Jason Witten	14%
6	Greg Olsen	28%	6	Desmond Clark	15%
7	Alge Crumpler	27%	7	Randy McMichael	16%
8	L.J. Smith	26%	8	Kellen Winslow Jr	17%
9	Marcedes Lewis	24%	9	Zach Miller	17%
10	Daniel Graham	24%	10	Ben Utecht	18%

Top 30 Risk Adjusted Tight Ends

Rank	Player	Risk Adjust
1	Chris Cooley	9.11
2	Heath Miller	8.17
3	Tony Gonzalez	7.61
4	Antonio Gates	7.32
5	Jason Witten	7.09
6	Desmond Clark	6.47

7	Randy McMichael	6.13
8	Kellen Winslow Jr	6.04
9	Zach Miller	5.89
10	Ben Utecht	5.65
11	Donald Lee	5.34
12	Dallas Clark	5.05
13	Chris Baker	4.91
14	Bo Scaife	4.86
15	Todd Heap	4.76
16	Jeff King	4.61
17	Vernon Davis	4.54
18	Jeremy Shockey	4.54
19	Tony Scheffler	4.51
20	Owen Daniels	4.43
21	Eric Johnson	4.39
22	David Martin	4.25
Rank	**Player**	**Risk Adjust**
23	Daniel Graham	4.20
24	Marcedes Lewis	4.13
25	L.J. Smith	3.85
26	Alge Crumpler	3.74
27	Greg Olsen	3.60
28	Benjamin Watson	3.54
29	Reggie Kelly	3.50
30	Marcus Pollard	3.38

Tight ends are a fairly cut-and-dried position. There are five players this year who are significantly better than the rest, and beyond that, the players are all very similar. Gates, Witten, Winslow Jr., Gonzalez, and Cooley may be worth a mid-round selection, but after that you will likely be able to pick up a slightly worse tight end much later in a draft.

This chart is informative in that it demonstrates what we already could have guessed — that many of the mid-level tight ends are extremely consistent but relatively few are explosive performers. Thus, when it comes time to either, select your backup (because you already have one of the top guys) or grab your first tight end very late after the good ones are gone. Heap or Daniels might be a good, safe value if they are still left and the top six have been taken.

SECTION V: Game Theoretic Approach to Drafting

We present a system for making optimal choices in a fantasy football draft. This system utilizes a game theoretic approach to model the draft. The resulting model is in the form of a game tree, to which we can then apply an artificial intelligence (AI) algorithm to calculate theoretically optimal strategies. Implementing the model and algorithm in software provides an “assistant” that can advise a fantasy football owner at each step of the drafting process. The advantage of using a system based in game theory is that it takes into account the likely behavior of other owners in the draft rather than simply using the expected point values of the players.

The basic operation of the system is as follows: At each round, the AI algorithm “thinks ahead” a certain number of rounds in the draft, considering all reasonable moves it might make and assuming that the other owners follow some pre-determined strategy. It then evaluates each of the possible outcomes and recommends choosing the player that leads to the strongest team relative to the other owners.

A key element of this approach is modeling the strategies of the other owners. This can be done either by manually assigning each owner to one of a few basic strategies or by automatically extrapolating their strategies from prior drafts in which they have participated.

Prerequisite Formalities

The information prerequisite to modeling a fantasy football draft will be represented by a collection of sets and functions. First, we present a single-year model, which accounts for all the information necessary to model a draft for a given year but does not allow for the representation of drafts in prior years (used for the Opponent Modeling methods, see below).

The participants in a draft are represented by O, the set of owners. One particular element of O, denoted o*, will be used when we want to distinguish a single owner as the "protagonist" of the modeling, i.e., the owner who is employing the techniques described here. The set of all other owners, O', would then be defined by set subtraction as

$$O' = O \setminus \{o^*\}.$$

Of course, the fundamental object in a fantasy football draft is a player in the NFL or some other football league. We let P denote the set of players, though these may not be literal "players," just as in the case of defensive/special-teams units, which are treated as single, atomic players, even though in reality they are not.

We represent the set of teams in the real league as T and represent the membership of players on teams by the function

$$t : P \to T.$$

For example, we might have t(DrewBrees) = NewOrleans.

Another property of a player is the position he plays. We let R represent the set of positions (or, more generally, roles). A typical value for R might be

$$R = \{QB, RB, TE, WR, DST\},$$

representing quarterback, running back, tight end, wide receiver, and defense/special teams, respectively.

The function assigning a position to each player is

$$r : P \to R$$

Continuing our example, we would have r(DrewBrees) = QB.

We would also like to know when each player will be playing, and potentially the team against which he will play. The set W will represent the weeks in the regular season, i.e.,

$$W = \{1, 2, ..., 17\}$$

for a standard NFL season. The function

$$c : T \times W \to T \cup \{bye\}$$

will give, for each team and each week, the opposing team for that week, or it will indicate that the team has a bye. For example, for the NFL 2008 season, c(SanDiego,14) = Oakland, and c(Indianapolis,4) = bye.

An effective strategy for fantasy sports drafting requires that we predict the number of points that each player will earn each week. Because points are calculated from the player's weekly statistics, we will represent the set of statistics that may be used in such a calculation as K. A typical value of K might be

K = {PassYds, Rush/RecYds, PassTDs, Rush/RecTDs, OffensiveInt/Fumbles, ExtraPts, FieldGoals, DefensiveTurnovers, Sacks, Safeties}.

Thus, the set of possible statistics a player might have is some subset of $\Re^K$. For example, the vector

$$\langle 0,191,0,2,0,0,0,0,0,0 \rangle$$

would represent Issac Bruce's performance in week 7 of the 1995 season: 191 receiving yards with two receiving touchdowns. As most players only attain numbers in a few categories, it is typical that such a vector would largely consist of zeros.

Each league will also have a scoring system that values a player's performance with a certain number of points. Most generally, this is a function

$$v : N^K \to \Re .$$

However, we can obtain a simpler model because most leagues assign a fixed number of points to each statistic, with the net point value of a player's performance being a simple linear combination of his statistics. For example, each passing touchdown might be worth 4 points. In this case, the scoring system could be represented as a dot product with the vector of values for each "accomplishment." The only difficulty is with systems that assign a point value to every discrete "chunk" of accomplishment, like giving 1 point for every 10 receiving yards. This challenge can be overcome by the proper definition of K, so that yardage is recorded in the proper unit. Thus, Issac Bruce's vector would instead be written as

$$\langle 0,19,0,2,0,0,0,0,0,0 \rangle .$$

(Note that 19.1 is rounded down to 19, as there are no partial points in fantasy football.)

With a definition of K that recorded passing yards in units of 25 yards and receiving yards in units of 10 yards, a typical scoring system would be represented by a dot product with the vector

$$\langle 1,1,4,6,-2,1,3,2,1,2 \rangle .$$

In this case, Isaac Bruce's performance would be valued at

$$\langle 0,19,0,2,0,0,0,0,0,0 \rangle \bullet \langle 1,1,4,6,-2,1,3,2,1,2 \rangle = 31 .$$

A scoring system with this property will be called a linear scoring system.

Estimating the Value of a Player

We will treat the estimation of the point value of a player as being based on a (possibly implicit) probabilistic model of the player's performance in which the statistic vector is a random variable. This model will produce an expected value for the statistic vector. If this model has no parameters (apart from the player himself), then we call it a static model. On the other hand, if the model depends on factors such as which team the player is facing or when in the season a game occurs, we call it a dynamic model.

For now, we shall consider static models of player performance. A "complete" model will give expected statistics for every player in the league, so a complete static model's estimations can be represented by function

$$e : P \to K .$$

Because the values of e are probabilistic expectations, in the case of a linear scoring system we can simply use the composite function

$$v \circ e$$

to calculate the expected point value of a player given a particular model and scoring system. Note that by linearity of expectation, this is exactly the probabilistic expectation that the model would generate directly.

We will assume the existence of a common-knowledge static model of player performance with estimations given by $\tilde{e}$. This is the model we will attribute to all other owners in the draft. We will also assume the existence of a separate and presumably superior model accessible only to us, with estimators given by e'.

In the more sophisticated approach to this problem, we will be building models of the drafting strategies of the other owners. These models will be based on the behavior of the owners in prior drafts, and so we will need to relate the decisions they made to their presumed player models *at the time*. For this purpose, we will need versions of $\tilde{e}$ from prior years. The estimator of the common-knowledge model for the year *y* will be denoted $\tilde{e}_y$.

A simple approach to valuing a potential draftee *p* would be to just determine the estimated points-per-game (ppg) he will earn according to your model. That is, calculating

$$v \circ e'(p).$$

Thus, one might make the goal of a drafting system be to maximize

$$\sum_p v \circ e'(p),$$

where the sum is over the first-string players you have drafted. (In considering the value of a player we would also want to consider the weeks in which they have byes, when second-string players would be used.) However, a more effective system can be designed by taking a game-theoretic approach and recognizing that maximizing the sum of estimated points-per-game may not be equivalent to maximizing success in the fantasy league. First, note that while a draft is a game with finite resources, the impact of one player gaining some portion of those resources is not distributed equally among all of his opponents. Second, in a head-to-head league, success in the league is determined by winning weekly match-ups. The actual objective should be to maximize final placing in the fantasy league (1st, 2nd, 3rd, etc.).

The method of determining the league champion, runner-up, etc. will vary from league to league. For simplicity's sake, we will assume that given an allocation of players to owners, represented by a function

$$a : O \to 2^P ,$$

we can simulate the entire season of the league (using the calendar, our estimator e', etc.) and obtain the ordinal ranking of owners. We can then assign a value to those ordinal rankings based either on cash prizes awarded in the league or some other valuation of success.

The ability to simulate the results of the league given the outcome of the draft means that we can now represent the draft as a complete-knowledge, *n*-player, extensive-form game, or, equivalently, as a deterministic Markov game.

Solving the Game

In order to make this Markov game feasible, we will build deterministic models of every other owner's behavior during the draft, and we will assume that an owner makes his decision at each step in the draft based only on the set of available players and the set of players he has accumulated. Such a drafting strategy can be represented by a function

$$\sigma : T \times T \to P,$$

where the first component represents available players, and the second represents the players selected by the owner thus far.

The assumption of deterministic opponents has reduced our problem to that of a simple game tree. Of course, we will want to take into account the actual decision of the other owners, who will almost certainly diverge from our models. Thus, at each decision point we will build a new game tree, taking into account decisions that have been made. Still, this is not a tractable problem---if we are drafting *n* players, the number of leaves of the tree when we make our *i*th decision will be of order

$$|P|^{(n-i)}.$$

To further reduce the complexity of the problem, we will restrict ourselves to considering, of the players available for each role (position) at each decision point, only the highest values of $v \circ e'(p)$. An example of this is that we would never consider Drew Brees if Tom Brady were available. This reduces the number of leaves to

$$\left(2|R|\right)^{(n-i)},$$

which is still exponential in *n*.

The key to making this problem tractable is to place a bound on how deep we will go in an exploration of the game tree. Once we reach a node at a particular depth, we will calculate the value of that node in some way. For example, we could fill in the remaining spots in each player's roster with "generic" players for each position, each of whom will have the estimated value of the best available player for his position. Another option (the one used in the current implementation) is to pick a simple (i.e., low computational cost) strategy for ourselves and simulate the rest of the draft to derive fictional rosters. With these fictional rosters, we can then simulate the fantasy league season and calculate a value for the node.

If our exploration bound is *b*, we will only have to perform this calculation

$$\left(2|R|\right)^{\min\{k,(n-i)\}}$$

times at our *i*th decision point. We can calculate an upper bound on the number of simulations as

$$\left(2|R|\right)^{k}$$

If we use the typical value of *R* given above, we have

$$|R| = 5,$$

and our upper bound for various values of *k* is:

k	$\left(2\|R\|\right)^{k}$
2	100
3	1,000
4	10,000
5	100,000
6	1,000,000
7	10,000,000

While there is an exponential time cost associated with higher values of k, the further we explore the tree, the more accurate our node value estimates will be.

Thus our algorithm is:

At each decision point:

- Build a game tree of depth no more than *k*, deterministically simulating the behavior of the other owners.
- Fill in the partial rosters at the leaves of the game tree with "generic" players in each position.
- Simulate a complete fantasy league season with each of those roster configurations to assign values to each leaf.
- Choose the player according to the first branch

Learning Models of Opponent Strategies

The most difficult aspect of this system is the development of models of your opponents' draft strategies. Ideally, we would have access to a large number of prior drafts in which our opponents previously participated. From this data, we would derive models to allow us to predict our opponents' behavior. However, even the previously mentioned set of limited strategy functions that are deterministic and ignore other owners' teams,

$$\sigma : T \times T \to P,$$

is enormously large.

The real problem with this approach is that prior drafts likely occurred during previous seasons and thus have different pools of players. Consequently, we want to abstract players to their relevant properties so that we can extrapolate from an owner's decisions with respect to a prior pool of players to his behavior with respect to the current pool of players.

We posit the existence of a reasonable property space of players, denoted Π. This might include the player's position and expected points-per-game according to the common-knowledge estimator $\tilde{e}_y$ for the appropriate year. That is,

$$\Pi = R \times \Re .$$

(To obtain a smaller property space, we might discretize the expected ppg. Alternatively, we could replace the expected ppg with the ordinal ranking of the player's ppg.)

In the current implementation opposing owners may be assigned one of two strategy models: Amateur or Expert. These strategies are identical to the specifications of Amateur and Expert in the Portfolio Theory section. To provide a reminder of the rules followed by Amateurs, Experts, and Us, the Drafting Strategy matrix is presented again below:

	Drafting Strategy Matrix	US	Amateurs	Experts	Ranking
1	Do not draft a kicker until the last round	X		x	10
2	If a Team Defense is among top 5 or TE is among top 6, they may be picked as early as round # N/2. If not, in last 3 rounds.	X	x	x	3
3	Do not draft a second TE unless you have a top 6 TE. Do not draft second TE until last 3 rounds.	X	x	x	1
4	Always draft a third RB prior to drafting a second QB	X	x	x	7
5	If you get Brady or Manning, don't draft a second QB until you already have 2 WRs, 2 RBs	X		x	9
6	If you get Peterson or Tomlinson or Westbrook, don't get a third RB until you have 2 WRs, 2 QBs	X		x	10
7	If you get Tomlinson, Peterson, or Westbrook, disregard rule #5	X	x	x	
8	Do not draft a second defense if you have a top 5 defense	X	x	x	5
9	Do not draft a second WR, QB, or RB until the 4th round. If a RB, QB, or WR is picked prior to the fourth round, do not pick another RB, QB, or WR until at least the 4th round	X			7
10	Always draft your (X+1)th WR prior to drafting an Xth QB or RB provided x>=3	X		x	1
11	Do not draft a backup QB, WR, or RB that has same bye week as your starter(s)	X	x	x	10
12	For 2nd backup WRs and RBs, take highest potential performance figure (80% Breakout Universe)	X			8
13	Pick lowest ADP #		x	x	10

The most feasible approach, taken above, is to generate a few basic and plausible models of drafting strategies and then assign each opposing owner to one of the strategies either by statistically fitting prior draft behavior, or simply making intuitive estimates. Another benefit of these rules is the reduction of the game tree to a manageable size. Specifically, the AI does not need to go through a virtually infinite series of possibilities in order to play out the entire draft at each decision point. Therefore, since it can now foresee every conceivable decision, it would be impossible to do better against the specified opponents while following the Drafting Strategy matrix.

Implementation

The current demonstration implements several elements of the model. The system is implemented as a web service written in the OCaml language (version 3.10) on top of the Ocsigen framework (version 1.1.0). In order to verify that the implementation was correct and error-free, an automated testing framework was developed. This framework simulated the operation of the system by having the AI algorithm participate in drafts with other (simulated) owners. The other owners were pitted against each other in a simulation meant to reflect a real draft. There were eight Amateurs, three Experts, and Us.

The selections made at each round of each draft were verified by the framework to be legal selections according to the appropriate draft rules. For each of the five league styles (nfl.com, ESPN, etc.), the number of owners in the simulation ranged from 4 to 12. This yielded 54 test scenarios, each of which was simulated and verified by the testing framework. The successful running of the error framework confirmed that all non-interface elements of the application (data management, draft simulation, AI algorithm) run without producing errors in the full range of real-world scenarios.

Additionally, manual testing was performed via the web interface to ensure that the algorithm produced intuitively reasonable suggestions and that the interface was error-free. A system for presenting error messages was incorporated into the web interface so that the user could continue using the system even if he provided invalid input. Examples of invalid input would be trying to add two owners with the same name or trying to select a player who would produce a roster that violates the league rules. When such an event occurs, an informative message is presented (e.g., "Illegal selection.") and the system state is left unchanged.

Summary Explanation of Game Theoretic Drafting Methodology

The above presentation of the drafting methodology is highly technical and explained from the point of view of practitioners able to understand set theory and game theoretic techniques. A more accessible presentation of the material may benefit the reader, so this section will review precisely what the Game Theoretic model of draft choices does at each stage. The web-based software application, hosted on a commercial server, is available through the following Web site: http://dev.5hut.net/draft/

The point of this web-based application is to generate a draft that an owner can use to maximize his or her expected points at the end of the season. The way that this application arrives at that list of players is through a thorough review of a number of factors at each stage of the decision-making process. First, the owner assigns a strategy to the opponents in the draft. The strategy choices are detailed above but will be summarized briefly. If your opponent is the type to simply pick the best player available, then the opponent would be assigned to the Amateur category. If your opponent is far better at drafting, and would consider the position highly in the draft and also be aware of the importance of having at least an Acceptable player as a starter for each position, the Expert category would be assigned.

At each decision point, your opponent picks who the strategy assigns (or who you input for your opponent). Then, when it is your turn, the algorithm suggests a pick based upon the strategies of the opposing players, the available players, the chosen players, and the results of a simulation of the season to conclusion. After considering all of those factors, the algorithm returns a pick for the next decision point that is game theoretically optimal. An example of the draft that results from such Artificial Intelligence reasoning is presented below.

Sample Draft using CBS League Style

In order to demonstrate the drafting mechanism, a CBS style, 12-owner league was simulated and all of the picks at each decision point were calculated. The three completed drafts were conducted with respect to all opponents following both strategies discussed above. The commentary on the drafts is presented below. The presentation of the rounds includes 12 players with their strategy indicated by Amateur, Expert, or <you>. The choices at each decision point follow from the strategy assigned to the opponents.

The Artificial Intelligence Algorithm

The completed draft based on CBS league rules:

Round	Choosing	Choice	Position	Points
1	#0-Amateur	LaDainian Tomlinson	rb	262.56
1	#1-Amateur	Adrian L. Peterson	rb	262.60
1	#2-Amateur	Brian Westbrook	rb	262.46
1	#3-Expert	Steven Jackson	rb	200.71
1	#4-Amateur	Joseph Addai	rb	230.41
1	#5-Amateur	Tom Brady	qb	374.30
1	<you>	Marshawn Lynch	rb	235.34
1	#7-Expert	Randy Moss	wr	225.70
1	#8-Amateur	Marion Barber III	rb	182.69
1	#9-Amateur	Frank Gore	rb	185.93
1	#10-Amateur	Larry Johnson	rb	220.79
1	#11-Expert	Clinton Portis	rb	191.54
2	#11-Expert	Reggie Wayne	wr	196.49
2	#10-Amateur	Ryan Grant	rb	156.43
2	#9-Amateur	Terrell Owens	wr	200.75
2	#8-Amateur	Willis McGahee	rb	169.69
2	#7-Expert	Peyton Manning	qb	349.06
2	<you>	Andre Johnson	wr	220.19
2	#5-Amateur	Braylon Edwards	wr	197.83
2	#4-Amateur	Laurence Maroney	rb	177.63
2	#3-Expert	Tony Romo	qb	316.95
2	#2-Amateur	T.J. Houshmandzadeh	wr	175.18
2	#1-Amateur	Brandon Jacobs	rb	207.75
2	#0-Amateur	Larry Fitzgerald	wr	212.91
3	#0-Amateur	Jamal Lewis	rb	204.02
3	#1-Amateur	Steve Smith	wr	155.59
3	#2-Amateur	Maurice Jones-Drew	rb	198.73
3	#3-Expert	Marques Colston	wr	203.92
3	#4-Amateur	Chad Johnson	wr	173.01
3	#5-Amateur	Reggie Bush	rb	194.99
3	<you>	Ben Roethlisberger	qb	315.29
3	#7-Expert	Torry Holt	wr	154.07

Round	Choosing	Choice	Position	Points
3	#8-Amateur	Drew Brees	qb	306.48
3	#9-Amateur	Ronnie Brown	rb	182.75
3	#10-Amateur	Earnest Graham	rb	160.58
3	#11-Expert	Plaxico Burress	wr	174.59
4	#11-Expert	Wes Welker	wr	154.29
4	#10-Amateur	Michael Turner	rb	36.32
4	#9-Amateur	Anquan Boldin	wr	166.50
4	#8-Amateur	Edgerrin James	rb	153.37
4	#7-Expert	Greg Jennings	wr	207.92
4	<you>	Dwayne Bowe	wr	183.84
4	#5-Amateur	Willie Parker	rb	143.45
4	#4-Amateur	Carson Palmer	qb	302.06
4	#3-Expert	Roy Williams	wr	181.24
4	#2-Amateur	Brandon Marshall	wr	149.50
4	#1-Amateur	Santonio Holmes	wr	173.10
4	#0-Amateur	Julius Jones	rb	76.28
5	#0-Amateur	Derek Anderson	qb	307.30
5	#1-Amateur	Darren McFadden	rb	141.00
5	#2-Amateur	Thomas Jones	rb	119.92
5	#3-Expert	Rudi Johnson	rb	137.49
5	#4-Amateur	LenDale White	rb	160.33
5	#5-Amateur	Calvin Johnson	wr	155.76
5	<you>	Marvin Harrison	wr	171.75
5	#7-Expert	Hines Ward	wr	125.43
5	#8-Amateur	Fred Taylor	rb	146.83
5	#9-Amateur	Matt Hasselbeck	qb	267.30
5	#10-Amateur	Jonathon Stewart	rb	138.00
5	#11-Expert	Matt Forte	rb	134.00
6	#11-Expert	Donovan McNabb	qb	315.05
6	#10-	Lee Evans	wr	141.30

Round	Choosing	Choice	Position	Points
	Amateur			
6	#9-Amateur	Selvin Young	rb	99.73
6	#8-Amateur	Donald Driver	wr	141.23
6	#7-Expert	Joey Galloway	wr	124.40
6	<you>	Jerricho Cotchery	wr	150.13
6	#5-Amateur	Roddy White	wr	131.86
6	#4-Amateur	Chris Chambers	wr	105.20
6	#3-Expert	Kevin Curtis	wr	119.95
6	#2-Amateur	Kevin Smith	rb	136.00
6	#1-Amateur	Rashard Mendenhall	rb	116.00
6	#0-Amateur	Laveranues Coles	wr	140.10
7	#0-Amateur	Antonio Gates	te	139.59
7	#1-Amateur	Jason Witten	te	127.78
7	#2-Amateur	Kellen Winslow Jr	te	125.06
7	#3-Expert	Tony Gonzalez	te	122.73
7	#4-Amateur	Dallas Clark	te	112.51
7	#5-Amateur	Chris Cooley	te	117.40
7	<you>	Chester Taylor	rb	153.18
7	#7-Expert	DeAngelo Williams	rb	110.21
7	#8-Amateur	Eli Manning	qb	253.27
7	#9-Amateur	Jay Cutler	qb	271.30
7	#10-Amateur	Marc Bulger	qb	204.18
7	#11-Expert	Reggie Brown	wr	122.60
8	#11-Expert	Deuce McAllister	rb	125.84
8	#10-Amateur	Anthony Gonzalez	wr	139.99
8	#9-Amateur	Justin Fargas	rb	137.19
8	#8-Amateur	Javon Walker	wr	126.61
8	#7-Expert	Ahman Green	rb	141.99
8	<you>	San Diego Defense	dst	203.47
8	#5-Amateur	Felix Jones	rb	105.00
8	#4-	Derrick Mason	wr	104.22

Round	Choosing	Choice	Position	Points
	Amateur			
8	#3-Expert	Bernard Berrian	wr	144.29
8	#2-Amateur	David Garrard	qb	301.29
8	#1-Amateur	Philip Rivers	qb	235.35
8	#0-Amateur	Bobby Engram	wr	107.28
9	#0-Amateur	Minnesota Defense	dst	188.28
9	#1-Amateur	Jerious Norwood	rb	92.78
9	#2-Amateur	Chicago Defense	dst	162.07
9	#3-Expert	Donte Stallworth	wr	93.25
9	#4-Amateur	New England Defense	dst	180.90
9	#5-Amateur	Aaron Rogers	qb	237.00
9	<you>	Reggie Williams	wr	136.80
9	#7-Expert	Green Bay Defense	dst	169.79
9	#8-Amateur	Darrell Jackson	wr	107.99
9	#9-Amateur	Sidney Rice	wr	114.34
9	#10-Amateur	Matt Schaub	qb	212.69
9	#11-Expert	Jake Delhomme	qb	289.00
10	#11-Expert	Jerry Porter	wr	105.68
10	#10-Amateur	Nate Burleson	wr	86.36
10	#9-Amateur	Nate Kaeding	k	135.47
10	#8-Amateur	Mark Clayton	wr	85.58
10	#7-Expert	Tatum Bell	rb	110.70
10	<you>	Kevin Jones	rb	143.16
10	#5-Amateur	Mason Crosby	k	123.02
10	#4-Amateur	Jon Kitna	qb	258.01
10	#3-Expert	Ladell Betts	rb	54.80
10	#2-Amateur	Vince Young	qb	224.72
10	#1-Amateur	Matt Leinart	qb	174.14
10	#0-Amateur	Stephen Gostkowski	k	131.41
11	#0-Amateur	Nick Folk	k	135.75
11	#1-Amateur	Shaun McDonald	wr	96.97
11	#2-Amateur	Adam Vinatieri	k	134.58
11	#3-Expert	Brett Favre	qb	276.98

Round	Choosing	Choice	Position	Points
11	#4-Amateur	Josh Brown	k	114.88
11	#5-Amateur	Isaac Bruce	wr	89.18
11	<you>	Kurt Warner	qb	250.77
11	#7-Expert	Jason Campbell	qb	240.21
11	#8-Amateur	Shayne Graham	k	131.89
11	#9-Amateur	Robbie Gould	k	114.94
11	#10-Amateur	Jason Elam	k	95.71
11	#11-Expert	Ron Dayne	rb	127.63
12	#11-Expert	Jeremy Shockey	te	90.35
12	#10-Amateur	Vernon Davis	te	87.89
12	#9-Amateur	Todd Heap	te	93.70
12	#8-Amateur	Heath Miller	te	97.50
12	#7-Expert	Benjamin Watson	te	96.84
12	<you>	Owen Daniels	te	96.89
12	#5-Amateur	Alge Crumpler	te	88.21
12	#4-Amateur	Tony Scheffler	te	89.54
12	#3-Expert	Baltimore Defense	dst	147.32
12	#2-Amateur	L.J. Smith	te	53.95
12	#1-Amateur	NY Giants Defense	dst	159.39
12	#0-Amateur	Greg Olsen	te	68.15
13	#0-Amateur	Jeff Garcia	qb	229.23
13	#1-Amateur	Pittsburgh Defense	dst	176.33
13	#2-Amateur	Mike Furrey	wr	101.92
13	#3-Expert	Dallas Defense	dst	174.23
13	#4-Amateur	Phil Dawson	k	124.55
13	#5-Amateur	Jacksonville Defense	dst	174.08
13	<you>	Kenny Watson	rb	132.84
13	#7-Expert	Shaun Alexander	rb	120.64
13	#8-Amateur	Indianapolis Defense	dst	185.84
13	#9-Amateur	Tennessee Defense	dst	170.49

Round	Choosing	Choice	Position	Points
13	#10-Amateur	Tampa Bay Defense	dst	163.38
13	#11-Expert	Seattle Defense	dst	175.43
14	#11-Expert	Rob Bironas	k	117.22
14	#10-Amateur	Neil Rackers	k	123.25
14	#9-Amateur	Muhsin Muhammad	wr	81.22
14	#8-Amateur	Carolina Defense	dst	138.03
14	#7-Expert	Matt Stover	k	100.49
14	<you>	Josh Scobee	k	124.16
14	#5-Amateur	Philadelphia Defense	dst	156.17
14	#4-Amateur	Patrick Crayton	wr	126.57
14	#3-Expert	David Akers	k	110.41
14	#2-Amateur	Jason Hanson	k	102.52
14	#1-Amateur	Kris Brown	k	124.18
14	#0-Amateur	Cadillac Williams	rb	112.96

In this model, the AI actively predicts what the other players (both the experts and the amateurs) will do with their picks, and then modifies its picks accordingly. Thus, some picks may seem strange at first but will often 'work out' in the end.

1. One such example is our decision to take Marshawn Lynch with the first pick. Marshawn Lynch was an extremely consistent rookie, who we project to score 235 points this year.
2. However, it may have seemed odd to take Lynch at #1 instead of Randy Moss (248), who we project as the top overall receiver, and who we also project to score 13 points higher than Lynch. Interestingly enough, however, the AI actually projected ahead and predicted that we would be able to get the #2 WR, Andre Johnson (220) with our second pick (which we do) and otherwise would have ended up with Willis McGahee (170 points) or Laurence Maroney (178 points) as our starting running back had we taken Moss first. Thus, by taking Lynch first, we gain about 60 points over the RB we would have gotten in the second round (McGahee or Maroney), while

taking Moss would have only netted us a 28 point gain over the WR that we ended up acquiring in the second round. So, by projecting ahead, the AI actually makes a brilliant move by avoiding the best player on the board in order to maximize long term point value.

3. Ben Roethlisberger is a great pick here. We held off picking a QB because we knew that there was a good chance that we could pick up a QB as good as Ben (the fourth rated QB overall) in the third round.
4. Dwayne Bowe adds a potentially great, fantastic young talent to the team. He should blossom in his second season.
5. If Bowe does not produce like we project him to, or if Bowe or Johnson get hurt, taking Marvin Harrison here is going to be a brilliant play. He will play at least two games (the bye) if not several more, and is very likely to score many fantasy points for our team. He is a consistent, proven veteran.
6. Jerricho Cotchery adds more depth at WR.
7. Chester Taylor adds depth at RB. Notice that since we got such a good RB and QB early, we have held off on drafting the second players at those positions for as long as possible to maximize our WR corps.
8. This is early to take a defense, but given that we have San Diego projected at #1 (although like we mentioned earlier, those rankings are highly volatile) this is a pretty solid pick that could pay off in the long run. Since this is a 'winner take all' draft, we are more willing to take risky picks such as this one.
9. Reggie Williams rounds out or WR corps.
10. Kevin Jones adds a third RB, giving us more options.
11. Kurt Warner will be a suitable second string QB – if he plays, the one-time NFL MVP could potentially put up great numbers.
12. Owen Daniels will provide an adequate solution at tight end. If he gets hurt or we need someone during his bye week, we will just release our worst (or non) performing backup and

pick up a second TE midway through the season. Given that we have such a great QB and a solid RB/WR corps and the top defense, we are just trying to lose as few points as possible as TE. It is a small sacrifice to maximize value in more important areas.

13. Kenny Watson finalizes our RB corps.
14. Josh Scobee is the best kicker available in the last round.

Once again, the amateurs make many mistakes. Player #0 takes a kicker too early, drafts three RB's in his first four picks, and takes a backup TE even though he already has Antonio Gates. Player #10 has a particularly horrific draft, blinded by the high point values of RBs; he takes RBs with his first five picks, finishes the draft with only three WRs, and also inexplicably chooses two kickers. Players #2, #4 and #9 also take two kickers, which while seemingly reasonable given their point production, does not make much sense from a game theory perspective, as one has no reliable way to project how many points a kicker will score in a given game and could likely have put that pick to much better use. Player #8 also suffers from the 'too many RBs' syndrome, taking four within his first five picks.

Of course, the experts do much better. Expert #7 goes for broke by getting Peyton Manning and stacking his WRs with the other five of his first six picks. He then makes a run on three RBs late, hoping that one of them will do decently, and then a top defense. He uses his last four picks on backups as well as a TE and a K. Experts #3 and #11 pursue a more balanced approach, getting a QB, two RBs and two WRs with their first five picks, building a traditional and very solid team across the board. After that, they grab the best backups available, and hold off on the less important players until late, although Expert #3 does realize the value in grabbing Tony Gonzalez in his particular draft order situation.

Overall, we have used our points projections and game theory rules and adjustments and also our ability to anticipate what our opponents will do and to act accordingly to draft the best team possible. Such a tool will be a fearsome weapon and a needed ally in any fantasy draft, and can really help a player look beyond the most obvious solution to a problem, finding more creative solutions with greater long-term benefits.

SECTION VI: Conclusion

This report provides the reader with an introduction to and explanation of a variety of concepts regarding novel and creative applications of various mathematical and statistical fields to the drafting of a Fantasy Football team. None of these tools is meant for independent use and each one has its strengths and weaknesses. With diligent study and an attentive mind, the reader and user of these tools should be able to create a winning fantasy team of NFL players based upon considerations ranging from portfolio theory and risk assessment to game theory and strategy analysis.