



A No-Tanking Draft Allocation Policy*

Martino Banchio[†] Evan Munro[‡]

January 30, 2020

1 Introduction

In the absence of intervention by the league, there is significant path dependence in a sports franchise's performance. Teams who perform well in one year have more opportunities to collect attendance and TV revenue. With this additional revenue, they can bid for better players. Over the long-term, without re-distributive mechanisms, leagues may end up with a handful of dominant teams and a set of weak teams that are not financially viable. For a business that relies on a large and engaged fan base for revenue, this is not a desirable outcome. As a result, to ensure that in the long-term there is some balance in team performance, sports leagues have instituted draft allocation mechanisms that favor the lowest ranked teams in the league. These mechanisms, however, can provide incentives for teams to purposely lose games in order to have a better chance of receiving a draft pick, a practice known as 'tanking'. If teams do tank as a result of league policies, then league revenue can be impacted through drops in game attendance and TV revenues.¹

Both the NHL and the NBA currently allocate top draft picks through a rank-order lottery, but have changed lottery probabilities frequently in attempt to minimize tanking while still favoring the lowest ranked teams. The NFL and the MLB simply allocate drafts based on reverse rank from the previous season. In all four leagues, which collectively produce over 30 billion dollars of revenue, there is frequent discussion in the media and among fans of teams purposely losing games in order to increase the expected value of their draft allocation for the year.² This shirking behavior has been documented empirically in Taylor and Trogdon (2002). We use NBA data for the applications in this paper, but the analysis is easily adapted to other sports leagues. Some of the general principles derived are also related to more general re-distributive settings, such as financial aid or food stamp allocation.

*We thank Edward Lazear, Mohammad Akbarpour, Andrzej Skrzypacz and Brad Ross for helpful comments and discussions. Code for the analysis in the paper is available at <http://github.com/evanmunro/draft-policy>.

[†]Graduate School of Business, Stanford University. Email: mbranchio@stanford.edu

[‡]Graduate School of Business, Stanford University. Email: munro@stanford.edu

¹For an example of this, see discussion of the Golden State Warriors in 2012: <https://bleacherreport.com/articles/1160424-golden-state-warriors-why-the-warriors-tanking-is-hurting-the-nba>

²Just in 2019, for example, commentators have argued frequently that the Golden State Warriors in the NBA and the Miami Dolphins in the NFL either are, or should be tanking: (<https://www.si.com/nba/2019/11/14/golden-state-warriors-tanking-steve-kerr>, <https://www.sbnation.com/nfl/2019/9/15/20861089/miami-dolphins-tanking-2019-draft-picks-cap-space>).



Currently, the NBA bases the draft probability lotteries on final rankings only. Our contribution highlights the importance of the dynamic aspects of the repeated two player games that occur during a sports season, in two different ways. First, we show that any re-distributive policy that relies on final rankings only to allocate draft picks can provide incentives for a team to lose in some possible history of a season. This negative results forces us to look at mechanisms that adapt the draft probabilities over time. We set up the optimization problem of the league as a central planner and we show that computing an optimal solution is not feasible. The objective is to maximize the probability that the first draft pick is assigned to the worst team, with the constraint that the mechanism cannot provide incentives for teams to tank. Then we propose a rule which, while not necessarily globally optimal, is optimal for a large subset of incentive compatible rules and is easy to implement. Our proposed rule adjusts the lottery draft probabilities after each game, to match the probability each team will be ranked last after all the games have been played, conditional on the games so far. This adjustment process is ended before incentive-compatibility would be violated if it continued. This means that in practice, the lottery probabilities are determined dynamically in each season, and depend on a restricted set of games. We show that not only does this process incentivize effort in every game, it is also approximately optimal in the sense that it maximizes the probability the draft pick is given to the worst team in any given season, conditional on a subset of results for that season.

There is existing work which has proposed alternative allocation systems. Gold (2010) proposes allocating the top pick to the team with the highest the number of wins after elimination from the playoffs while Lenten (2016) and Lenten et al. (2018) suggest the team that is eliminated first from playoff contention should receive the top pick. Though the authors of existing work have provided empirical evidence that their proposals would reduce tanking, there has not been a systematic theoretical evaluation of draft allocation policies. Our work provides the theory to determine if an optimal rule for draft allocation exists that satisfies the league's parity objectives while eliminating incentives for teams to lose. Under our framework neither rule is fully incentive compatible. Furthermore, Gold's rule would not favor the worst teams, so does not target our assumed league objective. Dagaev and Sonin (2018) model a different type of failure in incentive compatibility. When there are multiple tournaments played sequentially (e.g. pool play before playoffs), there can sometimes be incentives for teams to obtain a worse seed in the first round to obtain a better draw in the second round. In our model we disregard these types of incentive issues in favor of those arising from the misspecification of the draft allocation rule.

Our work is most closely related to the existing literature on tournament theory. Lazear and Rosen (1981) design a prize allocation system that incentivizes optimal effort for workers, where a worker's ranking in output, rather than their actual output, determines their prize allocation. The prizes decrease as rankings decrease. The workers have a single choice of effort, which depends on their ability and the the prize allocation system. Rosen (1986) investigates sequential elimination contests, where the prize structure must take into account how incentives change as a team progresses in the tournament. Our setting is different in a few important ways. As in Rosen (1986), we have a sequential series of games with repeated choices of effort and incentives that change over time. However, we do not have elimination, so even a team that loses games continues to play. At each game, teams know their past performance and make a choice to exert effort which is conditional on their past performance, their future schedule of games, and the prize allocation. Another important difference is that prizes are not increasing in rank. In the sports setting the league controls only the allocation of a secondary prize, a draft pick. The value of the secondary prize is taken as given by the planner, since it's determined by the skill of the newly eligible players participating in the draft that year. The maximization of



the league's objective requires that this secondary prize is allocated with higher probability to teams with the lowest rank. This structure, without a correct allocation policy that decouples prizes from the final ranking of the teams, will explicitly incentivize teams to intentionally lose games.

For readers who are not familiar with U.S. sports leagues, we provide some brief background of the draft allocation system in the NBA in Section 2. Section 3 defines the theoretical model of team decision making that we use to prove that the current allocation rule is not optimal and to introduce our draft allocation system. In Section 4, we show the results of our draft policy applied to simulated seasons and to real data from the NBA in the 1980s.

2 Background

Among the four U.S. sports leagues, the 'tanking' problem is considered most serious in the NBA.³ As a result, our empirical section focuses on the NBA, and we provide additional background for readers who are not familiar with league rules. In the current system, the NBA regular season begins in the last week of October and ends in the middle of April. There are thirty teams competing in each season, divided in two conferences. Each team plays eighty-two games in a single regular season, and each team plays every other team during the regular season, at least once at home and once away. At the end of the regular season, the teams are ranked by the number of wins. The top eight teams in each of the two conferences advance to the playoffs. The playoffs are an elimination tournament and the winner takes the championship. The remaining fourteen teams participate in the annual NBA draft lottery.

During the draft, teams can select players who are eligible and wish to join the league. An eligible player is at least nineteen years old and one year removed from their high school graduation date. The teams pick sequentially, according to a prescribed ordering, the player they value the most out of the remaining pool of eligible draftees. The first to fourth draft picks are considered the most valuable, and those are allocated using a lottery system. In Appendix A we include the draft probabilities for the first pick for systems from 1990-2019.

After the first four picks are allocated according to the outcome of the lottery, the remaining picks are allocated according to the following reverse order system: the lowest ranked team that didn't get one of the first four picks will get the fifth pick, the lowest ranked team that didn't get one of the first five picks will get the sixth pick, and so on. Once all thirty teams have received a draft pick in the first round, the second round begins and an additional thirty players are drafted according to the inverse ranking in the regular season. Players that go undrafted in the two rounds of draft picks are free to try out for any team. Teams are also allowed to sell their draft picks to other teams.

3 Designing a Draft Allocation Mechanism

3.1 Defining League and Team Objectives

First, we define some variables and notation. A season S is made up of a total of m games, each between two teams. There are a total of n teams in the season and team $i \in I = \{1, \dots, n\}$ has ability $\alpha_i \in [0, 1]$. $S_{t1} \in I$ denotes the home team for game t , and $S_{t2} \in I$ denotes the away team for game

³See, for example, <https://www.sbnation.com/2014/1/10/5266770/nba-draft-lottery-tanking-gm>



$t \in \{1, \dots, m\}$. A season is a pair of sequences $S = \{\{S_{t1}\}_{t=1}^m, \{S_{t2}\}_{t=1}^m\}$ that identifies the two contestants in every game. \vec{h}^t a t -length vector that encodes the history of the season up to and including game t . A history is defined recursively as follows:

$$\begin{cases} \vec{h}^0 &= \emptyset \\ \vec{h}^t &= \vec{h}^{t-1} \cup i & \text{if team } i \text{ wins game } t \end{cases}$$

Feasible full histories include any \vec{h}^m such that $\vec{h}_s^m \in \{S_{s1}, S_{s2}\} \forall s$. Let the set of feasible histories be \mathcal{H} . A season with m games will have $|\mathcal{H}| = 2^m$. For some history \vec{h}^t , the total number of wins for team i for that history is:

$$v_i(\vec{h}^m) = \sum_{s=1}^m \mathbb{1}(h_s^t = i)$$

The team with the k -th most wins after the season is played is $v^{(k)}(\vec{h}^m)$. For instance, the team with the most wins given a season's result is:

$$v^{(1)}(\vec{h}^m) = \arg \max_i v_i(\vec{h}^m)$$

and the team with the least wins given a season's result is :

$$v^{(n)}(\vec{h}^m) = \arg \min_i v_i(\vec{h}^m)$$

We assume that there is a primary prize V with value that is at least π^V for every team, and is allocated to the top k^* teams. For example, currently in the NBA, the top sixteen teams make the playoffs.⁴ In the NBA setting, π^V represents the minimum expected additional revenue from the playoffs that a team receives from playing in additional high-stakes games, and the related media exposure and merchandise sales. There is also a single secondary prize D . In the NBA setting it represents the first draft pick and has value of at most π^D , which represents the maximum increase in long-term expected revenue that a team expects to receive from drafting the top eligible player. We assume $\pi^D \leq \pi^V$. The secondary prize is allocated based on a league's policy $y_i(\vec{h}^t)$, which assigns a probability that team i gets the draft pick given a history \vec{h}^t . See Appendix B for a brief discussion of how to extend the model to the allocation of multiple draft picks.

In general, a draft allocation mechanism will be a rule

$$y: I \times \mathcal{H} \rightarrow [0, 1]^{m \times n}$$

such that $y_i(h^t)$ is the probability that team i will receive the draft pick given the history up to game t . The mechanism is restricted in the following ways. Since $y_i(h^t)$ represents the expected lottery probabilities conditional on information up to game t , the probabilities at time t must be dynamically consistent with the probabilities conditional on information up to game $t - 1$:

$$y_i(\vec{h}^{t-1}) = p_{it}(\alpha) y_i(\vec{h}^{t-1} \cup W) + (1 - p_{it}(\alpha)) y_i(\vec{h}^{t-1} \cup L) \quad \forall i, \vec{h}^{t-1} \quad (\text{DC})$$

⁴It is actually the top eight teams in each conference that make the playoffs, but we do not model divisions and conferences in this paper.



Moreover, the lottery probabilities at any history need to add up to one:

$$\sum_{i=1}^n y_i(\vec{h}^t) = 1 \quad \forall i, \vec{h}^t \quad (\text{PROB})$$

Let the space of all draft allocation mechanisms that satisfy DC and PROB be \mathcal{Y} . For all U.S. sports leagues currently, $y_i(\vec{h}^m)$ depends on a team's final ranking only. The final ranking is defined as follows:

$$r_i(\vec{h}^m) = k \text{ if } v^{(k)}(\vec{h}^m) = i$$

See Table A2 in Appendix A for the current probabilities for 2019, and Table A1 for a history of how the lottery mechanism has changed since 1966.

Each team makes a single strategic choice in each game, which is how much effort to exert. $e_{it} \in [0, 1]$ will denote the choice of effort by team i . A team exerting full effort in a game has $e_{it} = 1$, while a team that is tanking has $e_{it} = 0$. For our basic model specification, we assume that win probabilities for a team in game t against opponent j are generated by the following process:

- Each team independently draws a realization from the random variable z_{it} distributed according to

$$z_{it} \sim N(\alpha_i + e_{it}, 1)$$

- the probability that a team i wins against opponent j in game t is:

$$p_{it}(e_{it}, e_{jt}, \alpha) = Pr(z_{it} > z_{jt}) = 1 - \Phi\left(\frac{\alpha_j - \alpha_i + e_{jt} - e_{it}}{\sqrt{2}}\right) \quad (1)$$

where Φ is the standard normal cumulative distribution function

See Appendix B for a brief discussion of alternative methods for calculating win probabilities. In the current NBA system, a team makes the playoffs if it is in the top 16 teams at the end of the season. The probability a team makes the playoffs depends on what has happened so far, and the probability of every possible outcome in future games. For team i who plays against opponent j in game $t + 1$, $W = i$ and $L = j$ indicates team i wins. $W = j$ and $L = i$ indicates the opposing team wins and team i loses. We denote the probability of winning the primary prize in the season for team i given a history \vec{h}^t as $q_i(\vec{h}^t)$. These probabilities can be defined recursively.

$$\begin{cases} q_i(\vec{h}^m) = \mathbb{1}(r_i(\vec{h}^m) \leq k^*) \\ q_i(\vec{h}^t) = p_{i,t+1}(\alpha)q_i(\vec{h}^t \cup W) + (1 - p_{i,t+1}(\alpha))q_i(\vec{h}^t \cup L) \end{cases}$$

We suppress the dependence of $p_{i,t+1}$ on $e_{i,t+1}$ and $e_{j,t+1}$ because we focus on incentive compatible mechanisms for secondary prize allocation, where the effort chosen is always 1 and therefore the dependence on effort choice of each team disappears. In order to derive the mechanisms that disincentivize effort, we first explicitly define the team's and league's objectives:

Team Objective In game t , team $i \in \{S_{t1}, S_{t2}\}$ chooses an effort level e_{it} to maximize their expected payoff given the results in the games played so far, \vec{h}^{t-1} . For team i playing against team $j \in \{S_{t1}, S_{t2}\}$:

Problem 1.

$$\begin{aligned} \max_{e_{it}} & [p_{it}(e_{it}, e_{jt}, \alpha)(q_i(\vec{h}^{t-1} \cup W)\pi^V + y_i(\vec{h}^{t-1} \cup W)\pi^D) \\ & + (1 - p_{it}(e_{it}, e_{jt}, \alpha))(q_i(\vec{h}^{t-1} \cup L)\pi^V + y_i(\vec{h}^{t-1} \cup L)\pi^D)] \end{aligned}$$



No-Tanking Condition (NTC) The objective function can easily be rewritten as:

$$p_{it}(e_{it}, e_{jt}, \alpha) \left[(q_i(\vec{h}^{t-1} \cup W)\pi^V + y_i(\vec{h}^{t-1} \cup W)\pi^D) - (q_i(\vec{h}^{t-1} \cup L)\pi^V + y_i(\vec{h}^{t-1} \cup L)\pi^D) \right] + K$$

where K is constant with respect to the decision variable. Notice that the derivative of the probability of winning with respect to a team's own effort is positive. This means that if the term in the square brackets is positive, the team will choose the maximum e_{it} possible. Therefore, a necessary condition for team i to exert maximum effort in game t is that the following inequality is satisfied:

$$\left[q_i(\vec{h}^{t-1} \cup W) - q_i(\vec{h}^{t-1} \cup L) \right] \frac{\pi^V}{\pi^D} \geq \left[y_i(\vec{h}^{t-1} \cup L) - y_i(\vec{h}^{t-1} \cup W) \right] \quad (\text{NTC})$$

If this inequality is satisfied, it is optimal for team i to exert maximum effort in the game: the inequality implies that the derivative of the objective with respect to effort is increasing everywhere. NTC has a clear interpretation. If the team's increase in probability of receiving the secondary prize when they lose a certain game is less than the change in probability of receiving the primary prize, scaled by the ratio of the minimum value of winning the primary vs the secondary prize, then the team will exert maximum effort in a game ($e_{it} = 1$). If this condition does not hold for some team in some game, then the team will not exert maximum effort in that game ($e_{it} < 1$). We want to find a draft allocation mechanism that ensures NTC holds in every game. Note that we assume that if the team is indifferent between exerting effort and tanking, they will exert effort, since winning a game should be preferred to losing a game in the short-term, in the absence of long term incentives. In this setting, we don't need an explicit cost of effort. Instead, the cost of effort is implicit in the model and is a result of specific draft rules that make certain teams better off if they lose rather than win.

Another implication of NTC is that any allocation policy that is incentive compatible should not change the probability that a team receives the secondary prize in games where they are already out of the primary prize contest, since at that point the change in probability of receiving π^V from winning in all future games is always zero. That is not the case for any of the weighted lottery rules that the NBA has implemented that depend only on the final season ranking. This leads us to our first result:

Theorem 1. *The only secondary prize allocation policy that is a non-increasing function only of a team's final ranking $r_i(\vec{h}^m)$ and satisfies NTC for any history \vec{h}^m is a uniform lottery, which assigns equal probability to every team that does not receive π^V .*

Proof. Suppose that $y_i(\vec{h}^m)$ is chosen before the season starts, and assigns a fixed probability of getting the secondary prize to every position in the final ranking. Suppose that the probability of getting the secondary prize for team ranked $b + 1$ is greater than the probability for team ranked b . Consider a history such that before the very last game where they play each other, two teams are tied at position b and they are mathematically out of the contest for the primary prize. It is easy to see that in this scenario, NTC cannot hold. Since both teams are already eliminated from primary prize contention:

$$q_i(\vec{h}^{t-1} \cup W) - q_i(\vec{h}^{t-1} \cup L) = 0$$

However,

$$y_i(\vec{h}^{t-1} \cup L) - y_i(\vec{h}^{t-1} \cup W) \neq 0$$

If team i loses the game against team j and clinches the $b + 1$ -th ranking, then he receives a higher probability of getting the secondary prize than if he wins the game and is ranked b . The only way to



restore NTC for a lottery that is non-decreasing in rank is if the the secondary prize probabilities are equal for the b -th and $b + 1$ -th teams. Since b and $b + 1$ were chosen arbitrarily, this must hold for every position $b > k^*$, where teams in position 1 to k^* receive at least π^V . Therefore, only a uniform lottery satisfies NTC for every possible history in a season. □

This means that no matter how carefully the league chooses the weighted lottery used to assign the secondary prize, unless the weights are equal for all teams that did not win the primary prize in the season, there will sometimes be incentives for teams to tank once they have been eliminated from contention. We also briefly comment on a related incentive issue with the two proposed rules from Gold (2010) and Lenten et al. (2018). For those rules to satisfy NTC in every possible game in a season, they implicitly require that teams will always exert effort before they are eliminated from the playoffs. However, in our model, for teams that lose many games at the beginning of a season and have a high probability of elimination, NTC may be violated before the teams are actually mathematically eliminated.

In order to design a better rule, we need to define what the league’s preferred secondary prize allocation policy is. If it was not possible for teams to choose to lose games purposely, we assume the league would simply assign the secondary prize to the team with the lowest rank at the end of the season. This assumption is derived from qualitative statements made by sports leagues when draft policies have been introduced or modified; for example, the NFL directly justified the introduction of its draft in 1935 as a means of ensuring long-term financial viability of all franchises in the league and many other professional leagues followed.⁵ Without a re-distributive draft allocation system, the best new player each year was signing with the team that offered the most money. But the team that could offer the most money was generally the team that had won in previous years, reinforcing the gap between the best and the worst teams. In general, re-distributive policies enacted in sports leagues aim to maintain competitiveness of all teams over multiple seasons to increase fan engagement and revenues, see Coutinho da Silva and las Casas (2017) and Kendall and Lenten (2017). A model that incorporates fans that have a strong enough preference for year-over-year competitiveness in the league could justify a re-distributive draft as maximizing expected league revenue. However, the scope of this paper is limited to studying incentive issues, dynamics, and objectives within a single season. Where possible, we abstract from the interaction of fan preferences and long-term league revenues. In our setting the league only can influence its revenue through the secondary prize allocation. The league objective as stated below is justified as long as the leagues long-term expected revenue increases as the the rank of the team that receives the prize decreases.⁶ We also ignore other complementary policies that leagues might use to maintain competitiveness year-over-year, such as salary caps and lump sum transfers.

League Objective The league objective is to maximize the probability that it assigns the secondary prize to the team with the worst record at the end of the season, while maintaining NTC in every game.

Problem 2.

$$\max_{y \in \mathcal{Y}} \mathbb{E} \left[\sum_{i=1}^n \left(1[v^{(n)} = i] \cdot y_i(\vec{h}^m) \right) \right] \tag{2}$$

⁵<https://operations.nfl.com/the-players/the-nfl-draft/>

⁶When teams have heterogeneous abilities, the team that is the lowest rank is also the team with the lowest ability in expectation



subject to

$$\left[q_i(\vec{h}^{t-1} \cup W) - q_i(\vec{h}^{t-1} \cup L) \right] \frac{\pi^V}{\pi^D} \geq \left[y_i(\vec{h}^{t-1} \cup L) - y_i(\vec{h}^{t-1} \cup W) \right] \quad \forall i, \vec{h}^{t-1} \quad (\text{NTC})$$

$$\text{If } \alpha_i = C \forall i, y_i(\vec{h}^0) = \frac{1}{n} \forall i \quad (\text{FAIR})$$

The league wants to maximize the probability that the lowest performing team that year gets the secondary prize, while holding true the NTC condition so that no teams have an incentive to tank. Furthermore, if team abilities are assumed equal, the probability that any given team receives the secondary prize should be equal at the beginning of any given season (FAIR) and the probabilities should be dynamically consistent (DC), or equivalently, satisfy Bayes' rule.

This maximization problem is a convex program, and is theoretically solvable by backward induction. However, the optimal backward induction solution is not at all computationally feasible. For a standard NBA season the game tree is of size 2^m , where $m = 1,230$ and the decision problem requires a decision variable for each team at each node of the infeasibly large game tree. The recursive restrictions on probabilities add to the computational infeasibility. We propose a rule that is analytically feasible and satisfies NTC in any possible history while still targeting the problem (2) directly.

3.2 An Incentive-Compatible Allocation Mechanism

Theorem 1 indicates that a lottery that has a higher weight to teams with the worst records in a season will not satisfy NTC. We propose instead a weighted lottery based on teams' win-loss records for the first t^* games in a season; both the subset of games that count and the weights are determined dynamically as the season progresses. We call this mechanism the No Tanking Draft Allocation Rule (R-NTD).

Let $Pr(v^{(n)}(\vec{h}^m) = i | \vec{h}^t)$ be team i 's probability of having the lowest number of wins after all m games have been played, given the results from the first t games. The probability team i receives the secondary prize after game t in R-NTD is

$$y_i^{R-NTD}(h^t) = \begin{cases} Pr(v^{(n)}(\vec{h}^m) = i | \vec{h}^t) & \text{if } t \leq t^*(h^t) \\ Pr(v^{(n)}(\vec{h}^m) = i | \vec{h}^{t^*}) & \text{if } t > t^*(h^t) \end{cases} \quad (3)$$

which is the probability team i ends up last given the results of all games up to $t^*(\vec{h}^t)$. $t^*(\vec{h}^t)$ is the smallest $s \leq t$ such that inequality (4) holds true: given a history \vec{h}^s , NTC would be violated at time $s+1$ if the draft probabilities continued to be adjusted based on the probability that a team comes last. If the following condition does not hold true for any $s \leq t-1$, then $t^*(\vec{h}^t) = t$.

$$\begin{aligned} & \left[q_i(\vec{h}^s \cup W) - q_i(\vec{h}^s \cup L) \right] \frac{\pi^V}{\pi^D} \\ & < \left[Pr(v^{(n)}(\vec{h}^m) = i | \vec{h}^s \cup L) - Pr(v^{(n)}(\vec{h}^m) = i | \vec{h}^s \cup W) \right] \end{aligned} \quad (4)$$

Intuitively, this is how the rule works. When abilities are equal, a team's probability of receiving the secondary prize before the season begins is $\frac{1}{n}$ since each team has an equal probability of coming last. As wins and losses are recorded in each game, a team's probability of getting the secondary prize adjusts based on how their probability of coming last changes. For example, when a team loses multiple games early in the season, they increase their probability of coming last and their probability of receiving the secondary prize increases. This adjustment process permanently stops the first time NTC



would be violated for any team playing in any game, which is the first time any team might want to start tanking.

Let us now provide some intuition on why this mechanism is a good approximation to the optimum. With no information, if we didn't take into account the wins and losses from any games, then the policy that would maximize (2) in expectation when teams have equal ability is a uniform lottery. If we condition the allocation policy on more wins and losses, the expected value of the objective function increases. To illustrate this, if we ignore NTC, we can condition on the full history \vec{h}^m , which is the result of every game in the season. With a stopping time of m , R-NTD would allocate the draft pick to the lowest-ranked team with probability 1, since $Pr(v^{(n)}(\vec{h}^m) = i | \vec{h}^m) = 1$, which is the maximum possible value of the league objective. However, taking into account NTC, we end up in the middle between the no information case and the full information case; we maximize (2) in expectation, conditional on as much information as we can take into account without violating NTC.

Theorem 2. *R-NTD satisfies NTC, DC, FAIR and PROB. Moreover, the allocation rule maximizes the expected value of the league objective, conditional on the history up to time $t^*(\vec{h}^m)$.*

Proof. First, we show that the rule satisfies each of the constraints on the league objective.

- PROB is satisfied by definition, since $y_i(\vec{h}^t)$ are the probabilities that team i is ranked last in the season, which satisfy PROB.
- DC: The secondary prize probability change from time $t - 1$ to t is either:
 - $Pr(v^{(n)} = i | \vec{h}^{t-1}) = p_{it}(\alpha)Pr(v^{(n)} = i | \vec{h}^{t-1} \cup W) + (1 - p_{it}(\alpha))Pr(v^{(n)} = i | \vec{h}^{t-1} \cup L)$ if $t \leq t^*$
 - $Pr(v^{(n)} = i | \vec{h}^{t*}) = p_{it}(\alpha)Pr(v^{(n)} = i | \vec{h}^{t*}) + (1 - p_{it}(\alpha))Pr(v^{(n)} = i | \vec{h}^{t*})$ if $t > t^*$

After probabilities are fixed, DC holds since $(1 - p_{it}(\alpha)) + p_{it}(\alpha) = 1$. The first equality holds by definition; if draft probabilities do change, they correspond to an update in the conditional probability of the team ending up last in the season, and this conditional probability is DC.

- FAIR: In our model, at the start of the season if abilities are assumed equal, the prior probability that any team wins a given game is $\frac{1}{2}$. Before any games are played, each team has an equal chance of being ranked last at the end of the season, and $y_i(\vec{h}^0) = \frac{1}{n}$.
- NTC: Suppose by contradiction that NTC was violated at some time t . This is equivalent to inequality (4) being satisfied at time t . But then, R-NTD fixed the secondary prize probabilities at some time $s < t$, which implies that at time t , the secondary prize probabilities are no longer being adjusted, and the RHS of the inequality is zero. Since the LHS is always weakly positive, it cannot be that NTC was violated.

Now we turn to optimality. The statement follows after a careful examination of the objective function of the league. After t games have been played, the expected value of the league's objective is:

$$\mathbb{E} \left[\sum_{i=1}^n \left(1[v^{(n)} = i] \cdot y_i(\vec{h}^m) \right) \middle| \vec{h}^t \right]$$

Notice that this conditional expectation can be simplified to

$$\sum_{i=1}^n \left(Pr(v^{(n)}(\vec{h}^m) = i | \vec{h}^t) \cdot y_i(\vec{h}^m) \right)$$



Both $Pr(v^{(n)}(\vec{h}^m) = i|\vec{h}^t)$ and $y_i(\vec{h}^m)$ are probabilities and constrained to sum to 1. As a result, this objective is maximized when it is a sum of squares. The optimal secondary prize probability conditional on the history up to time t must have $y_i(\vec{h}^m) = Pr(v^{(n)}(\vec{h}^m) = i|\vec{h}^t)$, which is the value under R-NTD. R-NTD maximizes the expected value of the league objective conditional on any history up to when the adjustment process must stop to prevent violation of NTC. \square

We do not claim that this rule is globally optimal for the league objective. There may be other incentive compatible rules that improve the league objective; these might, for example, allow the league to change the secondary prize probabilities after $t^*(\vec{h}^m)$, or prescribe smaller changes in probabilities in the first games of the season. The way to find a global optimum is to optimize the league objective by an exhaustive backwards induction process, but this is infeasible as discussed previously. Our rule is optimal conditional on a limited history, and has theoretical guarantees that each team has the incentive to exert effort in every game. For certain histories, the game after which the rule stops taking wins and losses into account can be theoretically very early in the season. In the next section we show how the mechanism works in practice and find that in simulations and in real data, it takes a large proportion of the games in a season into account before freezing the draft lottery probabilities.

4 Examining the Rule in Practice

4.1 Simulations

We start with a simple setting of $N = 3$ teams which is the smallest number of teams such that there are incentive issues. With $N = 2$ teams, there are no incentive issues since a team that is eliminated from primary prize contention gets the secondary prize for sure. In the three team setting, if lottery weights are strictly decreasing in rank, the bottom two teams would prefer to be ranked last than second last, so they have an incentive to lose after they are eliminated from primary prize contention. Incentives to lose appear when winning a game increases the chance of a team of ending up in a middle rank that is inferior in terms of prize allocation to the lowest ranks.

We examine some of the incentive issues described more closely through simple simulations. First, we simulate a season with three teams, who play a total of 60 games. Each team plays 40 games, 20 against each opponent. Only the first place team receives π^V . The abilities of the team are assumed equal. We assume for the rest of this section that $\frac{\pi^V}{\pi^D} = 10$, so that the value of winning is at least 10 times the value of receiving the draft pick for every team. We randomly simulate a single season and show how the draft probabilities under R-NTD change for each team compared to their win records in Figure 1.

The final ranking of the season is Team 3 (23 wins), Team 2 (21), and Team 1 (16). Their respective probabilities of receiving the first draft pick according to R-NTD are 0%, 35%, and 65%. The lowest ranked team receives the draft pick with the highest probability, which is significantly higher than a uniform lottery, which would assign probabilities of 50% to Team 1 and Team 2. Examining Figure 1 shows how these probabilities are derived. At the beginning of the season, since abilities are treated as equal, each team has an equal probability of being ranked last after the season concludes, so their draft probabilities are equal. After 20 games, all three teams are tied, so draft probabilities are still close to equal. In the next 10 games, Team 1 has a long losing streak, which raises its draft probability to over 70%, as the chance they come third in the season increases. During another losing streak between 35 and 40 games, Team 1 is sufficiently behind Team 3 that they no longer have a meaningful probability of winning the primary prize. As a result, Team 1 no longer has an incentive to win and



would violate NTC if probabilities continued to be adjusted. The draft probabilities are fixed at this point going forward. Team 2 also has not performed well and has a meaningful chance of coming third at this point, so has a draft probability of 35%. Team 2 then has a lengthy winning streak at the end of the season that results in him ending the season with only 2 wins less than Team 3, but the losing streak does not affect his draft probability since it comes after the probabilities have been frozen.

Our draft allocation rule takes advantage of how incentives change over time during a season. Panel b) of Figure 1 shows the maximum amount that Team 1's draft probability can change at each game in the season without violating NTC. R-NTD adjusts the probability a team receives the draft during the initial period in the season when Team 1 still has a chance at winning the primary prize and the incentive to win is high, but game results still provide information about who will end up as the worst team after the season concludes. As the incentive to win decreases, the draft probability adjustment is eventually stopped to avoid giving Team 1 the incentive to lose intentionally.

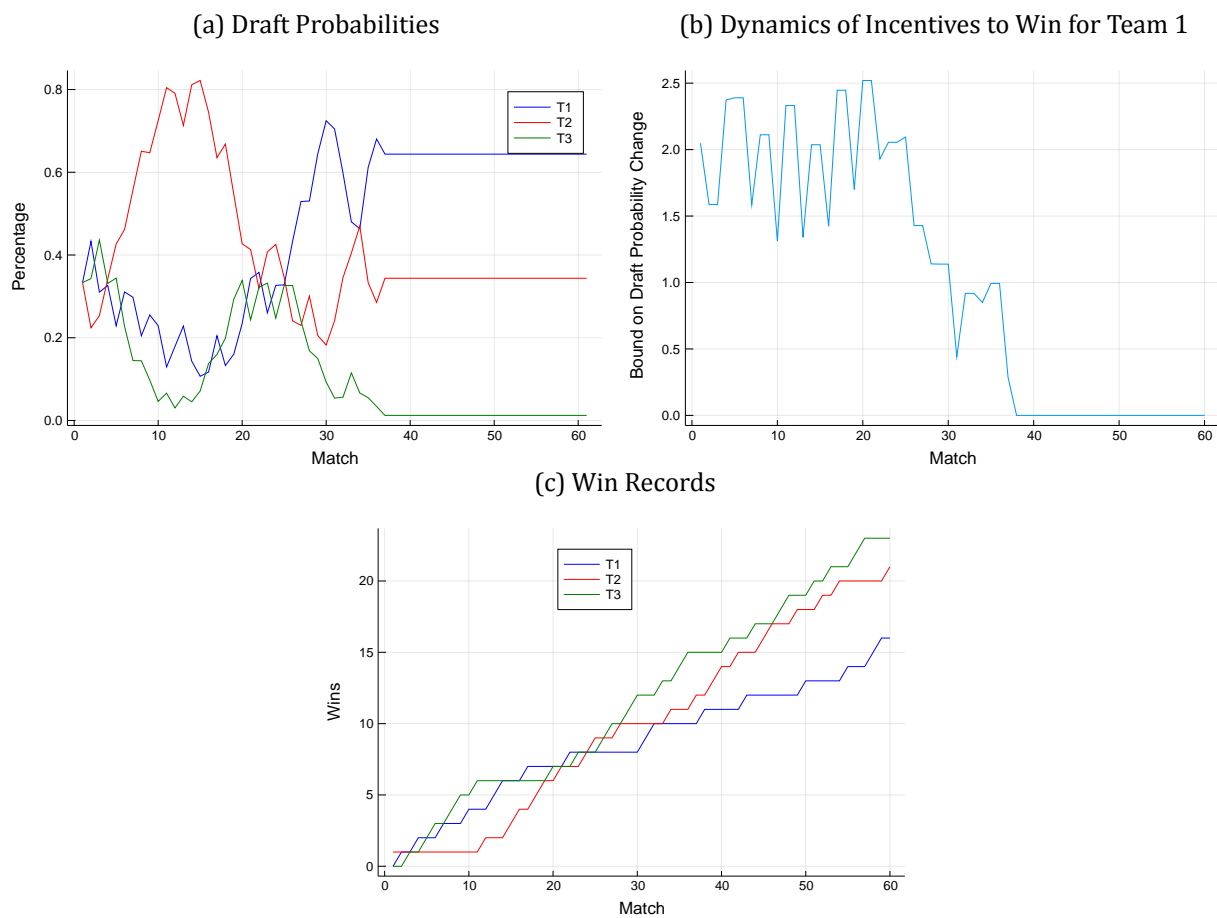


Figure 1: Allocation mechanism for simulation with $\frac{\pi^V}{\pi^D} = 10$

In Section 3, we proved that our draft allocation rule is optimal in a restricted sense; it maximizes the expected value of the league objective given game results up to the stopping time at which IC first fails. How late this stopping time is realized depends on the assumed lower bound for $\frac{\pi^V}{\pi^D}$ for any team. For example, if $\frac{\pi^V}{\pi^D} = 1$, then NTC will fail very early in the season, whenever a team has a game such that if they win, the probability of winning the primary prize increases by less than the probability of



winning the secondary prize increases if they lose. As $\frac{\pi^V}{\pi^D} \rightarrow +\infty$, then NTC will hold until the first team is mathematically eliminated from contention for the primary prize. Any team that can increase their probability of winning the primary prize by winning will have an incentive to win, no matter how small this increase is. To illustrate this, we calculate the expected gain of the league objective compared to a uniform lottery in a simple setting with $N = 3$ teams who play a total of 9 matches against each other. We enumerate all possible outcomes in a season, and calculate how the expected value of the league objective varies with $\frac{\pi^V}{\pi^D}$. This requires enumerating all possible outcomes of a season and calculating the draft rule for each possible outcome. Figure 2 shows how the expected stopping time t^* , as a percentage of all matches, and the expected improvement of the league objective over a uniform lottery depends on the ratio in prize values. If $\frac{\pi^V}{\pi^D} = 1$, then NTC binds very early in the season, so R-NTD cannot take into account more than one or two games, and does not improve meaningfully on the uniform lottery. As the ratio of the value of the primary prize compared to the secondary prize increases to 7, the expected gain increases rapidly, such that the lowest ranked player is expected to receive the secondary prize with a probability of over 75%. Illustrating the proof of Theorem 1, even as $\frac{\pi^V}{\pi^D}$ increases from 10 to 20, the expected gain does not approach 50% and the average percentage of games taken into account does not approach 100%. Even with a secondary prize of negligible value, there are always some histories such that NTC fails if every game is taken into account.

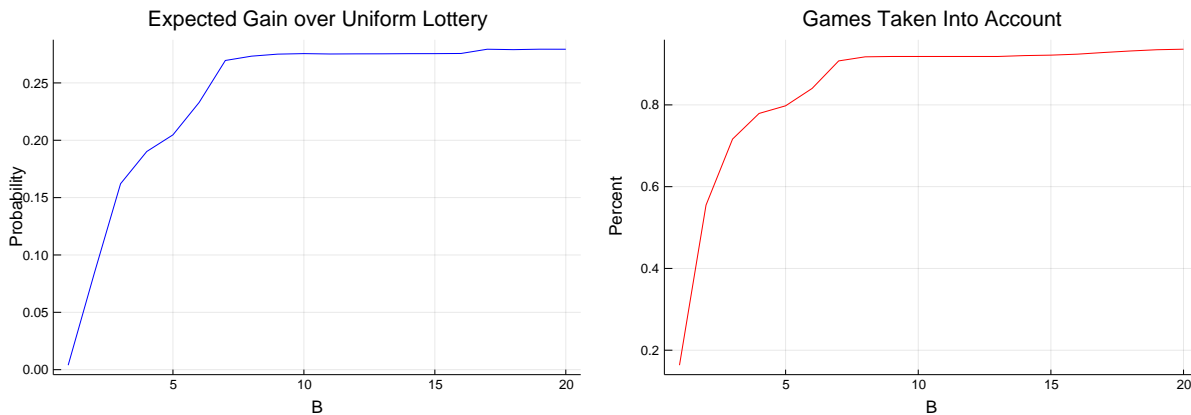


Figure 2: Gain over Uniform Lottery in Simulations Depends on $B = \frac{\pi^V}{\pi^D}$

The importance on the lower bound of $\frac{\pi^V}{\pi^D}$ also provides some insight on why the problem of ‘tanking’ might be more serious in basketball rather than baseball. In basketball, a top draft pick like LeBron James could turn a team into a title contender. In contrast, in baseball, the teams are much larger and a single player impacts a team’s prospects less. The model also explains why in some years, there is an increased incentive for teams to choose a tanking strategy. If an especially valuable draft pick is available in a certain year, the value of that draft pick might be greater even than the value of making the playoffs that year, with $\frac{\pi^D}{\pi^V} \leq 1$. In those years, R-NTD would indicate that a near to uniform lottery might be the only lottery that satisfies NTC.

4.2 Improving the NBA Lottery

From 1985 to 1989, the NBA had a uniform lottery system, where each non-playoff team received the first draft pick with equal probability. Under our model, teams under this system did not have an incentive to intentionally lose games. As a result, the league records in these seasons provide a real



world setting to examine the performance of our incentive-compatible rule and how it compares to a uniform lottery.

We calculate the draft probabilities based on our rule in 1987. In Table 1 we show the draft probabilities for our rule compared to the uniform lottery in place at the time. Teams are ordered by their final ranking. The draft probabilities in this season were not adjusted after the 266th game, which is when, based on simulations of the remainder of the season to approximate the win probabilities q_i , that we estimated that NTC would fail if lottery adjustment continued. The lowest ranked team, the L.A. Clippers, receives the pick with a 28.5% probability. This is much higher probability than in the uniform case, but the allocation mechanism does not violate NTC. The draft probabilities are mostly decreasing in the ultimate final ranking, except that the Kings have a slightly higher probability than the Spurs. This is because the probabilities were determined based on each team's probability of coming last conditional on the first 266 games of the season, when the Kings were ranked lower than the Spurs. Overall, though, in 1987 the teams that had performed badly after 266 games were largely the same teams that performed poorly over the whole season.

In this particular season, after 266 games, some of the teams that ultimately made the playoffs still had a non-zero probability of being last in the season. As a result, the mechanism described in Section 3 assigns small probability in the draft to teams that are playoff teams. We address this by adjusting the rule described in Section 3 without violating NTC; in Table 1 and Table 2 we redistribute the small probability assigned to teams that do make the playoffs uniformly to the teams that ultimately do not make the playoffs. This removes all probability from teams that make the playoff and increases the draft probability for each of the non-playoff teams by around 1%.

Rank	Team	Wins	NBA Lottery	R-NTD Lottery
23	Los Angeles Clippers	12	14%	28.5%
22	New Jersey Nets	24	14%	25.5%
21	New York Knicks	24	14%	21.0%
20	San Antonio Spurs	28	14%	9.0%
19	Sacramento Kings	29	14%	12.0%
18	Cleveland Cavaliers	31	14%	2.0%
17	Phoenix Suns	36	14%	2.0%

Table 1: Allocation Policy from 1987 Season, Adjustment Cutoff at 266th Game

We repeat the analysis for 1989, when there were two more teams in the NBA than in 1987, and summarize in Table 2. This time, the draft probabilities are adjusted until the 332nd game. The Miami Heat, the team ranked last in the season, had a poor enough start to the season that their draft probability is 45.5% when the NTC condition binds. Some other teams, though, performed better in the latter half of the season than the first half. As a result, the team ranked 24th has a lower draft probability than the team ranked 20th, since more of the Charlotte Hornets' losses occurred after draft probabilities were frozen compared to the Indiana Pacers' losses, which occurred closer to the beginning of the season.

For both the 1987 and the 1989 season, the proposed draft allocation results in a lottery that allocates the most probability to the team ranked last in the season: this probability is as high as 45% in the 1989 season. In both years, R-NTD favors the worst team more than the alternative incentive-compatible uniform lottery. The empirical results confirm our theoretical insight, which is that it is possible to design a draft lottery that favors the worst teams in a season, without providing incentives



Rank	Team	Wins	NBA Lottery	R-NTD Lottery
25	Miami Heat	15	11%	45.5%
24	Charlotte Hornets	20	11%	7.5%
23	San Antonio Spurs	21	11%	9.5%
22	Los Angeles Clippers	21	11%	3.5%
21	New Jersey Nets	26	11%	2.0%
20	Sacramento Kings	27	11%	12.0%
19	Indiana Pacers	28	11%	18.5%
18	Dallas Mavericks	38	11%	0.5%
17	Portland Trail Blazers	39	11%	0.5%

Table 2: Allocation Policy from 1989 Season, Adjustment Cutoff at 332nd Game

for teams to tank.

5 Conclusion

We show that any re-distributive draft lottery based on the final ranking will provide incentives for teams to lose intentionally, unless that lottery has equal weight on all non-playoff teams. Our main contribution is the proposed decoupling of the final ranking from the draft allocation rule. This would result in a rule that dynamically adjusts each season, taking into account teams' performance and incentives as the season evolves. It ensures that teams do not have an incentive to purposely lose games but is still optimal among all rules that take into account the results of games up to the stopping time t^* . This results in a draft allocation rule that maintains incentives to compete throughout the season while still redistributing draft picks to the lowest-ranked team with higher than uniform probability. Using simulations and empirical data on season records from the 1980s in the NBA, we show the outcome of our rule results in very sensible draft allocation policies in practice that improve on the uniform lottery without introducing incentive issues. To our knowledge, this is the first paper to examine the draft allocation problem using the theory of tournaments and incentives in a mathematical model of team decision making.

The insights gained in the more simple setting of sports leagues may also be helpful for understanding more general settings with incentive issues when there are allocations to the lowest performing individuals. The results in the sports setting suggest that in government welfare policy-making, if eligibility is determined based on an individual's cumulative performance before their application, then a re-distributive program will provide incentives to shirk. Kleven and Kopczuk (2011) indicate why having complexity in the application process can be optimal in social programs. This work indicates a certain kind of complexity that can be introduced to directly address incentive issues. For example, in college financial aid allocation, taking into account the path of an individual's bank account balances rather than the balance at time of application would reduce the incentive to spend or transfer assets immediately before college applications. Incentive issues that appear impossible to address in a static setting may be more simple to resolve in a dynamic setting. The sports setting indicates that IC constraints have more slack in earlier periods, where there is still substantial uncertainty about an individual's end period outcome, compared to later periods, where final outcomes are close to determined. When present, the dynamic incentives that appear in a contest need to be taken into account in order to implement a re-distributive prize structure.



References

- da Silva, Edson Coutinho and Alexandre Luzzi las Casas**, "Sport Fans as Consumers: An Approach to Sport Marketing," *British Journal of Marketing Studies*, 2017, 5 (4), 36–48.
- Dagaev, Dmitry and Konstantin Sonin**, "Winning by Losing: Incentive Incompatibility in Multiple Qualifiers," *Journal of Sports Economics*, 2018, 19 (8), 1122–1146.
- Gold, Adam M**, "NHL Draft Order Based on Mathematical Elimination," *Journal of Quantitative Analysis in Sports*, 2010, 6 (1).
- Kendall, Graham and Liam J.A. Lenten**, "When Sports Rules go Awry," *European Journal of Operational Research*, 2017, 257 (2), 377–394.
- Kleven, Henrik Jacobsen and Wojciech Kopczuk**, "Transfer Program Complexity and the Take-Up of Social Benefits," *American Economic Journal: Economic Policy*, February 2011, 3 (1), 54–90.
- Lazear, Edward P and Sherwin Rosen**, "Rank-order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 1981, 89 (5), 841–864.
- Lenten, Liam J.A.**, "Mitigation of Perverse Incentives in Professional Sports Leagues with Reverse-Order Drafts," *Review of Industrial Organization*, 2016, 49 (1), 25–41.
- Lenten, Liam JA, Aaron CT Smith, and Noel Boys**, "Evaluating an alternative draft pick allocation policy to reduce 'tanking' in the Australian Football League," *European Journal of Operational Research*, 2018, 267 (1), 315–320.
- Price, Joseph, Brian P. Soebbing, David Berri, and Brad R. Humphreys**, "Tournament Incentives, League Policy and NBA Team Performance Revisited," *Journal of Sports Economics*, 2010, 11 (2), 117 – 135.
- Rosen, Sherwin**, "Prizes and Incentives in Elimination Tournaments," *The American Economic Review*, 1986, 76 (4), 701 – 715.
- Rottemberg, Simon**, "The Baseball Players' Labor Market," *Journal of Political Economy*, 1956, 64 (3), 242 – 258.
- Taylor, Beck A. and Justin G. Trogdon**, "Losing to Win: Tournament Incentives in the National Basketball Association," *Journal of Labor Economics*, 2002, 20 (1), 23–41.



A Existing Draft Allocation Policies

Year	Mechanism
1966-1984	Coin flip between conference losers
1985-1989	Uniform lottery
1990-Present	Weighted lottery, modified in 1994, 2005, 2010, 2019

Table A1: History of the NBA Draft Allocation Policy

Rank	2019 -	2010 - 2018	2005 - 2009	1996 - 2004	1994	1990 - 1993
30	14.0	25.0	25.0			
29	14.0	19.9	17.8	22.5		
28	14.0	15.6	17.7	22.5		
27	12.5	11.9	11.9	15.7	25.0	16.7
26	10.5	8.8	7.6	12.0	16.4	15.2
25	9.0	6.3	7.5	8.9	16.4	13.6
24	7.5	4.3	4.3	6.4	16.3	12.1
23	6.0	2.8	2.8	4.4	9.4	10.6
22	4.5	1.7	1.7	2.9	6.6	9.1
21	3.0	1.1	1.0	1.5	4.4	7.6
20	2.0	0.8	0.9	1.4	2.7	6.1
19	1.0	0.7	0.7	0.7	1.5	4.6
18	1.5	0.6	0.6	0.6	0.8	3.0
17	0.5	0.5	0.5	0.5	0.5	1.5

Table A2: Draft Lottery Probabilities for the First Draft Pick

B Extensions

Multiple Secondary Prizes In our setting we assumed there is a single secondary prize π^D , and the objective of the league is to assign it to the worst performing team. In practice, the major sports leagues assign multiple secondary prizes, with value decreasing in rank. How does this change our mechanism? Suppose that there are J secondary prizes to be assigned to the worst j performing teams. The incentive problem looks similar, but there are more terms in the right-hand side of the NTC condition. The change in probability of getting one of the secondary prize has to be weighted by the value of that secondary prize, and the sum of those changes needs to be smaller than the change in probability of winning the primary prize weighted by the value of the primary prize. One way to set up the league problem is to add to the league objective some additional terms:

$$\max_y \mathbb{E} \left[\sum_{j=1}^J \sum_{i=1}^n \left(1[v^{(n-j)} = i] \cdot y_i^j(\vec{h}^m) \right) \right]$$

where $y_i^j(\vec{h}^m)$ is the probability that team i gets the secondary prize j at the final history \vec{h}^m . Moreover, we need to incorporate another constraint, in addition to the modified NTC, which is a version



of (PROB) for the additional secondary prizes. The problem is fairly similar, and a rule similar to R-NTD can be imagined such that it stops at the first instance of the season in which the modified NTC is violated. The simplest version would be a rule R-NTD-J such that each $y_i^j(\vec{h}^t)$ is adjusted according to the probability of ending up in position $n - j$ at the end of the season. There is no clear cut way to determine whether the expected t^* will be skewed towards the beginning or the end of the season, even when fixing a specific rule. To see this, suppose there were enough prizes such that at the end of the season every team receives one prize, whether it is the primary or a secondary prize. Then the incentive compatibility condition should get more slack, because losing a game will affect the expected final payoff less than in our benchmark case. Instead, with only two secondary prizes it's easy to imagine a situation in which the incentive compatibility binds earlier than in our benchmark, because the two secondary prizes have increased the incentives to lose disproportionately.

Alternative Models for Win Probability Estimation In the paper, we assumed abilities were known and generally treated as equal. If abilities are assumed equal, then each team, when both teams exert effort, wins with probability $\frac{1}{2}$, so every team has an equal chance of winning or losing the season before it starts. In practice, abilities are not likely to be known, especially at the beginning of the season before games are played. It is possible to use existing models to estimate win probabilities instead of the basic model described in the text. Leagues, betting websites, and others have a wide variety of sport-specific models for estimating team win probabilities based on their existing records.⁷ These may incorporate additional predictors of win probability for a team such as how many games they have played recently, whether the game is at home or away, and aggregated measures of player statistics. Our results on the optimal draft policy are robust to replacing a team's win probability $p_{it}(e_{it}, e_{jt}, \alpha)$ with any kind of sport-specific scoring model that estimates win probabilities and team specific factors α_i . The derivation of NTC depends only on the derivative of the win probability with respect to effort being positive, which must be true for any model. If NTC is satisfied by the draft rule, then effort can be ignored in the calculation of win probabilities since it is assumed equal for every team. Adjusting the draft rule and calculating NTC depends only on forecasts of a team's probability of winning different games, and their resulting probability of making the playoffs or coming last in the league. In the paper, we use the simple general model presented in Section 3 for such forecasts, but note that it is feasible to replace the win probability calculation with a variety of alternative situation-specific models.

⁷For example, the website FiveThirtyEight uses an adapted version of the ELO system, originally used to score chess players, to analyze team abilities in the NBA (<https://fivethirtyeight.com/features/how-we-calculate-nba-elo-ratings/>).