An Incentive-Compatible Draft Allocation Mechanism*

Martino Banchio† Evan Munro‡

April 20, 2020

Abstract

In the major U.S. sports leagues, the most valuable newly eligible player is allocated through a draft lottery that favors the lowest-ranked teams after each season. This causes some teams to exert low effort later in the season to secure a draft pick with higher probability. We derive a basic theoretical model of team decision making. We prove that under the current mechanism, it is impossible to both favor the lowest-ranked teams and eliminate all incentives for teams to lose intentionally. We design an incentive-compatible lottery that dynamically adjusts as the season progresses, favors the worst teams, and is optimal in a restricted class of mechanisms. We show that this policy performs well compared to alternative lotteries in simulations and using data from the National Basketball Association.

Keywords: Incentive-Compatible Mechanism, Tournaments, Competition Policy
JEL Codes: D82, Z28

*We thank Mohammad Akbarpour, László Csató, Edward Lazear, Michael Ostrovsky, Brad Ross, Andrzej Skrzypacz and Bob Wilson for helpful comments and discussions, as well as Daryl Morey, Mike Zarren and participants at the MIT Sloan Sports Analytics Conference 2020 and the Stanford Economic Theory Workshop. Code for the analysis in the paper is available at http://github.com/evanmunro/draft-policy.

†Graduate School of Business, Stanford University. Email: mbanchio@stanford.edu
‡Graduate School of Business, Stanford University. Email: munro@stanford.edu

1 Introduction

A central problem in economics is understanding whether it is possible to achieve some degree of redistribution without affecting individuals' incentives to exert effort. Results from Mirrlees (1971) and Epple and Romer (1991) indicate that redistribution must be limited by the distortions it places on incentives. Redistributive policies, however, are essential to reduce persistence in inequality, as described by Mookherjee and Ray (2003). In this paper, we examine the allocation of new players to teams in U.S. sports leagues, and its relation to teams' incentives to exert effort. The precise and tractable insights gained in this specific setting contribute to the understanding of more general dynamic settings with redistribution and relative performance evaluation.

The conventional wisdom in the sports industry is that there is significant path dependence in a franchise's performance in the absence of any centralized intervention. Teams who perform well collect more revenue and can recruit better players. Over the long-term, this hurts the competitiveness of the league and, ultimately, league revenues. As a result, sports leagues have instituted policies that support the worst-performing teams in various ways to increase fan engagement in the long-term. Perhaps the most important is the draft lottery, which determines the allocation of newly eligible players. This lottery, based on final rankings at the end of the season, favors the lowest-ranked teams but introduces perverse incentives for teams to lose intentionally, known colloquially as 'tanking'. Tanking is frequently and passionately discussed in the media and has been documented empirically in Taylor and Trogdon (2002) for the NBA and Fornwagner (2019) for the NHL. As a result, the optimal design of the allocation mechanism for new players is an important and open problem for the major sports leagues, which collectively produce over 30 billion dollars of revenue.

In order to design an improved draft mechanism, we formalize the league objective, which is to maximize the probability that the lowest ranked team receives the first draft pick, constrained by a no-tanking condition for each team in every game. This objective 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 poor-performing franchises in the league and many other professional leagues followed.[1] We set up a

[1] https://operations.nfl.com/the-players/the-nfl-draft/

simple model of team decision making, which intentionally abstracts from the complexity of real world coaching strategies in order to analyze only the effect of the draft policy design on incentives for teams to exert effort. We then derive two main contributions, which highlight the dynamic nature of incentives to win in a sports season. First, we show that any lottery based on final rankings that favors the lowest-ranked teams provides incentives for a team to lose in some possible history of a season. This negative results suggests we should look at mechanisms that adapt the draft probabilities over time. Second, we show that the globally optimal mechanism is the solution to a linear program with an infeasible number of constraints. So, we propose a rule that is optimal for a large subset of incentive compatible rules and is feasible to implement. Our mechanism updates draft lottery probabilities after each game so that they match the probability each team will be ranked last after all the games have been played, conditional on results from games played so far. This adjustment process ends as soon as incentive-compatibility would be violated. We show this process is approximately optimal since it maximizes the expected league objective, conditional on a subset of games in the season.

In the rest of this section, we provide a brief literature review and necessary background on the NBA draft. Section 2 defines the model of team decision making that we use to prove that the current allocation rule is not incentive-compatible and to introduce our draft allocation mechanism. In Section 3, we show that our rule performs nearly as well as the optimal solution in a small simulation, and we demonstrate that it performs well on real data from the NBA in the 1980s.

Literature Review. Our paper is closely related to the existing literature on tournament theory. Lazear and Rosen (1981) design a prize allocation system that incentivizes optimal effort for workers with a single choice of effort, where a worker's ranking in output, rather than their actual output, determines their prize allocation. The prizes decrease as rankings decrease. Rosen (1986) investigates sequential elimination contests, where the optimal prize structure takes into account how incentives change as a team progresses in the tournament. As in Rosen (1986), we have a sequential series of games with repeated choices of effort and incentives that change over time, although in a sports season there is no elimination. The first important difference is that prizes are not increasing in rank. In our setting we assume that there are two prizes, a primary prize that is assigned to the top k^* teams that make the playoffs, and a secondary prize, the first draft pick, that is assigned with the highest probability to the lowest ranked team. Without a correct allocation policy that decouples prize allocation from the final ranking of the teams, this prize structure will explicitly incentivize teams to intentionally lose games. Dagaev and Sonin (2018) model a different failure of incentive compatibility in repeated tournaments where a team may want to obtain a worse seed in the first round to obtain a better draw in the second round.

It is possible to describe the optimal draft mechanism as a solution to a simple optimization problem. However, it is impossible to compute that solution for settings of a reasonable size, so we propose an approximately optimal but computationally feasible mechanism instead. As described in Akbarpour et al. (2020), there are many other mechanism design settings where approximate solutions are necessary due to computational limits, including the optimal packing of cargo (Dantzig, 1957), radio spectrum allocation with interference constraints (Leyton-Brown et al., 2017) and computing the efficient allocation for combinatorial auctions (Lehmann et al., 2002).

There is a larger body of literature in sports economics which has proposed feasible allocation mechanisms in the absence of a theoretical model that explicitly describes the league objective and team decision making. Gold (2010) proposes allocating the top pick to the team with the highest 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. Under our framework neither rule is fully incentive compatible. Furthermore, Gold's rule would assign zero probability in the draft to a team that loses every game in the season, so it does not target our assumed league objective. The closest to our work is the concurrent paper by Kazachkov and Vardi (2020). They also set up a theoretical model of a sports season, and suggest running a draft based on reverse rankings at a stopping time earlier in the season. Rather than designing a fully incentive-compatible mechanism, they illustrate computationally the tradeoff between the prevalence of tanking and how much the draft benefits the lowest ranked team in expectation.

Background. Draft picks are considered more valuable in basketball compared to other sports, since teams are smaller and a single star player is more likely to have a major influence on a team's overall performance. As a result, we use NBA data for the application in this paper, but the analysis is easily adapted to other leagues. The NBA regular season begins in the last week of

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 1: Draft Lottery Probabilities for the First Draft Pick, NBA

October and ends in the middle of April. There are thirty teams competing, divided in two conferences. Each team plays eighty-two games in a single regular season, and each team plays every other team at least twice during the regular season. 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, in a prescribed ordering, the player they value the most out of the remaining pool of eligible draftees. In the NBA, the first four picks in the ordering are allocated by lottery. The remaining picks are based on reverse rank. In this paper, we focus on the problem of allocating the first pick, which is the most valuable, but provide some comments in Section 2 on extending the model to the allocation of multiple picks. Before 1985, the first draft pick was allocated based on a coin flip between the two conference losers. In response to accusations that teams were tanking in response to this system, the league switched to a uniform lottery over all non-playoff teams from 1985 to 1989. Due to concerns that the uniform lottery did not favor the worst teams, the league switched to a weighted lottery system starting in 1990. Table 1 describes the draft lotteries by rank from 1990-2019. The lottery has been changed frequently in response to complaints about tanking or competitive balance in the league, so the probability that the lowest ranked team receives the pick has ranged from 14% to 25% . The next section of the paper provides some insight into why the league has changed the system so often without finding a lottery that is satisfactory.

2 Designing a Draft Allocation Mechanism

2.1 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 $i \in I = \{1, \ldots, n\}$ 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 $t \in \{1, \ldots, m\}$. A season is a pair of sequences $S = \left\{\{S_{t1}\}_{t=1}^{m}, \{S_{t2}\}_{t=1}^{m}\right\}$ that identifies the two contestants in every game. A history is defined recursively as follows:

- $\vec{h}^0 = \emptyset$
- $\vec{h}^t = \vec{h}^{t-1} \cup i$ if team i wins game t

Feasible histories, including partial histories, are any t-length vectors $\vec{h}^t$ such that $t \leq m$ and $\vec{h}^t_s \in \{S_{s1}, S_{s2}\}$ $\forall s$. Let the set of feasible histories be $\mathcal{H}$. For some history $\vec{h}^t$, the total number of wins for team i is:

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

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 :

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

The ranking of a team i is defined as follows:

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

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.[2] In the NBA setting, π^V represents the minimum expected additional revenue from the playoffs that a team receives, in terms of media exposure and ticket sales. There is also a single secondary prize D. In the NBA setting it represents the first draft pick and has a 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 $\pi^D \leq \pi^V$.

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:

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

  where Φ is the standard normal cumulative distribution function.

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:

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

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 general, a draft allocation mechanism will be a rule

$$y \colon \mathcal{H} \to [0, 1]^n$$

[2] We abstract from the two conference structure in the NBA.

such that $y_i(\vec{h}^t)$ is the probability that team i will receive the secondary prize given the history up to game t. The mechanism is restricted in the following ways. Since $y_i(\vec{h}^t)$ 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 $t-1$:

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

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

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

A draft allocation mechanism satisfying DC and PROB is feasible, and the space of feasible mechanisms is $\mathcal{Y}$.

Each team makes a single strategic choice in each game, which is how much effort to exert. The choice of effort by team i in game t will be denoted by $e_{it} \in [0,1]$. In order to derive the mechanisms that always incentivize effort, we first explicitly define the teams' 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 $j \in \{S_{t1}, S_{t2}\}$:

Optimization Problem 1.

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

Maximizing the team objective, we derive a necessary condition for team i to exert maximum effort in game t:

No-Tanking Condition.

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

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 increase in the probability of receiving the secondary prize when losing compared to winning a certain game is less than the decrease in the probability of receiving the primary prize, scaled by the ratio of the prizes' values, then the team will exert maximum effort. 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.

A rule that is incentive-compatible is defined as any $y \in \mathcal{Y}$ such that NTC is satisfied for all teams $i \in I$ for all possible histories $h \in \mathcal{H}$. 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. *When $m > 2$, $n > 2$, and $k^* < n - 1$, the only secondary prize allocation mechanism y that*

1. *is a function only of a team's final ranking $r_i(\vec{h}^m)$,*
2. *for every $\vec{h}^m$ there exists a $k > k^*$ such that $y_{v^{(k)}(\vec{h}^m)}(\vec{h}^m) \leq y_{v^{(k+1)}(\vec{h}^m)}(\vec{h}^m)$,*
3. *satisfies NTC at every history $\vec{h}^t \in \mathcal{H}$,*

is a uniform lottery, which assigns equal probability to every team i with rank $r_i(\vec{h}^m) > k^*$.

Proof. Let y be a feasible mechanism satisfying conditions $1 - 3$. Let $i = S_{m1}$ and $j = S_{m2}$ denote the two teams playing in the last game of the season. There exists some partial history $\vec{h}_u^{m-1}$ such that

- team i and j are tied at position k, before the final game m is played.

$$r_i(\vec{h}_u^{m-1}) = r_j(\vec{h}_u^{m-1}) = k$$

By condition 2, this implies if team i loses game m, then its lottery probability is higher than if it wins game m:

$$y_i(\vec{h}_u^{m-1} \cup L) - y_i(\vec{h}_u^{m-1} \cup W) \geq 0$$

- Team i and j are both eliminated from playoff contention before this final game m is played. Since $k > k^*$, $k^* < n-1$, and $m > 2$ there is always some history where this is true.

$$q_i(\vec{h}_u^{m-1}) = q_j(\vec{h}_u^{m-1}) = 0$$

Then, we have for team i:

$$\begin{aligned} 0 =& \Big(q_i(\vec{h}_u^{m-1} \cup W) - q_i(\vec{h}_u^{m-1} \cup L)\Big)\frac{\pi_V}{\pi_D} \geq \\ & y_i(\vec{h}_u^{m-1} \cup L) - y_i(\vec{h}_u^{m-1} \cup W) \geq 0 \end{aligned}$$

where the first inequality is NTC and the last inequality is condition 2. This chain of inequalities is satisfied only with equality, making the only rule that satisfies conditions $1-3$ the uniform lottery over non-playoff teams. □

This means that no matter how carefully the league chooses the weighted lottery used to assign the secondary prize, if the lottery weights favor lower ranked teams and are not equal for all non-playoff teams, then there will always be incentives for teams to tank after they are eliminated from the playoffs in certain histories. A simple deduction from this theorem is that any incentive compatible draft allocation mechanism cannot take into account losses occurred after a team has been mathematically eliminated from playoff contention, as noted by Lenten et al. (2018).

Corollary 1. *If $q_i(\vec{h}^t) = 0$, any incentive-compatible rule must have $y_i(\vec{h}^s) = y_i(\vec{h}^{s+1})$ for all $s \geq t$ such that $i \in \{S_{s1}, S_{s2}\}$.*

In order to design a better rule, we need to define what the league's preferred secondary prize allocation policy is. We recognize that the true league objective is complex and may take into account incentives for teams to engage in a variety of unwanted strategies. We focus here only on the strategy of tanking within a single season, which is induced by the draft rule. Based on statements from the league about the purpose of the draft, we take the league objective with respect to the draft as given.

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.

Optimization Problem 2.

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

subject to

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

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

The league wants to maximize the probability that the lowest performing team gets the secondary prize, without providing incentives for teams to tank. If it was not possible for teams to exert low effort, the league would simply assign the secondary prize to the team with the lowest rank at the end of the season. 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).

This maximization problem is a linear program, but solving the global optimum is infeasible, since all possible histories of a season need to be enumerated to specify the NTC constraints and to calculate the expected value of the league objective over all possible histories. For a standard NBA season the game tree is of size 2^m, where $m=1,230$. This requires specifying more constraints than there are number of atoms in the known universe, so calculating the optimal solution via linear programming is not computationally feasible in a real-world setting. We propose a rule that is analytically feasible and satisfies NTC in any possible history while still targeting the objective (2) directly. In Section 3 we show how well our proposed rule approximates the optimal solution in a small simulation where specifying all constraints in the linear program is feasible.

2.2 An Incentive-Compatible Allocation Mechanism

Theorem 1 indicates that a lottery based on final rankings that favors the worst teams will not satisfy NTC. We propose instead a weighted lottery based on teams' win-loss records for the first $t^*(\vec{h}^m)$ games in a season. Both the subset of games that count and the lottery weights are determined dynamically as the season progresses. We call this mechanism the No-Tanking Draft 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. We define the probability that team i receives the secondary prize after game t in R-NTD as

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

which is the probability team i ends up last given the results of all games up to $t^*(\vec{h}^t)$. We next define the stopping time $t^*(\vec{h}^t)$. Consider the following quantity:

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

The quantity $B(i,s)$, derived from rearranging NTC, is the difference in utility for team i from losing game $s+1$ compared to winning game $s+1$. We cannot continue to adjust the draft probabilities according to the updated conditional probabilities that teams will be ranked last in the season when this benefit is positive. Let

$$t^*(\vec{h}^t) = \min\left\{s < t \colon \exists i \in I \text{ such that } B(i,s) > 0\right\}$$

The time $t^*(\vec{h}^t)$ is the smallest $s < t$ such that given a history $\vec{h}^s$, NTC would be violated at time $s+1$ if the draft probabilities continued to be updated based on the conditional probabilities that teams end up last in the season. When there is no such $s < t$, we let $t^*(\vec{h}^t) = t$.

The rule dynamically adjusts the lottery odds as the season progresses. 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 their updated conditional probability of ending up ranked last. 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 before NTC would be violated.

Theorem 2. *R-NTD is feasible and satisfies NTC and FAIR. 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 conditional probabilities.
- DC: The secondary prize probability change from time $t-1$ to t is either:
  - If $t > t^*(\vec{h}^m)$, there is no change, so DC holds
  - If $t \leq t^*(\vec{h}^m)$, then the draft probabilities correspond to updates in a conditional probability, so DC holds:

$$Pr(v^{(n)} = i|\vec{h}^{t-1}) = p_{it}(\alpha)Pr(v^{(n)} = i|\vec{h}^{t-1} \cup W) + (1-p_{it}(\alpha))Pr(v^{(n)} = i|\vec{h}^{t-1} \cup L)$$

- FAIR: If abilities are assumed equal, the model of win probabilities in (1) indicates the probability that any team wins a given game is $\frac{1}{2}$, so 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 $B(i,t)>0$. But then, R-NTD fixed the secondary prize probabilities at some time $s < t$. This implies the RHS of NTC is zero, and 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 s games have been played, the expected value of the league's objective is:

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

Notice that this conditional expectation can be simplified to

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

Both $Pr(v^{(n)}(\vec{h}^m) = i|\vec{h}^s)$ and $y_i(\vec{h}^s)$ 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 s must have $y_i(\vec{h}^s) = Pr(v^{(n)}(\vec{h}^m) = i|\vec{h}^s)$, which is the value under R-NTD for all $s \leq t^*(\vec{h}^m)$. R-NTD maximizes the expected value of the league objective conditional on any history up to the stopping time. □

We 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 over all teams. If we condition the allocation mechanism 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. R-NTD decouples the incentive constraints from the targeting of the objective function. The objective function is addressed by matching the draft probabilities to the optimal probabilities up to the stopping time, while incentives are taken care of separately by the stopping time t^*. This separation is the key to the feasible computation of the method, as incentives only affect the stopping time and not the adjustment of the weights.

The rule is optimal conditional on a limited history, and has theoretical guarantees that each team has the incentive to exert effort in every game. One drawback is that for certain histories it is possible that the stopping time is 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, under reasonable assumptions, it takes a significant proportion of the games in a season into account before freezing the draft lottery probabilities.

2.3 Practical Extensions

Our model is designed in a simplified environment, where team abilities are fixed, there is a single draft pick to be allocated, and team incentives are based only on prizes in a single season. We briefly comment on how the model might be extended to address environments where some of these simplifying assumptions are relaxed.

Incentives Across Seasons R-NTD solves the problem of teams choosing to exert low effort partway through a season, once they are near-eliminated from the playoffs and no longer have a good incentive to win within that season. It does not directly address a related issue, known as a rebuilding strategy, where a team makes a decision to intentionally lower their ability across multiple seasons by trading away good players. These teams intentionally limit their win probability across multiple seasons in order to accumulate a package of valuable draft picks. Our paper considers this issue to be secondary to the within-season effort problem. However, R-NTD could increase the incentives to engage in a rebuilding strategy, since in expectation it increases the probability the lowest ranked team receives the first draft pick compared to the existing system. To address this, it would be sufficient to modify the constraint on the league objective, taking into account cross-season incentives to intentionally lower a team's ability. Satisfying this constraint may require a modification of R-NTD, such as excluding teams that win the first pick from future draft lotteries for a certain number of years.

Multiple Secondary Prizes Our rule solved the problem of assigning the best draft pick, with value at most π^D, to the worst performing team. In practice a draft involves multiple picks. We briefly comment on how this changes the optimal mechanism. Suppose that there are J draft picks, and we directly extend the league objective so that the league would like to assign the j-th pick to the j-th worst team. A natural extension of our mechanism would hold a separate lottery for each pick. R-NTD would adjust J lotteries after each game rather than a single lottery. To determine the stopping time, NTC would be extended to take into account the total expected change in value of multiple secondary prizes when a team loses compared to wins.

Alternative Models for Win Probability Estimation In the empirical section of the paper, we use a general and simple structural model of win probabilities to calculate NTC and R-NTD. For the simulations in Section 3, we treat team abilities as fixed and equal, so that when both teams exert effort, each team wins with probability $\frac{1}{2}$. However, leagues, analysts, and betting websites have developed complex sports-specific models for forecasting the results in a season. These may adjust for unexpected events such as player injuries or trades, and incorporate additional predictors of win probability for a team such as how many games they have played recently and aggregated measures of player statistics. For example, the website FiveThirtyEight uses an adapted version of the ELO system, originally used to score chess players, to estimate win probabilities for teams in the NBA. The derivation of NTC requires only that the derivative of the win probability for a single game with respect to effort is positive, which must be true for any reasonable model. This and the absence of an explicit cost of effort produce a corner solution, where teams either exert full effort or no effort at all. Adjusting the draft rule and calculating NTC depends only on forecasts of a team's probability of winning the remaining games in the season. As long as the win probability forecasts from an alternative model can be feasibly calculated and are dynamically consistent, then we can replace the predictions of our simple model in (1) with a model tailored to a specific sport.

3 Examining the Rule in Practice

3.1 Simulations

We start with a simple setting of $N = 3$ teams, where teams have equal ability and only the top team receives the primary prize, 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 receives the secondary prize. 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 a dynamically-determined stopping time. How late this stopping time is realized depends on the assumed lower bound on $\frac{\pi^V}{\pi^D}$ for any team.

Under $N = 3$ teams in a season where there are $m = 15$ games, it is feasible, though computationally intensive, to enumerate all possible 2^{15} outcomes of the season and calculate the globally optimal incentive-compatible mechanism as the solution to the linear program in Optimization Problem 2. We calculate the league objective, which is the expected probability that the lowest-ranked team receives the first draft pick over all histories, for four different incentive-compatible mechanisms that also satisfy FAIR:

1. Ex-post uniform lottery: a uniform lottery over all teams that do not make the playoffs each year.
2. Ex-ante uniform lottery: a uniform lottery over all teams.
3. The globally optimal rule
4. R-NTD, the approximately optimal rule

Mechanisms 2-4 are always incentive compatible. Mechanism 1 is incentive compatible as long as $\frac{\pi^V}{\pi^D} \geq 1$. The forecasts of the season results require to compute NTC for R-NTD and the optimal rule are estimated by simulating the rest of the the season thousands of times.

Figure 1a shows how the value of the league objective varies with $\frac{\pi^V}{\pi^D}$. The ex-ante uniform rule gives the draft pick with expected probability 33% to the lowest ranked team. The ex-post uniform rule gives the draft pick with expected probability 50% to the lowest ranked team, but is only incentive compatible when $\frac{\pi^V}{\pi^D} \geq 1$. For low values of $\frac{\pi^V}{\pi^D}$, neither the globally optimal rule nor R-NTD does much better than the uniform rules. The incentive constraints bind early in the season, so the draft probabilities cannot be adjusted. However, as $\frac{\pi^V}{\pi^D}$ increases from 2 to 10, the value of the league objective increases rapidly for both R-NTD and the globally optimal rule, although the globally optimal rule increases faster. The rules take advantage of the fact that incentive constraints are more slack earlier in the season. Both rules converge to a constant value as $\frac{\pi^V}{\pi^D}$ increases beyond 10. The maximum possible value of the league objective is 1. As expected from Corollary 1, neither the globally nor the approximately optimal rule can ever give the draft with expected probability 1 to the lowest ranked team, no matter how large $\frac{\pi^V}{\pi^D}$ is. The globally optimal rule converges to a value of approximately 82% while R-NTD converges to a value of approximately 76%, so approximates well the global optimum. The performance of R-NTD is computationally feasible as m increases, whereas the globally optimal rule is not. Both are over 25 percentage points better than the ex-post uniform rule that was used in the NBA in the 1980s.

Figure 1b provides some intuition on what drives the good performance of the approximately optimal and globally optimal rule, and the differences between them. Figure 1b shows the average absolute value of the change in draft probability between match m and $m-1$, averaged over all histories and teams. This can be considered the average weight that each mechanism places on each game. Both rules place more weight on earlier games in the season, where incentive compatibility holds in more histories, since teams still are in contention for the primary prize. However, the globally optimal rule places more weight on games towards the end of the season compared to R-NTD. Our rule freezes the draft probabilities for all teams once incentives for a single team have been violated. The results for the globally optimal rule show there is some benefit in some histories of adjusting the draft probabilities for some teams after incentives have been violated for a single team. However, this benefit is quite small, since the two mechanisms converge to within 6 percentage points of each other as $\frac{\pi^V}{\pi^D}$ increases.

From these simulations we have indications that, for a reasonable assumption on $\frac{\pi^V}{\pi^D}$, the approximately optimal solution given by R-NTD is close to the globally optimal solution. We now turn to data on the NBA to examine how R-NTD performs on results from past sports seasons.

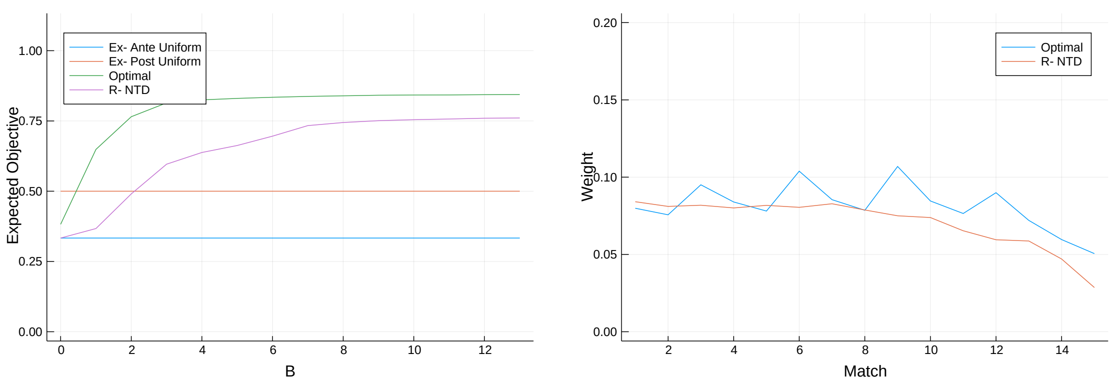

(a) League Objective, $B = \frac{\pi^V}{\pi^D}$

(b) Average Draft Probability Change by Match

Figure 1: Comparing Alternative Incentive-Compatible Rules in Simulations

3.2 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 and final rankings would not be expected to change under R-NTD. In more recent years, where the NBA lottery incentivizes tanking, team exertion of effort would have changed under our rule and expected final rankings would be different. As a result, we use records from 1985 to 1989, rather than more recent years, to examine the performance of our incentive-compatible rule on real data.

First, we calculate the draft probabilities based on our incentive-compatible rule for every season from 1985-1989 and examine the results from 1987 more closely. For these five years, we assign the draft to the lowest ranked team with a 38.6% probability on average. This is a large increase over the incentive-compatible ex-post uniform lottery, which gives the lowest ranked team a 14% probability from 1985-1988 and an 11% probability in 1989, when the league was expanded. The shortest stopping time is 353 games in 1989 and the longest is 527 games in 1985. On average, the draft probabilities are adjusted until 45% of the season has occurred.

Rank	Team	Wins	NBA Lottery	R-NTD Lottery
23	Los Angeles Clippers	12	14%	59.0%
22	New Jersey Nets	24	14%	5.1%
21	New York Knicks	24	14%	6.5%
20	San Antonio Spurs	28	14%	17.4%
19	Sacramento Kings	29	14%	7.4%
18	Cleveland Cavaliers	31	14%	0.9%
17	Phoenix Suns	36	14%	0.7%
< 17	Playoff Teams	N/A	0%	3.0%

Table 2: Allocation Policy from 1987 Season, Adjustment Cutoff at 371st Game

Table 2 shows the final draft probabilities for R-NTD in 1987 compared to the uniform lottery in place at the time. Teams are ordered by their inverse final ranking. The lowest ranked team, the L.A. Clippers, receives the pick with a 59% probability. 88% of the probability of receiving the first draft pick is concentrated among the four teams at the bottom of the final ranking. It is worth noting that probabilities are not necessarily always increasing as rank decreases; for example, the Spurs have a higher draft probability than the lower-ranked Knicks. This is because at the stopping time, the Spurs had a worse ranking than the Knicks, but improved their record by the end of the season. Also, there is a cumulative total of 3% probability assigned to the 16 teams that made the playoffs, but for individual teams this probability is negligible, and is only greater than than 0.2% for the Denver Nuggets, who were ranked 16th and were the last team to qualify for the playoffs.

Figure 2 shows how the probabilities were adjusted over the first 371 games for the Clippers, the Nets, and the Knicks, who were the worst 3 teams at the end of the season. Up until game 300, each had a similar win record, so each had a roughly equal probability of ending up last in the season. However, after game 300, the Clippers begin a lengthy losing streak; at first, the draft odds continue to adjust based on the rapidly increasing probability that the Clippers end up ranked last in the season. Early in the season, there are still incentives for the Clippers to win since there is still a chance they make the playoffs. After enough games have passed, our model indicates that the Clippers are increasingly certain that they will not make the playoffs, and incentives to lose increase enough that the draft probabilities are frozen after game 371.

From 1985-1989, R-NTD assigns the first draft pick with an average probability that is over 20 percentage points higher than the uniform lottery, while maintaining incentive compatibility. There are significant practical benefits to implementing a draft mechanism that is dependent not only on the cumulative total of wins and losses, but also when those wins and losses occur.

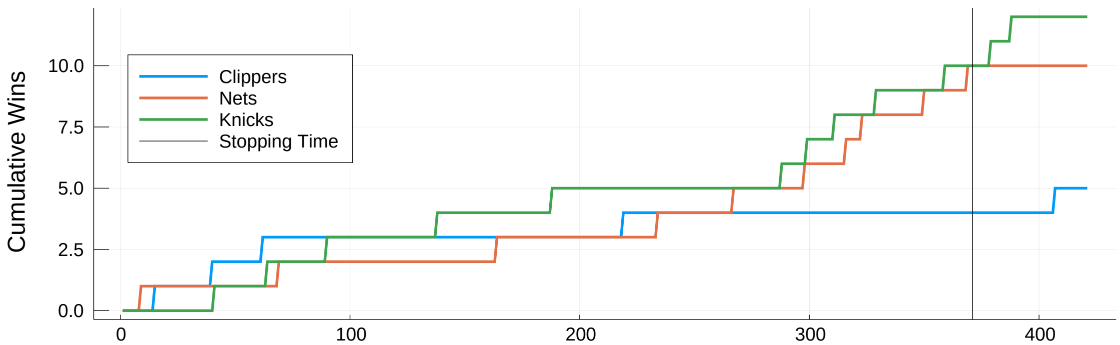

(a) Win Records

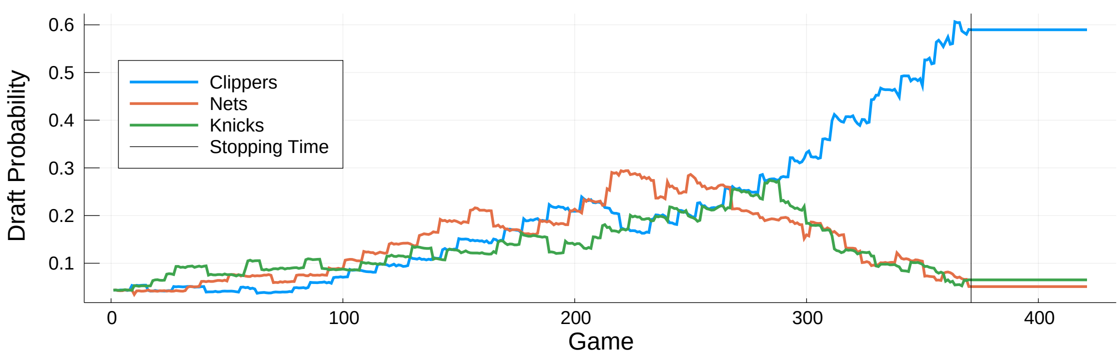

(b) Draft Probabilities

Figure 2: Dynamics of R-NTD for the 1987 season with $\frac{\pi^V}{\pi^D} = 10$

4 Conclusion

There has been substantial controversy over perverse incentives created by the draft lottery system in U.S. sports leagues. The NBA, for example, has repeatedly changed the lottery used to assign the draft pick without finding a satisfactory system that addresses the incentive issue while still favoring the worst teams in the league. Using a mechanism design approach, we identify what distortions are induced on incentives by the current policy, and prove that no incentive-compatible lottery based on final rankings can favor the worst teams.

We propose an alternative incentive-compatible lottery design that still subsidizes poor-performing teams and is approximately optimal for the league objective. Our rule directly targets the league objective by adjusting draft probabilities as wins and losses occur based on the updated conditional probability that each team ends up ranked last after the season concludes, up until a stopping time. The stopping time is dynamically determined in each season to ensure that the rule is incentive compatible in every possible history of a season. We show in a small simulation that our rule performs nearly as well as the globally optimal rule, which is the solution to a linear program that is not feasible in a larger, more realistic setting. We also show that on historical data from the NBA from 1985-1989, R-NTD significantly outperforms the league lottery in place at the time without introducing perverse incentives. We show that the dynamic nature of incentives in tournaments are important; incentive-compatibility constraints are more slack at the beginning of the season when uncertainty about final outcomes is higher, compared to the end of the season when final outcomes are close to determined.

The draft allocation problem is an example of how incentive issues that appear impossible to address in a static setting can be resolved when dynamics are considered. Our insights may be relevant to other settings, such as financial aid or welfare allocation, where the dynamic nature of incentives should be taken into account in the optimal design of incentive-compatible and re-distributive policies.

References

Akbarpour, Mohammad, Scott Duke Kominers, Shengwu Li, and Paul R Milgrom, "Investment Incentives in Near-Optimal Mechanisms," 2020. Working Paper.

Dagaev, Dmitry and Konstantin Sonin, "Winning by Losing: Incentive Incompatibility in Multiple Qualifiers," *Journal of Sports Economics*, 2018, *19* (8), 1122–1146.

Dantzig, George B, "Discrete-variable extremum problems," *Operations research*, 1957, *5* (2), 266–288.

Epple, Dennis and Thomas Romer, "Mobility and Redistribution," *Journal of Political Economy*, 1991, *99* (4), 828–858.

Fornwagner, Helena, "Incentives to lose revisited: The NHL and its tournament incentives," *Journal of Economic Psychology*, 2019, *75.*

Gold, Adam M, "NHL Draft Order Based on Mathematical Elimination," *Journal of Quantitative Analysis in Sports*, 2010, *6* (1).

Kazachkov, Alexander M. and Shai Vardi, "On Tanking and Competitive Balance: Reconciling Conflicting Incentives," 2020. Working paper.

Lazear, Edward P and Sherwin Rosen, "Rank-order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 1981, *89* (5), 841–864.

Lehmann, Daniel, Liadan Ita Oćallaghan, and Yoav Shoham, "Truth revelation in approximately efficient combinatorial auctions," *Journal of the ACM (JACM)*, 2002, *49* (5), 577–602.

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.

Leyton-Brown, Kevin, Paul Milgrom, and Ilya Segal, "Economics and computer science of a radio spectrum reallocation," *Proceedings of the National Academy of Sciences*, 2017, *114* (28), 7202–7209.

Mirrlees, J. A., "An Exploration in the Theory of Optimum Income Taxation," *The Review of Economic Studies*, 1971, *38* (2), 175–208.

Mookherjee, Dilip and Debraj Ray, "Persistent Inequality," *The Review of Economic Studies*, 2003, *70* (2), 369–393.

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.