

Bidding Strategies for Fantasy-Sports Auctions

Aris Anagnostopoulos^{1*}, Ruggiero Cavallo²,
Stefano Leonardi^{1*}, and Maxim Sviridenko²

¹ Sapienza University of Rome {aris,leonardi}@dis.uniroma1.it

² Yahoo Research {cavallo,sviri}@yahoo-inc.com

Abstract. Fantasy sports is a fast-growing, multi-billion dollar industry [10] in which competitors assemble virtual teams of athletes from real professional sports leagues and obtain points based on the statistical performance of those athletes in actual games. Users (team *managers*) can add, drop, and trade players throughout the season, but the pivotal event is the player draft that initiates the competition. One common drafting mechanism is the so-called *auction draft*: managers bid on athletes in rounds until all positions on each roster have been filled. Managers start with the same initial virtual budget and take turns successively nominating athletes to be auctioned, with the winner of each round making a virtual payment that diminishes his budget for future rounds. Each manager tries to obtain players that maximize the expected performance of his own team. In this paper we initiate the study of bidding strategies for fantasy sports auction drafts, focusing on the design and analysis of simple strategies that achieve good worst-case performance, obtaining a constant fraction of the best value possible, regardless of competing managers' bids. Our findings may be useful in guiding bidding behavior of fantasy sports participants, and perhaps more importantly may provide the basis for a competitive *auto-draft* mechanism to be used as a bidding proxy for participants who are absent from their league's draft.

1 Introduction

In fantasy sports, individuals compete against each other by becoming virtual managers of a team of professional athletes, choosing players and modifying rosters over the course of a season, competing based on the statistical performance of the athletes composing their respective teams. Athlete statistics from real-life games are converted into “fantasy points”, which are compiled and aggregated. Points may be manually calculated by a participant designated as “league commissioner” who coordinates and manages the overall league, or they may be compiled and calculated by online platforms tracking game results. Managers draft, trade, and cut athletes over the course of the season in response to changing evaluations of athlete potentials, analogously to real sports.

Fantasy sports are a multi-billion dollar industry [10]. According to the Fantasy Sports Trade Association (FSTA), in 2015 there were 56.8 million people

* This author's work was partly supported by the EU FET project MULTIPLEX no. 317532.

playing fantasy sports in the USA and Canada. On average, fantasy-sports participants (age 18+) spend \$465 on league-related costs, single-player challenge games, and league-related materials over a 12-month period [1]. Each fantasy-sports league can have up to 20 teams, although most managers prefer to have leagues of at most 12, presumably because large leagues require drafting lower-performance and lesser-known athletes.

There are two standard ways to pick (draft) a fantasy team: snake drafts and auctions. Even though the majority of fantasy leagues use the standard snake draft (taking turns choosing players), anecdotal evidence suggests that auction-style drafts—wherein managers take turns *nominating* players, who are ultimately allocated based on competitive bidding—are popular among the most experienced and engaged users. If one of the managers is not present during the draft, there is typically an *auto-draft* algorithm, which makes decisions for the absentee manager. Furthermore, in some cases the auto-draft system is employed as a practice tool for inexperienced users. The real-time fantasy auction draft involves sophisticated decision making and strategies [2, 7], and being able to practice such an auction in advance is important for one’s success during the actual draft.

This paper is motivated by the goal of designing an auto-draft proxy bidder good enough to keep a manager competitive—in some broad sense—with the rest of the league. Because this is a competitive and strategic environment, it is natural to take a game-theoretic approach and consider strategy profiles that form Nash equilibria. However, we find that plausible pure strategy equilibria often do not even exist, and in the cases where they do exist their computation would be completely impractical for non-computer-aided participants, thus removing any predictive value. We therefore instead focus on strategies with good *guarantees*, creating teams with a certain competitive value even when faced with optimal adversarial opponent bidders.

1.1 Description of the real auction

A fantasy sports league is composed of a set of k team *managers* (or *users*) u_1, \dots, u_k —where k usually ranges from 3 to 20—who form teams from a pool of n *athletes* (or *players*) P_1, \dots, P_n .³ Each fantasy team must be composed of m athletes. The number of athletes m depends on the sport and fantasy games provider; for example, Yahoo Daily Fantasy NFL team rosters have 9 slots: 1 quarterback, 2 running backs, 3 wide receivers, 1 tight end, 1 flexible position, and 1 slot for a defensive team.

As mentioned above, leagues are formed via snake draft or auction draft, with this choice determined by the initiator of the league. In snake drafts users successively take athletes, with no bidding or competition. But auction drafts, which are the focus of the current paper, are strategically complex. Each fantasy manager is given a fixed budget B to draft a team with a set number of

³ For instance, in the NFL each actual team consists of a 53-man roster (plus a 5-athlete practice squad). There are 32 NFL teams, making for a total of $n \geq 1,696$ athletes, all of whom are eligible to be drafted.

athletes. For instance, Yahoo provides each team manager with a \$200 budget, although the precise value is not extremely relevant. Managers take turns successively (in some pre-determined order) nominating athletes for bidding, which then proceeds via an *English auction* as follows. The default bid ascribed to the nominating manager is \$1, and it can be raised to any dollar value within the nominating manager’s budget. Managers are given a fixed amount of time (say 10 seconds) to place a higher bid; if any manager does so, the clock resets and users have another 10 seconds to place a yet higher bid, and so on. Managers are free to bid any amount as long as they have available funds. When bidding stops, the manager who submitted the high bid wins the athlete and his bid amount is subtracted from the funds available to him for future bidding.

There is no incentive to retain budget after the auction, because leftover money cannot be used for anything else. Managers are allowed to bid only to a dollar amount that leaves at least one dollar for each currently unfilled spot on the team. Further, each athlete comes with a fixed position (there are quarterbacks, running backs, wide receivers, tight ends, kickers, and defensemen for NFL fantasy athletes). Each team must meet a fixed distribution of positions that depends on the sport and fantasy-game provider (for example, an NFL manager has to acquire exactly 2 running backs). The goal of a manager is to fill her positions in a way that maximizes the overall value of the team according to the stats her players accrue during real play throughout the season.

1.2 The auction that we analyze

In this work we initiate the study of fantasy-sports auctions. We make some simplifying assumptions to make the problem amenable to a rigorous theoretical investigation. Our first simplifying assumption is that team managers agree on the *value* of every athlete. This assumption, clearly somewhat of a departure from reality, is made primarily to simplify presentation; in fact, if value estimates vary across managers, our main results continue to apply with respect to the manager’s subjective value estimates for the players in the league (see the Conclusion section for further discussion of this). Moreover, a rough correspondence in value estimates can be expected because previous season statistics (average number of fantasy points per game, etc.) and *next-season projections* are widely available on specialized services on the web (e.g., at rotowire.com), which many managers consult. Formally, we assume that each athlete P_i has an associated value v_i denoting the expected number of fantasy points he will earn throughout the season; v_i is a shared belief, common knowledge to all managers.

A second simplifying assumption is that each phase of the auction draft is a second-price sealed-bid auction in which managers are allowed to place arbitrary fractional bids (limited only by their remaining budget). Even though this also deviates from reality, the English auction described in the previous section yields similar results, in theory. If more than one manager announces the highest bid in our sealed bid auction, we allocate the athlete to each high-bidding manager with equal probability. The exception to this rule is when all managers place the minimal bid. In this case we assume that the athlete is won by the nominating

manager, deterministically. That is, no manager can impose an athlete on others simply by nominating him. This reflects the real-life policy: an athlete always goes to a nominating manager if there is no competing bid from other managers.

We assume that managers are allowed to bid anywhere between \$0 and their entire remaining budget for each athlete (recall that in reality bids are capped to ensure at least \$1 remains for every open roster position). If a manager expends the entire budget before filling her roster, she can pick up lesser-preferred athletes for free after all other managers have filled their rosters. If there are multiple managers with zero budget left at the end of the auction, the remaining athletes are chosen by managers according to the nomination order.

Having fixed the price of each athlete, we can assume that for each position type the player pool has exactly the number of athletes required to complete each team. For example, if there are 8 NFL managers, and each manager must hire 2 running backs, then there exist exactly 16 running backs overall; these will be the most valuable, as no manager will prefer a lower-ranked running back (and it is easy to show that no advantage can be gained by putting one up for bid). We can therefore assume that $n = km$.

Finally, it will be useful to define the value

$$V = \frac{1}{k} \sum_{i=1}^n v_i$$

This can be considered the *fair share* of total value each manager would get if all managers were equally competitive. Against optimally bidding adversaries, the best a manager can reasonably hope for is to draft a team of value not much worse than V .

1.3 Summary of results

After describing related work in Section 2, we begin in Section 3 by showing that pure strategy subgame perfect Nash equilibria do not generally exist in the game representing the fantasy auction draft. In fact, even for what is virtually the simplest example one can conjure (2 managers, 4 athletes to be drafted), unless the athletes are nominated in a particular pre-fixed order, there is no pure strategy Nash equilibrium. That entails—allowing the nomination choices to be in fact part of the strategy space—that there can be no pure-strategy subgame perfect equilibrium of the auction draft.

This, along with other practical critiques of an equilibrium-based analysis, motivates us to focus instead on worst-case analysis. In other words, can we describe a bidding algorithm that performs well no matter what opponent managers do? In Section 4 we provide an analysis of the simplified case where athletes are automatically nominated in non-increasing order of their values ($v_1 \geq v_2 \geq \dots \geq v_n$). For each athlete P_i we define his *fair price* as $c_i = \frac{v_i}{k} B$, and we use this fair price as the basis of the following simple (and nonadaptive) bidding strategy: letting b_i be the manager's current leftover budget before the i^{th} athlete is nominated, the manager always places bid $\min\{\alpha c_i, b_i\}$, with

$\alpha = 3/2$. We show that a manager playing this strategy is guaranteed to obtain value at least $V/3$, regardless of the other managers' bids. Moreover, we show that this is the best one can do for a strategy of this class, in that any alternative choice of α will yield a worst-case value no greater than $V/3$.

In Section 5 we analyze a more complex bidding strategy for the general case in which nominations are made in a general adaptive fashion according to manager strategies. We prove that the proposed bidding strategy is able to guarantee a final team with value at least $V/16$.

2 Related work

For a general reference to auction theory we refer to [15]. Dynamic auctions have been extensively studied in a literature that has recently been very active [3]. Note that there is no mechanism design component in the current paper, because the draft mechanism is determined by the platform and we are rather studying how to best engage with it. We focus on the strategic implications for users who in principle can play very sophisticated bidding strategies. Users can strategize on the order of nomination of the players and on the bids posted in each round of the auction. Existence of equilibria in dynamic auctions—subgame perfect equilibria [17]—indeed requires very strong rationality assumptions on agents [14].

In this work we do not attempt to analyze the structure of complex equilibria, which, when they even exist, are implausible to reach. We instead propose bidding strategies that achieve the formation of teams with a guaranteed share of the total value of the player pool, even in the presence of irrational opponents. Another difference from classical dynamic auctions is that users cannot participate in all rounds of the auction. Users pass if their budget is expired or if the role of the nominated player has already been filled on the team.

The imposition of financial constraints such as budgets is also known to alter the properties of even simple standard auctions [6]. For example, the VCG mechanism [20] does not retain all its glorious properties if payments are restricted by a budget [9]. Dynamic auctions with budgets have recently been investigated in the setting of ad auctions for sponsored search [12].

Although bidding and player nomination strategies are at the crux of our work, the reader may note some commonalities with the vast literature on the fair division of goods [4, 19, 16]. Given the symmetric, full-information setting we consider, the best outcome that a strategy played by all team managers can achieve is the proportional share of the total value of the players. Proportionality is also one of the main goals in the fair-division literature. A second commonality is the allotment of an equal budget to all users to play in all rounds of the auction. Competitive equilibrium from equal incomes (CEEI) [13, 5] also assumes the same budget given to each agent to acquire a set of indivisible goods. However, differently from CEEI, the prices of the goods in fantasy auctions are decided by iterative auction rounds rather than by a centralized pricing mechanism.

We also note the connection to an array of prior work analyzing bidding strategies and autonomous bidding across a diverse spectrum of competitive environments (see, e.g., [18, 8, 11]).

3 On the Absence of Equilibria

Given that our setting is one of strategic interaction amongst self-interested agents, at first blush game theory seems the most natural way to approach an analysis and evaluation of bidding strategies. However, there are issues in this approach. For any realistic version of the fantasy auction draft game, computing equilibria is virtually guaranteed to be computationally intractable. Thus, in a single-shot world where the participants (save for the one being served by our bidding proxy) are human rather than computational, it is highly implausible to ascribe predictive power to any given strategy profile merely because it constitutes a Nash equilibrium.

But there is an even more fundamental issue: generally there will *not even exist* any pure strategy subgame perfect equilibria in fantasy draft auctions. Consider the following simple example:

Example 1. There are two users with equal budgets; each team roster has two slots; there are four athletes eligible for nomination, two of unequal positive value, and two of value 0.

Claim. In Example 1, if the lower (nonzero) value athlete is nominated first, there exists no pure strategy Nash equilibrium forward from that point.

Proof. First note that in any equilibrium, regardless of what transpires in the first round, in the round in which the other nonzero athlete is nominated each agent must bid as much of his remaining budget as necessary to win (if possible). Then, if someone wins the first athlete for a nonzero price, he will lose the second (high-value) athlete; if he can bid instead to lose the first athlete in order to win the second, this is beneficial.

If bids for the first athlete are 0 and $\epsilon > 0$, the first manager has a beneficial deviation in instead bidding $\delta \in (0, \epsilon)$, losing the low-value athlete and taking the opposing manager out of contention for the high-value athlete.

If bids for the first athlete are $\delta > 0$ and $\epsilon > 0$, with $\epsilon > \delta$, the manager bidding ϵ has a beneficial deviation in instead bidding $\gamma \in (0, \delta)$, taking the opposing athlete out of contention for the high-value athlete, as above. Similarly, if the two managers bid the same value $\epsilon > 0$, they each have positive probability of winning, and thus each has a beneficial deviation in instead bidding $\gamma \in (0, \epsilon)$.

If bids for the first athlete are 0 and 0, then the nominating manager gets that athlete, and then can bid B in the second round and also win the high-value athlete with probability 1/2. The other manager has a beneficial deviation in bidding any $\epsilon > 0$ first, since he will then take the low-value athlete for free (with certainty), and still win the high-value athlete with probability 1/2.

This exhausts the space of possible bids in round 1, and none are consistent with an equilibrium strategy profile.

Since the low-value athlete being nominated first is a possible path of the full game, and since there is no equilibrium strategy profile forward from that subgame, we have the following corollary:

Corollary 1. *In Example 1 there exists no pure-strategy subgame perfect Nash equilibrium of the game as a whole.*

However, if the *high-value* athlete is nominated first, there is a straightforward pure strategy Nash equilibrium of that subgame: both users bid their entire budgets for each nonzero value athlete. Each user will win one nonzero value athlete: with probability 1/2 it will be the high-value athlete; deviating by underbidding in the first round will only ensure that the user gets the low-value athlete (or possibly nothing if he bids 0 in the first round).

Therefore, if we look at the game as a whole, recognizing that athlete nomination is a part of the strategizing, is there an equilibrium? There is, but it is not subgame perfect (which we knew must be the case from Corollary 1), and thus not very satisfying in any kind of predictive sense. Any equilibrium strategy profile must specify what happens off the equilibrium path, and there is no way to do so for the “lower-value athlete nominated first” path that is consistent with rational users (since there is no equilibrium there). Thus, the existence of pure strategy equilibria inherently depends on a model of user behavior that subverts the main rationale for considering equilibrium in the first place, that is, the idea that agents are rational. Other examples may have no equilibria whatsoever.

4 Analysis of Fair-Price Bidding

In this section we start our adversarial analysis by considering the case where athlete nomination does not form part of the strategy space. Instead, players are nominated in non-increasing value order, and the strategy space of each manager consists in bidding for these players. Recall that we have a second-price auction, and in the case of (positive bid) ties the athlete is allocated uniformly at random among the managers that placed the highest bid.

Recall that there are k team managers and m slots per team, that v_i is the value of the athlete with i^{th} highest value, and that $V = \frac{1}{k} \sum_{i=1}^{km} v_i$.

For each athlete P_i we define his *fair price* as:

$$c_i = v_i \frac{B}{V}.$$

We first consider the following natural nonadaptive strategy, which we call *simple fair-price bidding*. We will show that it may lead to a bad team, but can be made robust via a slight modification, motivating in this way our variant solution. Let b_i be the current leftover budget before the i^{th} athlete is nominated. In simple fair-price bidding the manager always places a bid equal to $\min\{c_i, b_i\}$. That is,

$$\text{SIMPLE-FAIR-PRICE-BID}(b_i, v_i) = \min \left\{ \frac{v_i}{V} B, b_i \right\}.$$

We will show that a manager following the simple fair-price bidding strategy could end up with a team of value arbitrarily close to V/k (which is bad). Consider the case where $v_1 = \dots = v_{k-1} = V(1 - \varepsilon)$ and $v_k = \dots = v_{km} = \frac{V(1+k\varepsilon-\varepsilon)}{km-k+1}$. Our manager bids $(1-\varepsilon)B$ for the first $k-1$ athletes; imagine that the other managers bid $(1-\varepsilon/2)B$. This means that after $k-1$ rounds our manager did not win any athletes and has his full budget B available for bidding, whereas the remaining $k-1$ managers have each won one athlete of value $V(1-\varepsilon)$ and have remaining budget $\varepsilon B/2$.

We assume that ε is chosen to be small enough that

$$\frac{V(1+k\varepsilon-\varepsilon)}{km-k+1} \cdot \frac{B}{V} > \varepsilon \frac{B}{2}.$$

In this case our manager bids the fair value of $\frac{(1+k\varepsilon-\varepsilon)B}{km-k+1}$ for each of the next k athletes and wins each of these bids. Therefore, the value of the final team for our manager, who uses the *simple fair-price bidding* strategy in the auction, is

$$\frac{(1+k\varepsilon-\varepsilon)Vm}{km-k+1},$$

which is arbitrarily close to V/k for large enough m and small enough ε .

This example motivates us to modify the simple fair-price bidding strategy. In the modified fair-price bidding strategy, with parameter $\alpha \geq 1$, the manager always places bid $\min\{\alpha c_i, b_i\}$. That is,

$$\text{FAIR-PRICE-BID}(b_i, v_i) = \min\left\{\frac{\alpha v_i}{V}B, b_i\right\}.$$

Theorem 1. *The expected value of the team generated by FAIR-PRICE-BIDDING with parameter $\alpha = 3/2$ is at least $V/3$, regardless of the other managers' bidding strategies.*

Proof. Assume we have $r \geq 0$ athletes with values $v_1 \geq \dots \geq v_r \geq V/\alpha$. Our manager bids her whole budget B for these athletes. After r rounds of the auction our manager gets an athlete of value $v_i \geq V/\alpha$ with probability $p \geq \min\{1, r/k\}$. If $r \geq k$ then our manager always gets one of these high-value athletes and there is nothing more to prove. So now assume that $r < k$.

Let $W = \sum_{i=1}^r v_i$. We claim that the expected value E' of the final team selected by our manager conditioned on the fact that she wins one of the high-value athletes is at least W/r . Let $p_i \geq \frac{1}{r-i+1}$ be the probability that our manager wins the i^{th} round of the auction conditioned on the fact that she did not win any of the previous rounds. Then we can estimate:

$$E' \geq v_1 p_1 + v_2 (1-p_1) p_2 + \dots + v_r \prod_{i=1}^{r-1} (1-p_i) p_r \geq \sum_{i=1}^r v_i / r$$

by using the standard majorization inequality (it can be proven easily by induction on t) and the fact that:

$$\begin{aligned} p_1 + (1 - p_1)p_2 + \cdots + \prod_{i=1}^{t-1} (1 - p_i)p_t &= 1 - \prod_{i=1}^t (1 - p_i) \\ &\geq 1 - \prod_{i=1}^t \left(1 - \frac{1}{r - i + 1}\right) = 1 - \frac{r - t}{r} = \frac{t}{r}. \end{aligned}$$

If our manager does not get a high-value athlete (it happens with probability $1 - p$), then the remaining $k - r$ managers with budgets B participate in the auction to acquire athletes with total value $kV - W$. Also, all the athletes auctioned in this case have values $v_i < V/\alpha$. When all the managers have spent their entire budgets or filled all the spots on their rosters, the remaining athletes are distributed at random in a fair way. We consider two cases:

1. There is a moment during the auction when our manager cannot bid the target price of αc_j for the current athlete P_j , that is, $b_j < \alpha v_j \cdot B/V$. We claim that the value of the athletes on our manager's roster at this moment is already at least $\frac{V}{2\alpha}$.

Assume to the contrary that the value of the athletes on our manager's roster at this moment is strictly less than $\frac{V}{2\alpha}$. Then $b_j > B/2$ and we derive that $v_j > \frac{V}{2\alpha}$. Therefore, all athletes nominated so far have values higher than $\frac{V}{2\alpha}$ and our manager did not win any of them (or we are done). Therefore $b_j = B$. But recall that $v_j < V/\alpha$ (the athlete is not one of the r high-value athletes), which means that our manager can bid the target price of αc_j , arriving at a contradiction. Therefore, the value of the athletes on our manager's roster is at least $\frac{V}{2\alpha}$. Let p' be the probability that the auction ends in this case.

2. Now take the case where we never reach a situation in which we cannot bid the target price, and our manager fills all the spots on her roster. Assume that p'' is the probability of this case, and note that $p + p' + p'' = 1$.

At any moment during the auction our manager either bids the target price of αc_j for the athlete or does not bid anything because she already filled positions on her roster of the same type as the position type of the athlete who is nominated now. Let X be the value of the team of our manager at the end of the auction. Consider manager u who is one of the remaining $k - r - 1$ managers. We divide the team of manager u into two sets of athletes $S^+(u)$ and $S^-(u)$, where $S^+(u)$ is the set of athletes that u won for price greater than or equal to the target price of αc_j and $S^-(u)$ is the set of athletes that u won for price strictly smaller than the target price.

We claim that $\sum_{i \in S^+(u)} v_i \leq V/\alpha$ and $\sum_{i \in S^-(u)} v_i \leq X$. To show the former, assume that $\sum_{i \in S^+(u)} v_i > V/\alpha$. Then u has paid at least

$$\sum_{i \in S^+(u)} \frac{\alpha v_i}{V} B = \alpha \frac{B}{V} \sum_{i \in S^+(u)} v_i > \alpha \frac{B}{V} \cdot \frac{V}{\alpha} = B,$$

leading to a contradiction, because the initial budget of manager u is B . The only way for u to win a bid for an athlete in $S^-(u)$ is if our manager (who bids the target value for each athlete she has a remaining slot for on her team) does not place a bid. That could only happen when our manager already filled the spots on its roster corresponding to the position of the athlete who is currently nominated. That means all athletes in $S^-(u)$ have smaller value than all athletes of the same type on our manager roster, that is, $\sum_{i \in S^-(u)} v_i \leq X$.

Analogously, for each manager u who got one of the first r athletes, the value of the remaining athletes on his team is at most X . Therefore,

$$kV \leq kX + (k - r - 1) \frac{V}{\alpha} + W,$$

which entails that

$$X \geq V - \frac{W}{k} - \frac{k - r - 1}{k} \cdot \frac{V}{\alpha}.$$

Overall, the expected value of the final team for our manager is lower-bounded by:

$$\begin{aligned} \frac{W}{r}p + \frac{V}{2\alpha}p' + Xp'' &\geq \frac{W}{r}p + \frac{V}{2\alpha}p' + \left(V - \frac{W}{k} - \frac{k - r - 1}{k} \frac{V}{\alpha} \right) p'' \\ &\geq \frac{W}{r} \cdot \left(p - \frac{r}{k} p'' \right) + \frac{V}{2\alpha} p' + \left(1 - \frac{k - r}{\alpha k} \right) V p'' \\ &\geq \frac{V}{\alpha} \cdot \left(p - \frac{r}{k} p'' \right) + \frac{V}{2\alpha} p' + \left(1 - \frac{k - r}{\alpha k} \right) V p'' \\ &= \frac{V}{\alpha} \cdot \left(p - \frac{r}{k} p'' \right) + \frac{V}{2\alpha} p' + \left(1 - \frac{k - r}{k} \right) \frac{V}{\alpha} p'' + \left(1 - \frac{1}{\alpha} \right) V p'' \\ &= \frac{V}{\alpha} p + \frac{V}{2\alpha} p' + \left(1 - \frac{1}{\alpha} \right) V p'' \geq \frac{V}{3}. \end{aligned}$$

We now give examples that show that our analysis of the algorithm is tight, in that for any choice of α , given $k > 2$ managers, there exists an example where our manager collects no more than a third of the fair value.

First off, for any $\alpha \leq 1$, the example given at the beginning of this section serves to demonstrate the claim. Now consider any $\alpha > 3/2$ and the following example: there are two types of athletes, k managers, and team composition constraints where we need to choose exactly one athlete of type one and $m - 1$ athletes of type two. The athletes of type one have value $v_1 = \dots = v_k = \frac{1}{3}V(1 + \varepsilon)$ (k athletes). There are $2k - 2$ athletes of type two with value $v_{k+1} = \dots = v_{3k-2} = \frac{1}{3}V$. The remaining athletes are also of type two but have value

$$v_{3k-1} = \dots = v_{km} = \frac{kV - kv_1 - (2k - 2)v_{k+1}}{km - k - 2(k - 1)} \approx \frac{2V}{3km},$$

where the approximation is close for large enough m and small enough ε .

During the first bidding round our manager bids $\min\{\alpha B(1 + \varepsilon)/3, B\} > B(1 + \varepsilon)/2$ for the first athlete. The other managers let our manager win with this bid and force her to pay $B(1 + \varepsilon)/2$ (by bidding this amount), and they get the remaining $k - 1$ athletes of the first type for free (one athlete per each manager). During the next $2k - 2$ bidding rounds our manager bids her whole remaining budget of $B(1 - \varepsilon)/2$ whereas the other managers bid $B/2$ and win each of these rounds. After that our manager wins the next $m - 1$ bidding rounds and obtains a team of value $v_1 + (m - 1)v_{3k-1} \approx V/3$.

We now consider the case where $\alpha \leq 3/2$. Consider the example where there is only one type of athlete, there are k managers, and athletes have the following values: $v_1 = \dots = v_{k-1} = \frac{2}{3}V(1 - \varepsilon)$ and $v_k = \dots = v_{km} = \frac{kV - (k-1)v_1}{km - k + 1}$. Our manager bids $2\alpha B(1 - \varepsilon)/3 < B$ for each of the first $k - 1$ athletes and loses all of those rounds to other managers who bid their whole budgets B . After that our manager wins m bids and ends up with a team of value

$$m \cdot \frac{kV - (k - 1)v_1}{km - k + 1},$$

which is arbitrarily close to $V/3$ for large enough m and small enough ε .

5 Arbitrary Nomination Order

We now consider the more general (and realistic) setting in which each manager is allowed to nominate an arbitrary athlete when his turn comes up, with the order of nominators pre-determined arbitrarily.

5.1 Algorithm description

Our algorithm for this expanded version of the problem, which we call SELECTIVE-FAIR-BIDDING, will depend on three parameters: α , β , and γ , with $\beta > \alpha \geq 1$ and $\gamma \geq 1$. We define two groups of athletes: $L = \{i : v_i \geq V/\beta\}$ and $S = \{i : v_i < V/\beta\}$.

If $\sum_{i \in L} v_i \geq \sum_{i \in S} v_i$, our manager ignores (bids 0 for) all athletes with indices in the set S and only bids for athletes with indices in L . The bid value is always the whole budget B .

Now we describe the strategy in the case that $\sum_{i \in L} v_i < \sum_{i \in S} v_i$. Letting t be the number of distinct types of athletes that each team must have, let d_j be the number of athletes of type $j = 1, \dots, t$ the team must have, with $\sum_{j=1}^t d_j = m$. All athletes in the set S are partitioned into t sets S_1, \dots, S_t by type. Let

$$A_j = \frac{\sum_{i \in S_j} v_i}{kd_j}$$

be the average value for the athlete of type $j = 1, \dots, t$ (among the athletes in group S) per available spot on the team roster.

Our manager only bids for an athlete P_i if either $i \in L$, or $i \in S_j$ and $v_i \geq A_j/\gamma$ for some $j \in \{1, \dots, t\}$. We will call such athletes *desirable*. The bid value for desirable athletes is $\min\{\alpha c_i, b_t\}$ where b_t is the remaining budget in the current period t of the bidding procedure and c_i is the fair value of athlete i . The athletes from the set S of value less than A_j/γ will be called *undesirable*, and our manager bids 0 for all of them.

Our manager always nominates the highest-value available athlete that can fill one of the positions on her roster.

5.2 Analysis

Theorem 2. *The expected value of the team generated by SELECTIVE-FAIR-BIDDING with $\alpha = 16/3$, $\beta = 8$, and $\gamma = 2$ is at least $V/16$.*

Proof sketch. The proof can be divided into two cases, the second of which we only sketch here due to limitations on space (the proof is otherwise complete).

Case 1. Assume that $\sum_{i \in L} v_i \geq \sum_{i \in S} v_i$. Recall that our manager bids her whole budget B for each athlete in L until she either wins one or there are no more athletes in L left. There are $|L|$ rounds of the auction that are relevant to our manager (as she bids zero in the others); let $p_j \geq 1/k$ be the probability that our manager wins an athlete during the j th round our manager bids in, for $j = 1, \dots, |L|$, conditioned on the fact that she did not win an athlete during the previous rounds. And let $v_{L(j)}$ be the value of the athlete nominated in the j th round our manager bids in. Then the expected value of the athletes won by our manager is at least

$$E' = v_{L(1)}p_1 + v_{L(2)}(1-p_1)p_2 + \dots + v_{L(|L|)} \prod_{i=1}^{|L|-1} (1-p_i)p_{|L|}.$$

If $\prod_{i=1}^{|L|} (1-p_i) \leq 1/2$, then

$$E' \geq \left(1 - \prod_{i=1}^{|L|} (1-p_i)\right) \min_i v_{L(i)} \geq \frac{V}{2\beta}.$$

Otherwise, that is, if $\prod_{i=1}^{|L|} (1-p_i) > 1/2$, then

$$E' > \sum_{i \in L} \frac{v_{L(i)}p_i}{2} \geq \sum_{i \in L} \frac{v_{L(i)}}{2k} \geq \sum_{i=1}^n \frac{v_{L(i)}}{4k} = \frac{V}{4}.$$

Case 2. Assume that $\sum_{i \in L} v_i < \sum_{i \in S} v_i$. This is the more difficult of the two cases, and due to space constraints we must omit much of the analysis. However, it can be established that in this case:

$$E' \geq V \left(\frac{1}{2} \min \left\{ \frac{1}{\gamma}, 1 - \frac{1}{\gamma} \right\} - \frac{1}{\alpha} \right),$$

If we choose $\gamma = 2$, combining bounds for the various cases, the expected value of our manager’s team at the end of the auction is at least:

$$V \cdot \min \left\{ \frac{1}{4} - \frac{1}{\alpha}, \frac{1}{\alpha} - \frac{1}{\beta}, \frac{1}{4}, \frac{1}{2\beta} \right\}.$$

Now choosing $\alpha = 16/3$ and $\beta = 8$, we derive the lower bound of $V/16$ on the expected value of the final team.

6 Conclusion

In this paper we initiated the study of fantasy auction drafts, which play an important role in the large and growing market of fantasy sports. We abstracted the problem by defining a simple but realistic auction, which models the real-life process. We studied pure-strategy Nash equilibria and showed that even for two players, there are no pure-strategy subgame perfect Nash equilibria. We thus turned our attention to worst-case outcomes and designed a deterministic algorithm for bidding in fantasy auction drafts, which guarantees the creation of a team with a total value that in expectation (over the random choices of the allocation mechanism in case of ties in the highest bids) is at least a constant (1/16) approximation of the optimal possible.

Throughout the paper we assumed that for each athlete P_i there exists a universal value v_i ; yet our results are more general. Our worst-case bounds hold also for the setting where each manager u_j can have an idiosyncratic value $v_{i,j}$ for each athlete P_i . Our algorithm can guarantee the creation of a team that has a value of at least $V_j/16$, where

$$V_j = \frac{1}{k} \sum_{i=1}^{km} v_{i,j}.$$

Consideration of equilibrium outcomes was a nonstarter for our purposes quite apart from concerns about collusion and the like. However, one can show rather easily that if other managers do collude, it can have an adverse effect on the quality of the team our manager ends up with. Yet our algorithm is *collusion resistant* with respect to the worst case: even if other managers collude, the guarantees remain unchanged.

The main open theoretical question regards closing the gap between what we guarantee with our strategy and what is *possible* to guarantee. We have given an algorithm that provides a constant approximation (1/16) to the best one can hope for. Can the constant be improved? Can one show any upper bound?

Finally, there are empirical angles that we hope will be pursued. The fact that the worst-case guarantee of the algorithm we propose is $V/16$ does not at all indicate that it will not be competitive in practice. If such bidding strategies are ultimately implemented via auto-draft proxies in actual fantasy auctions, this will yield an intriguing dataset that can be studied and perhaps used as the basis for an empirically grounded iteration on the work we initiated here.

References

1. FSTA industry demographics. <http://fsta.org/research/industry-demographics/>.
2. Michael Beller. Fantasy football 2015 draft preview: Auction strategies. *Sports Illustrated*, August 20, 2015.
3. Dirk Bergemann and Maher Said. Dynamic Auctions: A Survey. Cowles Foundation Discussion Papers 1757R, Cowles Foundation for Research in Economics, Yale University, March 2010.
4. S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
5. E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium. *Journal of Political Economy*, 119:1061–1103, 2011.
6. Y.-K. Che and I. Gale. Standard auctions with financially constrained bidders. *The Review of Economic Studies*, 65:1–21, 1998.
7. Tristan H. Cockcroft. Fantasy football auction strategy. http://espn.go.com/fantasy/football/story/_/id/13084199/fantasy-football-auction-strategy-fantasy-football, Aug 18, 2015.
8. Janosz A. Csirik, Michael L. Littman, Satinder Singh, and Peter Stone. Faucs: An FCC spectrum auction simulator for autonomous bidding agents. In *Proceedings of the Second International Workshop on Electronic Commerce*, pages 139–151, 2001.
9. Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. In *FOCS*, pages 260–269, 2008.
10. Brian Goff. The \$70 billion fantasy football market. *Forbes*, August 20, 2013.
11. A. Greenwald and J. Boyan. Bidding under uncertainty: theory and experiments. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pages 209–216, 2004.
12. R. Gummadi, Peter B. Key, and Alexandre Proutiere. Optimal bidding strategies in dynamic auctions with budget constraints. In *49th Annual Allerton Conference on Communication, Control, and Computing, Allerton Park & Retreat Center, Monticello, IL, USA, 28-30 September, 2011*, page 588, 2011.
13. A. Hylland and R. Zeckhauser. Efficient allocation of individuals to positions. *Journal of Political Economy*, 87:293–314, 1979.
14. Krishnamurthy Iyer, Ramesh Johari, and Mukund Sundararajan. Mean field equilibria of dynamic auctions with learning. In *Proceedings of the 12th ACM Conference on Electronic Commerce, EC '11*, pages 339–340, New York, NY, USA, 2011. ACM.
15. Paul Robert Milgrom. *Putting auction theory to work*. Cambridge University Press, 2004.
16. Herve Moulin. *Fair Division and Collective Welfare*. MIT Press, 2004.
17. M. J. Osborne. *An Introduction to Game Theory*. 2004.
18. Daniel M. Reeves, Michael P. Wellman, Jeffrey K. MacKie-Mason, and Anna Osypayshvili. Exploring bidding strategies for market-based scheduling. *Decision Support Systems*, 39:67–85, 2005.
19. H. Steinhaus. The problem of fair division. *Econometrica*, 16, 1948.
20. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.