Salary Cap Drafts in Fantasy Sports:

An Analysis of Budget-Constrained Sequential Auctions

Connor Larson

ECO 325: Organization and Design of Markets

Professor Jakub Kastl

December 8, 2020

This paper represents my own work in accordance with University regulations.

/s/ Connor Larson

Table of Contents

1. **Introduction**

In this paper, I set out to explore the mechanism, strategy, benefits, and shortcomings of budget-constrained sequential auctions as implemented in fantasy sports drafts. I begin with an overview of the market and an explanation of the auction procedure. Next, I explore the connections between these auctions and existing work in the field, before providing brief commentary on factors affecting individual preferences and valuations. I then propose a bidding strategy using data from a scaled-down example, which I relate to some more foundational themes in auction theory as well as larger-scale trends from FantasyPros. Finally, I illustrate some adverse behavior through a few specific examples, and end with some concluding remarks on limitations and opportunities for future exploration.

2. **Background**

2.1. Overview

Fantasy sports are a popular type of game in which several individuals, known as "managers," compete against each other by leading a virtual team of athletes over the course of a season. Though born as a pen-and-paper affair tabulated from daily newspapers, modern fantasy sports are almost entirely organized online, where managers have access to sophisticated tools to research players and manage their rosters.

Based on their statistics and performance in real life, athletes earn "points" for their teams, which are totaled over the course of each week or the whole season to determine winning and losing teams. Over the course of many months, managers have the opportunity to sign new players, bench or cut rostered players, and trade with each other, just as a manager would do in real life. At the end of the season, winners and losers are declared, frequently earning a betting pot from their league and bragging rights, or conversely having to feel shame until the fun starts again the following year.

2.2. Market Sizing

According to the Fantasy Sports and Gaming Association, almost 60 million Americans – over 1 in 6 – played a fantasy sport in 2017, up from 54 million three years prior.[1] Of those, a large majority spend money on their leagues. In addition to intra-league bets and cash prizes, over 73% spend money on an additional non-cash prize like a trophy, 84% spend money on a draft party, and 68% spend money on a losing punishment.[2] Per a 2013 study referenced in *Forbes*, across 32 million fantasy football players, estimated average annual spending was $467 per person, for a whopping $15 billion total. All in, the author estimated the market to be worth over $70 billion.[3]

Figure 1: "The League," a network TV show focused on fantasy football which has run for 7 seasons on FX[4]

In addition to the traditional season-long format, a new type of fantasy sports game has recently exploded: "daily" and "weekly" competitions. Taking advantage of a legal loophole defining fantasy sports as a game of skill rather than a form of sports betting, sites like FanDuel and DraftKings allow users to create and place bets on new teams every day – even in states that prohibit online gambling. In these competitions, even more-so than in a traditional league, one event plays an outsize role in determining outcomes: the initial draft.

[1] The Fantasy Sports and Gaming Association, "Industry Demographics."
[2] The Fantasy Sports and Gaming Association.
[3] Brian Goff, "The $70 Billion Fantasy Football Market."
[4] "The League."

2.3. The Draft

At the beginning of a fantasy sports season, managers must use a draft system to select an initial allocation of players for their team. This is one of the most important events of a season – though managers can and do edit their rosters on an ongoing basis, a bad initial allocation can sink a team from the start.

There are two main formats for drafts. The most common is the "snake" draft, where managers take turns filling positions from the overall pool of players. This proceeds similarly to a rookie or expansion draft in many pro leagues – a random order is selected, and teams select in alternating descending order $(A, B, C, D; D, C, B, A \dots)$ until all rosters are filled. While this is an interesting allocation mechanism in its own right – and shares many similarities with the serial dictatorship mechanism we discussed in class – it is not the main focus of this paper.

The other common format to allocate players is the "salary cap" draft system, otherwise known as an "auction" draft. In this format, every manager is given a fixed budget to spend on their players. The mechanism proceeds via a series of sequential English auctions, wherein managers take turns nominating players by placing $1 initial bids. Then, any manager may raise the bid by any increment of $1 provided they have ample budget and space for the player on their roster.

After each successive bid, a timer descends (usually between 10 and 30 seconds), and at its expiration the highest bidder will win the player for their team. Finally, once every manager has filled their teams, any leftover money expires worthless. Thus, managers are incentivized to spend their entire budget. There is one additional constraint: managers must reserve enough budget to fill their entire roster, so in a league with 22-man rosters and $200 budgets, for instance, the highest initial bid would be $\$200 - 21 * \$1 = \$279$.

3. Auction Theory

3.1. Comparison to Studied Formats

Each English-style auction, when considered in isolation, bears many similarities to the second-price mechanism studied in class. If the minimum bid ratio $\frac{\$1}{B}$ is small, the winner pays only slightly above the final bid of the last competitor to drop out; thus, in the absence of externalities, there is no incentive to bid non-truthfully. Indeed, by overbidding, a buyer needlessly exposes themselves to negative utility if the second-highest bidder is above their true value, while underbidding decreases the odds of winning without affecting the final price.

However, the broader context of the system adds several nuances to the auction mechanism. Each auction does not actually occur in isolation, but rather in a dynamic sequence. This also sets sports auctions apart from the simultaneous multiunit and combinatorial mechanisms – such as for treasury bills and lots of goods – as only one item is up for purchase at a time. Furthermore, the budget constraints and direct competition add externalities to each auction, broadening the scope of each bidder's payoff function.

3.2. Relevant Academic Research

While work regarding auction theory is quite extensive, few papers have been published specifically on fantasy sports. Those that exist largely focus on algorithmic analysis in a computer science context and on regression analysis using a behavioral economics lens. However, these insights are still useful to build intuition for bidding strategy and to predict shortcomings over larger samples.

Motivated by the goal of building a better auto-draft algorithm,[5] Anagnostopoulos et al.'s 2016 paper contains extensive analysis on the presence of equilibria and determines – through a heavily constrained example – that they do not always exist.[6] Boudreau and Shunda take a

[5] Most fantasy sports draft sites include an algorithm to "auto-pick" in case a player disconnects from the internet. These are also often used by beginners to gain a feel for the system.

[6] Anagnostopoulos et al., "Bidding Strategies for Fantasy-Sports Auctions."

different approach; by analyzing the relationships between winning bids and position rank, the authors find that broadly, participants overbid for players at the start and end of the draft, while underbidding for players in the middle.[7]

3.3. Bidding, Projections, and Signals

As there is no real money exchanged and no "seller" seeking revenue, maximizing team values in a competitive context is the primary goal of the draft. Accordingly, each manager has their own set of beliefs about the players at auction. To gain insights, it is helpful to break these beliefs down into three components: public signals, private signals, and dynamic elements that change as the auction progresses.

The first of these three is typically the largest in magnitude and most consistent; vast amounts of public data exist on player performance, so projected season totals are a natural choice of signal. Several services provide these projections based on players' previous year statistics, team dynamics, schedules, injury likelihood, and countless other factors. Yahoo and ESPN – two of the most popular fantasy platforms – have these projections directly integrated into the draft display. Another popular choice is FantasyPros, a third-party service which excels in providing large, downloadable tables of data – these will form the basis of my analysis.

The other facets of managers' beliefs are complex yet necessary because they add a stochastic element to the auction – without them, every bid would be symmetric. Managers may have initial private signals which influence their beliefs for any number of reasons. For instance, most managers have favorite real-life teams, and likely value their hometown players over those of a rival. Alternatively, managers often have personal beliefs that rookies are either significantly over- or under-projected, as less data is available to predict their performance.

---

[7] Boudreau and Shunda, "Sequential Auctions with Budget Constraints."

Though highly complex in real life, one approach might be to model these signals as uniform random variables scaled to some fraction of a standard deviation; by comparing last season's draft week projections to final results, this is easy to calculate. As a frame of reference, this calculation for quarterbacks is displayed in *Table 1* below.

# of QBs	Average Difference	Standard Deviation
40	-0.0031	0.2214

Table 1: Difference Between Draft Week Projections and 2019 Results

The final components which affect player valuations are dynamic and arise over the course of the sequential auction. While harder to pin down, one example of these would be the common phenomenon of "handcuff" running backs, the backup RBs on NFL teams. Usually, these players will not start and therefore have no value. However, if the starting RB gets hurt, the backup will take their place and accrue many of their points. Thus, a synergy exists for the manager who owns the starting RB on a given NFL team; their specific handcuff is more valuable to them than their league-mates because it ensures continuity in case of an injury.

4. <u>Data and Analysis</u>

4.1. The VORP Bidding System

Though projected points are the most logical scale for player value, a challenge arises when modeling the auction: how do points translate to dollar-value bids consistent with each player's starting budget?

To solve this problem, we seek some function which maps each manager's set of signals S (which are the composite of the statistical projections and private beliefs for every player) onto a set of dollar-values V. In other words, we seek some function $f: S \longmapsto V$.

The model I propose is based on a metric called $VORP$, which stands for Value Over Replacement Player. This is a commonly-used statistic to evaluate trades in fantasy sports and is based on the idea that players generally have decreasing marginal returns – as exhibited in

Figure 2 below, the difference between, for instance, the 1st and 2nd players is significantly greater than between the 11th and 12th. To find a rostered player's VORP in-season, one typically measures the difference between their projections and the best available "replacement" – a free agent at their position, who is not claimed or rostered by any other team. In order to adjust this metric in a draft context, I make some adjustments to produce an "A-VORP" as detailed below.

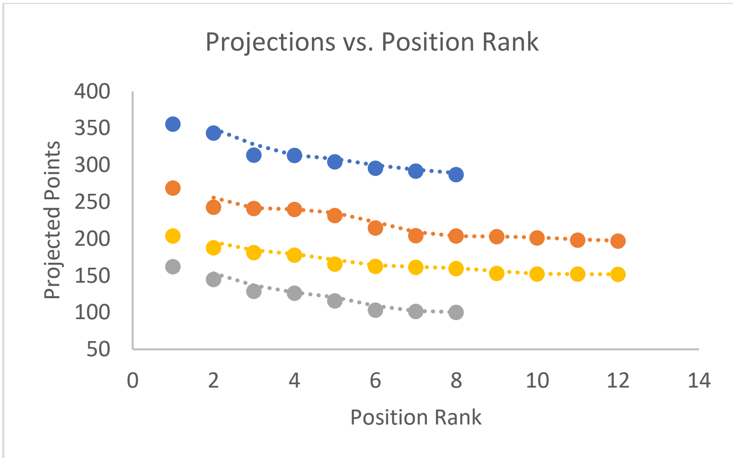

Figure 2: Declining Marginal Differences - Note the Decaying Trend[8]

4.2. Simplifications and Adjusted VORP

In order to model efficiently, it is important to make some simplifying assumptions. First, it is helpful to assume that the set of all drafted players is the same for every manager; in other words, for a roster position with n slots, the top n^{th}-preferred players are the same for every manager (note that this does not constrain preferences within the top n). This also implies that, especially at the lower end of player rankings, the differences in personal beliefs are small.[9]

Next, to reduce the number of variables, it is helpful to use a smaller draft setting; this also exacerbates the differences in player values by the theory of diminishing marginal difference above, as considering few players implies larger differences. Thus, while a standard draft might consist of 14 teams filling 18 roster slots, I chose to focus on 4 teams filling 12 slots as follows:

[8] "Fantasy Football Projections - Draft Week."

[9] This is generally a safe assumption; when managers have strong personal beliefs about players, it is usually the top few on their team or at a given position. See section *Limitations* for further commentary.

1 (+1) QB, 2 (+1) RB, 2 (+1) WR, 1 (+1) TE, 1 DST, and 1 K, where the parentheses indicate bench slots. This hints at the final assumption – generally, bench slots are unrestricted. However, I allocated specific numbers to positions in order to guarantee that the set of drafted players would remain consistent.

Lastly, I used the constraints of the sequential auction mechanism to inform my adjustments to VORP. In particular, because roster allocations are fixed, the last player at every position will always sell for $1 if the managers use an optimal strategy, as there is no possible competition. Thus, with n teams for a position with k slots, the value $v_{pos}^{(nk)}$ must be the "replacement" value and be associated with a bid $1.

4.3. Accounting for Injuries

To adjust for the value of bench players, I chose to scale their projections down by the expected portion of games they would start in a fantasy league with k starters at their position. To do so, I utilized data on the average number of game starts for starters at each position, sourced from Pro Football Logic.[10] I also added in a factor of $\frac{k}{18}$ to account for the starters' bye weeks:[11]

$$AVORP = VORP_o \cdot \mathbb{P}(Starter\ Injured) = VORP_o \cdot \left(1 - \left(\frac{Avg.Starts}{16}\right)^k + \frac{k}{18}\right)$$

Finally, to convert each $AVORP$ back to a dollar value, I summed the total budgets across all teams (excluding the minimum reserve of $1 per position) and divided this total by the sum of players' $AVORP$s across all positions to find a dollars-per-$AVORP$ factor. Finally, I multiplied

---

[10] Michael Gertz, "NFL Injury Rate Analysis."

[11] In the NFL, each team plays 17 out of 18 weeks; on a starter's bye week, fantasy football managers start the backup. Also note that the denominator of 16 in the left hand term was a result of the source data; Pro Football Logic quoted average starts per 16 games.

this by each player's $AVORP$ and added back the minimum bid of $1 to find each player's final baseline value v. An example of this spreadsheet for running backs is shown in *Figure 3* below.

$$v = AVORP \cdot \frac{n \cdot (B - \#_{roster})}{\sum AVORP} + 1$$

#	RB	Points	VORP		Avg. Games		AVORP	Value
1	Christian McCaffrey	268.8	72.1		13.3		72.1	$ 65.17
2	Ezekiel Elliott	242.8	46.1				46.1	$ 42.03
3	Saquon Barkley	240.7	44		# Starters		44	$ 40.16
4	Derrick Henry	239.4	42.7		2		42.7	$ 39.01
5	Dalvin Cook	231.4	34.7				34.7	$ 31.89
6	Alvin Kamara	214.2	17.5		BE Factor		17.5	$ 16.58
7	Clyde Edwards-Helaire	204	7.3		0.42		7.3	$ 7.50
8	Nick Chubb	203.4	6.7				6.7	$ 6.96
9	Josh Jacobs	202.6	5.9				2.5	$ 3.21
10	Joe Mixon	200.8	4.1				1.7	$ 2.53
11	Miles Sanders	197.9	1.2				0.5	$ 1.45
12	Aaron Jones	196.7	0				0.0	$ 1.00

Figure 3: AVORP Calculations for 2 (+1) RB in a 4-Team League

4.4. Strategic Analysis

Given this strategy, two questions naturally arise: can one do better, and does this guarantee at least a somewhat competitive team? The former answer is almost certainly "yes," as examples in the next section show. However, I also believe that the second answer is a "yes," at least if manager's set of values V are somewhat accurate.

The intuition for this result closely tracks with that of the second-price auction: if other teams underbid for a player under my calculated set of valued v, I have the opportunity to buy that player at a profit $v - b_{others}$ while if other players overbid, they will be left with less money in the system to spend. Because the total amount of money is zero-sum, and my values v were calculated based on the total amount in the system, overspending by anyone else automatically guarantees a net surplus for me under my set of values V. So, the worst I can do by bidding truthfully under my own system is with my fair share of total value $\frac{\sum V}{k}$, which seems like a good result – players who draft conservatively and who's initial value set V is not off the mark should hit the ground to a level playing field.[12]

4.5. Comparison to FantasyPros Suggested Salaries

To get a sense of the form of this solution, it is also helpful to check against the salaries suggested by FantasyPros (based upon a proprietary algorithm) for their standard auction drafts. These are plotted below for four positions:

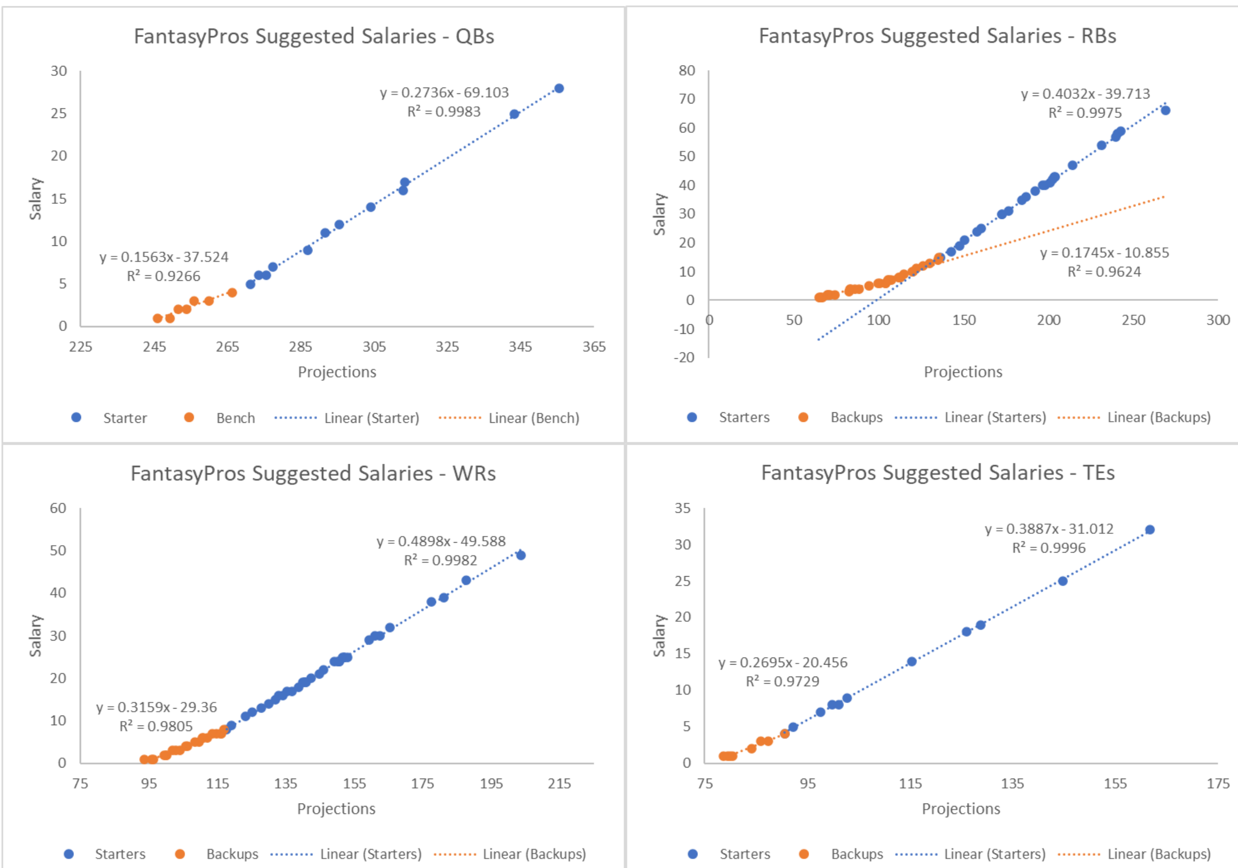

Figure 4: FantasyPros Suggested Salaries for a Standard 12-Team League

[12] While this result holds in the generalization when n is relatively large and $\frac{1}{B}$ relatively small, there might be some examples where it does not at small n. For instance, consider filling a position with a player undervalued by3, only for a player at the same position to be undervalued by 10 in the next auction. However, because you have filled the position, you cannot bid. In this case, the missed surplus instead goes to whoever won it at 10 – in a 4-team league, as things are zero-sum, this is shared as a loss of 3.33 by the other players, resulting in a net loss of 0.33.

Though FantasyPros' suggestions are based on a much larger draft and dataset, the apparent linear form of their solution is an encouraging litmus test for the AVORP method detailed above. Indeed, the core trend – a linear increase from a non-zero intercept, with a shallower slope (lower multiplied dollars-per-point) for bench players – concurs exactly with the results of the AVORP model.

4.6. Externalities and Competitive Behavior

Though the AVORP is elegant, the simplifying constraints on the problem remove the opportunity to gain competitive advantage by utilizing externalities. Indeed, a key point that distinguishes every auction in the sequence from standalone auctions is that, even when a bidder loses for an item, they can still have a positive profit if their opponent overpays. More generally, since the budget is zero-sum, for every dollar a competitor overspends, everyone else shares the same amount in profit equally split. This allows us to model certain scenarios by using calculus to find the bidding function which maximizes an individual's expected profit; below are some examples.

5. An Example: Robbie the Ravens Fan

5.1. Confident Robbie

Robbie is a Ravens fan. Though he uses the VORP system above to convert his beliefs to prices like the three other managers in his league, he has biased beliefs about the players on his own team. On draft day, Robbie exclaims, "Lamar is going to put up 400 points this season – I *know* it!" Here, Robbie is confident; let us assume he has a signal $s = 400$ and uses the VORP system to obtain a value $v_R = \$96$.

On the other hand, Adam, Bob, and Charlie are all unbiased and believe the FantasyPros projections are spot on. Adam and Bob are simple; they will purely bid up to their VORP value of $v_o = \$62$, no less, no more. However, Charlie is clever – he realizes that if he bids up to $v_R - 1 = \$95$, Robbie will spend \$96 instead of \$63, an overspend of \$34 in the eyes of everyone else! And because the others are in direct competition with Robbie and face his team an equal number of times over the season, they all share the profit $\pi = \$11.33$ evenly.

5.2. Uncertain Robbie

Robbie is still a Ravens fan, but this time, he has less conviction in his beliefs. "I think Lamar's projections are too low, but I'm not sure he'll put up 400 points," he ponders. Here, let's assume that he has some signal s such that his value $v_R \sim U[62,96]$ follows a uniform random variable. Once again, the others believe that the FantasyPros projections are accurate.

Adam and Bob are still simple, and Charlie is still clever – however, this time he must balance his desire to drive up Robbie's spending with his own risk of overspending. If Charlie decides to bid up to a boundary $b > \$62$ and ends up winning, he will bear the loss up to whenever Robbie drops out at v, netting a profit $\pi = -(v - 62)$. However, if he drops out, Robbie will instead pay b and Charlie (and each of the others) will profit $\pi = \frac{1}{3}(b - 62)$. Thus, Charlie's expected profit (assuming continuous bidding for simplicity) is:

$$\mathbb{E}[\pi|b] = \int_{62}^{96} \pi(b, v) \cdot \mathbb{P}(v \in dv)\, dv$$

Here, we can split the integral based on the mutually exclusive winning and losing profit conditions, and substitute in the uniform p.d.f. $\mathbb{P}(v \in dv) = \frac{1}{96-62} = \frac{1}{34}$:

$$\mathbb{E}[\pi|b] = \int_{62}^{b} (62 - v) \cdot \left(\frac{1}{34}\right)\, dv + \int_{b}^{96} \frac{1}{3} \cdot (b - 62) \cdot \left(\frac{1}{34}\right)\, dv$$

Finally, Charlie will choose his bid to maximize his expected profit, which can be found by finding the zero(es) of the derivative w.r.t. b:

$$\hat{b} = \arg \max_{b \in [62,96]} -\frac{(b - 62)^2}{2 \cdot 34} + \frac{(96 - b) \cdot (b - 62)}{3 \cdot 34}$$

$$\frac{d}{db}(\dots) = \frac{344 - 5b}{102} = 0$$

Thus, we see that Charlie should bid up to a bound of $b = \frac{344}{5} \approx \$69$, and will expect to gain a profit of $\mathbb{E}[\pi] = -\frac{7^2}{68} + \frac{27 \cdot 7}{102} \approx 1.13$ from his strategy (as will both Adam and Bob).[13]

5.3. Charlie's Coalition

Though Charlie is happy with his increased profits, he does not like that he currently bears all the risk in his strategy while sharing the reward equally with Adam and Bob. "If we all bid Robbie up together," he convinces them, "our expected profits will be higher." Charlie proposes a mixed strategy: with probability $p = \frac{1}{3}$, the three will randomly alternate raising their bids up to a boundary b. The math proceeds as above, with slight alteration to the leftmost term:

$$\mathbb{E}[\pi|b] = \int_{62}^{b} \frac{1}{3} \cdot (62 - v) \cdot \left(\frac{1}{34}\right) dv + \int_{b}^{96} \frac{1}{3} \cdot (b - 62) \cdot \left(\frac{1}{34}\right) dv$$

$$\hat{b} = \arg \max_{b \in [62,96]} \frac{(62 - b) \cdot (3b - 254)}{204}$$

$$\frac{d}{db}(\dots) = \frac{110}{51} - \frac{b}{34} = 0$$

Thus, the three will randomly alternate bidding up to a bound of $b = \frac{220}{3} \approx 73$, and will expect to raise their profits to $\mathbb{E}[\pi] = \frac{(62-73)\cdot(3\cdot 73-254)}{204} \approx 1.89$ each.

5.4. Robbie Responds

All three above examples assume that Robbie still follows the truthful mechanism – accurately reporting and bidding according to his preferences. However, they also open the door to another possibility – what if Robbie chooses to be deceitful himself? By outwardly feigning

---

[13] Note that the coefficient on the order-one term in the first derivative is negative; therefore, the second derivative is negative for all b and thus this is indeed a maximum.

confidence that Lamar's value should be higher than he truly believes, Robbie may be able to fool his competitors into bidding him up in the same way they did to him – only to drop out and leave them with the burden of overpaying.

6. <u>Discussion</u>

6.1. Results

In reality, all four above scenarios are heavily simplified. A more nuanced model, for instance, might expand the stochastic element to every auction participant; an attempt at modelling even a simple case of this with two symmetrical players, included in *Appendix: A General Case* below, illustrates how quickly this becomes mathematically complex.

However, even these highly constrained examples shed light on more broadly applicable results of budget constraints and competition. The existence of a profit incentive to bid above one's value brings the theoretical model of bids away from the truthful second-price framework, and closer towards (although not the same as) the first-price auction model. Furthermore, the ability of players to increase their profits by forming coalitions provides valuable insight into budget-constrained competitions even beyond the scope of fantasy sports; while teaming up against a friend in the league certainly does not seem like a good outcome, cartels in the business world is undoubtedly a more concerning possibility.

6.2. Limitations

Aside from the myriad simplifications and constraints discussed throughout this paper, I believe the largest limitation to this analysis was a lack of large-scale human data. Unfortunately, as the football season is nearing the end, the online draft lobbies which I had aimed to use to collect additional data and test the AVORP framework were largely closed and empty. For instance, both Yahoo and ESPN, two leaders in the field, had taken their systems offline by the time I was ready to collect data, and will not reopen them until next year.[14][15]

While this development was disappointing, the Boudreau and Shunda study referenced earlier does shed some light on the behavioral side of the analysis; though there dataset centered on fantasy basketball instead of football, their analysis included two strong trends: the top few players, with star-player status on their teams, generally were the most overbid on (presumably by individuals like Robbie the Ravens Fan); and, the overall trend on over/underbidding was downward start-to-finish. This latter result also seems to echo a result discussed in *Milgrom*, where-in actual prices in art and wine auctions followed a similarly decreasing pattern.[16] This certainly merits further exploration.

6.3. Stability and Other Concerns

One final area of exploration which is compelling, but outside the scope of my analysis, is the stability and efficiency of budget-constrained sequential auctions as an allocation mechanism relative to the "snake" draft also common in fantasy sports. Because fantasy sports offer a secondary market – the ability to trade players as soon as seconds after the draft – insight on different quantities of trading could be a fascinating indication of stability difference. Likewise, with a larger dataset, comparisons of the averages and standard deviations of net team projections might lend valuable insights into the efficiency of both algorithms to maximize the number of good players in starting positions.

7. Conclusions and Further Applications

Though fantasy football is "just a game" with a heavily constrained auction method using fake money, the observed phenomena of idealized second-price behavior marred by competitive bidding, dis-truthful action, and coalition-forming in practice provides valuable insight into both

[14] "Mock Draft Lobbies Are Closed."
[15] "Mock Draft Auction | Fantasy Football | Yahoo! Sports."
[16] Milgrom, *Putting Auction Theory to Work*.

the advantages and perils of sequential budget-constrained auctions in the broader context. For better or for worse, the framework – though with real money and likely on the scale of months or years instead of seconds and minutes – of competitive, sequential auctions is certainly present in the broader business, government, and sports world. Whether the focus is airline slots, telecom licenses, or real-life free-agent negotiations, there certainly remains a vast amount of complexity and nuance to discover and explore.

8. Bibliography

Anagnostopoulos, Aris, Ruggiero Cavallo, Stefano Leonardi, and Maxim Sviridenko. "Bidding Strategies for Fantasy-Sports Auctions." In *Web and Internet Economics*, edited by Yang Cai and Adrian Vetta, 10123:102–15. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016. https://doi.org/10.1007/978-3-662-54110-4_8.

Boudreau, James W., and Nicholas Shunda. "Sequential Auctions with Budget Constraints: Evidence from Fantasy Basketball Auction Drafts." *Journal of Behavioral and Experimental Economics* 62 (June 2016): 8–22. https://doi.org/10.1016/j.socec.2016.03.002.

Brian Goff. "The $70 Billion Fantasy Football Market." Forbes, August 20, 2013. https://www.forbes.com/sites/briangoff/2013/08/20/the-70-billion-fantasy-football-market/.

Fantasy Pros. "Fantasy Football Projections - Draft Week." Database, September 10, 2020. https://www.fantasypros.com/nfl/projections/.

Michael Gertz. "NFL Injury Rate Analysis." Sports Analytics. Pro Football Logic, February 22, 2016. https://www.profootballlogic.com/articles/nfl-injury-rate-analysis/.

Milgrom, Paul R. *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge, UK ; New York: Cambridge University Press, 2004.

Yahoo Sports. "Mock Draft Auction | Fantasy Football | Yahoo! Sports," December 8, 2020. https://football.fantasysports.yahoo.com/f1/mock_closed.

ESPN.com. "Mock Draft Lobbies Are Closed.," December 8, 2020. https://fantasy.espn.com/football/mockdraftlobby.

The Fantasy Sports and Gaming Association. "Industry Demographics." FSGA. Accessed December 8, 2020. https://thefsga.org/industry-demographics/.

"The League." FX Network, October 29, 2009.

9. Appendix: A General Case

Consider two bidders, A and B, both with values $v_A, v_B \sim U[0,1]$. As in section 5, bidders gain profit/loss $v_{me} - b_{them}$ from winning and profit/loss $b_{me} - v_{me}$ from losing. Thus, for player A, expected profit is:

$$\pi_A = \mathbb{P}(b_A > b_B) \cdot (v_A - b_B) + \mathbb{P}(b_A < b_B) \cdot (b_A - v_A)$$

Applying inverse functions and the uniform c.d.f. find:

$$\pi_A = b_B^{-1}(b_A) \cdot (v_A - b_B) + \left(1 - b_B^{-1}(b_A)\right) \cdot (b_A - v_A)$$

Using the first order condition:

$$\frac{2v_A}{b_B'\left(b_B^{-1}(b_A)\right)} - \frac{b_B}{b_B'\left(b_B^{-1}(b_A)\right)} - \frac{b_A}{b_B'\left(b_B^{-1}(b_A)\right)} - b_B^{-1}(b_A) + 1 = 0$$

Now assume that $b_{A,B}$ are symmetric so that $b_B'\left(b_B^{-1}(b_A)\right) = v_A$, and remove subscripts for simplicity:

$$2v - 2b + b' + v \cdot b' = 0$$

This is a linear ODE with solution:

$$b(v) = c \cdot (v+1)^2 + 2v + 1$$

The only c that both makes this monotonically increasing and non-negative is $c = -\frac{1}{2}$.

And, under this function, expected long-run profits for each individual are $\frac{1}{6}$, which exactly matches the result from the second-price auction. This result certainly deserves more exploration.