

Targeting in Tournaments with Dynamic Incentives*

Martino Banchio Evan Munro[†]

October 26, 2020

Abstract

We study the problem of a planner who wants to reduce inequality by awarding prizes to the worst contestants in a tournament without incentivizing shirking. We prove that no ex-post targeting mechanism eliminates perverse incentives and show that the optimal dynamic rule is computationally infeasible. We design an approximately optimal, incentive-compatible mechanism that targets low-ranked contestants based on the tournament’s history up to an endogenous stopping time. We describe applications to eligibility for remedial education, retraining benefits for the unemployed, and draft lotteries in sports. Using data from the NBA, we show how our mechanism aligns incentives and improves targeting.

Keywords: Incentive-Compatible Mechanism, Tournaments, Targeting

JEL Codes: D82, H23, 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 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

In this paper we study the incentive-compatible design of mechanisms that target ‘equalizing’ prizes to the lowest ranked agents in a tournament. There are a variety of tournaments with non-monotonic prize structures, with awards for the top-ranked contestants, but where a planner also targets some secondary resources to the lowest-ranked contestants. Non-monotonic prize structures result in incentive issues in later stages of a tournament where contestants who are eliminated from the competition for top awards compete for the secondary prizes by intentionally underperforming. We show that it is possible to implement a non-monotonic prize structure in expectation without introducing incentive issues. Our first main result in this setting is that an incentive-compatible targeting mechanism must rely on the total history of the tournament, not just final scores. Second, we show that the mechanism should adapt the amount of targeting of secondary prizes based on the history of the tournament. Improved targeting is possible in histories of a tournament where competition is close and incentive compatibility constraints are slack compared to histories where the performance of contestants diverge early and incentive compatibility constraints are tight.

This analysis is motivated by a variety of settings in education, labor, and sports where prizes are awarded both to the highest and lowest-ranked contestants. In the major U.S. sports leagues, the top prize of the regular season tournament is qualification for the playoffs, which comes with media exposure and additional ticket sales. The draft lottery allocates the best newly eligible players to the bottom ranked teams, but introduces perverse incentives for teams to lose intentionally. Teams that intentionally lose in order to receive a better draft pick are colloquially known as ‘tanking’. Tanking is frequently and passionately discussed in the media and has been documented empirically in Taylor and Trogdon (2002) for the National Basketball Association (NBA) and Fornwagner (2019) for the National Hockey League (NHL). 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.

The allocation of an equalizing prize to the lowest-ranked contestants in a tournament, however, is not specific to the sports industry. The objective to support lowest-ranked competitors and the associated incentive issues also appear in a variety of social policies. In classroom settings, educators and administrators often devote extra resources to bottom-ranked students or under-performing schools. Examples of these resources include turnaround funding for schools, extra tutoring programs available in a variety of U.S. public school districts and valuable regrade options offered to failing students. In the U.S., the Adult Program of the Workforce Innovation and Opportunity Act (WIOA) provides funding for a vocational or other training course to a limited number of unemployed individuals that show the least potential of obtaining

full-time employment. This can provide incentives to perform poorly on skills tests administered by career centers in order to become eligible for the retraining benefit.

To address the optimal design of the draft lottery in sports leagues and other related problems in the design of eligibility criteria for social programs, we formalize the general problem of allocating prizes to the lowest-ranked individuals in a tournament without elimination. We formulate the designer objective, which is to maximize the probability that workers ranked below a cutoff receive an equalizing prize, constrained by an incentive-compatibility condition for each agent in every period in the tournament. We set up a simple model of worker decision-making that provides a framework for analyzing the effect of the prize allocation design on incentives for workers to exert effort. We then derive two main contributions which highlight the dynamic nature of incentives to exert effort. First, we show that any ex-post lottery based on final rankings that favors the lowest-ranked workers, such as the lotteries in the NBA, or the grading system in schools, provides incentives for a worker to shirk in some possible history of the tournament. This negative result suggests we should look at mechanisms that adapt the allocation probabilities over time. Second, we show that the globally optimal mechanism is the solution to an optimization program with an infeasible number of constraints. We propose a rule, which we call T-IC, that converges to the optimum and is feasible to implement. Our mechanism updates the secondary prize allocation probabilities after each period so that they match the probability each contestant will be ranked below a certain cutoff at the end of the tournament, conditional on results from periods played so far. This adjustment process ends as soon as incentive-compatibility is violated for any contestant. T-IC is computationally feasible since it separates incentive compatibility and targeting; the stopping time satisfies the incentive constraints so that they do not affect the adjustment of the allocation probabilities.

We show this process is approximately optimal since it minimizes a version of the designer objective that is conditional on all information up to the stopping time. Furthermore, this approximate objective converges to the true designer objective as the stopping time converges to the end of the tournament. In a small simulation, where it is possible to enumerate all possible outcomes of a tournament and solve the non-linear program directly, we show that this approximately optimal rule performs nearly as well as the globally optimal rule. In the simulation, T-IC reaches 73.5% of the performance of the globally optimal rule, as measured by the gain in targeting of the lowest-ranked contestant compared to the ex-post uniform lottery.

We then describe how our theoretical framework applies to three practical settings: remedial programs in education, retraining programs for the unemployed in the U.S., and the allocation of new players to teams in the NBA. Our first theorem indicates that

remedial programs with eligibility criteria based on ex-post cumulative performance can result in incentive issues for students and school administrators. We argue that using performance over time to partially restrict eligibility for remedial programs could reduce incentive issues while still targeting struggling students and schools. In contrast, the eligibility requirements of the WIOA already take into account a partial work and skill achievement history of workers, which qualitatively matches the process suggested by our allocation mechanism T-IC. 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 Theorem 1, was incentive-compatible. As a result, in our model, team effort and records would not change if T-IC, another incentive compatible rule, was in place instead of the ex-post uniform lottery. 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 is without trading off any violation in incentive compatibility for any team.

Before introducing the model, the next section provides a brief review of the related literature in tournament theory, social program eligibility design, and sports economics.

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).

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, where a worker's ranking in output, rather than their actual output, determines their prize allocation. As a result, the optimal prize structure is decreasing in ranking. Rosen (1986) investigates sequential elimination contests, where the optimal prize structure takes into account how incentives change as a contestant progresses in the tournament. As in Rosen (1986), we have a sequential series of periods with repeated choices of effort and incentives that change over time, however in our setting there is no elimination. The main point of departure from standard tournament theory is that the designer does not fully control the prize distribution. In our setting we assume that there are

two prizes: a primary prize that has a fixed value outside of the designer’s control and is assigned to the top v^* contestants, and a secondary prize that the designer would like to assign with the highest possible probability to the bottom d^* contestants. 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 contestant 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 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 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, 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).

We consider two kinds of applications of the theory: the first application provides insights on incentive-compatible design of eligibility requirements for social programs, and the second proposes a new design for the NBA draft allocation mechanism. There is related work 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. This includes the literature on tagging, see Akerlof (1978) and Allcott et al. (2015), and the optimal complexity of the application process, see Kleven and Kopczuk (2011).

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. Furthermore, Gold’s rule would assign zero probability in the draft to a team that loses every game in the season; thus, it does not target our assumed league objective. 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.

3 Designing a Targeting Mechanism

3.1 Defining Designer and Worker Objectives

First, we define some variables and notation. A tournament is made up of a total of T periods. In the sports setting, each period is a game between two teams, while in more general settings in any period some subset of workers can compete in a task with a certain probability of successful completion. Even if the task is not a sports game, we will refer to a worker successfully completing a task as a winning worker, and any unsuccessful workers will be called losing workers. There are a total of n workers in the tournament and worker $i \in I = \{1, \dots, n\}$ has ability $\alpha_i \in [0, 1]$.

A history is a sequence $S = \{S^t\}$ where S^t is a n -dimensional vector such that $S_i^t \in \mathbb{N}$ is the score of worker i at the end of time t . The score is simply the cumulative number of tasks successfully completed by the worker, but it can be interpreted more broadly as any measure of performance used to rank employees, students, or teams. Feasible histories need to satisfy some general constraints:

- Score Consistency: $S_i^t \geq S_i^u$ for any $u \leq t$.
- Initial Condition: $S_i^0 = 0$ for every $i \in I$.

The function $w : \{1, \dots, T\} \rightarrow 2^I$ selects the subset of workers called to play in each period. The possible outcomes in each period are a set of vectors $\mathcal{O}_t \subset \{0, 1\}^n$, where $O_{it} \in \{0, 1\}$ indicates whether worker i successfully completed the task in period t . For example, in the sports setting, $w(t)$ selects the two teams that play against each other in game t . \mathcal{O}_t is then composed of two vectors, with $O_{it} = 1$ and $O_{jt} = 0$ or $O_{it} = 0$ and $O_{jt} = 1$. $\mathcal{O}_t = \{(0, \dots, 0, 1, \dots, 0), (0, \dots, 1, 0, \dots, 0)\}$, since only one of the two teams can win each game. Let the set of feasible histories be \mathcal{S} : a tournament is a tuple $(\mathcal{S}, w, \{\mathcal{O}_t\}_{t=1, \dots, T})$.

We assume that there are v^* primary prizes with value π_i^V for worker i , and they are allocated to the best v^* performers. 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. There are also d^* indivisible secondary prizes. In the NBA these represent the draft picks and have value of at most π_i^D , which represents the increase in long-term expected revenue that a team expects to receive from drafting the top eligible player. While the tournament primary prize is taken as given, the planner would like to target the d^* lowest-ranked performers with the secondary prize, perhaps in order to reduce long-term inequality. The assumption

that primary prizes are out of the control of the planner reflects for example market-determined prizes, as in the labor market interpretation of our model. We assume $0 < \pi_i^D \leq \pi_i^V$ and that the incentive problem is non-trivial, so $v^* + d^* < n$: workers ranked between v^* and $n - d^*$ are worse off than workers above v^* or below $n - d^*$. We assume that both vectors π^V and π^D are common knowledge. Where it is unrealistic to assume that contestant-specific values for the prizes are common knowledge, we can instead assume that a lower bound $\bar{\pi}^V$ on the value of the primary prize and an upper bound $\bar{\pi}^D$ on the value of the secondary prize is common knowledge.

The workers, once faced with a task, can choose how much effort $e_{it} \in [0, 1]$ to exert. The task is completed with probability $p_{it}(\alpha_{w(t)}, e_{w(t),t})$, where $\alpha_{w(t)}$ and $e_{w(t),t}$ are vectors of size $|w(t)|$. α represents the contestants' abilities, which may be heterogeneous. Heterogeneity in α reflects anticipated differences in performance: we can normalize initial scores to 0 and adjust ability values to accommodate for handicaps. Notice that $p_{it}(\alpha_{w(t)}, e_{w(t),t}) = \mathbb{P}(O_{it} = 1 | \alpha_{w(t)}, e_{w(t),t})$. We do not take a stance on the functional form of the underlying probability law on the outcome space, except for assuming that the probability p_{it} is increasing in both α_i and e_{it} . Additionally, we will assume that, conditional on the outcome of worker i , her effort does not affect the probability of success of other workers. This assumption is weaker than unconditional independence and is satisfied by every environment we describe. At any point in the tournament, the probability that a worker will be ranked high enough at the end of the contest to obtain the primary prize depends on what has happened so far, and the probability of every possible outcome in the future. The ranking of a worker i is defined as follows, where the worker with the maximum score has rank 1:

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

Define $r(S^t)$ as the vector $(r_1(S^t), \dots, r_n(S^t))$. For worker i who competes in period t , $W_{it} \in \mathcal{O}_t$ is a random vector recording a successful task for worker i and stochastic outcomes for every other worker. The distribution of W_{it} will depend on the choices of effort of each worker in $w(t)$. Similarly, L_{it} is a random vector recording an unsuccessful task for worker i . For instance, if each worker competes in each period, W_{it} is simply a random vector in $\{0, 1\}^n$ that always takes value 1 in position i . In the NBA setting, where only two teams compete against each other, W_{it} is deterministic and is a vector with value 1 in position i and value 0 everywhere else. We denote the probability of winning the primary prize in the tournament for team i given a history S^t as $q_i(S^t)$. These probabilities can be defined recursively:

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

If a worker is called to play in period t , using the tower property, probabilities take the form:

$$q_i(S^{t-1}) = p_{it}(\alpha_{w(t)}, e_{w(t),t})\mathbb{E}[q_i(S^{t-1} + W_{it})] + (1-p_{it}(\alpha_{w(t)}, e_{w(t),t}))\mathbb{E}[q_i(S^{t-1} + L_{it})]$$

where the expectations are taken with respect to the distributions of W_{it} and L_{it} .

In general, we consider a secondary prize allocation mechanism of the form:

$$y: \mathcal{S} \rightarrow \underbrace{\Delta^n \times \cdots \times \Delta^n}_{d^*}$$

such that $y_i^k(S^t)$ is the probability that worker i receives the secondary prize targeted to the k -th ranked player 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 $t - 1$:

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

which implies

$$y_i^k(S^{t-1}) = p_{it}(\alpha_{w(t)}, e_{w(t),t})\mathbb{E}[y_i^k(S^{t-1} + W_{it})] + (1 - p_{it}(\alpha_{w(t)}, e_{w(t),t}))\mathbb{E}[y_i^k(S^{t-1} + L_{it})] \quad (\text{DC})$$

Moreover, the lottery probabilities at any history need to add up to 1:

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

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

Each worker makes a single strategic choice in each period, which is how much effort to exert. Since we assume in the following section no explicit cost of effort, then the efficient level of effort has every worker exerting maximum effort in every period. In order to derive the efficient mechanisms that always incentivize effort, we first explicitly define the workers' and planner's incentives.

Worker Objective In period t , worker $i \in w(t)$ chooses an effort level e_{it} to maximize his expected payoff given the results so far, S^{t-1} . We are interested in equilibria where each player exerts full effort in each period: this allows us to suppress the de-

pendence of $p_{it}(\alpha_{w(t)}, e_{w(t),t})$ for ease of exposition. Following the one-shot deviation principle, worker i in period t in such an equilibrium faces the following objective:

Optimization Problem 1.

$$\begin{aligned} \max_{e_{it}} \quad & p_{it} \left(\mathbb{E}[q_i(S^{t-1} + W_{it})] \pi_i^V + \sum_{k=n-d^*+1}^n \mathbb{E}[y_i^k(S^{t-1} + W_{it})] \pi_i^D \right) \\ & + (1 - p_{it}) \left(\mathbb{E}[q_i(S^{t-1} + L_{it})] \pi_i^V + \sum_{k=n-d^*+1}^n \mathbb{E}[y_i^k(S^{t-1} + L_{it})] \pi_i^D \right) \end{aligned} \quad (1)$$

Maximizing the worker's objective, we derive a necessary condition for worker i to exert maximum effort in period t :

Incentive Condition.

$$\begin{aligned} \mathbb{E} \left[q_i(S^{t-1} + W_{it}) - q_i(S^{t-1} + L_{it}) \right] \frac{\pi_i^V}{\pi_i^D} \geq \\ \sum_{k=n-d^*+1}^n \mathbb{E} \left[y_i^k(S^{t-1} + L_{it}) - y_i^k(S^{t-1} + W_{it}) \right] \end{aligned} \quad (\text{IC})$$

If this inequality is satisfied, it is optimal for worker i to exert maximum effort in the period: the inequality implies that the derivative of the objective with respect to effort is increasing everywhere. IC has a clear interpretation. If the increase in the probability of receiving the secondary prize when losing compared to winning a certain task is less than the decrease in the probability of receiving the primary prize, scaled by the ratio of the prizes' values, then the worker will exert maximum effort. Note that we assume that if the worker is indifferent between exerting effort and shirking, they will exert effort, since 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 to win or lose. 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 allocation rules that make certain workers better off if they lose rather than win. We discuss in Appendix A how results are affected by the assumption of an explicit cost of effort.

A rule $y \in \mathcal{Y}$ is incentive-compatible if IC is satisfied for all workers $i \in I$ for all possible histories $S \in \mathcal{S}$. Another implication of IC is that any allocation policy that is incentive compatible should not change the probability that a worker receives the secondary prize in periods 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 periods is always zero. This leads us to our first result:

Theorem 1. When $T > 2$ and $n > 2$, the only secondary prize allocation mechanism y that

1. is a function only of a worker's final ranking $r_i(S^T)$,
2. there exists a $l > v^*$ such that

$$\sum_{k=n-d^*+1}^n y_{i(l)}^k(S^T) \leq \sum_{k=n-d^*+1}^n y_{i(l+1)}^k(S^T)$$

where $i(j)$ is such that $r_{i(j)}(S^T) = j$

3. satisfies IC at every history $S^t \in \mathcal{H}$,

is a uniform lottery, which assigns equal probabilities to every worker i with rank $r_i(S^T) > v^*$.

Proof. Let y be a feasible mechanism satisfying conditions 1 – 3. There exists some partial history S^{T-1} such that

- worker i and j are tied at position l , before the final period T is played.

$$r_i(S^{T-1}) = r_j(S^{T-1}) = l$$

and $i \in w(T-1)$. By condition 2, this implies if worker i loses period T , then his lottery probability is (in expectation) higher than if he wins period T :

$$\sum_{k=n-d^*+1}^n \mathbb{E}[y_i^k(S^{T-1} + L_{it}) - y_i^k(S^{T-1} + W_{it})] \geq 0$$

- worker i is eliminated from playoff contention before this final period T is played. Since $l > v^*$, $v^* < n-1$, and $T > 2$ there is always some history where this is true.

$$q_i(S^{T-1}) = q_j(S^{T-1}) = 0$$

Then, we have for worker i :

$$0 = \mathbb{E}[q_i(S^{T-1} + W_{it}) - q_i(S^{T-1} + L_{it})] \frac{\pi_i^V}{\pi_i^D} \geq \sum_{k=n-d^*+1}^n \mathbb{E}[y_i^k(S^{T-1} + L_{it}) - y_i^k(S^{T-1} + W_{it})] \geq 0$$

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 workers ranked lower than v^* . \square

This means that no matter how carefully the planner chooses the weighted lottery used to assign the secondary prize, if the lottery weights favor lower ranked workers and are not equal for all the losers, then there will always be incentives for workers to shirk after they are eliminated from primary prize contention in certain histories. A simple deduction from this theorem is that any incentive compatible targeting rule cannot take into account losses occurred after a worker has been mathematically eliminated from primary prize contention, as observed by Lenten et al. (2018) in the NBA context. In order to design a better rule, we need to define what the planner’s preferred secondary prize allocation policy is. We recognize that the true planner objective in different environments is likely complex and may take into account incentives for workers to engage in a variety of unwanted strategies. The common thread in all our applications, however, is that the planner would like to target the equalizing secondary prize to the lowest-ranked contestants, and as a result we choose a simple objective that reflects that. Justifying this objective formally is beyond the scope of the paper. Informally, 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. This also applies to policymakers who may want to set education and unemployment policy to maximize long-term competitiveness in the labor market. Given this justification for the planner objective through repeated tournaments, we focus on how to satisfy the objective within a single tournament.

Planner Objective The planner’s objective is to maximize the probability that the lowest-ranked performers obtain the secondary prize. To that end, we define the random variable $\mathbb{1}[r_i(S^T) = k]$. It is a random variable since S^T is a random variable, and takes value 1 whenever the final score is such that worker i is the k -th performer. The conditional probability of this event, $Pr(r_i(S^T) = k | S^t)$, evolves as the scores evolve over time. We assume that the planner would like to construct a dynamic allocation rule ex ante that will adjust in a way that minimizes the distance $|\mathbb{1}(r_i(S^T) = k) - y_i^k(S^T)|$ in expectation for each i and k , while maintaining IC in every period.

Optimization Problem 2.

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

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

$$\begin{aligned} & \mathbb{E} \left[q_i(S^{t-1} + W_{it}) - q_i(S^{t-1} + L_{it}) \right] \frac{\pi_i^V}{\pi_i^D} \\ & \geq \sum_{k=n-d^*+1}^n \mathbb{E} \left[y_i^k(S^{t-1} + L_{it}) - y_i^k(S^{t-1} + W_{it}) \right] \end{aligned} \quad (\text{IC})$$

The planner wants to maximize the probability that the lowest-ranked performing workers receive the secondary prize, without providing incentives for workers to shirk. If it was not possible for workers to exert low effort, the planner would simply assign the secondary prize to the workers with the lowest rank at the end of the season.

This maximization problem is a linear program, but 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 $T = 1,230$. This requires specifying more constraints than there are atoms in the known universe, so calculating the optimal solution via linear programming is not computationally feasible. In fact, it is problematic to even make sure that a given rule is incentive compatible and feasible: it requires checking an exponential number of constraints. We propose a rule that is computationally feasible. 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 estimated in each period without enumerating all possible future histories. We show our rule satisfies IC in every possible history while targeting an approximate version of the objective (2). In Section 3.3 we show how well our proposed rule approximates the optimal solution in a small simulation where specifying all constraints in the linear program and minimizing the objective directly is feasible.

3.2 An Incentive-Compatible Targeting Mechanism

Theorem 1 indicates that a lottery based on final rankings that favors the worst workers will not satisfy IC. We propose instead a weighted lottery based on workers' records

for the first $t^*(S^T)$ periods in the tournament. Both the subset of periods accounted for and the lottery weights are determined dynamically as the tournament progresses. We call this mechanism the Incentive Compatible Targeting Rule (T-IC).

Since $Pr(r_i(S^T) = k | S^t)$ is worker i 's probability of having the k -th ranked score after all T periods, given the results from the first t periods, we define the probability that worker i receives the secondary prize after game t in T-IC as

$$y_i^{k,T-IC}(S^t) = \begin{cases} Pr(r_i(S^T) = k | S^t) & \text{if } t \leq t^*(S^t) \\ Pr(r_i(S^T) = k | S^{t^*}) & \text{if } t > t^*(S^t) \end{cases} \quad (3)$$

which is the probability worker i ends up in the k -th position given the results of all games up to $t^*(S^t)$. We next define the stopping time $t^*(S^t)$. Consider the following quantity:

$$B(i, s) = \sum_{k=n-d^*+1}^n \left[Pr(r_i(S^T) = k | S^s + L_{it}) - Pr(r_i(S^T) = k | S^s + W_{it}) \right] \pi_i^D - \left[q_i(S^s + W_{it}) - q_i(S^s + L_{it}) \right] \pi_i^V$$

The quantity $B(i, s)$, derived from rearranging IC, is the difference in utility for worker i from losing compared to winning in period $s + 1$. We cannot continue to adjust the targeting probabilities according to the updated conditional probabilities that workers will be ranked in the last d^* contestants when this benefit is positive. Let

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

The time $t^*(S^t)$ is the smallest $s < t$ such that given a history S^s , IC would be violated at time $s + 1$ if the targeting probabilities continued to be updated based on the conditional probabilities that workers end up last in the tournament. When there is no such $s < t$, we let $t^*(S^t) = t$.

The rule dynamically adjusts the lottery odds as the tournament progresses. When abilities are equal, a worker's probability of receiving a secondary prize before the tournament begins is $\frac{1}{n}$ since each worker has the same probability of ending up in position k . If instead there is an initial level of inequality, the ex-ante probability takes this into account. As wins and losses are recorded in each period, a worker's probability of getting a secondary prize adjusts based on their updated conditional probability of ending up k -th ranked. For example, when a worker loses in multiple periods early in the tournament, they increase their probability of ending with a low rank and their probability of receiving the secondary prize increases. This adjustment process permanently stops before IC would be violated.

In Theorem 1 we showed the planner will not be able to use all the information revealed during the tournament. Our incentive compatible targeting rule instead takes into account only a fraction of the information revealed. We will prove in what follows that our rule is optimal for an approximate version of the league objective limited to a smaller information set and that, for $t^* \rightarrow T$, T-IC approaches the global optimum¹.

Theorem 2. *T-IC is feasible and satisfies IC. Moreover, for $t^* \rightarrow T$ we have that*

$$y_i^{k,T-IC}(S^T) \rightarrow \begin{cases} 1 & \text{if } r_i(S^T) = k \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

which is the global minimum of the planner objective, and for $t^* \rightarrow T$

$$\mathbb{E} \left[\sum_{k=n-d^*+1}^n \sum_{i=1}^n \left| \mathbb{1}(r_i(S^T) = k) - y_i^{k,T-IC}(S^T) \right| \middle| S^{t^*} \right] \rightarrow 0 \quad (5)$$

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

- PROB is satisfied by definition, since $y_i^{k,T-IC}(S^t)$ are conditional probabilities.
- DC: The secondary prize probability change from time $t - 1$ to t is either:
 - If $t > t^*(S^T)$, there is no change, so DC holds
 - If $t \leq t^*(S^T)$, then the draft probabilities correspond to updates in a conditional probability, so DC holds:

$$\begin{aligned} Pr(r_i(S^T) = k | S^{t-1}) &= p_{it}(\alpha_{w(t)}, e_{w(t),t}) Pr(r_i(S^T) = k | S^{t-1} + W_{it}) + \\ &+ (1 - p_{it}(\alpha_{w(t)}, e_{w(t),t})) Pr(r_i(S^T) = k | S^{t-1} + L_{it}) \end{aligned}$$

- IC: Suppose by contradiction that IC was violated at some time t . This is equivalent to $B(i, t) > 0$. But then, T-IC fixed the secondary prize probabilities at some time $s < t$. This implies the RHS of IC is zero, and since the LHS is always weakly positive it cannot be that IC was violated.

With respect to the limit, note that the random variable $\mathbb{1}(r_i(S^T) = k)$ is in L^1 . Once we observe that the S^T are the terminal histories, we can apply Levy's zero-one law to the filtration generated by the histories S^t to obtain

$$\lim_{t^* \rightarrow T} Pr(r_i(S^T) = k | S^{t^*}) = Pr(r_i(S^T) = k | S^T) = \mathbb{1}(r_i(S^T) = k)$$

¹While $t^*(S^T)$ depends on the history, for notational simplicity we will often omit the dependence.

Since by definition $y_i^{k,T-IC}(S^{t^*}) = Pr(r_i(S^T) = k | S^{t^*})$, we proved the statement in (4).

Now we turn to approximate optimality. The designer's objective is ex-ante, and is an expectation over all possible histories, and infeasible to compute and to optimize. His optimization is consistent with information arrival, in the sense that the following expectation is minimized conditional on some observed information S^{t^*} :

$$\mathbb{E} \left[\sum_{k=n-d^*+1}^n \sum_{i=1}^n \left| \mathbb{1}(r_i(S^T) = k) - y_i^k(S^T) \right| \middle| S^{t^*} \right] \quad (6)$$

Let us focus on this objective and consider only rules that are constant after S^{t^*} . In general this problem is still infeasible, since it requires enumeration of all the possible continuation histories. However, it's simple for the seller to minimize the following problem:

$$\sum_{k=n-d^*+1}^n \sum_{i=1}^n \left| Pr(r_i(S^T) = k | S^{t^*}) - y_i^k(S^{t^*}) \right| \quad (7)$$

It is clear that this objective is set to 0 by T-IC. Now, notice that

$$\begin{aligned} & \lim_{t^* \rightarrow T} \sum_{k=n-d^*+1}^n \sum_{i=1}^n \mathbb{E} \left[\left| \mathbb{1}(r_i(S^T) = k) - y_i^k(S^T) \right| \middle| S^{t^*} \right] - \\ & \quad - \left| Pr(r_i(S^T) = k | S^{t^*}) - y_i^k(S^{t^*}) \right| = \\ & = \sum_{k=n-d^*+1}^n \sum_{i=1}^n \lim_{t^* \rightarrow T} \mathbb{E} \left[\left| \mathbb{1}(r_i(S^T) = k) - y_i^k(S^T) \right| \middle| S^{t^*} \right] - \\ & \quad - \lim_{t^* \rightarrow T} \left| Pr(r_i(S^T) = k | S^{t^*}) - y_i^k(S^{t^*}) \right| = \\ & = \sum_{k=n-d^*+1}^n \sum_{i=1}^n \mathbb{E} \left[\left| \mathbb{1}(r_i(S^T) = k) - y_i^k(S^T) \right| \middle| S^T \right] - \\ & \quad - \left| Pr(r_i(S^T) = k | S^T) - y_i^k(S^T) \right| = 0 \end{aligned}$$

where for the second-to-last equality we applied Levy's zero-one law twice. In this sense then we know that, as $t^* \rightarrow T$, the tractable objective (7) approaches the original objective (6). Combining this with the fact that T-IC sets objective (7) to 0 yields the result. \square

The analysis of this specific rule helps identify some features of any rule that satisfies incentive compatibility. First, Theorem 1 shows that any incentive compatible mechanism cannot rely only on ex-post rankings, as many targeting mechanisms in tournament-like settings do, see Section 4. It is also not enough to depend on dynamic

rankings rather than ex-post rankings, as we prove in the Appendix. Any incentive compatible rule must depend on the full history of scores in the tournament. This property of incentive compatible mechanisms highlights the complexity of solving problem (2). The rule we propose takes care of the complex restrictions by decoupling the incentives from the optimality of the rule. The objective function is addressed by matching the targeting probabilities to the conditional ranking probabilities, while incentives are taken care of separately by 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 weighted lottery over all workers. Conditioning the allocation mechanism on more wins and losses decreases the value of the objective function. If we ignored the restrictions posed by IC, we could condition on the full history S^T . With a stopping time of T , T-IC would allocate the secondary prizes to the bottom d^* workers with probability 1, as shown in Theorem 2. This yields the minimum possible value of the planner objective. However, the incentive compatibility requirements places us somewhere 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. 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 some small simulations showing how T-IC compares to the globally optimal, ex-ante, and ex-post IC rules.

3.3 Simulations

We use a simple setting of $n = 3$ workers, where workers have equal ability, the top worker receives the primary prize, and there is only one secondary prize available. This is the smallest number of workers such that there are incentive issues. With $n = 2$ workers, there are no incentive issues since a worker that is eliminated from primary prize contention receives the secondary prize. We proved that our allocation rule is optimal in a restricted sense; it maximizes a version of 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 worker.

We simulate a tournament with $n = 3$ workers, $T = 12$ periods and $|w(t)| = 2$ for all t : additionally, we constrain the possible histories so that only one team out of the two playing wins each game. This models a small sports tournament in which teams compete pairwise against each other in a total of 12 games. In this tournament it is feasible, though computationally intensive², to enumerate all possible 2^{12} outcomes of the tournament and calculate the globally optimal incentive-compatible mechanism as the solution to the non-linear program in Optimization Problem 2. We report the designer objective divided by two, which is equivalent to the total variation distance between allocating the lowest-ranked team the secondary prize with probability 1 and the probabilities given by the allocation rule, averaged over all possible histories. We calculate this for four different incentive-compatible mechanisms :

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

Mechanisms 2-4 are always incentive compatible. Mechanism 1 is incentive compatible as long as $\frac{\bar{\pi}^V}{\bar{\pi}^D} \geq 1$.

Figure 1a shows how the value of the designer objective varies with $\frac{\bar{\pi}^V}{\bar{\pi}^D}$. The ex-ante uniform rule gives the equalizing prize with expected probability 33% to the lowest-ranked agent. The ex-post uniform rule gives the equalizing prize with expected probability 50% to the lowest-ranked agent, but is only incentive compatible when $\frac{\bar{\pi}^V}{\bar{\pi}^D} \geq 1$. For low values of $\frac{\bar{\pi}^V}{\bar{\pi}^D}$, neither the globally optimal rule nor T-IC does 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 from 2 to 10, the designer objective decreases rapidly for both T-IC and the globally optimal rule, although the globally optimal rule decreases faster. The rules take advantage of the fact that incentive constraints are more slack earlier in the tournament. Both rules converge to a constant value as $\frac{\bar{\pi}^V}{\bar{\pi}^D}$ increases beyond 10. The minimum possible value of the planner objective is 0. As expected from Theorem 1, neither the globally nor the approximately optimal rule can ever give the secondary prize with expected probability 1 to the lowest-ranked worker, no matter how large $\frac{\bar{\pi}^V}{\bar{\pi}^D}$ is. The globally optimal rule converges to a value of approximately 0.16 while T-IC converges to a value of approximately 0.25, so approximates well the global optimum compared to the uniform

²We use the Julia optimization package JuMP and the open source solver Ipopt to perform this optimization

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.

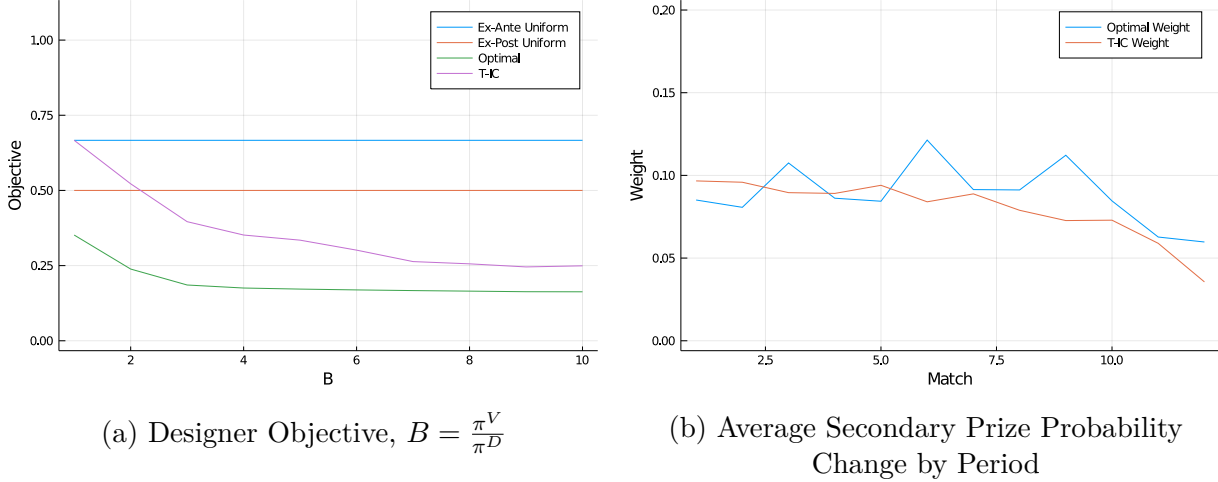


Figure 1: Comparing Alternative Incentive-Compatible Rules in Simulations

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 secondary prize allocation probability between period t and $t - 1$, averaged over all histories and workers. This can be considered the average weight that each mechanism places on each period. Both rules place more weight on earlier periods in the tournament, where incentive compatibility holds in more histories, since workers still are in contention for the primary prize. However, the globally optimal rule places more weight on periods towards the end of the tournament compared to T-IC. Our rule freezes the secondary prize probabilities for all workers once incentives for a single worker have been violated. The results for the globally optimal rule show there is some benefit in some histories of adjusting the secondary prize probabilities for some workers after incentives have been violated for a single worker. However, this benefit is quite small, since the two mechanisms have performance that is quite close as $\frac{\bar{\pi}^V}{\bar{\pi}^D}$ increases.

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

4 Policy Design

Our theory indicates that there exist computationally feasible mechanisms that can target low-ranked contestants in a tournament without affecting incentives to exert effort. However, they requires a dynamically-adjusted policy, as any ex-post targeting will result in incentive issues. In this section, we examine three different economic environments where our theory applies: education policy, labor policy, and sports policy. With the insights we obtained so far, we analyze the design of remedial education programs and unemployment retraining programs. Then, we provide a quantitative analysis of the performance of T-IC when it is applied to the design of the draft lottery for the NBA.

4.1 Targeting of Low-Performing Students and Schools

There are a variety of examples in education where additional resources are given to the lowest ranked contestants in a tournament-like setting. Schools are ranked on a variety of metrics including standardized test scores, dropout rates, and truancy rates. The Renewal program in New York City targeted low-ranked schools with extra funding for longer school hours and integrated social and health programs. The Turnaround Program in Nevada targets low-ranked schools with additional funding over a period of two years. The Bagrut program in Israel, analyzed in Lavy and Schlosser (2005), provided schools with a low matriculation rate with funding for remedial education. Remedial education programs are another example of targeting in tournaments, providing students with the lowest grades with additional support such as personalized tutoring.

We can apply the theory developed in this paper to any of these settings. Grades and school rankings are usually derived from continuous assessment that makes it possible to condition allocation on dynamic performance rather than ex-post rankings. Furthermore, there is a non-monotonic prize structure such that contestants just above the cutoff for additional support are worse off than contestants just below the cutoff. As a result, making eligibility for student remedial or school turnaround programs dependent only on ex-post rankings will introduce incentive issues. Our results suggest that by examining the history of a school or student's performance the planner can target the worst-performing contestants without introducing incentive issues.

As an example, consider the retake option offered by IIT Delhi. If a student fails a class at IIT Delhi, the failing grade does not affect their GPA, and the student can repeat the same class for credit. This policy introduces non-monotonicity: the best students improve their GPA with an A, the worst students leave their GPA unchanged through the retake option, while students receiving marginal passing grades see their

GPA drop. The retake option can be considered as the secondary prize in our model. The results from Theorem 1 indicate that if the retake option is based only on final grades, determined from relative standings at the end of the semester, there will be some students with an incentive to deliberately reduce effort on assignments later in the term. Theorem 2 also holds: it is possible to design an incentive-compatible, computationally-feasible allocation of the retake option. An incentive compatible rule would assign more weight to results in midterms and earlier assignments, when incentives constraints have more slack, when granting the retake option.

4.2 Labor Policy: Workforce Innovation and Opportunity Act

Labor markets offer another prominent example of targeting of low-ranked individuals. For example, the WIOA provides funding for career and training services for unemployed individuals through local job centers in the United States. Basic career services are offered to anyone who enters a job training center, and include information provision and some basic skills assessments. Individualized career services include more comprehensive skills assessments interviews with a career counselor. Access to training services, unlike basic career services, is limited to targeted groups, which includes laid off workers through the Dislocated Worker Program, and disadvantaged groups with poor work histories through the Adult Program. Budget constraints mean that access to training slots is a scarce resource allocated by staff. Training services include a wide range of programs, including programs operated by the private sector, occupational skills training that can lead to a post-secondary degree, skills and entrepreneurial training, and adult literacy activities. There is some discretion given in the eligibility requirements for the Adult Program, which means that the requirements vary depending on location. In San Diego, an individual qualifies if they are deemed to be unable to obtain employment with only basic career services, and their household income must be lower than 70% of the Lower Living Income Standard Level. Additional priority is given to individuals who receive public assistance or are basic skills deficient, as defined by either staff assessments or by lack of a high school diploma.³

This program is another example of a situation with relative performance evaluation and dynamic incentives. Our model provides a simplified representation of the economic environment and the structure of T-IC provides some guidance on what kinds of eligibility requirements can be less likely to provoke incentive issues in this environment. A tournament is the period of unemployment before an individual might apply to a job training center for assistance. The workers are the set of individuals in the

³See the San Diego Workforce Partnership Operations Manual (July 2020) for additional details

local labor market who are struggling to find work and are under-skilled. During their period of unemployment, individuals can complete a series of task with some probability of success in order to increase their chances of full-time employment; these include applying for interviews, volunteering, applying for part-time work, or networking. The relative ranking of workers in terms of success in these tasks is important in that we assume the top v^* workers are hired for full-time work without intervention, receiving π_i^V . Scarcity of job offerings and the tournament structure induce competition. The U.S. Department of Labor would like to target the bottom d^* workers with retraining programs, perhaps in order to prevent low-skilled individuals from leaving the labor force entirely. This leads to incentive issues: if the retraining program for a more in-demand industry is valuable enough, a worker who is struggling to find full time work after a 5 month period of unemployment might have the incentive to perform poorly on further tasks, such as skill assessments, to improve their chances of receiving retraining funding.

These concerns can be addressed while still targeting the lowest-skilled workers by examining the dynamics of the incentives. The key is to leverage the longer history available, and to weigh early negative performance (i.e. unsuccessful job applications or failed certification courses) more than performance at the end of the period to determine eligibility. The theory suggests why the WOIA eligibility requirements already include longer-history metrics in their evaluation process. These additional metrics not only reduce noise in assessment but also are less susceptible to incentive issues.

4.3 Sports Policy: Improving The NBA Draft Lottery

Our final example is the design of the draft lottery in the U.S. sports leagues, where there is frequent and vigorous discussion in the media on the incentive issues of the current ex-post lotteries. When teams lose intentionally towards the end of a season in order to receive a higher draft probability, it is known as ‘tanking’. We show that it is possible to apply T-IC directly to allocate the first draft pick in an incentive-compatible way and address the tanking issue. 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 and show that the rule performs well.

We use NBA data for this application, but the analysis is easily adapted to other leagues. The NBA tournament occurs from October to April. There are 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 fourteen teams participate in the draft lottery for the first draft pick.

In the notation and language of our model, $n = 32$, the tournament is a season and workers 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 easily 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

For the years from 1985-1989 when the NBA had a uniform lottery system, teams did not have an incentive to intentionally lose games and final rankings would not be expected to change under T-IC. 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 T-IC on real data. 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 $y_i^n(S^t)$: The probability any given team will be ranked last conditional on their record after each game $t = 1, \dots, t^*$:

$$Pr(r_i(S^T) = n | S^t)$$

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% probability⁴.

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 , where W_{it} is a vector with 1 for team i and 0 for all other teams if team i wins against opponent j , and L_{it} is a vector with 1 for their opponent team j and 0 for all other teams if team i loses against opponent j . S^s is the n -dimensional cumulative wins for each of the n teams for the first s games of the season.

$$B(i, s) = \left[Pr(r_i(S^T) = n | S^s + L_{it}) - Pr(r_i(S^T) = k | S^s + W_{it}) \right] \bar{\pi}^D - \left[q_i(S^s + W_{it}) - q_i(S^s + L_{it}) \right] \bar{\pi}^V$$

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

⁴It 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.

season from 1985-1989 and then 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	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

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 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.


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.



(a) Win Records



(b) Draft Probabilities

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

5 Conclusion

There are a variety of settings with relative performance where the planner is interested in assigning an equalizing prize to lower-ranked agents. In a static model with a single choice of effort, any kind of non-monotonic prize structure will affect incentives and reduce effort. We prove that in a dynamic model with multiple choices of effort, while no ex-post re-distributive policy can be incentive 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. Our rule targets an approximate version of the planner objective by adjusting allocation probabilities as wins and losses occur based on the updated conditional probability that each agent ends up ranked last after the tournament concludes, 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.

This design directly leads to an improved draft allocation mechanism for the NBA, where the lottery has been changed repeatedly over time without finding a satisfactory system that addresses ‘tanking’ while still supporting the worst teams in the league. We also show that on historical data from the NBA from 1985-1989, T-IC significantly outperforms the league lottery in place at the time without introducing perverse incentives. Furthermore, our results apply more generally to the optimal design of eligibility requirements for social programs. As technology and data analysis improves, it is increasingly feasible for government programs to use a more comprehensive set of data about individuals in order to determine eligibility. We clarify how using performance data on a contestant over time that is predictive of their final outcome, rather than purely their final rankings, can allow targeting in social programs without affecting incentives.

Incentive-compatibility constraints are more slack at the beginning of a period when uncertainty about final outcomes is higher, compared to the end of a tournament when final outcomes are close to determined. As a result, the dynamic nature of incentives should be taken into account in the optimal design of policies that target low-performing groups.

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.
- 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.
- 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.
- Lavy, Victor and Analía Schlosser**, “Targeted remedial education for underperforming teenagers: Costs and benefits,” *Journal of Labor Economics*, 2005, *23* (4), 839–874.
- 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 O’callaghan, 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.

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:

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

Then the worker's problem becomes

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

Worker i exerts maximum effort in period t if

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

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 2 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 secondary prize 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 Ranks are not sufficient

Lemma 1. *A feasible incentive compatible rule $y_i^k(S^t)$ such that there exists a function $g_i^k: I \rightarrow [0, 1]$ such that $g_i^k(r(S^t)) = y_i^k(S^t)$ for every S^t must be the ex-ante uniform rule.*

Proof. We proceed by contradiction: fix an incentive compatible rule $y_i^n(S^t)$ such that there exist a g_i^n as described above. Consider a history S^T such that the rank of player i is n , but the score of his predecessor j is $S_j^T = S_i^T + 1$ in history 2. Additionally, assume that i is the only worker called to play in period T . By assumption $y_i^n(S^{T-1}) = y_i^n(S^T)$ because the score at the end of period $T - 1$ was the same as in period T . This coupled with (DC) implies that also $y_i^n(S^{T-1} + W_{iT})$ is equal to $y_i^n(S^T)$. But then the secondary prize probability must be the same when i and j are tied in the second to last position. Iteration of this procedure shows that the secondary prize probability has to be the same also when i is ranked above j , and all the other players: the only rule that satisfies these requirements is the ex-ante uniform rule. \square