Time-bundled Contracts and Effort Decisions: Evidence from the NFL

Ivan Li

December 14, 2020

Abstract

We are interested in the effects that multiple levers have on effort incentives in a time-bundled contract. In particular, our setting provides multiple primary rewards for high performing individuals and an alternative reward for poorly performing individuals. Standard assumptions about zero effort and eliminating deadlines are not feasible in our setting because of the competitive nature of sports. We estimate a single-agent dynamic structural model of hidden effort choice in the context of the National Football League (NFL). The model builds upon current dynamic models in providing an identification technique for a hidden policy function using observed heterogeneity in a sports setting. We find evidence that confirms general intuition about tanking: less talented teams that have little chance of making the playoffs and more talented teams that are likely to make the playoffs tend to exert lower effort. Counterfactual results show that modifying the length of the contract, timing of rewards, and thresholds of the rewards have ambiguous effects on effort incentives. Removing the alternative reward improves effort, especially at lower states, but might be antithetical to the league's objective.

Keywords: Time-bundled contracts, Effort incentives, deadline effects, effort distortion, dynamic decision making, unobserved actions

1 Introduction

A hotly debated topic among executives and officials in the major US sports leagues has been the effect of the draft pick allocation mechanism on the effort exerted by teams. The amateur draft is the primary way that teams obtain new talent each season, generating fan and media interest especially in the NBA and NFL. Currently, all leagues implement some sort of mechanism which gives the worst performing teams in a season the top picks in the upcoming draft. These mechanisms are founded on the belief that the draft can redistribute talent to improve the worst teams. Otherwise, the rich get richer since better performing teams will drive more revenues and be able to afford higher quality talent. In the long run, leagues fear that such path and state-dependent effects will cause interest in their sport to deteriorate in certain markets.

However, the problem with draft mechanisms is that a single player may have a disproportionate impact on the trajectory of a franchise. Occasionally, teams will intentionally lose ("tank") to secure a draft position to select a highly coveted amateur player. Leagues find tanking problematic because tanking teams can be uncompetitive, ruining the quality of the league's product and the fan experience. The most famous example of tanking comes from the NBA, where in the mid 2010s the Philadelphia 76ers intentionally traded away all of their talent in order to collect top draft picks over multiple seasons. This behavior, famously dubbed "The Process", was very successful as the Sixers obtained three top draft picks over a span of four years and eventually reached the playoffs just a few years later. Unfortunately, "The Process" came with a glaring flaw; attendance at Sixers games was abysmally low and media attention of the team was extremely negative. NBA league officials were displeased with the optics surrounding "The Process" and forcibly removed Sixers General Manager Sam Hinkie from his duties midway through one of the tanking seasons.

Given the interest from both fans and league officials in tanking, we study the tanking problem in this paper through a single-agent dynamic structural model of discrete hidden effort choice in the context of the National Football League (NFL). Teams choose effort levels which impacts their probability of winning and making the playoffs. These decisions are made under uncertainty, since exerting effort does not guarantee a win and not exerting effort does not mean the team will certainly lose. A model with hidden action should represent the true effort allocation process of teams. Unlike Sam Hinkie and the 76ers, most teams do not want to face retribution from fans or the league by openly admitting to tanking. One of the goals of this paper is to measure the extent of tanking in sports and particularly, in the NFL, where draft picks are allocated in reverse order of standings at the end of the season. The measurement problem is important because it informs the importance of the mechanism design problem in economics. Papers such as Banchio and Munro (2020) are more concerned with the theoretical incentive-compatible draft mechanisms for the sports leagues and less focused on documenting the empirical existence of tanking.

Our paper is related to the marketing literature through three avenues. First, we contribute to the literature of dynamic structural models with hidden actions most frequently seen in the sales-force compensation literature. We model the effort decision differently, as a binary discrete choice rather than a continuous variable. This allows for a slightly different identification strategy than the ones in the literature. Our paper also addresses the issue of deceleration in time-bundled incentive contracts, like in rewards programs. In a general time-bundled contract, deceleration is usually observed when individuals cannot reach a specified threshold before the end of the contract. Our unique setting provides us multiple levers to modify effort, not just a deadline and a reward. For example, our setting naturally provides an *alternative reward* for the worst performing individuals. Having various adjustable parameters of the contract gives the contract designer more flexibility in shaping the effort profile and addressing deceleration. Lastly, we tackle the problem of product/program design. Time-bundled contracts generally act as a compensation mechanism to elicit effort. While our setting requires compensation, effort is a fundamental part of the product in sports competition, providing sports leagues another reason to try and maximize it.

Formally, our research question is concerned with studying the effects that multiple levers in time-bundled incentive contracts have on eliciting effort. We frame the NFL regular season as a time-bundled contract, and the parameters that define the league (number of games played, number of playoff teams, draft mechanism, etc.) can be thought of as the endogenous levers. We approach this question assuming that the league's objective is to maximize teams' effort. We largely ignore the talent redistribution objective from the league's perspective, but its existence is recognized by teams and they gain utility from obtaining a higher draft pick. To maximize teams' effort, the league must set the parameters of the season to optimally incentivize effort while minimizing effort deceleration. In other settings studied in this literature, it is possible to linearize the time-bundled contract as a means of minimizing deceleration. Often, this means removing any threshold which a reward is tied to. However, it is difficult to do so in sports because of the exogenous nature of competition. The competitive nature of sports necessitates the existence of a playoff or a championship game. It also requires that determining the participants of a playoff cannot be arbitrary. Therefore, our counterfactual simulations only change some parameters of the NFL regular season, keeping intact the exogenous competitive structure which the league cannot feasibly change.

Our results provide an empirical first step towards understanding the effects of multiple levers in time-bundled contracts. We present descriptive evidence documenting the dynamic effects that different states and talent levels have on the probability of winning. Our parameter estimates imply that an average team facing an average quality draft will exert effort about 60% of the time in the first game of the season. Teams tend to almost always exert effort in the last week of the season if they are on the threshold of making the playoffs or obtaining a first-round bye. Weak teams exert effort about a quarter of the time if they are eliminated from playoff contention and there are strong amateur prospects in the upcoming draft. Conditional on making the playoffs but not able to reach any another thresholds, teams exert effort about 40% of the time. Our first counterfactual tests the implementation of the new NFL league parameters in 2021-2022. This counterfactual changes the threshold rules and the length of the season. We find that these changes ambiguously affect effort; increasing the length of the season and allowing more playoff teams creates leeway in certain states of the world while incentivizing other states. Our second counterfactual replaces the current draft mechanism with a purely random lottery, essentially eliminating the alternative reward. Unsurprisingly, eliminating the reward for the worst performing teams increases effort across the board. Nevertheless, this counterfactual does not address the talent redistribution problem that the league faces. Our third counterfactual implements a "tanking deadline", shifting the realization of the alternative reward from the last period T to an earlier period $t < T$. Predictably, poorly performing teams will further decrease effort prior to the tanking deadline, but will increase their efforts post tanking deadline relative to the status quo. Our results seem to follow general intuition and anecdotal evidence about competitiveness and tanking in sports.

The paper proceeds as follows: section 2 dives deeper into the literature related to our study. Section 3 gives background of our setting, summarizes our dataset, and provides descriptive evidence supporting tanking. We introduce a general model of dynamic discrete choice effort allocation in section 4. Section 5 details the functional form assumptions and the estimation procedure of our model. Identification of our model is discussed in section 6. Section 7 presents the parameter estimates of our model and the results of our counterfactual simulations. We provide concluding remarks in section 8.

2 Literature Review

In this section we discuss other research relevant to our study. While the "time-bundled contract" has been studied in many empirical contexts, we first see this phrase used in a behavioral economics paper by Aggarwal et al. (2020). The authors run a field experiment in India attempting to incentivize exercise in diabetic patients. In one experimental group, the reward for reaching the exercise goal is "bundled" over time; instead of receiving rewards for every day they reach the goal, participants must reach the goal at least four or five times a week in order to receive their reward. This experiment uncovers a crucial attribute of the time-bundled contract: bimodality in the distribution of effort at the extremes. In other words, subjects under the time-bundled contract were most likely to reach their daily exercise goal either seven times a week or zero times a week. The time-bundled compensation was just as effective as daily compensation in incentivizing exercise, but the time-bundled contract comes at a lower cost to the policymaker. Our analysis differs because we care about why the contract is creating the bimodality in effort, not how it is creating the most efficient compensation scheme.

The study of time-bundled contracts also exists in the marketing and economics literatures. Loyalty programs are one setting where these contracts are quite popular. Hartmann and Viard (2008) study switching costs in a buy n get one free setting, where the threshold n and a deadline creates switching costs for individuals short of the threshold. Kopalle et al. (2012) find a "points pressure" created by deadlines in loyalty programs with different rewards tiers. Another marketing setting where time-bundled contracts are studied is sales force compensation. Chung et al. (2014) recover unobserved effort in a sales force setting, using the nonlinearity induced by the quota-based compensation design for identification. After obtaining estimates of sales force agents' preferences, they simulate counterfactuals to find optimal compensation plans. Our setting provides a unique challenge since the competitive nature of sports makes smoothing out the nonlinearity in compensation impossible. Studies such as the one done by Misra and Nair (2011) find optimal counterfactual policies (and in this particular paper, compensation policies implemented in the field) by removing compensation tied to thresholds such as quotas. Instead, we are restricted to using multiple endogenously defined compensation levers, such as the draft pick mechanism, to shape effort.

We provide an identification technique of single agent dynamic discrete choice models with hidden actions. The standard dynamic discrete choice models (cf. Rust 1987, Hotz and

Miller 1993, etc) all require that actions are observed in order to identify the state transition densities. More recent papers have been concerned with unobserved state variables, such as unobserved heterogeneity (Arcidiacono and Miller, 2011). The aforementioned sales force literature deals with challenges in identifying hidden actions most similar to ours. Most recently, Chung et al. (2019) identifies a hidden-action dynamic model of effort with beta-delta present bias preferences. There are two key differences between our paper and the sales force ones: first, we use a discrete effort choice rather than a continuous choice. Second, our identification technique does not rely on an assumption about zero effort or maximum effort. Because effort is a discrete choice, the unobserved state variable ensures that the probability of exerting effort is never exactly zero or 100%. A zero-effort state is not feasible in our setting because of the inherent nature of competition; professional athletes are not likely to face situations where they would all give up. Instead, we rely on observing some measure of true ability to identify the payoffs and cost of effort.

In a sense, we are related to the literature focused on the attributes of product design. The most noticeable work in this field is the work on conjoint analysis (Green and Srinivasan, 1978) which helps firms set product attributes using a series of revealed preference designs. Sports offers entertainment; its key attribute is the competition aspect. Therefore, leagues will want to create a contract structure that can maximize competitive effort. We explore how multiple levers in a time-bundled contract can affect effort, which inherently affects the degree of competition in sports.

Finally, we are related to the literature on sports economics and the growing literature on optimal draft allocation. Banchio and Munro (2020) propose a dynamic, incentive compatible draft allocation mechanism which they implement on data from NBA games. Their mechanism satisfies the objective of allocating the best draft pick to the worst team while addressing tanking. Lenten (2016) and Lenten et al. (2018) suggest that the team first eliminated from playoffs should receive the best pick in the draft. Our work is less concerned with the efficient allocation of draft picks and is more focused on empirically documenting the existence of tanking.

3 Data and Descriptive Evidence

3.1 Setting

The National Football League (NFL) is the premier professional American football league in the United States. There are 32 teams, split evenly into two conferences, the American Football Conference and the National Football Conference. Each conference is further split into four equal sized divisions - North, South, East, and West. Every team plays a 16-game regular season schedule over a span of 17 weeks starting in late September. Through the 2019-2020 NFL season, a total of twelve teams qualified for the playoffs - six from each conference, split into four division winners and two "wild card" teams that had the best remaining records. Starting in the 2020-2021 NFL season, the number of playoff teams increased to 14, with each conference adding a third wild card team. Our data only runs through the 2019-2020 NFL season, so we will proceed using the state of the league from that year.

In each conference, there are three single-elimination playoff rounds: the wild card, divisional, and conference rounds. The division winning teams with the two best regular season records get a "bye" in the wild card round and skip directly to the divisional round. The two winners of the conference rounds then face off in a winner-take-all Super Bowl to determine a champion of that season.

For teams that do not make the playoffs, their season ends after they have played their sixteenth and final regular season game. In between seasons, teams are able to modify their player personnel through a few means. The first is through the seven round NFL Draft, where teams are able to select amateur prospects. The order of the NFL Draft is determined by the inverse of the regular season record. That is, the team with the worst record will select first. Another way of adding personnel is through free agency - players whose contracts have expired are able to sign contracts to play for other teams. Finally, teams can trade their players for other players or draft picks. The ability to trade also exists in the regular season up until the trade deadline (Week 8 of the NFL season).

3.2 Data

Our data consists of information about all NFL games played from the 2009-2010 through the 2019-2020 NFL seasons. In this dataset, we observe all teams, scores, dates, locations, weather and betting spreads of every game played. We pool the observations from these 11 regular seasons and treat each team-year combination as an individual team. In all, there are 5632 observations of 352 team seasons. To construct α_i from the data, we use a Quarterback-adjusted ELO[1] from FiveThirtyEight. Note that this isn't a perfect measure of ability; there tends to be stickiness in talent year to year, so the ELO prior to week 1 is most reflective of the performance of the teams as of the last game they had played in the prior season. Our construction also also implicitly assumes that the most talented team will have a win probability of $p(0)$ if they do not put in effort.

Additionally, we model observed draft strength ds as a number on the interval $[0, 1]$. To create this, we take projected prospect grades from ESPN's archive of NFL drafts from 2009-2019 and convert the prospect grade to a number ds. We use a weighted calculation of the best quarterback prospect rating and the best overall prospect rating in the draft. The quarterback prospect rating is heavily weighted because according to almost all NFL pundits, it is single most important position in the NFL. These draft ratings give us an ex-ante view of what teams believe the quality of the draft will be.

3.3 Descriptive Evidence

Below, we provide some descriptive evidence that different effort levels are exerted by different NFL teams. Figure 1 shows the probability of winning the next game in weeks 2, 5,

[1] A rating system used in chess (Wikipedia Link)

12, and 17. Each colored line represents a different talent quartile given by FiveThirtyEight's NFL QB-adjusted ELO. The $x-$axis in each grid is the win percentage entering the week, which should be interpreted as the state of the teams. In week 2 (top left), ELO quartile seems to explain all of the variation in win percentages and the state of the team does not appear to matter much. As soon as week 5, ELO becomes less predictive of the win probability, and the state of the world gains much more explanatory power. At this point in time, every team still has a mathematical chance of making the playoffs; even so, the talented teams at lower states seem to have given up already. We see this upwards relationship between current win percentage and probability of winning the next game continuing through the season, including in week 12. In week 17, talented teams at higher states seem less likely to win. The anecdotal evidence suggests that these teams have already clinched the playoffs or a first-round playoff bye, giving them less incentive to put in effort. Teams are most likely going to win if they have won between $40-60\%$ of their games heading into week 17; i.e, the group of teams most likely to win are the teams who have not been eliminated from playoff contention in the last period. The main takeaway from these plots is that a dynamic context is necessary to explain the patterns in the data. There seems to be a change in effort at different time periods as well as in different states of the data. We introduce our a general version of our dynamic model in the next section.

4 Model

There are n players, $i \in I = \{1, \ldots, n\}$. Each player has an observed heterogeneity parameter, α_{it} which affects their individual performance. There are $T-1$ periods where effort can be invested. In these periods players can make a discrete effort decision $e_{it} \in \{0, 1\}$, and an output $w_{it}(e_{it}, \alpha_{it})$ is realized conditional on effort and observed heterogeneity. Players realize a final, state-contingent payoff at $t = T$.

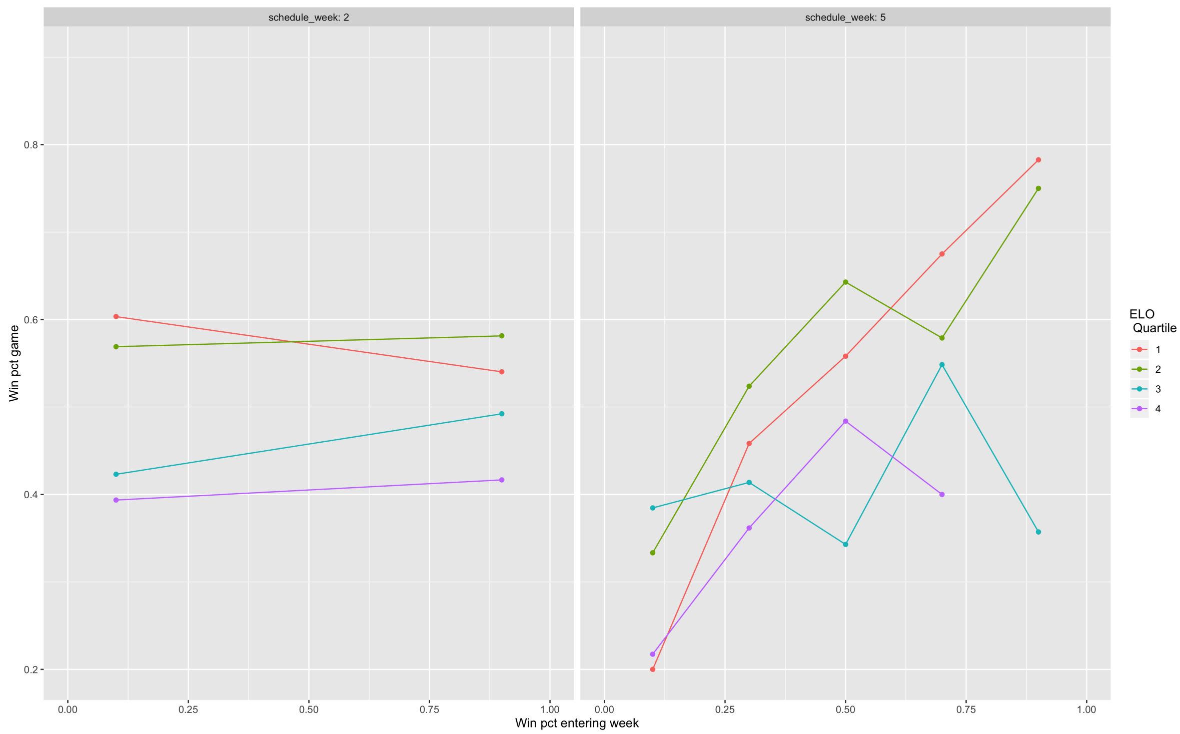

(a)

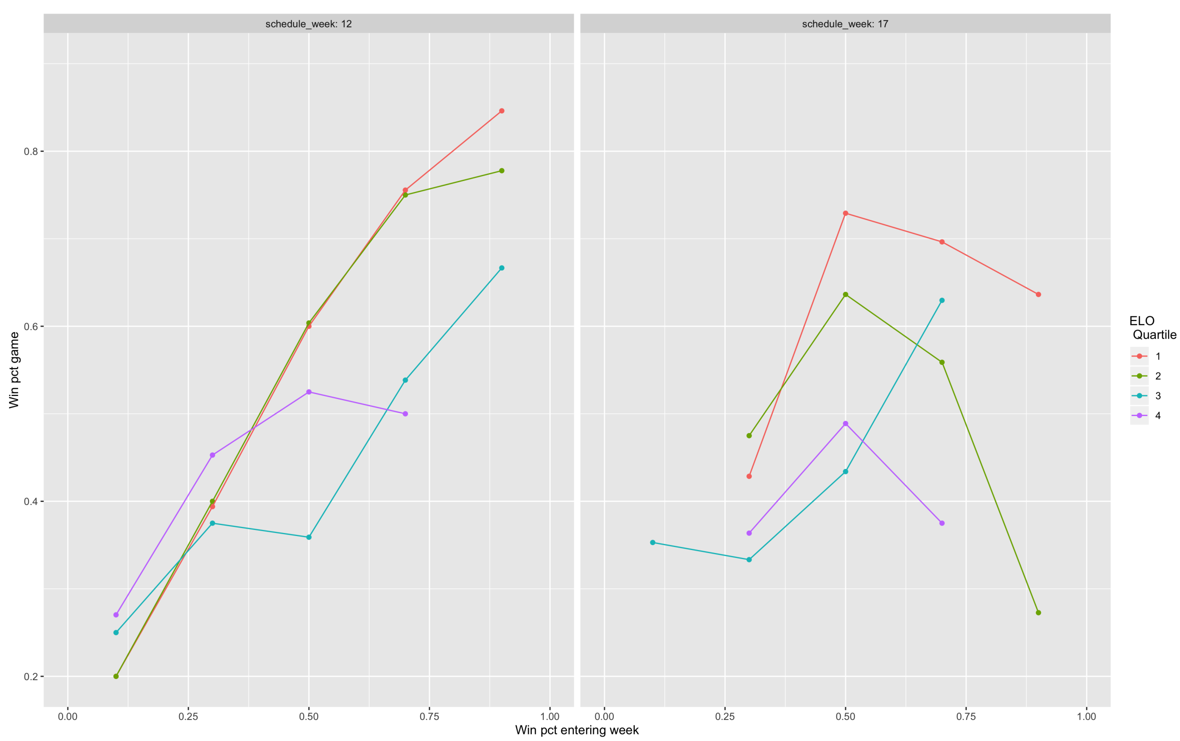

(b)

Figure 1: Win percentage by ELO quartile, weeks $= 2, 5, 12, 17$

4.1 State Variables and State Transitions

The state variables are the observed heterogeneity parameter, α_{it}, and the accumulated output through up until time t, W_{it}.

The output $W_{it} \in W$ realized by time t follows the transition law:

$$W_{it} = \begin{cases} W_{i,t-1} + w_{i,t-1}, & t > 1 \\ 0, & t = 1 \end{cases} \tag{1}$$

We also have observed heterogeneity, $\alpha_{it} \in \mathcal{A}$. We assume it does not change over time, so $\alpha_{it} = \alpha_i$. Overall, we write the state variable $s_{it} = (\alpha_i, W_{it}) \in \mathcal{A} \times W$

4.2 Per-period and Last Period Utility

In every period but the last one, we specify a common per-period utility function. Let $s_{it} = (\alpha_{it}, W_{it})$ be our observed state variables. Player i in period t derives per period utility based on the effort e_{it} exerted:

$$U(e_{it}|s_{it}) = -C(e_{it}) + \epsilon_{it}(e_{it}) \tag{2}$$

The per-period utility is simply the cost of effort plus an error term $\epsilon_{it}(e_{it})$, which can be interpreted as a state that is unobserved by the econometrician. We assume that these error terms follow a Type-I extreme value distribution (T1EV) and is distributed i.i.d across players and choices over time.

Next, we define the payoffs in period $t = T$. In this period, no effort choices are made. Instead, the players simply realize their state and rewards:

$$U_{iT} = \underbrace{\beta_p \cdot 1\{W_{iT} \geq \bar{w}\}}_{\text{Primary Reward}} + \underbrace{\beta_d \cdot \mathcal{D}(W_{iT})}_{\text{Alternative Reward}} \tag{3}$$

Where $\mathcal{D}$ is a decreasing, convex function satisfying $\mathcal{D}' \leq 0$, $\mathcal{D}'' \geq 0$. $\bar{w}$ is the threshold of output that a player needs to reach in order to realize the primary reward. For now, we assume that this is the same fixed number for each player. This assumption implies that the effort exerted by one player does not affect the threshold needed achieve the reward for another player. The assumption is feasible as long as there are minimal competition effects.

There are two rewards that are realized. The primary reward is the reward for players who have crossed the threshold $\bar{w}$. In the context of the NFL, this can be thought of as the utility gained from making the playoffs. The alternative reward is decreasing in the state W_{iT}, so players who have performed worse will receive a larger alternative reward. This is role of the NFL Draft, where the better draft picks (conditional on not making the playoffs) are assigned in the reverse order of the number of games won.

In summary, we have the following ex-ante utility functions:

$$U_{it}(e_{it}, s_{it}) = \begin{cases} \beta_p \cdot 1\{W_{iT} > \bar{w}\} + \beta_d \mathcal{D}(W_{iT}) & t = T \\ -C(e_{it}) + \epsilon_{it}(e_{it}) & t < T \end{cases} \tag{4}$$

4.3 Optimal Actions

Given the per-period utility function in equation 4 and the state transitions above, we can formulate a player's problem as choosing the optimal effort to maximize the present-discounted utility in each period. A player's present-discounted utility under the optimal effort policy can be represented by the value function

$$\tilde{V}_t(e_t, s_t; \theta, \mu) = \max_{e_t, e_{t+1}, \ldots} \mathbb{E}\left[\sum_{\tau=t}^{T} \delta^{\tau-t} U_\tau(e_\tau, s_\tau; \theta, \mu)\right] \tag{5}$$

The player's ex-ante (integrated) value function from choosing effort e_t in period t is:

$$V_t(s_t) = \int \max_{e_t} \left\{ U(e_t, s_t) + \delta \sum_{s_{t+1}} V_{t+1}(s_{t+1}) f_t(s_{t+1}|e_t, s_t) \right\} dG_\epsilon(\epsilon_t) \tag{6}$$

Where $f_t(s_{t+1}|e_t, s_t)$ is a column vector of transition probabilities. Note that the value functions and transition probabilities must be indexed by t because we are in a finite horizon setting. Let $u(\cdot)$ be the flow utility net the epsilon shock, $U(\cdot) - \epsilon_t$. Our corresponding choice-specific value function is:

$$v_t(e_t, s_t) = u_t(e_t, s_t) + \delta \sum_{s_{t+1}} V_{t+1}(s_{t+1}) f_t(s_{t+1}|e_t, s_t) \tag{7}$$

The optimal effort policy is:

$$e_t = e(s_t) = \begin{cases} 1 & v_t(1, s_t; \theta, \mu) > v_t(0, s_t; \theta, \mu) \\ 0 & \text{else} \end{cases}$$

which maximizes the present-discounted utility conditional on the current state variables.

5 Estimation

Our goal is to recover the effort profile $P_t(e = 1|s_t)$, which is akin to recovering the conditional choice probability in a dynamic discrete choice problem. We use a nested fixed point approach similar to Rust (1987) where in the inner algorithm we take θ and μ as given and compute the value function using backwards induction. In the outer algorithm we search over the parameter space to maximize the log likelihood.

To begin, we adapt the general model in Section 4 to our specific application involving the NFL. There are $n = 32$ teams and $T = 17$ periods. There are $G = T - 1 = 16$ games in a season, and a final period where the playoff and the draft rewards are realized. In each period, effort is chosen before the realization of the outcome of the game. Effort can be thought of as a collective team effort: from a General Manager's (GM), effort is spent constructing the team, actively scouting players, or seeking out trades. GMs can mandate that coaches make certain personnel decisions, like giving more playing time to less experienced players. Coaches spend effort preparing gameplans for the next opponent, watching film, and running practice. Players exert effort watching film, practicing, and playing in the actual football game. One must note that tanking is not necessarily incentive compatible to players and coaches; if a player performs poorly, they can get benched. Additionally, tanking potentially allows the team to replace the poorly performing player with a more talented replacement from the draft. Coaches can be fired if the team performs poorly.

Define W_{it} as the total number of wins that a team has heading into period t. Let w_{it} be the indicator that a team wins in period t. It follows the distribution:

$$w_{it} \sim Binom(1, p(q(e_{it}, \alpha_{it}; \mu))) \tag{8}$$

Where $q(\cdot) : \mathcal{E} \times \mathcal{A} \to \mathbb{R}$ maps the effort choice and individual heterogeneity to a some quality score in $\mathbb{R}$. We assume that $q(\cdot)$ is additively separable in e_{it} and α_{it}:

$$\begin{aligned} q(\cdot) &= \mu \cdot (\alpha_i, e_{it}) \\ &= (\mu_1, \mu_2) \cdot (\alpha_i, e_{it}) \end{aligned}$$

$p(\cdot) : \mathbb{R} \to [0, 1]$ is a strictly increasing function that maps the quality score to the win probability of each game; in our estimation, we use $p(\cdot) = \Phi(\cdot)$, the standard normal CDF. μ is a vector of parameters that indexes $q(\cdot)$. We obtain α_i by transforming the ELO of a team at the beginning of time period t to a number on $[-1, 0]$. These assumptions on α_i and $p(\cdot)$ impose a 50% win probability for $q = 0$ teams. In particular, the win rate of the most talented team $\alpha = 0$ is normalized to 50% if they do not exert effort[2]. Note that the probability of winning depends on two factors: the innate ability of the team and the effort choice.

We assume the functional form of $C(\cdot)$:

$$C(e_{it}) = \beta_c e_{it}$$

where $\beta_c > 0$ denotes the team's cost of exerting effort.

The timing is as follows: at the beginning of time $t = 1$, each team realizes its talent, α_i. Given the parameters θ and μ, the team computes its value function, realizes an independent shock $\epsilon_{i1}(e_{i1}) \sim T1EV$ for each effort choice, and then chooses its effort $e_{i1} \in \{0, 1\}$. After choosing e_{i1}, the team realizes wins w_{i1} according to the distribution given in equation 8. We then advance to the next period where the team again computes its value function and makes an optimal effort choice.

5.1 Inner algorithm

We restate some of the assumptions[3] mentioned earlier and introduce new ones. We assume that teams are risk neutral, so that $\gamma = 0$. Furthermore, we specify a functional form for the draft utility function, $\mathcal{D}(W_{iT})$:

$$\mathcal{D}(W_{iT}|ds) = \begin{cases} ds \cdot (G - W_{iT})^2 & W_{iT} < \bar{w} \\ ds \cdot (G - \bar{w})^2 & W_{iT} >= \bar{w} \end{cases} \tag{9}$$

The function is conditional on ds, which is the strength of the draft class observed in the data. Within a season, ds is assumed to be constant. We obtain variation on the draft

[2]Empirically, this may not be a bad assumption. If we look at the regular season records of teams which won the Super Bowl the year prior, the worst performing teams still obtained 7 wins in 16 games. See Appendix for details

[3]Aguirregabiria and Mira (2010) gives a list of common assumptions for the single agent dynamic model that we use if not mentioned here. Noticeably, we assume the two conditional independence assumptions.

payoff across seasons. This functional form is convex in wins to reflect the idea that the earlier draft picks are more valuable because more talented amateur prospects are available to be selected. After crossing the playoff threshold $\bar{w}$, the draft payoff is constant. This is in accordance with NFL policy; conditional on making the playoffs, draft picks are assigned by a team's performance in the playoffs and not in the regular season.

We assume that the threshold for making the playoffs is $\bar{w} = 10$ wins. This is a realistic threshold that teams aim for, as 90% of the teams since the playoff expansion in 1990 have reached the playoffs conditional on achieving exactly 10 wins[4]. Lastly, we add a third reward for the first round playoff bye. This reward is achieved if a team reaches 13 or more wins[5]. Anecdotally, teams do not tend to rest their players once they qualify for the playoffs. Instead, they often compete until they have clinched a first round bye or cannot mathematically achieve a bye. These assumptions modify the per period utility in 2 and 4 to be:

$$U_{it}(e_{it}, s_{it}) = \begin{cases} \beta_p \cdot 1\{W_{it} > 10\} + \beta_b \cdot 1\{W_{it} > 13\} + \beta_d \mathcal{D}(W_{it}) & t = T \\ -\beta_c * e_{it} + \epsilon_{it}(e_{it}) & t < T \end{cases} \tag{10}$$

Under the assumption that the unobserved state variable is distributed T1EV, the integrated value function in equation 6 has the closed form:

$$V_t(s_t; \theta, \mu) = \log \left(\sum_{e=0}^{1} \exp \left(u(e_t, s_t) + \delta \sum_{s_{t+1}} V_{t+1}(s_{t+1}) f_t(s_{t+1}|e_t, s_t) \right) \right) \tag{11}$$

We can solve for $V_1, V_2, \ldots, V_T$ using backwards induction, and then compute the condi-

[4]Historical data from pro-football-reference.com

[5]Empirical data shows that this cutoff is usually between 12 and 14 wins

tional choice probability of exerting effort using the conditional logit formula:

$$P_t(e=1|s_t,\theta,\mu) = \frac{\exp(v_t(1,s_t))}{\exp(v_t(0,s_t)) + \exp(v_t(1,s_t))} \tag{12}$$

5.2 Outer algorithm

Given the computation of the conditional choice probabilities $P_{it}(\cdot)$ at every t, we can the result into a likelihood function. The likelihood of observing wins $\hat{w}_t$ in period t for individual i is:

$$L_{it}(\theta,\mu) = \sum_e Pr_t(e_{it}|\theta) f_t(\hat{w}_t|e_{it},\mu) = \sum_e Pr_t(e_{it},\hat{w}_t|\theta,\mu) \tag{13}$$

$$L_i = \prod_{t=1}^{T} L_{it} \tag{14}$$

$$\log(L) = \sum_{i=1}^{N} \log(L_i) \tag{15}$$

where $f(\cdot)$ is the probability mass function of a binomial distribution as described in equation 8. We then use an optimization routine to search for the (θ,μ) that maximizes this likelihood function.

6 Identification

We observe the following data:$\{W_{it}, w_{it}, G_{it}, \alpha_{it}, ds_t, : \ i = 1,\ldots,N; t = 1,2\}$. Unlike classic models of dynamic discrete choice, we do not observe the effort choice; we must infer the probability of putting in effort given the observed number of wins in the period, w_{it}. The formal identification argument is tricky and has not been fully fleshed out before the deadline of this paper. The main challenge is untangling the benefits of exerting effort, μ_2, from the cost of exerting effort, β_c. Ideally, we need to find an exclusion restriction that shifts one of either the cost or the benefit of effort without shifting the other one. We proceed with a brief discussion about the intuition surrounding the identification of the parameters.

We need to be able to identify the following parameters: $\theta = (\beta_p, \beta_d, \beta_c, \beta_b), \mu = (\mu_1, \mu_2)$. To do so, we rely on the following identifying assumption:

Assumption 1 *The probability of winning, $P_i(s_t)$, does not depend on the observable state W_t when we condition on α_t, e_t. W_t affects the probability only through the choice of effort e_t.*

We begin by first considering the identification of μ_1, the parameter that affects the win probability through α_i. We are able to identify this parameter by focusing on teams at states $W \geq \bar{b} = 13$. Assumption 1 allows us to ignore W_t and focus on how effort e_t would be chosen in these states. These teams have exhausted all incentives; there are no more thresholds to reach and the draft reward crucially does not depend on regular season record conditional on making the playoffs. Hence, all teams in these states should have the same effort incentives, so variation in win percentage such teams in the data must be attributed to difference in talent, α. Analyzing the win probabilities of teams in states $W \geq \bar{b}$ while varying α gives us identification of μ_1.

We can separately identify the effect of effort on winning, μ_2, and the cost of effort, β_c. Doing so relies on the assumption that $p(\cdot)$ is not linear in μ_2; the effect of exerting effort on win probability is not constant but relies on α and μ_1. If we had an exclusion restriction, we can isolate each effect individually as well.

The difference in behavior for teams that have a chance at making the playoffs versus teams that are mathematically eliminated from the playoffs can help us recover β_p, the utility of making the playoffs. A similar argument allows us to identify β_b, using the threshold of the playoff bye. We can separately identify β_d because the draft utility acts almost as a second cost of effort. Teams at lower win states will face a higher marginal cost of effort because of the convex nature of $\mathcal{D}(\cdot)$. Year-to-year variation in the quality of the draft alters the payoff, so we can also compare performance of teams with low wins and similar quality across seasons to identify β_d. We end up assuming risk neutrality for all teams, so $\gamma = 0$.

Lastly, we cannot identify the discount factor so we assume that $\delta = 0.98$.

7 Results

7.1 NFL Season Estimation Results

Table 1 reports the parameter estimates of our model estimated on the NFL regular season data:

	Win function $p(\cdot)$
	Normal CDF
β_p	18.85***
	$[1.75, 56.11]$
β_b	52.27***
	$[3.42, 684.2]$
β_d	0.07***
	$[0.02, 0.35]$
β_c	0.48
	$[-0.70, 1.09]$
μ_1	0.82***
	$[0.68, 2.11]$
μ_2	0.73***
	$[0.59, 2.92]$

Note: *p<0.1; **p<0.05; ***p<0.01
Bootstrap SEs clustered at the season level.

Table 1: Estimated parameters, NFL 2009-2019 regular seasons

At first glance, the parameter estimates do not seem to make much intuitive sense. Why would the utility achieved from a playoff bye be three times as large as the utility achieved from making the playoffs? Why is the cost of effort not statistically significant when considering the 95% confidence interval? Some context about the setting may help

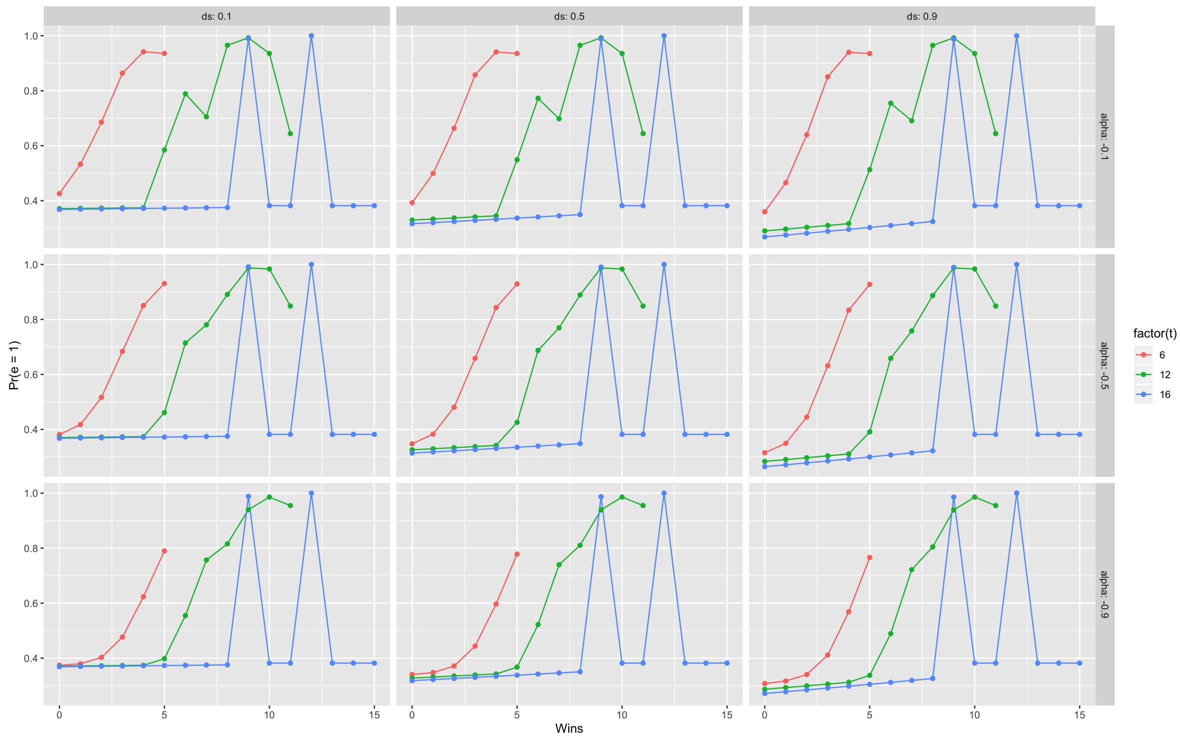

(a)

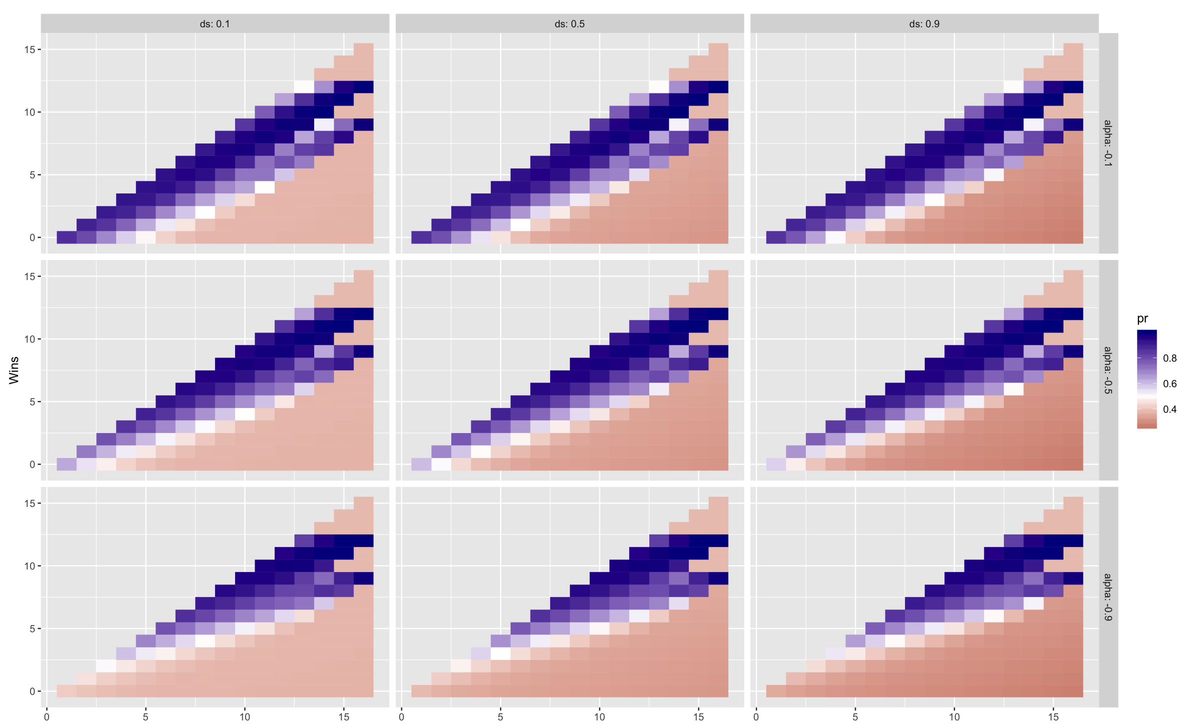

(b)

Figure 2: Implied probability of exerting effort, NFL 2009-2019

provide some intuition. For the top performing teams, simply making the playoffs is usually not their only goal; many have aspirations to win the Super Bowl. A first round playoff bye is valuable - it is equivalent to winning a playoff game since the team advances to the next round of the playoffs without needing to play an opponent. A bye allows teams more time to rest up and prepare for whomever their next opponent will be. Therefore, we do not see the absolute best teams rest many starters or shirk on effort even when they have already locked up a playoff birth. The large values of β_b and, to an extent, β_p are a result of effort almost always expended when a playoff birth or a playoff bye are in reach. The insignificance of the cost of effort parameter is a possible result of the components of effort discussed in section 5. While it may be in the best interests of the front office and GM to tank, coaches and players are fighting for their jobs and may not choose to tank.

Figure 2 displays the effort profile given different combinations of ds and α. The top figure shows the effort profile in time periods $t = 6, 12$, and 16. In each row, we compute the effort profile while holding observed talent fixed and increasing the strength of the draft prospects from left to right. Within each column, we hold the draft strength fixed while increasing the observed talent from bottom to top. Within each grid, the x-axis represents the number of wins W that each team has entering period t; the y-axis represents the probability that they will exert effort that period. The bottom figure is a heatmap that shows the effort profile across all weeks and all possible win states. Each row and column in the bottom figure has the same interpretation as in the top figure. Within a grid, the 45 degree line represents an undefeated team in a period t. To convert between the two figures, consider the following example: The blue lines in a grid in the top figure correspond to the rightmost columns in a grid in the bottom figure. The green lines correspond to the 12th columns from the left in each grid, and so forth.

At a lower time period such as $t = 6$, the effort profiles follow a shape that seems intuitive. The probability of exerting effort seems to be monotonically increasing in the win state, almost matching the shape of the descriptive win-probability plots in Figure 1. More talented teams become more likely to exert effort faster as their state improves. Amongst 3-win teams, a team with high talent $\alpha = -0.1$ exerts effort almost 90 percent of the time while a team with low talent $\alpha = -0.9$ will barely exert effort 40 percent of the time. At $t = 6$, no teams have mathematically been eliminated from the playoffs, so we would expect effort to be increasing monotonically. At a later period in time, such as $t = 12$, we start to see lapses in the monotonicity of effort. Highly talented teams exert effort more when they have six wins compared to when they have seven wins. The probability of exerting effort declines across all talent and draft levels when we reach undefeated teams (team with 11-wins), although the decline is much greater for talented teams. These teams have already made the playoffs, and only need to win 2 of their remaining 5 games to achieve a first round bye. Teams do not exert as much effort when they have earned some leeway for themselves. In the final period where effort can be exerted $t = 16$, we see that teams one win below the thresholds will exert effort with probability close to 1, while everyone else tends to not exert effort. Using the parameters estimated in Table 1, we assess the our model fit in Figure 3 by simulating outcomes. Our model slightly underperforms the lower tail of the data, but is able to capture the bimodality of the data. We perfectly capture the spike at 7 wins, and slightly overpredict the spike at 10 wins (which is like an artifact of the threshold set at $\bar{w} = 10$). As a final check, we use our model parameters to simulate the outcomes in the 2009-2019 NFL seasons 100 times. On average, a team will win 8.01 games per year, which is just a sliver larger than the theoretical maximum of 8 games.[6] Thus, we have confidence that our model is not highly miscalibrated and gives reasonable predictions.

Overall, we document the existence of tanking in the NFL. While effort isn't always exerted, the propensity to exert effort is much lower when teams are eliminated from playoffs than in any other state. Less talented teams give up faster; compared to more talented teams, they exert much less effort in lower win states that are not mathematically eliminated. Low skill teams do not fully put in effort until their prospectus of reaching the threshold is quite

[6] Since each team plays 16 head-to-head games a year and there is one winner per game, in the absence of ties the average number of games won by any team will be 8.

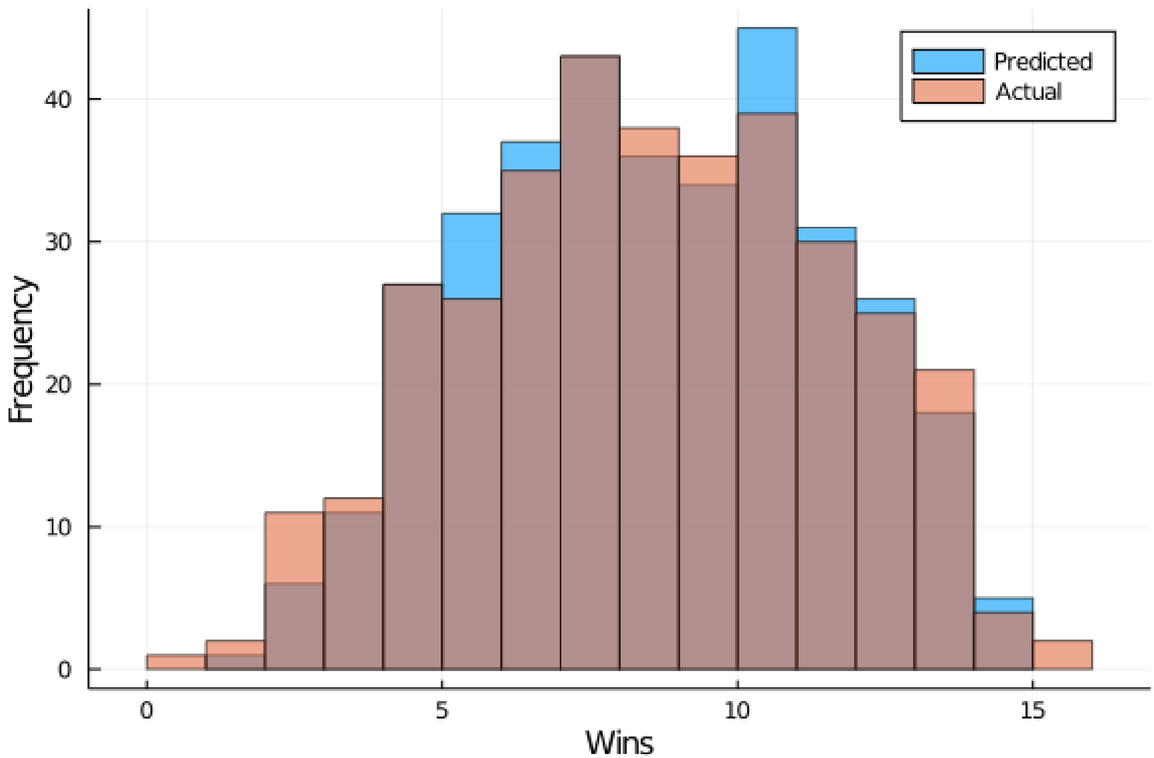

Figure 3: Predicted vs Actual Wins

good. High skill teams are also different from low skill teams because they are more likely to shirk on effort at higher states. Their raw talent can carry them to the next threshold reward without exerting a lot of effort.

7.2 Counterfactual Results

We run three counterfactuals, all of which adjust different parameters of the NFL season. The first change mirrors the rules implemented for the 2021-2022 NFL season. The number of games played increases to 17 and two additional teams can make the playoffs. Implementing this counterfactual is interesting because contract lengths and thresholds are changing. Increasing the length of the season allows for more uncertainty to be resolved before crucial effort decisions need to be made. Allowing additional teams into the playoffs may incentivize higher effort.

The second counterfactual we run replaces the current draft allocation mechanism by randomly assigning a draft pick to every team, regardless of their season record. We implement this by making the draft allocation function $\mathcal{D}(W)$ equal to a constant for every realization of the final number of wins, W. This counterfactual demonstrates the importance of the reward structure in linking together the individual decisions of teams across time periods. Furthermore, it provides insight into an interesting sports problem that has plagued league commissioners recently.

The third counterfactual is similar to the second, but instead of randomly assigning draft picks we determine the draft order at some deadline $t < T$, changing the *timing* of the alternative reward. This is the recommendation of Banchio and Munro (2020), so it will be interesting to see if this policy change elicits a more "optimal" effort profile. Specifically, we choose $t = 6$ as one of the "tanking deadlines" because under our model, every team can still qualify for the playoffs. We also choose $t = 10$ as an alternative deadline because at $t = 6$, not enough uncertainty might have been resolved to determine if the lowest win state teams are actually the worst teams.

7.2.1 Proposed Future NFL Changes

We first simulate the scenario where the proposed changes for the 2021-2022 NFL season and beyond are implemented. To recap the changes, the number of regular season games increases from 16 to 17 and the number of playoff teams increases from 12 to 14. Increasing the number of playoff games decreases the number of teams receiving first round byes from 4 to 2.

To model the 2021-2022 proposed changes in our counterfactual, we increase the number of games played in a season, G, to 17. Since we need a final period to realize end-of-season payoffs, we increase T to 18. We leave the threshold number of wins to make the playoffs, $\bar{w}$, at 10. Since there is an additional regular season game, there are more opportunities to reach 10 wins and more teams will do it. Increasing the number of playoff teams also decreases the number of first round byes, so the threshold for a playoff bye must be more competitive.

Hence, we raise the threshold to obtain a playoff bye from 13 to 14 wins. We believe that the number of additional teams that will reach 10 wins under this counterfactual scenario will reflect the updated playoff structure.

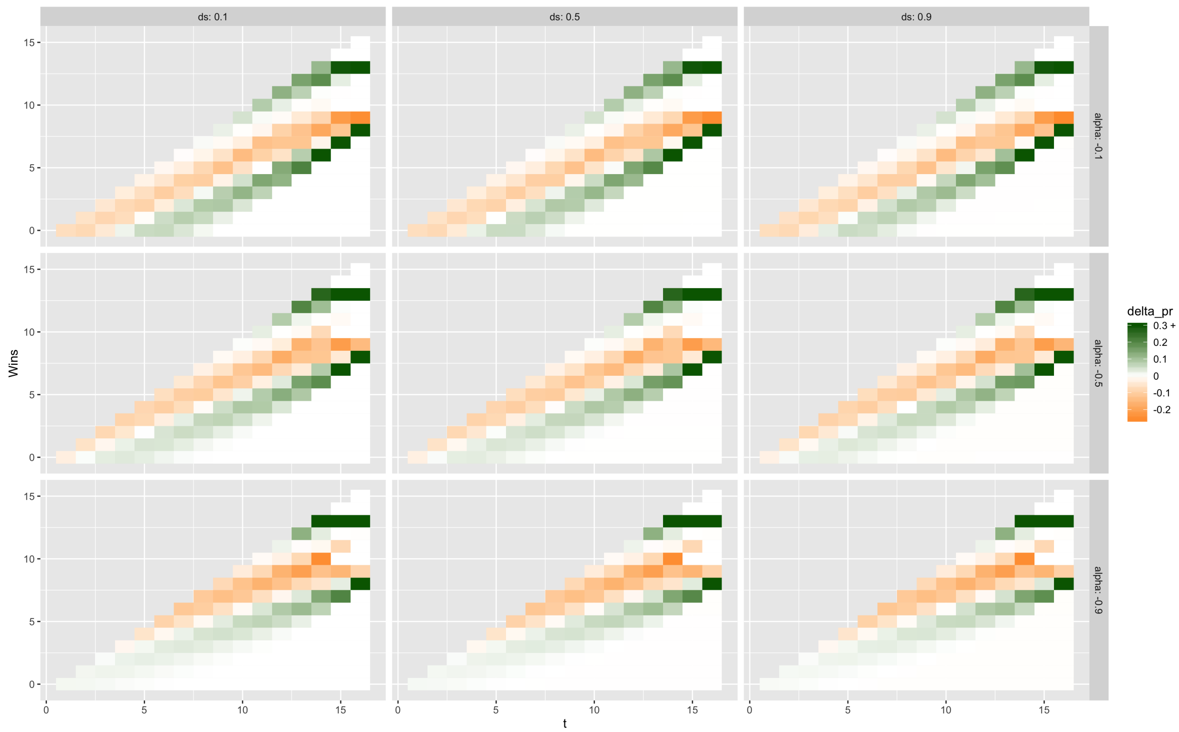

Figure 4: Counterfactual effort profile: Implementing 2021-2022 policy changes, change in effort

The effects of the 2021-2022 regular season changes on the effort profile are displayed in Figure 4.[7] We see that increasing the number of games and playoff teams increases the probability of effort being exerted when teams have won half of their games. Effort decreases when teams have won a little more than half of their games, reflecting the additional leeway that a 17th game provides in making the playoffs. Effort also tends to increase at very high win states in later time periods; it is likely that these changes are a results of the increase in the playoff bye threshold. These patterns hold regardless of the number of games played and the draft strength. Interestingly, the change in season structure seem to have varying impact on teams with different α. Less talented teams seem to increase their effort at $t = 1$,

[7]We do not display changes in effort at $t = T' = 17$, since we only see a shifted copy of the $t = 16$ heatmap columns from Figure 2

while more talented teams decrease their effort.

Because our model is single-agent and lacks competition effects, we need to check if, on average, we predict a feasible number of wins. We do so by running 100 simulations using our estimated parameters and the counterfactual changes. We then obtain the average number of wins that our model predicts for any team. In a 17 game season, the average number of games won per team cannot realistically exceed 8.5, since for every winner of a game there is a loser[8]. Our model predicts an average of 8.47 wins per team, a realistic average.

7.2.2 Randomized Draft Order

We next visit the scenario where draft order is randomized. In order to conduct this counterfactual, we need to set $\mathcal{D}(W_{i3})$ equal to some constant. Given the functional form specification for $\mathcal{D}(\cdot)$ is section 5.1, we can set $\mathcal{D}(W_{i3}) = \mathbb{E}\mathcal{D}(W_{i3}), \forall\ W_{i3}$. Note that the expectation of the function is not equivalent to the function value of the expected number of wins because of the convex nature of the draft function. Holding everything else fixed, this situation results in the change-in-effort profile displayed in Figure 5.

We see a large increase in the probability of exerting effort at the lower win states, and almost no change in probability of exerting effort in the higher win states. No teams decease their probability of exerting effort. It appears as if teams contending for the playoffs do not let the alternative reward affect their decision making. The largest increase in the probability of exerting effort is around 10%, which occurs in the lower win-states at $t = T$. Stronger drafts lead to larger increases. This is unsurprising because the payoff for performing poorly is much lower in the counterfactual.

Simulations predict that the teams will win on average about 8.17 out of 16 games in a season. The biggest gains in predicted wins are among teams with the lowest α. For example, the 2009 Detroit Lions ($\alpha = -1$) win on average 4.8 games in the simulation under current NFL rules, whereas they win on average 5.0 games in this counterfactual simulation.

[8] We can have slightly fewer than 8.5 wins because of ties

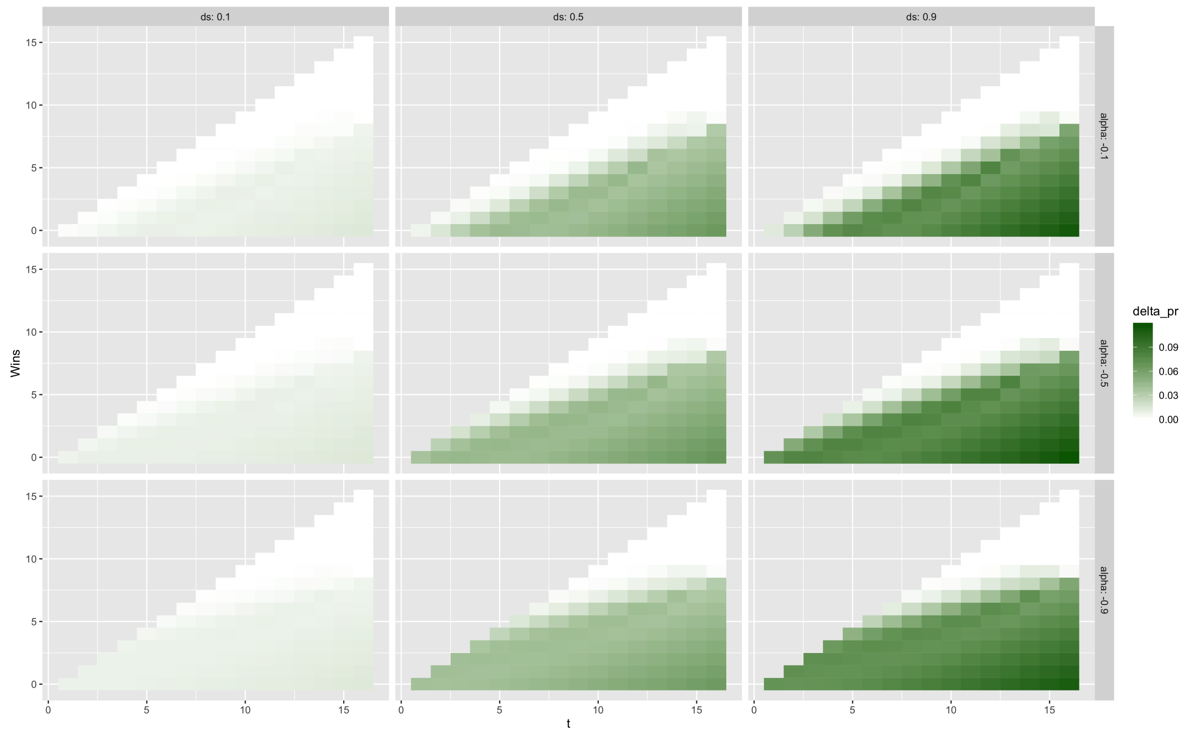

Figure 5: Counterfactual effort profile: Random Draft Order, change in effort

As discussed previously, the average number of wins of all teams in a season can not exceed realistically 8. So while effort levels may indeed increase with a randomized draft order policy, they may not increase to the extent that the model predicts since we are missing the competition element that a dynamic games model would bring.

7.2.3 Determining Draft Order Midseason

Our last counterfactual scenario implements the recommendation from Banchio and Munro (2020) to determine the draft order using the standings of teams at a time $t < T$. In particular, they recommend determining the draft order before any teams are mathematically eliminated from playoff contention. Our assumption of setting the playoff threshold at 10 wins means the latest we can set the playoff determination deadline is week 6. However, six games played seems like too small of a sample to determine if teams are truly bad or if they are simply unlucky. We proceed with conducting the counterfactual at two levels - a "tanking deadline" at $t = 6$ and another at $t = 10$.

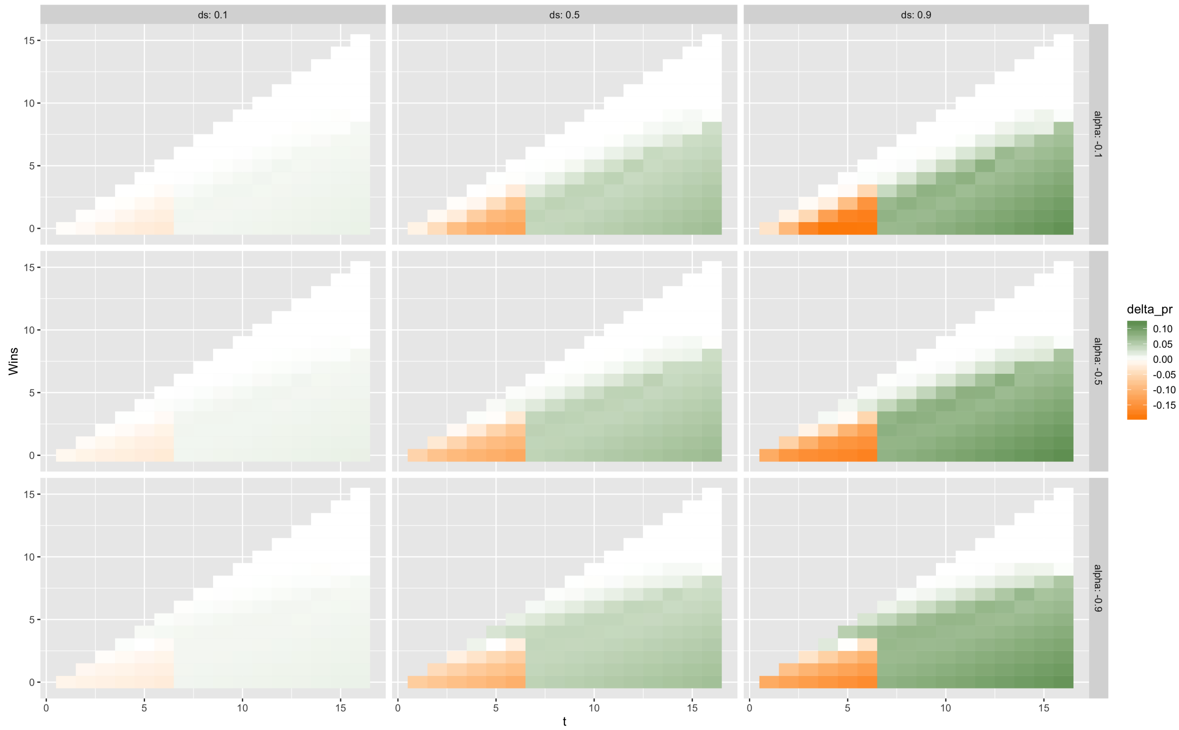

Figure 6: Counterfactual effort profile: Draft order determined week 6

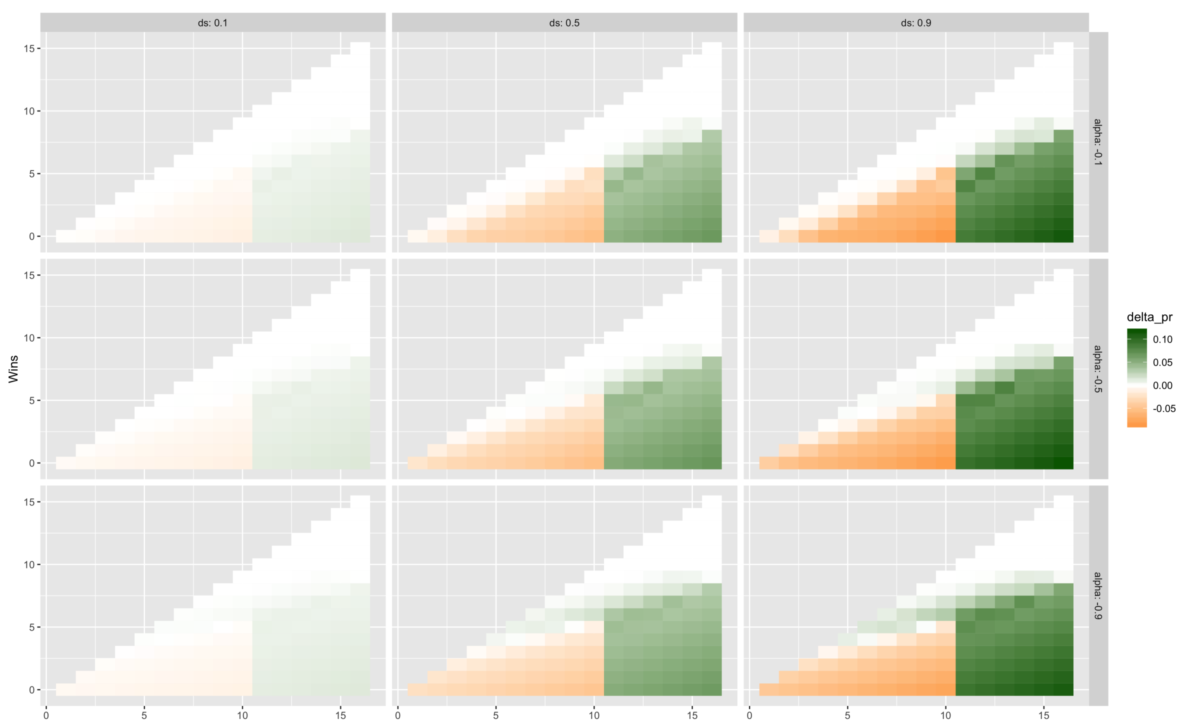

Figure 7: Counterfactual effort profile: Draft order determined week 10

Figure 6 shows the change in effort heatmap when the tanking deadline is set at $t = 6$ and Figure 7 shows the same thing for $t = 10$. We obtain predictable deadline effects regardless of when it is set - effort decreases in lower states prior to the tanking deadline, but increases afterwards. Draft strength affects the amount of bite that the deadline has; for weaker drafts, the change in effort is very small compared to the change for a stronger draft. Interestingly, the tanking deadline increases the effort of lower skill teams at high win states before both deadlines. These teams would have otherwise decreased their effort because they were unlikely to win enough games to make the playoffs. However, the early tanking deadline makes it more worthwhile to pursue the playoffs rather than a better draft pick.

For the $t = 6$ deadline, the largest decrease in probability of exerting effort is around 20%. The magnitude of decreasing effort is larger than the magnitude of increasing effort, although effort increases in more states. Our simulations predict an average of 7.98 wins per team. For $t = 10$, the decrease in effort is milder, but the decreasing effort occurs in more states. Simulations predict 8.01 average wins per team.

7.3 Discussion

Our three counterfactuals explored the following changes to the NFL season structure:

1. Increased the number of periods and increased one of the reward thresholds
2. Effectively removed the alternative reward
3. Changed the timing of the alternative reward

Under the objective where the league wants to maximize effort and allow for talent redistribution, it appears as if the status quo policy is somewhat optimal. These counterfactual results mainly demonstrated that the league can only shape the effort of teams in a limited manner. If the league is willing to trade off effort at certain state-time period combinations - for example, if they want more competitive games among bad teams late into the season -

then some of our proposed counterfactual policies can be adopted. If the NFL doesn't mind the potential media optics of a "tanking deadline", or the negative effort effects it would bring before the deadline, then adopting the third counterfactual policy would increase effort post deadline for struggling teams.

The first counterfactual policy of increasing the number of time periods and the reward threshold somewhat achieves the goal of incentivizing worse-performing teams to try harder. In the real world, we have seen such a counterfactual policy been implemented by airline rewards programs. Because travel was restricted for most of 2020 by the Coronavirus pandemic, many airlines extended the timeframe for points and status accumulation to the end of 2021. Our counterfactual predicts that changing the length of the season T and making the playoff reward easier to achieve will lead to mixed changes in effort. On one hand, the additional playoff spot in conjunction with a longer season incentivizes unlucky teams to keep putting in effort since there is more time and margin of error to recover from bad luck early in the season. On the other hand, the changes give additional leeway to teams on pace to make the playoffs, which decreases their effort. Additionally, these changes do nothing to address the absolute worst performing teams for whom tanking is still optimal.

If the league did not care about talent redistribution, then they could implement our second counterfactual policy: replace the current draft mechanism with a completely random draft lottery. In this scenario, the league would essentially be removing the alternative reward and only leaving the primary rewards of playoffs and byes. Because effort is costly, the increase in effort is not drastic; at best, teams may increase their propensity to try by 15%. Removing the alternative reward actually creates a uniform effort profile among teams mathematically eliminated from the primary rewards (playoffs). This counterfactual teaches us that an alternative reward almost always decreases and distorts effort whenever the primary reward has not been achieved. In a non-sports example, if a sales-force agency wanted to implement a training program for just their worst-performing salespeople, our counterfactual tells them to expect even more shirking by the poor performers than in the status quo.

8 Conclusion

In this paper, we have studied the effects of multiple levers on effort in a time-bundled contract. Previous literature has documented i.) effort deceleration in the presence of a deadline and a threshold reward and ii.) the deceleration mostly goes away when the threshold is removed. Our sports setting precludes us from removing any thresholds, but provides more levers to alter effort behavior. We simulate three counterfactual worlds; the first changes the number of time periods and the threshold for one of the primary rewards. The effect on effort is ambiguous due to simultaneous incentivizing and disincentivizing effects acting on different states of the world. Our second counterfactual eliminates the alternative reward that awards individuals in low states. Predictably, effort increases in every state and especially in the lower states. The final counterfactual alters the timing of the alternative reward, resulting in yet another ambiguous effect on effort. Effort pre-deadline declines for poor performing teams but increases for high performing teams. Effort post-deadline increases across the board. Overall, our work provides possible alternative contract structures for a contract designer with our specific constraints and very particular effort shifts in mind.

We have taken quite a few liberties with respect to assumptions in this paper. Most noticeably, we model our problem as a single agent dynamic discrete choice problem and do not directly allow for competitive effects in a sports setting. Further work can be done in modelling our setting as a dynamic game; such a direction will face similar challenges in identifying the latent action and dealing with a large, sparse state space. Furthermore, we do not take a strong stance on the objective of the league, making it difficult to determine what an optimal time-bundled contract would be. Additional work can be done to address what an "optimal" contract would be. Ideally, this would include studying possible effects on a demand side, such as TV viewership or game attendance.

References

Aggarwal, S., Dizon-Ross, R., and Zucker, A. D. (2020). Incentivizing behavioral change: The role of time preferences. Technical report, National Bureau of Economic Research.

Aguirregabiria, V. and Mira, P. (2010). Dynamic discrete choice structural models: A survey. *Journal of Econometrics*, 156(1):38–67.

Arcidiacono, P. and Miller, R. A. (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6):1823–1867.

Banchio, M. and Munro, E. (2020). An incentive-compatible draft allocation mechanism.

Chung, D. J., Kim, B., and Park, B. G. (2019). The comprehensive effects of sales force management: A dynamic structural analysis of selection, compensation, and training.

Chung, D. J., Steenburgh, T., and Sudhir, K. (2014). Do bonuses enhance sales productivity? a dynamic structural analysis of bonus-based compensation plans. *Marketing Science*, 33(2):165–187.

Green, P. E. and Srinivasan, V. (1978). Conjoint analysis in consumer research: issues and outlook. *Journal of consumer research*, 5(2):103–123.

Hartmann, W. R. and Viard, V. B. (2008). Do frequency reward programs create switching costs? a dynamic structural analysis of demand in a reward program. *Quantitative Marketing and Economics*, 6(2):109–137.

Hotz, V. J. and Miller, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529.

Kopalle, P. K., Sun, Y., Neslin, S. A., Sun, B., and Swaminathan, V. (2012). The joint sales impact of frequency reward and customer tier components of loyalty programs. *Marketing Science*, 31(2):216–235.

Lenten, L. J. (2016). Mitigation of perverse incentives in professional sports leagues with reverse-order drafts. *Review of Industrial Organization*, 49(1):25–41.

Lenten, L. J., Smith, A. C., and Boys, N. (2018). Evaluating an alternative draft pick allocation policy to reduce 'tanking'in the australian football league. *European Journal of Operational Research*, 267(1):315–320.

Misra, S. and Nair, H. S. (2011). A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quantitative Marketing and Economics*, 9(3):211–257.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033.

9 Appendix

9.1 Empirical Evidence for Assumptions

9.1.1 Assumption about quality score 0 teams

We assume that the most talented team $\alpha = 0$ exerting effort $e = 0$ will win 50% of their games. To empirically test this, we look at the outcomes of all Super Bowl winners the season after they win the championship. The idea is that talent is sticky in the NFL, so the previous Super Bowl winners should still be highly talented in the following season. Out of the last 20 Super Bowl Winners, only one has finished the following season with worse than a $8-8$ record (2003-2004 Tampa Bay Buccaneers) and even then, they finished only 1 win short of winning 50% of their games. If anything, this evidence shows that lazy but very talented teams might have a winning probability of more than 50%.

9.2 Some Identification Arguments

9.2.1 Setup

Utility functions

$$U_{it}(e_{it}, W_{it}) = -\beta_c e_{it} + \epsilon_{it}(e_{it}) \tag{16}$$

$$U_{iT} = \underbrace{\beta_p \cdot 1\{W_{iT} \geq \bar{w}\}}_{\text{Playoffs reward}} + \underbrace{\beta_d \cdot \mathcal{D}(W_{iT})}_{\text{Draft reward}} \tag{17}$$

Win distrubution:

$$w_{it} \sim Binom(G_{it}, p(q(e_{it}, \alpha_i; \mu))) \tag{18}$$

For the identification argument, assume that $G_{it} = 1$ for every i, t. The above expression is therefore also a a Bernoulli distribution with parameter $p(q(\cdot))$. The function $q(\cdot)$ is a linear

Super Bowl	Winning Team	Next Year's Record	Change in Wins (%)
XII	Cowboys	12-4	0* (-.107)
XIII	Steelers	12-4	-2 (-.125)
XIV	Steelers	9-7	-3 (-.188)
XV	Raiders	7-9	-4 (-.250)
XVI	49ers	3-6	-10** (-.479)
XVII	Redskins	14-2	+6** (-.014)
XVIII	Raiders	11-5	-1 (-.063)
XIX	49ers	10-6	-5 (-.313)
XX	Bears	14-2	-1 (-.063)
XXI	Giants	6-9	-8*** (-.475)
XXII	Redskins	7-9	-4*** (-.296)
XXIII	49ers	14-2	+4 (+.250)
XXIV	49ers	14-2	0
XXV	Giants	8-8	-5 (-.313)
XXVI	Redskins	9-7	-5 (-.313)
XXVII	Cowboys	12-4	-1 (-.063)
XXVIII	Cowboys	12-4	0
XXIX	49ers	11-5	-2 (-.125)
XXX	Cowboys	10-6	-2 (-.125)
XXXI	Packers	13-3	0
XXXII	Broncos	14-2	+2 (+.125)
XXXIII	Broncos	6-10	-8 (-.500)
XXXIV	Rams	10-6	-3 (-.188)
XXXV	Ravens	10-6	-2 (-.125)
XXXVI	Patriots	9-7	-2 (-.125)
XXXVII	Bucs	7-9	-5 (-.313)
XXXVIII	Patriots	14-2	0
XXXIX	Patriots	10-6	-4 (-.250)
XL	Steelers	8-8	-3 (-.188)
XLI	Colts	13-3	+1 (+.063)
XLII	Giants	12-4	+2 (+.125)
XLIII	Steelers	9-7	-3 (-.188)
XLIV	Saints	11-5	-2 (-.125)
XLV	Packers	15-1	+5 (+.313)
XLVI	Giants	9-7	0
XLVII	Ravens	8-8	-2 (-.125)
XLVIII	Seahawks	12-4	-1 (-.063)
XLIX	Patriots	12-4	0
L	Broncos	9-7	-3 (-.188)
LI	Patriots	13-3	-1 (-.063)
LII	Eagles	9-7	-4 (-.250)
LIII	Patriots	12-4	+1 (+.063)

Table 2: Post Super Bowl Performance

combination of e and α:

$$q(e, \alpha; \mu) = (\alpha, e) \cdot (\mu_1, \mu_2)$$

we also assume that $p(\cdot)$ is invertible and known. We have the following parameters to identify:

$$\mu_1, \mu_2, \beta_p, \beta_c, \beta_d$$

assuming that the teams are risk neutral.

In summary, the following assumptions are made:

1. The probability of winning, $P_i(w_t)$, does not depend on the observable state W_t when we condition on α_t, e_t. W_t affects the probability only throuagh the choice of effort e_t.

2. $p(\cdot)$ is increasing in its argument and invertible

3. $q(\cdot)$ is a linear combination of (α, e)

4. Teams are risk neutral ($\gamma = 0$)

We can think of the μ parameters as the ones that affect utility through the probability of winning, and the β parameters as the ones that affect utility directly. Using knowledge of our setting, we can specify some constraints for our parameters:

- Putting in effort should increase the probability of winning, so $\mu_2 > 0$. Howver, no team should be guaranteed to win if they put in effort. Therefore, $\mu_2 < p^{-1}(1)$

- More talented teams should be more likely to win, $\mu_1 > 0$. Similar to above, no team should be guaranteed to win if they are the most talented. So $\mu_1 < p^{-1}(1)$

- Playoffs and drafts should matter $\beta_p, \beta_d > 0$

- There should no temporary benefits to effort, $\beta_c \geq 0$

9.2.2 Non-parametric identification

We observe $\mathbb{E}[w|s]$, which is the marginal probability of winning. We do not observe effort. Because effort is discrete, this expression is equivalent to:

$$\mathbb{E}[w|s] = \mathbb{E}[w|e=1,s] \cdot Pr(e=1|s) + \mathbb{E}[w|e=0,s] \cdot Pr(e=0|s) \tag{19}$$

If we assume that there exists a state such that $Pr(e=1|s) = 1$, that is, effort is exerted with certainty, then $\mathbb{E}[w|e=1,s]$ can be identified. If $\mathbb{E}[W|e,s]$ takes the same functional form regardless of e, then it is identified.

Because of the binary nature of the effort choices, we only have $Pr(e=1|s)$ to identify, as $Pr(e=0|s) = 1 - Pr(e=1|s)$. We can easily identify this after recovering $E[w|e,s]$; we use observations where the effort choice will not be $e=1$ with certainty.

9.2.3 Parametric Identification

Equations at $T-1$

Consider the following argument for the identification at the last period $t = T-1$: we have the following choice specific value functions (dropping i, t indices for expositional clarity):

$$\begin{aligned} v(e,s) &= u(e,s) + \delta[U_T(s_0) * P(w=0) + U_T(s_1) * P(w=1)] \\ &= -\beta_c e + \delta[U_T(s_0) * (1 - \mathbb{E}[w|e,s]) + U_T(s_1) * \mathbb{E}[w|e,s]] \end{aligned}$$

where s_0 is the state in the final period if the team does not win in $t = T-1$ and s_1 if the team does win. We can group the $\mathbb{E}[w|e,s]$ arguments and rewrite the CSVF:

$$v(e,s) = \kappa\mathbb{E}[w|e,s] - \beta_c e + \delta U_T(s_0) \tag{20}$$

$$\kappa = \delta[U_T(s_1) - U_T(s_0)] \tag{21}$$

The conditional choice probability is:

$$\begin{aligned}Pr(e=1|s) &= \frac{\exp(v(1,s))}{\exp(v(0,s))+\exp(v(1,s))}\\ &= \frac{\exp(\kappa\mathbb{E}[w|1,s]-\beta_c+\delta U_T(s_0))}{\exp(\kappa\mathbb{E}[w|0,s]+\delta U_T(s_0))+\exp(\kappa\mathbb{E}[w|1,s]-\beta_c+\delta U_T(s_0))}\\ &= \frac{\exp(\kappa\mathbb{E}[w|1,s]-\beta_c)}{\exp(\kappa\mathbb{E}[w|0,s])+\exp(\kappa\mathbb{E}[w|1,s]-\beta_c)}\end{aligned}$$

From above, we observe $\mathbb{E}[w|s]$:

$$\mathbb{E}[w|s] = \mathbb{E}[w|e=1,s]\cdot Pr(e=1|s)+\mathbb{E}[w|e=0,s]\cdot Pr(e=0|s) \tag{22}$$

Substituting the first equation above into the observed win probability equation, we obtain:

$$\begin{aligned}\mathbb{E}[w|s] = \mathbb{E}[w|1,s]&\frac{\exp(\kappa\mathbb{E}[w|1,s]-\beta_c)}{\exp(\kappa\mathbb{E}[w|0,s])+\exp(\kappa\mathbb{E}[w|1,s]-\beta_c)} +\\ \mathbb{E}[w|0,s]&\left(1-\frac{\exp(\kappa\mathbb{E}[w|1,s]-\beta_c)}{\exp(\kappa\mathbb{E}[w|0,s])+\exp(\kappa\mathbb{E}[w|1,s]-\beta_c)}\right)\end{aligned}$$

Assumption 1 allows us to write the expectation as the probability $p(q(\cdot))$, and assumption 3 introduces linearity in the μ parameters:

$$\begin{aligned}\mathbb{E}[w|s] = p(\mu_2+\alpha\mu_1)&\frac{\exp(\kappa p(\mu_2+\alpha\mu_1)-\beta_c)}{\exp(\kappa p(\alpha\mu_1))+\exp(\kappa p(\mu_2+\alpha\mu_1)-\beta_c)} +\\ p(\alpha\mu_1)&\left(1-\frac{\exp(\kappa p(\mu_2+\alpha\mu_1)-\beta_c)}{\exp(\kappa p(\alpha\mu_1))+\exp(\kappa p(\mu_2+\alpha\mu_1)-\beta_c)}\right)\end{aligned} \tag{23}$$

Case when $\alpha = 0$

If we have many observations of $\alpha = 0$, then equation 23 simplifies greatly:

$$\mathbb{E}[w|W, \alpha = 0] = p(\mu_1)\frac{\exp(\kappa p(\mu_1) - \beta_c)}{\exp(\kappa p(0)) + \exp(\kappa p(\mu_1) - \beta_c)} + \\ p(0)\left(1 - \frac{\exp(\kappa p(\mu_1) - \beta_c)}{\exp(\kappa p(0)) + \exp(\kappa p(\mu_1) - \beta_c)}\right)$$

At win states $W \geq \bar{w}$, any additional win will not change the payoff in the final period. This relies on the assumption that the alternative reward payoff function, $\mathcal{D}$, is flat at state $W \geq \bar{w}$. In the NFL, draft order after you make the playoffs depends more on performance in the playoffs and less on regular season record, so the assumption is not egregious. This means that $U_T(s_1) = U_T(s_0)$, which implies that $\kappa = 0$:

$$\mathbb{E}[w|s] = p(\mu_2)\frac{\exp(-\beta_c)}{1 + \exp(-\beta_c)} + \tag{24}$$

$$p(0)\left(1 - \frac{\exp(-\beta_c)}{1 + \exp(-\beta_c)}\right) \tag{25}$$

However, we cannot disentangle the effects of μ_1 from β_c here. We need some sort of exclusion restriction. The exclusion restriction must satisfy:

- Modify payoff to effort while keeping cost of effort fixed OR
- Modify cost of effort while keeping payoff of effort fixed
- Must affect teams who have clinched playoffs (or home-field advantage) and have $\alpha = 0$

Case when $\alpha \neq 0$, $W \geq \bar{w}$

Using the same assumptions from above, we simplify $U_T(s_1) = U_T(s_0)$ implying $\kappa = 0$. We rewrite equation 23 as:

$$\mathbb{E}[w|s] = p(\alpha\mu_1 + \mu_2)\frac{\exp(-\beta_c)}{1+\exp(-\beta_c)} + \\ p(\alpha\mu_1)\left(1 - \frac{\exp(-\beta_c)}{1+\exp(-\beta_c)}\right) \tag{26}$$

We observe the left hand side of equation 26 for different α. Thus, variation in win percentages in this state must be attributed to α because the probability of exerting effort is only a function of β_c. Hence, μ_1 is identified conditional on μ_2 and β_c being identified.

Identification of β_d

We use $\mathbb{E}[w_{i,T-1}|\alpha, \bar{w}-2], \mathbb{E}[w_{i,T-1}|\alpha, \bar{w}-3], \ldots$ to identify β_d. The idea is the following: In these states, no matter what the team does, they cannot make the playoffs. Thus, final period payoffs only come from the draft utility, $\beta_d\mathcal{D}(\cdot)$. In different seasons, we observe variation in $\mathcal{D}(\cdot)$ because the quality of the amateur prospects differs. Hence, performance of teams in this state-time combination across different seasons will vary only as a result of the draft quality varying. So β_d is identified.