Targeting in Tournaments with Dynamic Incentives *

Martino Banchio Evan Munro[†]

April 12, 2021

Abstract

We study the problem of an institution that targets aid to its needy members, where need is determined by observed performance. This introduces incentives to shirk to qualify for aid. Examples include colleges allocating financial aid to students and U.S. sports leagues assigning newly eligible players to teams. We prove that no targeting mechanism based on cumulative performance can avoid introducing perverse incentives. We show that the optimal allocation rule is computationally infeasible. We design a simple, incentive-compatible, and dynamic mechanism that targets low-ranked agents based on their performance history up to an endogenous stopping time. We prove the mechanism is optimal among a large subset of allocation rules. Using data from the NBA, we show how our mechanism aligns incentives and improves targeting.

Keywords: Incentive-Compatible Mechanism, Institutional Design, Targeting
JEL Codes: D82, H23, Z28

*We thank Mohammad Akbarpour, Susan Athey, László Csató, Edward Lazear, Michael Ostrovsky, Brad Ross, Andrzej Skrzypacz and Bob Wilson for helpful comments and discussions, as well as Connor King, 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. A previous version of the paper was circulated under the title *An Incentive-Compatible Draft Allocation Mechanism.*

[†]Graduate School of Business, Stanford University. 655 Knight Way, Stanford CA 94305. Banchio: mbanchio@stanford.edu, Munro: munro@stanford.edu

1 Introduction

Numerous private and public institutions help their members in need with targeted aid. The business schools of Stanford and Harvard University, for example, award financial aid packages to incoming students with the lowest income and assets. When a new container ship arrives at West Coast ports, the International Longshore and Warehouse Union assigns priority to workers with the lowest number of hours worked in the current quarter.[1] In the annual draft lottery, the NHL, NBA and the NFL assign the most coveted newly eligible players with a higher probability to the teams with the worst win records in the previous year. In the curious West Point "Goat" tradition, the lowest-ranking graduating cadet receives a cash prize and a rousing ovation.[2] In all of these examples, individuals are targeted based on performance measures.

However, targeting aid based on performance measures introduces perverse incentives to intentionally under-perform in order to qualify for aid. In the NHL and the NBA, the practice of teams intentionally losing near the end of a season in order to receive a better draft pick is colloquially known as 'tanking'. This strategy is frequently and passionately discussed in sports media and has been empirically documented in the NHL by Fornwagner (2019) and in the NBA by Taylor and Trogdon (2002). Tanking-like strategies appear outside of sports leagues as well. For example, college consultants have devised strategies for sheltering assets on the Free Application for Federal Student Aid (FAFSA), including taking unpaid leaves of absence and purchasing non-reportable luxury goods.[3] Reyes (2008) provides empirical evidence that eligibility criteria for need-based financial aid introduces incentives that substantially reduce family savings. Major Rittenburg, a Class of 1973 West Point graduate, reports that "The Goat candidates had to know the material perfectly so they would know exactly how many questions to answer wrong... not failing, but barely passing and still below their competitors."[4]

The negative effect of tanking is two-fold: not only does it reduce the overall effort exerted by the members of an institution, it also hurts the effectiveness of targeting. Poor performance is no longer a good signal of need in the presence of incentives to tank. Institutions have devoted significant resources to adjusting eligibility criteria in attempt to improve targeting and minimize perverse incentives. In addition to the FAFSA, many private universities collect a wider set of information from parents using the CSS Profile, including business income and medical expenses, to reduce incentives to engage in asset sheltering strategies. The NHL and NBA have changed the design

[1] See Cole (2018)

[2] https://www.wsj.com/articles/SB10001424127887324352004578131262893535452

[3] See, for example https://finaid.org/financial-aid-applications/maximize/ and https://www.cappex.com/articles/money/how-to-shelter-assets-on-the-fafsa

[4] https://www.west-point.org/family/bicent/goat.html

of the draft lottery repeatedly in the last 30 years in order to address concerns about competitive balance and tanking, sometimes increasing the probability with which the lowest-ranked teams are allocated the top draft pick, and sometimes decreasing it. Targeting aid without encouraging tanking is a challenging problem.

In this paper, we study formally how a planner, such as a manager of an institution, can target aid based on relative performance criteria without introducing incentives to under-perform. We first build a framework that captures the incentives introduced by aid targeting mechanisms. We choose to model this problem as a contest between multiple agents.[5] The tournament consists of several periods, where in each period agents engage in competition with a certain probability of success. Successes and failures are recorded in a cumulative score. The probability of success is affected by the agent's choice, which is whether or not to exert effort. Agents maximize the value of the prizes that they expect to receive at the conclusion of the tournament, and prizes are allocated based on the cumulative scores. The key difference in this paper compared to a classic model of incentives in tournaments (see Lazear and Rosen (1981) and Rosen (1986)) is that there is a non-monotonic prize structure. While an exogenously-determined market prize is assigned to the top performers, the planner decides how to target the bottom-ranked agents with a limited number of aid packages. This non-monotonicity introduces an endogenous cost of effort in the model. Effort is costly only when exerting it reduces the agents' chances of receiving an aid package but does not sufficiently increase their chances of receiving the market prize.

We model the choice of aid allocation system as a mechanism design problem in the contest model described. The manager's objective is to maximize the probability that agents ranked below a cutoff receive aid, constrained by an incentive-compatibility condition requiring that agents play the efficient equilibrium, where every agent exerts effort in each period. Maximizing this objective leads to a choice of mechanism that targets the lowest-ranked agents without introducing any incentive to reduce effort. The first question we address is what kind of mechanisms satisfy the incentive-compatibility constraints. Most existing aid mechanisms based on relative performance use an ex-post cumulative measure of performance. The NBA determines draft lottery weights based on final rankings after the season concludes and the FAFSA collects asset information at the date of application. Our first result is that any ex-post mechanism based on final rankings that favors the lowest-ranked agents provides incentives to shirk in some possible history. This negative result suggests we should look at mech-

[5] Aid is often allocated to the neediest members of an institution on a fixed budget. For this reason, we analyze the design of an optimal targeting rule in a contest model. However, the contest structure is fairly general: it captures single-person decision problems where an agent qualifies for aid based on a static threshold, as in many government programs. See Appendix A.3 for details: in this sense, the model accommodates both relative and absolute performance evaluation.

anisms that adapt the allocation probabilities over time based on the performance history of each agent. In our applications, this implies that using data on historical performance can improve targeting while reducing incentives to tank. In line with our insight, rather than taking a snapshot of information, the Stanford Graduate School of Business recently began using three years of asset and income information from incoming MBA students to determine need-based aid packages, which reduce incentives to shelter assets immediately before enrolling in the program.

The challenge to address is how to adjust the allocation probabilities dynamically in an optimal and computationally feasible way. We show that the designer's problem is a convex program with an infeasible number of constraints. In Section 3.2 we address the computational issues by first solving a relaxed version of the designer's problem that does not include the IC constraints. A solution to this relaxed problem is a simple dynamic mechanism that in each period targets agents based on the conditional probability they rank below the aid cutoff. We then address the incentive constraints separately by freezing targeting probabilities at an endogenous stopping time. We refer to this incentive compatible targeting rule as T-IC. We show this process is constrained-optimal, since it minimizes the designer objective conditional on all information up to the stopping time. Furthermore, this conditional objective monotonically decreases with additional information, which guarantees that our rule outperforms any policy that leverages less information. We show that by leveraging the dynamics of the tournament it is possible to target aid to the worst performers without introducing any incentives to shirk.

Though computing the globally optimal rule is not computationally feasible in general, in Section 3.3 we set up a toy model with three contestants and six periods where it is possible to enumerate all 2^{18} possible histories. In this simulation, it is possible to compute the value of the globally optimal rule as the solution to a convex program and compare it with T-IC. We are then able to provide some interesting comparative statics. For example, we show that, as the ratio of the value of the market prize to that of the aid package increases, the value of T-IC closely approximates the value of the globally optimal rule. In addition, we highlight how the dynamic nature of incentives in this problem constrains the optimal rules. We show that in the absence of incentive issues, an optimal allocation mechanism would weight performance in periods closer to the end of a tournament higher than those earlier in the tournament, since performance in those periods is most relevant in determining final rankings. However, earlier in the tournament incentives are more slack since final rankings are uncertain and agents are still competing for the top market prizes. Towards the end of the tournament, incentives are more likely to bind as some agents who have lost repeatedly in earlier stages of the tournament have given up on pursuing a market prize and may compete for an aid package instead. As a result, when incentives are taken into account, the weights that both the optimal mechanism and T-IC place on each period are hump-shaped. Both allocation mechanisms place higher weights on intermediate periods where performance is increasingly indicative of final rankings but incentive constraints are still slack enough to allow for meaningful adjustment of the allocation probabilities.

We then describe in Section 4 how T-IC applies to the practical setting of designing the NBA draft lottery. For the NBA there are detailed data available on the histories of every season, which allows a quantitative analysis of our mechanism. From 1985-1989 the NBA used an ex-post uniform lottery which, according to our first theorem, was incentive-compatible. As a result, we can estimate the counterfactual performance of T-IC in those years without observing effort. We simulate how T-IC would have adjusted draft probabilities over time up until a season-specific stopping time; we show that our proposed mechanism, while stopping well in advance of the end of the season, targets the lowest-ranked team with a 39% average probability, compared to 14% from 1985-1988 and 11% in 1989 for the uniform lottery. This gain in targeting does not trade off any violation in incentive compatibility for any team.

2 Literature Review

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 the optimal level of redistribution is limited by the distortions it places on incentives. Redistributive policies, however, are essential to reduce persistence in inequality, as described theoretically by Mookherjee and Ray (2003).

There is an extensive literature on improving targeting of social programs while minimizing perverse incentive effects, but to our knowledge these issues have not been analyzed in a tournament setting. Kanbur (1987) suggests that welfare targeting based on observed characteristics introduces incentives to change behavior, as in our setting where the characteristics are a performance score. Of the various directions taken by the literature, one of the most popular is tagging, see Akerlof (1978) and Allcott et al. (2015), which relies on characteristics that are not manipulable. Another direction is deterrence or complexity-based mechanisms, as in Besley and Coate (1992) and Kleven and Kopczuk (2011). Our paper focuses on how dynamic information can improve targeting, without adding complexity or deterrence to the application process.

There is a wide literature studying various incentive problems in tournaments, beginning with work from Lazear and Rosen (1981) and Rosen (1986) who studied efficient prize structures in tournaments as optimal labor contracts. Dagaev and Sonin (2018) study incentive compatibility violations in tournaments with multiple qualification rounds. Brown (2011) shows that inequality in ability can result in effort reduction in a tournament setting and Brown and Minor (2014) show how dynamic strategies in multiple stage competitions affect the probability the strongest player wins.

In sports economics, there is a body of literature proposing draft allocation mechanisms but most of it lacks 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. Concurrent work from Kazachkov and Vardi (2020) sets up a theoretical model of a tournament and also suggest the NBA draft would be improved by running a lottery at a stopping time earlier in the tournament. Rather than designing a fully incentive-compatible mechanism, they computationally illustrate the trade-off between the prevalence of tanking and how much the draft benefits the lowest-ranked team in expectation.

It is possible to describe the optimal incentive-compatible allocation mechanism as a solution to a simple optimization problem. However, it is impossible to compute that solution for even reasonably sized settings, so we propose a constrained-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, such as 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).

3 Designing a Targeting Mechanism

3.1 Designer and Agent Objectives

A tournament is made up of a total of T periods. There are n agents in the tournament. Each agent accumulates a performance metric, which we call a score. Examples of a score include cumulative assets in the financial setting or cumulative victories in the sports setting. In every period agents have an opportunity to increase their score. They choose how much effort $e_{it} \in \{0,1\}$ to exert, knowing that it affects the probability of increasing their score. In every period, a random vector of outcomes $O_t \in \mathcal{O}_t \subseteq \{0,1\}^n$

is realized according to a distribution μ_t.[6] The outcome $O_{it} = 1$ indicates agent i successfully completed her task in period t. We define the score of agent i as the integer $S_{it} = S_{i,t-1} + O_{it}$. We allow for heterogeneous initial scores $S_{i0} \geq 0$ to model inherited differences. A t-history is a sequence $S^t = \{S_s\}_{s \leq t}$ where S_s is the vector of scores at time s. A full history is a T-history. The set of all histories is denoted by $\mathcal{S}$. A tournament is a tuple $(n, T, S^0, \{\mu_t, \mathcal{O}_t\}_{t=1,\ldots,T})$. Finally, define W_{it} as the random vector in $\mathcal{O}_t$ with distribution $Pr(W_{it}) = Pr(O_t | O_{it} = 1)$. W_{it} records a successful period for agent i and stochastic outcomes for every other agent. Similarly, the random vector L_{it} is distributed according to $Pr(L_{it}) = Pr(O_t | O_{it} = 0)$.

We assume that there is a market prize with value π_i^V for agent i, which is allocated to the best v^* performers. There are also d^* aid packages with value π_i^D, and the planner would like to allocate them to the d^* lowest-ranked performers. We assume $0 < \pi_i^D \leq \pi_i^V$ and that the incentive problem is non-trivial, so $v^* + d^* < n$.[7] This generates a trade-off for the planner, as the first-best prize structure is non-monotonic: agents ranked between v^* and $n - d^*$ receive no prize, while agents above v^* or below $n - d^*$ do. The assumption that market prizes are out of the control of the planner reflect exogenously-determined rewards to good performance. In the NBA setting, v^* is the cutoff for making the playoffs and π_i^V represents the expected additional revenue from the playoffs that a team receives, in terms of media exposure and ticket sales. Aid represents the draft picks, whose value can be interpreted as the increase in long-term expected revenue that a team expects to receive from drafting the top eligible player. In the financial aid setting, aid packages may include tuition reductions as well as housing and dining discounts given based on demonstrated need.

The agent succeeds at a task with probability $p_{it}(e_t)$,[8] defined by

$$p_{it}(e_t) = Pr(O_{it} = 1 | e_t) = Pr(W_{it} | e_t)$$

We assume that the probability is non-decreasing in e_{it}. If the probability p_{it} is constant 0 across effort levels, agent i doesn't face any task during that period.[9] Additionally, we will assume that an agent's effort affects other agents' success probabilities only through her own outcome.[10] The ranking of an agent i is defined as follows, where the

[6] Formally, let $\mathcal{O}_t \subseteq \{0,1\}^n$ be a measurable space for every t. The space of measures over $\mathcal{O}_t$ is denoted $P(\mathcal{O}_t)$. The distribution of O_t is $\mu_t(e_t)(O_t)$, where $\mu_t \colon \{0,1\}^n \to P(\mathcal{O}_t)$ is a functional mapping vectors of effort choices to probability measures over the set of outcomes.

[7] In settings where it is unrealistic to assume that agent-specific values for the prizes are common knowledge, we can instead assume that a lower bound $\bar{\pi}^V$ on the value of each market prize and an upper bound $\bar{\pi}^D$ on the value of each aid package is common knowledge.

[8] The dependence of p_{it} on the identity of the agent reflects anticipated differences in performance.

[9] See Appendix A.3 for a discussion of contests with inert agents.

[10] This assumption is weaker than unconditional independence and is satisfied by every environment we describe.

agent with the maximum score has rank 1:

$$r_i(S^t) = 1 + \sum_{j \in I \setminus i} \mathbb{1}(S_j^t > S_i^t)$$

Ties are broken at random. Denote by $R_j(S^t)$ the identity of the agent ranked j, that is, $R_j(S^t) = i$ if and only if $r_i(S^t) = j$. We denote the probability of receiving the market prize in the tournament for agent i given a history S^t as $q_i(S^t)$. These probabilities can be defined recursively:

- $q_i(S^T) = \mathbb{1}(r_i(S^T) \leq v^*)$
- $q_i(S^{t-1}) = \mathbb{E}_{t-1}[q_i(S^{t-1} + O_t)]$ for every $t \leq T$

where the subscript $t-1$ represents the conditional expectation on information S^{t-1}.

Consider the set V_{d^*} of all the vectors $v \in \{0,1\}^n$ such that there are exactly d^* coordinates with $v_j = 1$. We define $d^*\Delta^n$ as the affine n-simplex over a fixed maximal subset of linearly independent vectors of V_{d^*}. We consider a targeting mechanism of the form:

$$y \colon \mathcal{S} \to d^*\Delta^n$$

where $y_i(S^t)$ represents the probability that agent i receives the aid package given the history up to period t. The mechanism is restricted in the following ways. Since $y_i(S^t)$ represents the expected allocation probabilities conditional on information up to time t, the probabilities at time t must be dynamically consistent with the probabilities conditional on information up to time $t-1$:

$$y_i(S^{t-1}) = \mathbb{E}_{t-1}[y_i(S^{t-1} + O_t)] \qquad \text{for every } i, t \leq T \tag{DC}$$

Moreover, the lottery probabilities at any history need to add up to d^*:

$$\sum_{i=1}^{n} y_i(S^t) = d^* \qquad \forall S^t \tag{PROB}$$

A targeting allocation mechanism satisfying DC and PROB is feasible, and the space of feasible mechanisms is $\mathcal{Y}$. A simple linear algebra argument shows that any feasible targeting mechanism induces a lottery over deterministic aid allocations, see Lemma 6 in Appendix A.2.

Each agent makes a single strategic choice in each period, which is how much effort to exert. Besides literal interpretations of effort, the agent's action may reflect a savings commitment, or any other agent's choice that affects the expected score. Since in what follows we assume no explicit cost of effort, the efficient action for every agent is to exert effort in every period. We are interested in mechanisms that implement the efficient equilibrium.

Agent Objective. For a given targeting mechanism, in period t agent i chooses an effort level to maximize her expected payoff given the results so far. Following the one-shot deviation principle and applying the tower property of conditional expectations, agent i in period t in such an equilibrium faces the following objective.

Optimization Problem 1.

$$\begin{aligned} \max_{e_{it}} \quad & p_{it}(e_t)\Big(\mathbb{E}_{t-1}[q_i(S^{t-1}+W_{it})]\pi_i^V + \mathbb{E}_{t-1}[y_i(S^{t-1}+W_{it})]\pi_i^D\Big) \\ +(1-p_{it}(e_t))&\Big(\mathbb{E}_{t-1}[q_i(S^{t-1}+L_{it})]\pi_i^V + \mathbb{E}_{t-1}[y_i(S^{t-1}+L_{it})]\pi_i^D\Big) \end{aligned} \tag{1}$$

Notice that this optimization is relevant only when the probability is not constant across effort levels. When an agent is indifferent between effort levels, we assume that she exerts maximum effort.[11] Since probability is assumed non-decreasing in effort, maximizing the agent's objective we derive a necessary condition for efficiency:

Incentive Condition (IC)*.*

$$\begin{aligned} \mathbb{E}_{t-1}\Big[q_i(S^{t-1}+W_{it}) - q_i(S^{t-1}+L_{it})\Big]\frac{\pi_i^V}{\pi_i^D} \geq \\ \mathbb{E}_{t-1}\Big[y_i(S^{t-1}+L_{it}) - y_i(S^{t-1}+W_{it})\Big] \end{aligned} \tag{IC$_{(i,t)}$}$$

If this inequality is satisfied, it is optimal for the agent to exert maximum effort. (IC) has a clear interpretation. If the increase in the probability of receiving aid when unsuccessful is less than the decrease in the probability of receiving the market prize, scaled by the ratio of the prizes' values, then the agent will exert effort. Notice how in our model without an explicit cost of effort, an implicit cost arises as a result of allocation rules targeting the worst performers. The non-monotonicity of the reward structure makes the agent's problem non-trivial. The results do not hinge on this cost-of-effort assumption, as discussed in Appendix A.1.

The common thread in our applications is that a planner would like to target aid to the lowest-ranked contestants. For example, the U.S. sports leagues have indicated that maximizing competitiveness in the league over repeated tournaments requires allocating the valuable first round draft picks to the lowest ranked teams. Universities promote long-run equality of opportunity through aid policies. Our simple planner minimization reflects these observations.

[11] Successfully completing a task should be preferred to losing in the short-term, in the absence of long term incentives. For example, in the NBA, teams receive a benefit in terms of fan engagement from winning a game and would likely exert effort as a result of this, even in the absence of other external incentives. Parents would likely choose to accumulate assets absent incentives to lower their bank account balance.

Planner Objective. The first-best policy would allocate the aid packages to the bottom ranked agents at the end of the tournament. Since in general the first-best policy will not be incentive compatible, we focus on finding an optimal second-best policy. The planner's objective is to find an allocation rule y that satisfies IC and minimizes the mean squared deviation from the first-best mechanism. Let the cutoff $n - d^*$ be denoted by k^*: the planners aims to target agents whose rank falls below k^*.

Optimization Problem 2.

$$\min_{y \in \mathcal{Y}} \mathbb{E}\left[\sum_{i=1}^{n} \Big(\mathbb{1}(r_i(S^T) > k^*) - y_i(S^T)\Big)^2\right] \tag{2}$$

subject to $\forall i, t, k, S$

$$\mathbb{E}_{t-1}\Big[q_i(S^{t-1} + W_{it}) - q_i(S^{t-1} + L_{it})\Big]\frac{\pi_i^V}{\pi_i^D} \geq \mathbb{E}_{t-1}\Big[y_i(S^{t-1} + L_{it}) - y_i(S^{t-1} + W_{it})\Big] \tag{IC$_{(i,t)}$}$$

We begin by examining allocation mechanisms that rely on final rankings only, disregarding the full history of performance. We first restrict our attention to tournaments where the implicit cost of effort affects agents' choices.

Definition 1. A tournament $(n, T, S^0, \{\mu_t, \mathcal{O}_t\}_t)$ is non-degenerate if for every $l > m \geq v^*$ there exists an agent i and a history S^{t-1} such that

$$\mathbb{E}_{t-1}\Big[q_i(S^{t-1} + W_{it}) - q_i(S^{t-1} + L_{it})\Big] = 0 \tag{3}$$

but

$$\begin{aligned}&\mathbb{E}_{t-1}\Big[\mathbb{1}(r_i(S^{t-1} + L_{it}) > l) - \mathbb{1}(r_i(S^{t-1} + W_{it}) > l)\Big] > \\ &\mathbb{E}_{t-1}\Big[\mathbb{1}(r_i(S^{t-1} + L_{it}) > m) - \mathbb{1}(r_i(S^{t-1} + W_{it}) > m)\Big] \geq 0\end{aligned} \tag{4}$$

Non-degenerate tournaments require the existence of some histories where agents cannot increase their probability of receiving the market prize but can decrease their expected ranking.

Our first result shows that any incentive compatible rule that targets lower-ranked agents cannot depend only on the final rankings. It provides a rationale for turning our attention to mechanisms that account for the dynamics of the tournament.

Theorem 1. *If a tournament $(n, T, S^0, \{\mu_t, \mathcal{O}_t\}_t)$ is non-degenerate, the only targeting mechanism y such that:*

1. *The mechanism is a function only of an agent's final ranking $r_i(S^T)$,*

2. *The mechanism targets lower-ranked agents with a higher probability than higher-ranked agents. Formally, there exists a* $l > m \geq v^*$ *such that*

$$y_{R_m}(S^T) \leq y_{R_l}(S^T)$$

3. *The mechanism satisfies (IC) at every history* $S^t \in \mathcal{S}$*,*

is a uniform lottery, which assigns equal probabilities to every agent i *ranked lower than* v^* *.*

Proof. Let y be a feasible mechanism satisfying conditions $1-3$. Under the assumptions, there always exists a history S^t such that such that:

- Agent i's chances of ranking in the first v^* positions are constant:

$$\mathbb{E}_{t-1}\left[q_i(S^{t-1} + W_{it})\; -\; q_i(S^{t-1} + L_{it})\;\right] = 0$$

- Agent i's chances of ranking in the bottom positions can be improved by losing in period t:

$$\begin{aligned}&\mathbb{E}_{t-1}\left[\mathbb{1}(r_i(S^{t-1} + L_{it}) > l) - \mathbb{1}(r_i(S^{t-1} + W_{it}) > l)\;\right] > \\ &\mathbb{E}_{t-1}\left[\mathbb{1}(r_i(S^{t-1} + L_{it}) > m) - \mathbb{1}(r_i(S^{t-1} + W_{it}) > m)\;\right] \geq 0\end{aligned}$$

Then, we have for agent i:

$$\begin{aligned}0 =&\mathbb{E}_{t-1}[q_i(S^{t-1} + W_{it}) - q_i(S^{t-1} + L_{it})]\frac{\pi_i^V}{\pi_i^D} \geq \\ &\mathbb{E}_{t-1}[y_i(S^{t-1} + L_{it}) - y_i(S^{t-1} + W_{it})] \geq 0\end{aligned}$$

where the first inequality is (IC) 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 agents ranked lower than v^*. □

The theorem shows that independently of how carefully the planner chooses the targeting lottery, if the lottery favors lower ranked agents, then there will always be incentives for agents to shirk after they are eliminated from market prize contention. Theorem 1 provides a compelling reason to turn our attention to the dynamics of performance.

The planner optimization problem (2) is a convex program. However, solving for the global optimum is infeasible. Such a calculation would require all possible histories of a tournament to be enumerated to specify the (IC) constraints and to calculate the planner objective which is an expectation over all possible histories. For example, in a standard NBA season the game tree is of size 2^T, where $T = 1,230$. We propose a simple rule that is computationally feasible and provides us with clear intuition. It satisfies the (IC) constraints through a history-dependent stopping time, so that incentive compatibility does not affect the adjustment of the weights in periods before the stopping time. This means that the allocation probabilities can be determined in each period without enumerating all possible future histories.

3.2 An Incentive-Compatible Targeting Mechanism

Theorem 1 indicates that a lottery based on final rankings that favors the worst agents will not satisfy (IC). Let us instead turn our attention to dynamic rules, which adjust the targeting probabilities as the score progresses. We can write a relaxed version of the planner objective as a dynamic programming problem:

$$\min_{y \in \mathcal{Y}} V_y(S^0) \tag{5}$$

subject to

$$\begin{aligned} V_y(S^T) =& \sum_{i=1}^{n} \Big(\mathbb{1}(r_i(S^T) > k^*) - y_i(S^T)\Big)^2 \qquad \forall S^T \\ V_y(S^t) =& \mathbb{E}\left[V_y(S^{t+1})|S^t\right] \qquad \forall S^t \end{aligned}$$

Following Bellman's principle of optimality, if y^* solves the problem above then its truncation after a history S^t is tail-optimal, meaning it solves problem $\min_{y\in\mathcal{Y}} V_y(S^t)$ for every S^t. This implies that problem (5) is equivalent to a relaxation of problem (2) disregarding the incentive compatibility condition, and y^* is the optimal control of problem (2). Problem (5) may have arbitrarily many solutions, since the intermediate values of the control do not affect the value of the objective. Because every solution yields a minimum value of zero, we select one simple minimizer which we know to be feasible. This solution to the relaxed problem assigns to each agent her conditional probability of ending the tournament below rank k^* as her targeting probability. We call this solution $\overline{y}$. We first prove that this is indeed a solution of the relaxed problem:

Lemma 2. *The rule $\overline{y}$, defined as $\overline{y}_i(S^t) = Pr(r_i(S^T) > k^*|S^t)$ for all i, t, is a feasible solution of the unconstrained problem (5).*

Proof. To prove feasibility, note that $\sum_{i=1}^{n} \overline{y}_i(S^t) = \sum_{i=1}^{n} Pr(r_i(S^T) < k^*|S^t) = d^*$, and by the law of iterated expectations

$$\overline{y}_i(S^t) = Pr(r_i(S^T) > k^*|S^t) = \mathbb{E}[\mathbb{1}(r_i(S^T) > k^*)|S^t] =$$
$$= \mathbb{E}\left[\mathbb{E}[\mathbb{1}(r_i(S^T) > k^*)|S^{t+1}]\middle|S^t\right] = \mathbb{E}\left[\overline{y}_i(S^{t+1})\middle|S^t\right]$$

which proves dynamic consistency. To see that $\overline{y}$ is optimal, it is sufficient to notice that $Pr(r_i(S^T) > k^* \ |S^T) = \mathbb{1}(r_i(S^T) > k^*)$. □

The rule $\overline{y}$ is simple and intuitive: it dynamically adjusts the lottery odds as the tournament progresses. When agents are equal, the probability of receiving aid before the tournament begins is $\frac{d^*}{n}$. If instead there is an initial level of inequality, the ex-ante probability is higher for disadvantaged agents. As successes and failures are recorded in each period, an agent's probability of receiving aid adjusts based on their updated conditional probability of ending up ranked below k^*. For example, when an agent fails in multiple periods early in the tournament, the probability of ranking low increases and the planner increases their targeting probability accordingly. However, this rule may not be incentive compatible. Intuitively, we expect (IC) to hold earlier in the tournament, when incentive constraints are slack since scores are equal. We know that it will bind towards the end, when competition for the market prize is diminished due to accumulated inequality in scores. We follow this intuition and propose an incentive compatible targeting rule T-IC, denoted by y^{T-IC}, that satisfies incentive compatibility and still targets the bottom of the rank distribution. The rule y^{T-IC} coincides with the global optimum $\overline{y}$ until the incentive compatibility condition starts to bind. We denote the first period where (IC) binds along a particular history S^T as $t^*(S^T)$. Along any history following S^{t^*} the value of y^{T-IC} is constant. This ensures that (IC) is satisfied in every period.

Formally, we define the stopping time t^* as

$$t^*(S^t) = \min\left\{s \le t \colon \exists i \in I \text{ s. t. } (\text{IC})_{i,s} \text{ is violated under rule } \overline{y}\right\}$$

When there is no such $s \le t$, we let $t^*(S^t) = t$. The targeting policy y^{T-IC} then takes the following form:

$$y_i^{T-IC}(S^t) = \begin{cases} Pr(r_i(S^T) > k^* \ |S^t) \text{ if } t \le t^*(S^t) \\ Pr(r_i(S^T) > k^* \ |S^{t^*}) \text{ if } t > t^*(S^t) \end{cases}$$

Lemma 3. *The rule y^{T-IC} is feasible and incentive compatible.*

Proof. As in Lemma 2, the feasibility constraint (PROB) is satisfied by definition. To see that dynamic consistency is satisfied, note that:

- in period $t \leq t^*$ we can apply the same reasoning as in Lemma 2
- in period $t > t^*$ the targeting probabilities are constant, hence

$$y_i^{T-IC}(S^t) = \mathbb{E}[y_i^{T-IC}(S^{t+1})|S^t]$$

To show that (IC) is satisfied, we proceed by contradiction. Suppose that some (IC)$_{(i,t)}$ was violated. This is equivalent to

$$\begin{aligned}\mathbb{E}_{t-1}\Big[q_i(S^{t-1}+W_{it}) - q_i(S^{t-1}+L_{it})\Big]\pi_i^V \\ < \mathbb{E}_{t-1}\Big[y_i(S^{t-1}+L_{it}) - y_i(S^{t-1}+W_{it})\Big]\pi_i^D\end{aligned}$$

But T-IC fixed the targeting probabilities at some time $s \leq t$. This implies the RHS is zero, and since the LHS is always weakly positive it cannot be that (IC) was violated. □

By stopping the probability updating process we are able to satisfy incentive compatibility in every period, but we can no longer target the lowest-ranked agent in every history. Nonetheless, the performance of y^{T-IC} has an optimal property.

Theorem 4. *The rule y^{T-IC} is optimal among the rules that are constant after period t^*.*[12]

Proof. We exploit the dynamic programming characterization of the designer's problem. Let $\mathcal{Z}_t$ be the subspace of $\mathcal{Y}$ of feasible policies constant after period t, and z a generic element of $\mathcal{Z}_t$. The value of problem (5) for any z is $V_z(S^0) = \mathbb{E}[V_z(S^t)]$. The period with the relevant minimization is period t: the information arriving afterwards cannot be incorporated in the control. Then our problem reduces to minimizing the following:

$$\min_{z\in\mathcal{Z}_t} \mathbb{E}\left[\sum_{i=1}^{n}\Big(\mathbb{1}(r_i(S^T) > k^*) - z_i(S^t)\Big)^2\Big|S^t\right] \tag{6}$$

where $z_i(S^t)$ are non-random constants. Relaxing the feasibility constraint, we simply minimize

$$\min_{z_i\in\mathbb{R}} \mathbb{E}\left[\Big(\mathbb{1}(r_i(S^T) > k^*) - z_i\Big)^2\Big|S^t\right]$$

for each $i = 1, \ldots, n$. The value of z_i that minimizes this conditional average squared deviation of the random variable $\mathbb{1}(r_i(S^T) > k^*)$ is simply the conditional mean of the random variable,

$$\mathbb{E}[\mathbb{1}(r_i(S^T) > k^*)|S^t] = Pr(r_i(S^T) > k^*|S^t) = \overline{y}_i(S^t)$$

[12]We omit the dependence of t^* on the history for ease of notation

Therefore, we can construct a $z \in \mathcal{Z}_t$ that solves minimization (6):

$$z_i(S^s) = \begin{cases} \overline{y}_i(S^s) \text{ if } s \leq t \\ \overline{y}_i(S^t) \text{ if } s > t \end{cases}$$

Setting $t = t^*$ proves the Theorem's claim. $\square$

From the proof it transpires another property of the T-IC policy. Under T-IC, in each period until t^*, the value of the designer's objective monotonically decreases.

Corollary 5. *For any $s \leq t$,*

$$\min_{z \in \mathcal{Z}_t} V_z(S^0) \leq \min_{z \in \mathcal{Z}_s} V_z(S^0)$$

Proof. The claim follows from the following observation: a policy that is constant after period t is also constant after period $t+1$. Therefore the sets $\mathcal{Z}_t$ are nested: $\mathcal{Z}_1 \subseteq \mathcal{Z}_2 \subseteq \cdots \subseteq \mathcal{Z}_T = \mathcal{Y}$. $\square$

Any incentive compatible rule must depend at least partially on the history of scores in the tournament, as shown in Theorem 1. This property of incentive compatible mechanisms highlights the complexity of solving the planner's problem. The rule we propose takes care of the complex restrictions by decoupling the incentives from the optimality of the rule. We address the optimality requirements by solving a relaxed program, and we take care of incentives separately by means of the stopping time t^*. This separation is the key to the computational feasibility of the method, as incentives only affect the stopping time and not the adjustment of the weights. In the globally optimal solution computed in Section 3.3 for a small tournament, the incentives affect both the stopping time and the weights and as a result the optimal weights require enumerating all possible histories of the tournament.

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 record from any periods, then the policy that would minimize Optimization Problem 2 in expectation would be an ex-ante lottery over all agents. Conditioning the allocation mechanism on additional outcomes progressively decreases the value of the objective function. If we ignored the restrictions posed by (IC), we could condition on the full history S^T. Without a stopping time, T-IC would allocate aid to the bottom d^* agents with probability 1, as shown in Lemma 2. This yields the minimum possible value of the planner objective. However, the incentive compatibility requirements places us strictly in between the no information case and the full information case; we minimize problem (2) conditional on as much information as we can take into account without violating (IC). The first binding incentive compatibility constraint determines how far from the optimum our solution will be. One drawback is that for certain histories where incentive constraints bind early for a contestant in the tournament, the stopping time will occur quite early. If this is the case, the approximate objective that T-IC satisfies, which is conditional on a limited history up to the stopping time, may be far from the true objective.

We next provide a small simulation showing how T-IC compares to the globally optimal, ex-ante, and ex-post IC rules, by explicitly computing the stopping time t^* for every history.

3.3 Simulation

Consider a simple setting with 3 agents, where agents have equal ability, the top agent receives the market prize, and there is only one aid package available. This is the smallest number of agents such that there are incentive issues.

We simulate a tournament with $n = 3$ agents and $T = 6$ periods. This simulation can be considered a toy model of the financial aid setting. For each of 6 years, all three agents choose whether or not to be thrifty and try to save a discrete sum. If they do attempt to save, they are successful with a 50% probability.[13] The goal is to allocate an aid package to the agent with the lowest savings at the end of 6 years, without introducing incentives to intentionally reduce savings in order to receive aid. It is feasible, though computationally intensive[14], to enumerate all possible 2^{18} outcomes of the tournament and calculate the globally optimal incentive-compatible mechanism as the solution to the convex program in Optimization Problem 2. We report the designer's value averaged over all possible histories for four different incentive-compatible mechanisms:

1. Ex-post uniform lottery: a uniform lottery over all agents that do not receive the market prize.
2. Ex-ante uniform lottery: a uniform lottery over all agents.
3. The globally optimal rule
4. T-IC, the approximately optimal rule

All mechanisms are incentive compatible as long as $\frac{\bar{\pi}^V}{\bar{\pi}^D} \geq 1$.

We proved that our allocation rule is constrained-optimal; it maximizes the planner objective conditional on results only up to a dynamically-determined stopping time. How late this stopping time is realized depends on the assumed lower bound $\frac{\bar{\pi}^V}{\bar{\pi}^D}$ for any agent. Figure 1a shows how the value of the designer objective varies with $\frac{\bar{\pi}^V}{\bar{\pi}^D}$.

[13] Failing to save can be interpreted as a transient negative income shock.

[14] We use the Julia optimization package JuMP and the commercial solver Mosek to perform this optimization

The ex-ante uniform rule gives the aid package with expected probability 33% to the lowest-ranked agent. The ex-post uniform rule gives aid with expected probability 50% to the lowest-ranked agent. For low values of $\frac{\bar{\pi}^V}{\bar{\pi}^D}$, T-IC does not do much better than the uniform rules. The incentive constraints bind early in the tournament, so the allocation probabilities cannot be adjusted. However, as $\frac{\bar{\pi}^V}{\bar{\pi}^D}$ increases, the designer objective decreases rapidly for T-IC and decreases slowly for the globally optimal rule. Additionally, both rules converge to a constant value as $\frac{\bar{\pi}^V}{\bar{\pi}^D}$ increases. The value of the planner objective under the first-best policy is 0. As expected from Theorem 1, neither the globally nor the approximately optimal rule can ever achieve the first best, no matter how large $\frac{\bar{\pi}^V}{\bar{\pi}^D}$ is. The globally optimal rule converges to a value of approximately 0.038 while T-IC converges to a value of approximately 0.067, so approximates well the global optimum compared to the uniform incentive compatible rules. The performance of T-IC is computationally feasible as T increases, whereas the globally optimal rule is not, since it requires enumerating all possible histories.

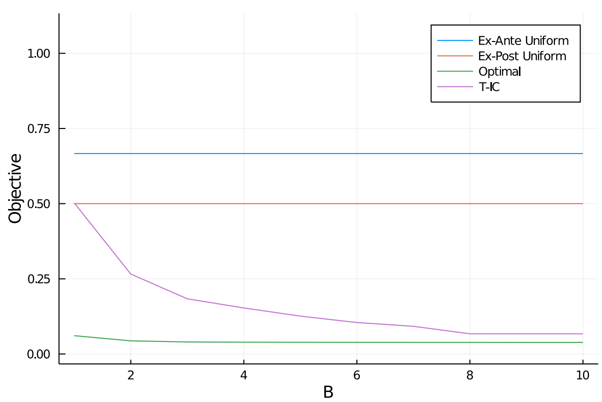

(a) Designer Objective, $B = \frac{\pi^V}{\pi^D}$

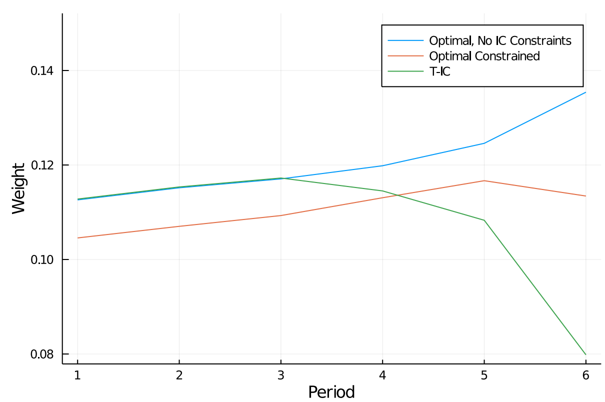

(b) Average Aid Allocation Probability Change by Period

Figure 1: Comparing Alternative Rules in Simulations

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 aid allocation probability between period t and $t-1$, averaged over all histories and agents, when we assume that $\frac{\bar{\pi}^V}{\bar{\pi}^D} = 10$. In the the absence of incentive constraints the weight placed on each period is higher later in the tournament, when it is increasingly clear who will have the lowest savings at the end of 6 years. However, when constraints are added, both T-IC and the optimal rule show a turning point where they place decreasing weights on later periods. In earlier periods of the tournament, incentive compatibility is more likely to hold, because every agent is still in contention for the market prize. In later periods, incentive compatibility constraints are tighter since certain agents may have failed to save repeatedly in the first few periods and have given up competing for the top prize.

T-IC follows the optimal rule without IC constraints for the first three periods. However, T-IC freezes the aid allocation probabilities for all agents once incentives for a single agent have been violated. As a result, after period 3 the weight that T-IC places on each period decreases. The globally optimal rule enumerates all possible histories of the tournament and can take into account the IC constraints when adjusting aid allocation probabilities. As a result, the optimal rule adjusts less in earlier periods, but can place a higher weight on later periods compared to T-IC. The benefit of this optimal adjustment process is quite small, since the two mechanisms have close performance as $\frac{\bar{\pi}^V}{\bar{\pi}^D}$ increases.

From these simulations we have indications that, for reasonable values of $\frac{\bar{\pi}^V}{\bar{\pi}^D}$, the constrained-optimal solution given by T-IC is close to the globally optimal solution and substantially better than the uniform rules.

4 Improving the NBA Draft Lottery

In the sports setting the dynamic performance data and predictive modeling necessary for estimating T-IC directly are publicly available. After providing some background on the history of the draft allocation lottery we estimate the performance of T-IC on empirical data on the NBA from 1985-1989. The NBA tournament occurs from October to April. Though the number of teams have changed over the years, there are currently thirty teams competing, divided in two conferences. Each team plays eighty-two games in a single 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 teams participate in the draft lottery for the first draft pick. In the notation and language of our model, the tournament is a season and agents are teams. π_i^V is the value of making the playoffs for the top $v^* = 16$ teams, which we assume has lower bound $\bar{\pi}^V$. $\bar{\pi}^D$ is the upper bound on the value of the first draft pick to any given team, which is targeted to the bottom-ranked team, so $d^* = 1$.

During the draft, teams 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 empirical example, we focus on the problem of allocating the first pick, which is the most valuable, but our framework extends to the allocation of multiple picks to the teams who do not make the playoffs. 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 intentionally losing 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%. Theorem 1 explains why the league has changed the system so often without finding a lottery that is satisfactory; it is not possible to have an ex-post incentive compatible lottery that also favors the worst-ranked teams.

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

When estimating our model on this data, one complication arises: we do not observe effort exerted by teams. We need to account for the changes in teams' effort choices arising from the alternate incentives for our counterfactual estimate to be valid. We can address this issue by exploiting specific institutional details of the draft lottery. Between 1985 and 1989, when the NBA had a uniform lottery system, teams did not have an incentive to intentionally lose games. According to our model, teams played the efficient equilibrium where no one had an incentive to reduce their effort. Since our rule also operates in an efficient equilibrium, we can directly estimate the performance of T-IC on the seasons from 1985 to 1989. However, we cannot use more recent years for this counterfactual exercise. Since the current NBA lottery incentivizes tanking, team exertion of effort would have changed under our rule and expected final rankings would be different. For each of the years from 1985-1989, in order to calculate the evolution of the draft probabilities up until the stopping time t^* as well as the stopping time itself, we need to calculate the following quantities for each game played in the season:

1. In order to adjust $y_i(S^t)$: The probability any given team will be ranked last conditional on their record after each game $t = 1, \ldots, t^*$:

$$Pr(r_i(S^T) = n \mid S^t)$$

   This probability is approximated by simulating the rest of the season, based on the simplifying assumption teams have equal ability, so win a game with 50% probability[15].

2. In order to determine the stopping time t^*: the incentives to win for each of the two teams that plays in every game t. This requires simulating the rest of the season for each game conditional on the results so far, assuming team i wins and assuming team i loses. This approximates the changes in the teams' probabilities of ending up ranked last and their probabilities of making the playoffs conditional on winning versus losing, which determines their incentives to exert effort in our model. We also assume for the purposes of determining the stopping time that $\frac{\bar{\pi}^V}{\bar{\pi}^D} = 10$.

We calculate the draft probabilities based on our incentive-compatible rule for every season from 1985-1989 and then examine the results from 1987 more closely, when there were 23 teams in the NBA. 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, over a total of roughly 1000 games per season. On average, the draft probabilities are adjusted until 45% of the season has occurred.

Table 2 shows the final draft probabilities for T-IC 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

[15]It is possible to replace this simple simulation with a more sophisticated forecast model. For example, the website FiveThirtyEight uses a version of the chess scoring system ELO to calculate win probabilities and forecast the results of the remaining games in the NBA season.

Rank	Team	Wins	NBA Lottery	T-IC 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, Stopping Time at 371st Game

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. Another potential drawback of the mechanism is that depending on the stopping time, it can assign a small probability to teams who marginally make the playoffs. In 1987, however, there is less than 4% total probability that any playoff team also gets the first draft pick.

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, T-IC 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 ex-post cumulative total of wins and losses, but also when those wins and losses occur.

5 Conclusion

We document various settings with relative performance where the planner is interested in assigning aid to under-performing agents. We prove that in a dynamic model with multiple choices of effort, while no ex-post re-distributive policy can be incentive

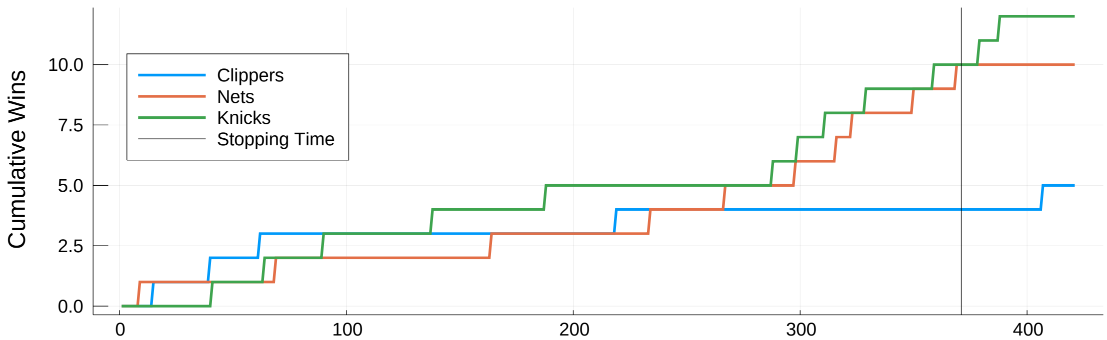

(a) Win Records

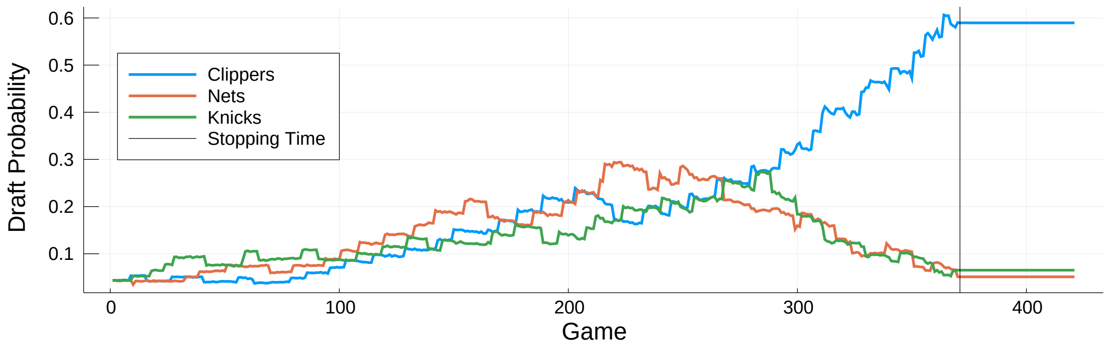

(b) Draft Probabilities

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

compatible, it is possible to target individuals that are likely to be ranked lower at the end of the tournament by adjusting allocation probabilities over time, up until a stopping time. The stopping time is dynamically determined in each tournament to ensure that the rule is incentive compatible in every possible history of a tournament. We show in a small simulation that our rule performs nearly as well as the globally optimal rule, which is the solution to a non-linear program that is not feasible in a larger, more realistic setting.

Our results have implications for the design of a variety of aid programs. It indicates that financial aid targeting is best improved by collecting a longer history of a simple set of financial metrics, rather than an increasingly complex snapshot of financial information at application time. As discussed, this sort of change has been recently made in the Stanford Graduate School of Business in order to reduce incentives to strategically reduce reportable wealth. In an institutional sports setting, the design of T-IC directly leads to an improved draft allocation mechanism for the NBA. Their lottery has been changed repeatedly over time without reaching a satisfactory system that addresses tanking while still supporting the worst teams in the league. We show using historical data from the NBA from 1985-1989 that T-IC significantly outperforms the league lottery in place at the time without introducing perverse incentives.

Our results apply more generally to the optimal design of eligibility requirements for aid and social programs, including in models that are not based on tournaments. As technology and data analysis improves, it is increasingly feasible for institutions and managers to use historical data about individuals in order to determine eligibility. We provide insight on how using performance data on an applicant over time can allow targeting without affecting incentives.

References

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

Akerlof, George A, "The economics of" tagging" as applied to the optimal income tax, welfare programs, and manpower planning," *The American economic review*, 1978, *68* (1), 8–19.

Allcott, Hunt, Christopher Knittel, and Dmitry Taubinsky, "Tagging and targeting of energy efficiency subsidies," *American Economic Review*, 2015, *105* (5), 187–91.

Besley, Timothy and Stephen Coate, "Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs," *The American Economic Review*, 1992, *82* (1), 249–261.

Brown, Jennifer, "Quitters never win: The (adverse) incentive effects of competing with superstars," *Journal of Political Economy*, 2011, *119* (5), 982–1013.

_ **and Dylan B Minor**, "Selecting the best? Spillover and shadows in elimination tournaments," *Management Science*, 2014, *60* (12), 3087–3102.

Cole, Peter, *Dockworker Power : Race and Activism in Durban and the San Francisco Bay Area.*, Urbana: University of Illinois Press, 2018.

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).

Kanbur, Ravi, "Transfers, Targeting and Poverty," *Economic Policy*, 1987, *2* (4), 112–147.

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

Kleven, Henrik Jacobsen and Wojciech Kopczuk, "Transfer program complexity and the take-up of social benefits," *American Economic Journal: Economic Policy*, 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.

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 tankingin 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.

Reyes, Jessica Wolpaw, "College financial aid rules and the allocation of savings," *Education Economics*, 2008, *16* (2), 167–189.

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.

Appendix A

A.1 Explicit Cost of Effort

In our model the cost of effort arises endogenously as a function of the v-shaped prize structure. However, effort could be intrinsically costly. In this section we analyze whether assuming effort costs impact our results. Suppose effort is binary, $e \in \{e^l, e^h\}$, and costly with $c(e^h) > c(e^l)$. Additionally, probabilities are increasing in effort:

$$p_i(\alpha_{w(t)}, e_{w(t)-i,t}, e^h_{it}) - p_i(\alpha_{w(t)}, e_{w(t)-i,t}, e^l_{it}) = \Delta_p \geq 0 \quad \forall e_{w(t)-i,t}$$

Then the agent's problem becomes

$$\begin{aligned}
\max_{e_{it}} \quad & [p_{it}(\alpha_{w(t)}, e_{w(t),t})\mathbb{E}[(q_i(S^{t-1} + W_{it})]\pi_i^V + \mathbb{E}[y_i(S^{t-1} + W_{it})]\pi_i^D) \\
& +(1 - p_{it}(\alpha_{w(t)}, e_{w(t),t}))(\mathbb{E}[q_i(S^{t-1} + L_{it})]\pi_i^V + \mathbb{E}[y_i(S^{t-1} + L_{it})]\pi_i^D)] - c(e_{it})
\end{aligned}$$

Agent i exerts maximum effort in period t if

$$\begin{aligned}
\mathbb{E}\,[q_i(S^{t-1} + W_{it}) - & \, q_i(S^{t-1} + L_{it})]\pi_i^V \geq \\
& \geq \mathbb{E}\,[y_i(S^{t-1} + L_{it}) - y_i(S^{t-1} + W_{it})]\pi_i^D + \frac{c(e^h) - c(e^l)}{\Delta_p}
\end{aligned}$$

This inequality corresponds to (IC), with the last additional term reducing the slack in the incentives. It should be obvious how Theorem 1 holds under these assumptions. A more interesting question is whether Theorem 4 holds. Since incentives only affect the stopping time but not the objective of the planner, the results are maintained under this specification. However, for any specific tournament the optimal rule T-IC will stop updating the targeting probabilities at an earlier period t^* in this model: the last term reduces slack in the incentives by a constant amount, therefore incentive compatibility will in general bind earlier in the tournament. While the left-hand side of the IC constraint is the benefit of exerting effort, the whole right-hand side is the cost: it's the sum of explicit cost of effort appropriately scaled and the implicit cost of effort due to the non-monotonic prize structure.

A.2 Targeting Mechanisms Induce Feasible Allocations

Lemma 6. *Every feasible targeting mechanism $y \in \mathcal{Y}$ induces a probability distribution over deterministic allocations of the equalizing transfers.*

Proof. Consider a vector $x \in d^*\Delta^n$. Any element of the affine simplex $d^*\Delta^n$ can be decomposed as a convex combination of the vertices of the simplex. The convex combination's coefficients form a probability distribution with the usual interpretation. But then, notice that evaluating a feasible targeting mechanism at any S^t yields a vector $y(S^t)$ in $d^*\Delta^n$. We can decompose each $y(S^t)$ as a probability distribution over the vertices of the simplex, that is, deterministic allocations of the d^* transfers. $\square$

A.3 Dummies

As mentioned in the introduction, our model accommodates single-person decision problems, which model numerous governmental aid programs. In certain environments it is useful to construct *dummy* agents:

Definition 2. A dummy agent is an agent i such that $p_{it}(e_{-i,t}, 0) = p_{it}(e_{-i,t}, 1)$ for every t.

A dummy agent's probability does not depend on his own effort in any history. If we construct dummies whose probabilities are either 0 or 1 in every period, we can interpret them as static thresholds against which other agents effectively compete. Their final score is determined at the very beginning of the tournament, and it is common knowledge. Incidentally, the generality of the model allows us to capture absolute performance evaluation by having an agent compete against static dummies. Absolute performance evaluation corresponds to a tournament with n agents where $n-1$ of them are dummies. In this setting, there is only one agent making choices that affect the equilibrium outcomes. She is competing against possibly stochastic thresholds, whose distribution is known at the beginning of the tournament. If we ask for dummies with probabilities either 1 or 0, we obtain an agent competing against deterministic thresholds, which models many classical settings of absolute performance evaluation. All the analysis of the paper applies: particularly, the stopping time will be determined exclusively by the only active player, since dummies' incentive constraint will never bind strictly. Our rule performs well and provides the same insight into the dynamics of incentives.