Fundamentals and Institutional Choice in US Sports Leagues*

Martín Gonzalez-Eiras

University of Copenhagen†

Nikolaj A. Harmon

University of Copenhagen‡

Martín A. Rossi

Universidad de San Andrés§

January 16, 2020

Abstract

To shed light on the relation between fundamentals and adopted institutions we examine institutional choice across the "Big Four" US sports leagues. Despite having very similar business models and facing the same economic and legal environment, these leagues exhibit large differences in their use of regulatory institutions such as revenue sharing or salary caps. We show, theoretically and empirically, that these differences can in part be rationalized as optimal responses to differences in one fundamental characteristic of the sports being played: how strongly win probabilities respond to hired talent. Our results thus show evidence that in professional sports existing institutions are tailored to fundamentals.

JEL Classification: D02, L10, L83, O17

Keywords: Regulations, institutional choice, sport economics, win probability

*We thank Sebastian Barfort, Ferrán Elías Moreno, Alejandro Manelli, Dirk Niepelt, Christian Ruzzier, and participants at seminars and conferences for useful comments and suggestions.

†Øster Farimagsgade 5, 1353 Copenhagen, Denmark. E-mail: mge@alum.mit.edu.

‡Øster Farimagsgade 5, 1353 Copenhagen, Denmark. E-mail: nikolaj.harmon@econ.ku.dk.

§Vito Dumas 284, (B1644BID) Victoria, Argentina. E-mail: mrossi@udesa.edu.ar.

1 Introduction

Whether comparing countries, firms or other economic entities, it is well established that institutions typically vary across different settings. The reason for these differences is less clear. On the one hand, institutional differences may simply arise through a combination of random chance and path dependence. On the other hand, it may be that the observed differences in institutions across settings reflect optimal responses to differences in the underlying preferences, technology or other fundamentals. Distinguishing between these possible explanations is of first order importance for providing policy prescriptions. In particular, if differences in institutions are largely caused by random chance, we can expect large gains from identifying and transplanting successful institutions from one context to another. Understanding the relationship between fundamentals and institutions is a challenging task, however, as institutions are complex objects and fundamentals often differ in a myriad of ways across economic settings.

To study the interplay between fundamentals and adopted institutions, this paper analyzes the specific case of institutional choice across the "Big Four" US sports leagues: the National Football League (NFL), the National Basketball Association (NBA), the National Hockey League (NHL), and Major League Baseball (MLB).[1] Each of these leagues currently consists of about 30 teams that make hiring decisions and play games against each other to generate revenue from fans and media contracts. To regulate this process each league decides on a set of institutions. In particular the league chooses the extent to which it will regulate teams' hiring of players, either through the use of direct constraints such as a salary cap, or through redistributive policies such as revenue sharing or various forms of payroll taxes.

The US sports leagues are particularly well suited as a case study for understanding the relationship between institutions and fundamentals. All leagues face very similar fundamental conditions regarding revenue generation. They all draw most of their fans from the U.S. population, negotiate broadcasting contracts with the same networks, and operate under the same legal system.[2] At the same time, the four US sports leagues exhibit large differences in their choice of institutions. At one end, the NFL features both a hard cap on salaries and extensive revenue sharing across teams. At the other end, the MLB has significantly less revenue sharing and only a modest payroll tax. These differences in institutions have periodically led league officials and commentators to argue

[1]Except for the NFL, all these leagues have at least one team in a Canadian city so a more precise description might be "North American" sports. For brevity we refer to them as US sports however.

[2]An outlier is the NHL which has seven teams in Canadian cities (NBA and MLB each have only one Canadian team).

that institutions should be transplanted across leagues.[3] Importantly, however, the four leagues differ in terms of one fundamentally basic premise: the rules of the sports that they play. In this paper we show that the observed differences in institutions across sports leagues can in part be rationalized as reflecting optimal responses of the leagues to differences in the fundamental rules of the underlying sports.

Our analysis starts from a standard off-the-shelf league model from the sports literature. In the model, each team in the league invests in a one-dimensional input, called skill or talent, which increases its chance of winning a given game, and thus the championship. Given this investment, the team will attract fans and viewers, from whom it will collect revenue. Relative to the standard model, we introduce two simple modifications. First, to allow leagues' institutions to influence the total amount of skills hired by teams, we do away with the assumption of a fixed supply of talent. Second, we allow for heterogeneity in the rules of the game reflected in one parameter, the elasticity of the odds of winning a match with respect to the relative skills of the playing teams. We denote this parameter the *productivity of skills*.

We analyze teams' individual hiring decisions and compare the competitive allocation of talent to the choice made by a planner that maximizes total league profits. Externalities arise because each team's talent level affects the revenue and win probabilities of other teams. As a result teams will in general not hire the efficient amount of skills.[4] These externalities thus provide a rationale for leagues to introduce regulatory and redistributive institutions. When teams have an incentive to hire more skills than the efficient amount, revenue sharing can increase aggregate profits by depressing teams' hiring incentives. Moreover, the strength of these externalities depends on the productivity of skills, as the incentives to hire increase when skills have a stronger effect on win probability. As a result, it is optimal for leagues that play sports with a higher productivity of skills to introduce a higher level of revenue sharing and other regulatory institutions that affect hiring incentives.

We use recent data on match outcomes and team payrolls to provide estimates of the productivity of skills across the four major US sports leagues. We find large differences across the four leagues. The elasticity of the odds of winning a match with respect to the relative skills ranges from about 1.54 in the NBA, to about 0.16 in MLB. Comparing these estimated productivity of skills with the institutions actually observed in the different leagues, we find that the main prediction from our theoretical model fit the data well:

[3] As an example of the former see Levin et al. (2000). Examples of the latter are Rogers (2015) and Gordon (2012).

[4] We refer to efficiency with respect to the objective of maximization of league-wide profits.

leagues where the productivity of skills is higher have more revenue sharing and other regulatory institutions relative to leagues with lower productivity of skills. The only slight exception is the NFL which has the highest level of revenue sharing and arguable the strongest regulatory institutions but where the estimated productivity of skills is only the second highest of the four sports. As an additional validation exercise, we show that our model makes an additional prediction regarding the relationship between the productivity of skills and the dispersion of payrolls, which is also in line with what we observe in the data.[5] We conclude that the observed differences in institutions across sports leagues can be explained as optimal responses to differences in the fundamentals of the underlying sports.

Our paper contributes to the existing literature on the economics of sports (for reviews of this literature see Downward and Dawson 2000; Andreff and Szymanski 2006; Downward, Dawson, and Dejonghe 2009). Relative to this literature, our model highlights the importance of a parameter, the productivity of skills, that has not been studied so far. Its main advantage is that, being related to match level data, it represents an intrinsic characteristic of a given sport that is largely unaffected by league institutions and regulations. In contrast, most existing comparisons across sports consider outcome measures that respond endogenously to institutions, such as the probability of winning the championship. Thus, the productivity of skills provides a microfoundation for the observed pay-performance sensitivity across sports highlighted in Szymanski (2003).

There is an academic literature that rationalizes some of the differences in revenue redistribution policies between leagues. For example, Peeters (2015) argues that U.S. leagues should not adopt the same revenue-sharing policies because the underlying markets characteristics are not homogenous. In particular, local sharing should be higher if teams in the league serve more homogeneous local markets. Salaga et al (2014) focus on the relative attractiveness of sharing media revenues versus stadium revenues. They show that revenue sharing policies differ between leagues because of differences in revenue generation mechanisms and magnitudes. In contrast, our focus is on explaining the observed differences in revenue sharing arrangements between leagues based on the differences in the marginal productivity of skills, which is an intrinsic characteristic of the sport. Thus, to emphasize our proposed mechanism, we assume revenue generation and underlying

[5] In Gonzalez-Eiras et al. (2018) we also allow the league to choose season length, and find that the model's implication regarding the relation between productivity of skills and season length is consistent with the data: In the model leagues with low productivity of skills choose longer season length. MLB has the lowest productivity of skills, and its teams play 162 games per season, significantly more than in the other sports.

markets' characteristics to be the same across leagues.[6]

At the broader level, our paper also contributes to the literature on the optimal design of institutions by highlighting the crucial role of fundamentals in shaping institutions, at least in the context of professional sports. In particular, our results suggest that the common practice of advocating the transplant of institutions across contexts should be taken with a pinch of salt.[7] Similar arguments have been made in the existing literature and debate. For example, much of the recent critique of the Washington Consensus can be viewed as arguing for more focus on differences in fundamentals across countries.[8]

The paper continues as follows. Section 2 describes the context and institutions across the four major US sports leagues. Section 3 develops the model, while Section 4 presents the data and reports the results. Section 5 concludes.

2 US sports leagues

Professional sports in North America are dominated by the so-called "Big Four" leagues: NFL, NBA, NHL and MLB.[9] The first professional league, for baseball, was founded in 1875, and the last, for basketball, was founded in 1946. Average attendance (and yearly revenues), as of 2015, is 17,500 (3.7 billion dollars) in the NHL, 17,800 (5.2 billion dollars) in the NBA, 30,500 (9.5 billion dollars) in the MLB, and 68,200 (13 billion dollars) in the NFL. In North America, the sports market was worth \$63.5 billion in 2015, and the sports industry contributes approximately 500,000 jobs.[10]

Each of the four sports leagues comprise a stipulated number of teams, also known as franchises. Currently, the NFL has 32 teams, while the NBA, NHL, and MLB have 30 teams each. Although franchises are corporate entities separated from their leagues, they operate only under league auspices. The formal structure of the leagues is the cooperative

[6] In contrast to Peeters (2015), our model predicts that an increase in local market heterogeneity should lead to an increase in revenue sharing. This can explain why a secular increase in revenue sharing has taken place alongside league expansions in recent decades.

[7] Attempts to transplant institutions across contexts is a common policy prescription. As examples of actual reform from above we have the dissemination of the English, French and German legal traditions through conquest, colonization or imitation in the 19^{th} century. Japan in 1945 imported many of US institutions, and in 2004 Dubai adopted common law for its International Finance Centre. Transplanting institutions was in large part the essence of the Washington Consensus and, more recently, suggestions that the US should adopt some Scandinavian institutions have surfaced several times in the 2016 presidential primaries.

[8] See for example Stiglitz (2008). Relatedly, Acemoglu et al. (2017) uses a similar argument to criticize the idea of exporting Scandinavian institutions to other countries.

[9] It should be noted that the MLB is actually composed of two leagues, NL and AL, that have slightly different rules.

[10] See Heitner (2015).

association of team owners. All strategic questions of league-wide relevance are decided by majority voting, and only franchise owners are allowed to vote (Downward and Dawson 2000). The Big Four leagues have franchises placed nationwide, and all leagues grant territorial exclusivity to their owners, precluding the addition of another team in the same area unless the current team's owners consent. All four major leagues have strict rules regarding who may own a team, and generally do not allow anyone to own a stake in more than one franchise, to prevent the perception of being in a conflict of interest.

Viewed as economic entities, the activities of the Big Four leagues are very similar. Once a year a season takes place in which the teams of a league play games against each other in the respective sports. Teams make decisions about player hiring and coaching, most of them in-between seasons, which influence the outcome of the games. At the end of the season the most successful team is crowned as the winner (champion) of the league.[11] This process generates revenue for teams. Major revenue sources are admissions and tickets (35% of revenue), television and broadcasting rights fees (30%), advertising, sponsorships, and endorsement fees (10%), and concessions and merchandise sales (5%).[12]

All leagues use institutions that regulate the behavior of teams in various regards. In this paper we focus on regulatory institutions that affect teams' incentives to hire players, either through redistribution or direct constraints on team behavior. The simplest example of such a regulatory institution is "revenue sharing", which is essentially the sports league version of redistributive taxation. Assuming (reasonably) that hiring better players increases team revenue, revenue sharing depresses teams' hiring incentives.

The Big Four leagues all use revenue sharing, though they do so to very different degrees. In the NFL teams share close to 61% of all league related revenues. For instance, all the revenue generated from broadcasting deals is shared equally among all teams, and a significant share of net gate income goes to the visiting team. Even licensing deals, such as income generated from jerseys and posters, is shared. In the NBA, all teams contribute annually a fixed percentage of their total local revenue, roughly 50%, into a revenue-sharing pool. Each team then receives an allocation equal to the league's average team payroll for that season from the revenue pool. In the MLB all teams share national broadcasting revenue equally and contribute 34% of their local TV revenue into

[11] All major sports leagues use a similar type of regular season schedule with a playoff tournament after the regular season ends. The best teams in the regular season reach the playoffs, and the winner of the playoffs is crowned champion of the league

[12] National TV rights are sold collectively by the league, and all Big Four leagues have launched a network of their own, NBA TV in 1999, the NFL Network in 2003, the NHL Network in Canada in 2001, and in the U.S. in 2007, and the MLB Network in 2009. In all leagues but the NFL, individual teams negotiate with local broadcasters to air most of their games (NFL teams do not negotiate local broadcast contracts, but are allowed to negotiate their own television deals for pre-season games).

a shared fund which is divided equally among all teams, and teams can keep all other revenue for themselves (this leads to an estimate of teams sharing approximately 15% of all revenues).[13] The NHL has a revenue sharing program that allocates 6% of total league revenue, primarily away from the top 10 revenue-generating teams, to financially struggling teams. Thus, the effective revenue sharing tax would be above 6%.[14]

Closely related to revenue sharing is the "luxury tax" (sometimes called a competitive-balance tax). This is a payroll tax that typically only kicks in when the total payroll of a team exceeds a predetermined threshold, thus there is a tax levied on money spent above a predetermined limit set by the corresponding sports league.[15] For every dollar a team spends above the tax threshold, those exceeding the limit must also pay some fraction to the league, with the tax increasing with the number of times in which the club exceeds the threshold. The money derived from this tax is distributed among the teams with smaller payrolls. The first luxury tax in professional sports was introduced in 1996 by MLB as part of its Collective Bargaining Agreement (CBA). This luxury tax forces MLB teams with high payrolls to pay a dollar-for-dollar penalty. These funds go into a central MLB fund and it is used for marketing programs. In 1999, the NBA also introduced a luxury tax.

While revenue sharing and luxury taxes affect teams incentives to hire players, a more forceful institution that affects hiring is a "salary cap". This is a limit on the amount of money a club can spend on players' salaries that is negotiated in CBAs between players' unions and team owners. The cap is usually defined as a percentage of average annual revenues and limits a club's investment in playing talent. In 1984, the NBA became the first league to introduce salary cap provisions.[16] Salary caps can be either hard or soft. Under a hard salary cap, the league sets a maximum amount of money allowed for player salaries, and no team can exceed that limit. The NFL and the NHL currently have hard salary caps. A soft salary cap sets a limit to players' salaries, but there are exceptions that allow teams to exceed the cap. In the NBA, for example, teams can exceed the salary cap when keeping players that are already on the team (Dietl et al., 2010). MLB has no salary cap.

[13]NFL revenue sharing was obtained from http://money.cnn.com. NBA revenue sharing was obtained from http://sportsbusinessdaily.com. MLB revenue sharing was calculated using data from http://awfulannouncing.com and http://forbes.com.

[14]A maximum of 50% of the redistribution commitment is drawn from the top 10 highest-grossing teams based on pre-season and regular season revenue. Each team's contribution is based on how much they earn over and above the 11th-ranked team (implying that the teams in the 8-10 spots contribute less than the top three). NHL revenue sharing was obtained from http://ontheforecheck.com.

[15]See Dietl et al. (2010).

[16]For details, see Fort and Quirk (1995), Szymanski (2003), and Vrooman (1995)

Finally, another common institution is the "draft". This is a process used to allocate certain players to teams. In a draft, teams take turns selecting from the pool of new players that want to start playing in the league, with the order being determined (partly) by teams performance last season so that worse performing teams go first. A draft could therefore affect teams incentives to hire good players because good performance one year affects draft order the next year.[17] Together with the "player reservation" system, the draft can potentially perform similar functions as revenue sharing by distributed income from strong to weak teams through selling players. Free agency, which was first introduced in 1976 in MLB and shortly afterwards in the other sports, however, has greatly undermined the impact of the draft as an institution that affects the allocation of talent across teams. [18]

Table 1 summarizes the different regulatory institutions in the Big Four US sports leagues. Hard salary caps were binding for between 6 and 18 teams in NFL, and for between 5 and 13 teams in NHL between 2011 and 2015. In that period between 5 and 6 teams in NBA and between 1 and 4 teams in MLB paid luxury taxes. A fairly clear ranking of the four leagues arises in terms of the amount of regulatory institutions affecting teams hiring of players. The NFL has the strongest regulatory institutions, followed by the NBA and the NHL. MLB has the least regulatory institutions.[19]

[17]The effect of a draft on teams incentive to hire skills depends on what is assumed about wage setting. If wage setting is competitive, the order in which teams get to choose the specific players is unimportant because better players will also receive higher wages. If wages are not competitive, however, the draft order may matter, which will affect teams incentives to hire good players because good performance one year affects draft order the next year. Whether the draft increases or decreases incentives to hire good players is still unclear, however, and depend on whether early or late draft picks offer "better deals" relative to the wages they are paid. Conventional wisdom seems to assume that picking early in the draft allows teams to get better players at a lower cost, however, empirical work by Massey and Thaler (2012) have suggested that in the NFL draft, players drafted later actually offer greater value relative to the salaries they are paid.

[18]The reserve clause was devised in the National Baseball League in 1879 to restrict competition on hiring of superstar players. Each team in the league was allowed to "reserve" five of its players, implying that other owners could not attempt to hire them in the end-of-season market. On the revenue sharing implications of the reservation and draft system see Fort and Quirk (1995).

[19]As a check on this ranking, we informally polled several known sports economists regarding the regulatory rank of the Big Four leagues. The answers we received confirmed the ranking, with the only slight point of disagreement being the relative standing of the two "in-between" leagues, NBA and NHL.

Table 1: Regulations across US sports leagues

League	Revenue shared (%)	Salary cap	Luxury tax
NFL	61	hard	no
NBA	50	soft	yes
NHL	6	hard	no
MLB	15	no	yes

3 The Model

3.1 Teams and competition

The league consists of N ($N \geq 2$) teams that compete in a tournament. Teams play against each other once, which determines the number of games, $(N-1)$. We take N as exogenous, and most of the analysis will consider $N = 2$.[20]

The role of teams in a league is to hire, and coach, players. We model this by assuming that teams invest in a one-dimensional input called skill or talent, S, that will influence their probability of winning as well as the entertainment value of each game. Teams choose S to maximize expected profits, taking league rules and institutions, and other teams' hiring decisions as given.

3.2 Team revenue

Following the sports literature we assume that team i's revenue is proportional to the (exogenously given) size of its local fan base or population, F_i, and to the number of games played, $(N-1)$, while also depending on the team's probability of winning the tournament, w_i, as well as the quality of the league, captured by the aggregate amount of talent of competing teams, $\bar{S} \equiv \sum_{i=1}^{N} S_i$, through some function $R(\bar{S}, w_i)$.[21] Total team revenue is thus $(N-1)F_i R(\bar{S}, w_i)$. We make the following standard assumptions on the revenue function, R:

[20]In Gonzalez-Eiras et al. (2018) we allow the league to choose season length and thus have teams playing against each other n times. See Christenfeld (1996) for a critical comparison of season length across several sports.

[21]The probability of winning the tournament, w_i, is a function of $(S_1, \ldots, S_N)$. We describe this in section 3.4.

Assumption 1.

$$\frac{\partial R}{\partial \bar{S}} \geq 0, \quad \frac{\partial^2 R}{\partial \bar{S}^2} \leq 0, \quad \lim_{\bar{S}\to\infty} \frac{\partial R}{\partial \bar{S}} = 0$$
$$\frac{\partial R}{\partial w_i}|_{w_i=\frac{1}{N}} > 0, \quad \frac{\partial^2 R}{\partial w_i^2} < 0, \quad \frac{d^2 R}{dS_i^2} < 0.$$

This reduced-form revenue function can be rationalized as capturing two sources of income to teams, as in Vrooman (1995). First, there is the prize that the winner of the championship earns, and therefore increases with win probability. This could include direct rewards as well as the monetary value of qualifying for the playoffs. Second, there is broadcasting revenue. This increases with the overall quality of the league which is why revenue increases with $\bar{S}$.[22]

The broadcasting revenue, however, also depends on the "suspense" of the outcome. As w_i gets too large, suspense obviously goes down, which is why the function is strictly concave in win probability.[23] We allow this second effect to possibly dominate, and thus only require that the marginal effect of win probability be positive when team i's probability of winning, w_i, is average, $\frac{1}{N}$. To guarantee a unique solution to teams' optimization problem we require that the revenue function be concave in own skills.

For tractability we assume the following revenue function to derive some results.[24]

Assumption 2.

$$R(\bar{S}, w_i) = k + \log(w_i).$$

Where $k = \sum_i \log(F_i)$ is chosen such that revenue per game is positive for the efficient allocation. Since revenues under assumption 2 do not directly depend on the level of skills, $\bar{S}$, we further assume that in order to generate positive revenues it must be the case that $\bar{S} \geq \bar{S}^*$, with $\bar{S}^* > 0$. The rationale for this additional assumption will become clear when deriving the efficient allocation from the perspective of the league in 6.1.

[22]See Falconieri et al. (2004).

[23]Several empirical studies have documented that the demand for match tickets peaks at a point in which the home team is about twice as likely to win as the away team. See Szymanski (2003) for a survey of these results, as well as inconclusive evidence on seasonal uncertainty. In a recent contribution, Bizzozero et al. (2016) use data from professional tennis and report that both suspense and "surprise" positively affect live TV audience figures.

[24]We will explicitly point out which results require assumption 2. Importantly, our main results in proposition 5 do not require it. Note that this revenue function satisfies $\frac{d^2R}{dS_i^2} < 0$, as required in assumption 1.

3.3 The supply of talent

The sports literature usually assumes that there is a fixed exogenous supply of skills. We deviate from this assumption for two reasons. First, it seems highly questionable in modern sports, especially in the long run. The share of people participating in competitive sports has increased significantly over time, suggesting that the total supply of skilled players is not fixed over time.[25] Moreover, the observation of a large number of foreign players in many sports leagues suggests that the pool of talent is not fixed at the national league level in the short run.[26] Second, for studying the choice of league institutions, it is important that different institutions have an effect on the total amount of talent hired by teams, as this is a significant determinant of aggregate profits. This is of course precluded if the supply of talent is exogenously fixed.[27]

Instead of a fixed supply of skills, we will here assume a perfectly elastic supply such that teams can hire as much talent as they want at a constant per game cost of c. To be clear, this assumption ignores many features of real world labor markets. In practice, the labor market for sport talent has many complications, as players are obviously indivisible and wages typically being set through complicated bargaining with individual players and/or player unions. While the assumption of a constant marginal cost of talent simplifies the analysis, it is not essential to the main mechanisms at work in our model or the empirical results we present later. The only requirement we need is that the complex hiring process is in the aggregate well approximated by a constant cost per unit of talent across teams at a given point in time. Since all the sports we study have adopted free agency, we view this as a reasonable approach.[28]

3.4 Win probability

The second deviation of our model from the standard assumptions in the sports literature is on how teams' talents are translated into winning probability. Usually the probability of winning the *championship* is related to the distribution of skills across teams. Since we

[25] See Rossi and Ruzzier (2018).

[26] The proportion of foreign players is approximately 30% in the MLB and the NBA, 25% in the NHL (considering Canadian players as nationals), and 3.5% in the NFL.

[27] We also note that El-Hodiri and Quirk (1971) and Fort and Quirk (1995) show that, with an exogenous supply of skills, revenue sharing has no impact on the *distribution* of skills, and only reduces players' salaries.

[28] A number of studies show that, under free agency, players are paid their marginal productivity. See, for example, Scully (1974) and Rosen and Sanderson (2001). Moreover, note that a salary cap restricts the total level of talent that a team can hire, but does not affect the fact that the former is proportional to payroll for all teams.

are interested in examining how differences in the fundamental rules of different sports may affect the choice of institutions, we will instead model the probability of winning *individual games.*

In particular, we assume that the probability that team i defeats team j in a match is given by:

Assumption 3.

$$w_{ij} = \frac{S_i^\alpha}{S_i^\alpha + S_j^\alpha}.$$

This formulation has been widely used in the contest literature.[29] Skaperdas (1996) shows that assumption 3 is the only formulation that depends on the ratio of skills and satisfies a series of basic axioms, including independence of irrelevant alternatives. The parameter α reflects how much skills matter for the win probability. We therefore refer to it as the *productivity of skills.* More formally, since the odds of winning is given by the ratio of the probability of winning to the probability of losing, α corresponds to the elasticity of the odds of winning to the ratio of skills $\frac{S_i}{S_j}$.

The parameter α will depend on the fundamental rules of the sport that teams are playing. As an extreme and trivial example, if the sport being played simply involved flipping a coin to declare a winner, α would equal 0 as the win probability would always be a half. Conversely, if the sport involved comparing players average height, α would equal infinity as the team with more talent (taller players) would always win for sure. In reality the level of α depends on many details of the rules, such as the number of times each team gets to be on the offensive in a given game and the amount of randomness involved in scoring. In section 4.3 we estimate this elasticity for each of the major four North American sports and show that it differs significantly, going from 0.16 for baseball to 1.54 for basketball.

It is worth noting that the previous literature on comparing different sports usually focuses on ex post outcomes, such as the probability of winning the championship, or the dispersion of wage bills across teams. These measures are clearly influenced by the institutions adopted by the leagues and therefore do not represent a *fundamental* characteristic of the corresponding sports.[30] In contrast, the elasticity of the odds of winning to the ratio of skills describes a characteristic of sports that is practically invariant to their corresponding institutions.

[29] For applications in the sports literature for the probability of winning the championship see Szymanski's (2003) analysis of pay-performance sensitivity.

[30] For example, the dispersion of wage bills is affected by the degree of revenue sharing. The probability of winning the league is also a function of season length (see Christenfeld (1996)).

Given the above formulation for the probability of winning a game, we assume that the team that wins the championship is simply the team that wins the most games each season.[31] The probability of winning the tournament, which determines revenue, is then given by a complex relation between the skills of team i, skills of other teams, and parameter α. For now, we simply represent this relation by[32]

$$w_i = W(S_i, \vec{S}, \alpha).$$

3.5 Team's hiring decision without regulatory institutions

In the "laissez-faire" case, when there are no regulatory institutions in place, each team decides how much talent, S_i, to hire by maximizing profits, $\pi_i(S_i)$, taking as given the amount of talent hired by the other teams, S_j:

$$\max_{S_i} \ \pi_i(S_i) = (N-1)F_i R(\bar{S}, W(S_i, \vec{S}, \alpha)) - c(N-1)S_i.$$

Under assumption 2, the first-order condition for team i is

$$\frac{1}{w_i}\frac{dw_i}{dS_i} - \frac{c}{F_i} \leq 0. \tag{1}$$

This first-order condition holds as an equality if $S_i > 0$.

3.6 League's objective and efficient allocation

We assume the league wants to maximize aggregate expected profits.

$$\sum_i \pi_i(S_i).$$

If a planner were able to perfectly dictate teams' choice of skills $\vec{S}$, it would solve the

[31]In practice of course, the leagues we consider use a play-off format. We abstract from this and consider qualifying to the playoffs as the reward to "winning" teams.

[32]We use vector notation such that $\vec{S} \equiv (S_1, S_2, \ldots, S_N)$.

following problem:

$$\max_{\vec{S}} \quad \sum_{i=1}^{N} \pi_i\left(S_i\right)$$
$$= (N-1)\sum_{i=1}^{N} F_i R(\bar{S}, W(S_i, \vec{S}, \alpha)) - c(N-1)\bar{S}.$$

Note that we can reparameterize the planner's problem as choosing $\bar{S}$, and $N-1$ win probabilities, w_i. We now characterize the profit maximizing league-wide allocation of talent.

Proposition 1. Under mild concavity assumptions on total revenue (a sufficient condition being $\frac{\partial^2 R}{\partial \bar{S} \partial w_i} = 0$) there is a unique profit maximizing outcome that solves the planner's problem, $\vec{S}^P$. Furthermore $\bar{S}^P \equiv \sum_{i=1}^{N} S_i^P$, and w_i^P do not depend on α.

Proof. All proofs are in the appendix. □

3.7 Decentralized equilibrium allocation

Henceforth we will restrict the analysis to a two-team league, $N = 2$. Note that with $w_1 = w_{12}$ given by assumption 3, $w_2 = 1 - w_1$, $\frac{\partial w_1}{\partial S_1} = -\frac{S_2}{S_1}\frac{\partial w_1}{\partial S_2}$, and $\frac{\partial w_1}{\partial S_1} = -\frac{\partial w_2}{\partial S_1}$.

We now turn to comparing the decentralized outcome to the allocation that maximizes league-wide profits. We start by showing existence and uniqueness of the Nash equilibrium.

Proposition 2. Under assumption 2, there exists a threshold $\underline{\alpha}$ such that sports with $\alpha < \underline{\alpha}$ do not generate revenues. If $\alpha \geq \underline{\alpha}$, there exists a unique Nash equilibrium.

A useful way of understanding whether and how the decentralized equilibrium differs from the profit maximizing league-wide allocation is to examine the first-order conditions for the planner and team 1 with respect to S_1. The first order condition for the planner is given by:

$$F_1 \frac{\partial R(\bar{S}, w_1)}{\partial \bar{S}} + F_2 \frac{\partial R(\bar{S}, w_2)}{\partial \bar{S}} + F_1 \frac{\partial R(\bar{S}, w_1)}{\partial w_1}\frac{dw_1}{S_1} - F_2 \frac{\partial R(\bar{S}, w_2)}{\partial w_2}\frac{dw_1}{S_1} = c.$$

The first order condition for team 1 is given by:

$$F_1 \frac{\partial R(\bar{S}, w_1)}{\partial \bar{S}} + F_1 \frac{\partial R(\bar{S}, w_1)}{\partial w_1}\frac{dw_1}{S_1} = c.$$

Comparing the first-order conditions for the team's problem to the first-order conditions for the planner's problem, we see that they are the same, except that the marginal profit from hiring an additional unit of skill includes two additional terms in the planner's case: $F_2 \frac{\partial R(\bar{S}, w_2)}{\partial \bar{S}}$ and $-F_2 \frac{\partial R(\bar{S}, w_2)}{\partial w_2} \frac{dw_1}{S_1}$. These two terms reflect two externalities that teams fail to incorporate when choosing skill investments. The first is positive and reflects the fact that when team 1 increases its skill level, this increases team 2's revenue through its effect on total league talent. This creates a *common pool* problem. The failure of team 1 to internalize this externality implies that it will tend to hire *less* skill than what is optimal for the league as a whole. The second term is negative and reflects the fact that when team 1 employs more skill and raises its own win probability this lowers the win probability of team 2 because winning is a *zero sum* game. The failure of team 1 to internalize this externality implies that it will tend to hire *more* skill than what is optimal for the league.

Importantly, the strength of the zero sum externality, and thus a team's incentive to hire skills, depends crucially on the productivity of skills, α. In particular we have the following proposition.

Proposition 3. There exists $\underline{\alpha}(F_1, F_2) > 0$, with $\lim_{F_2 \to \infty} \underline{\alpha}(F_1, F_2) = 0$, such that for $\alpha > \underline{\alpha}(F_1, F_2)$ the equilibrium levels S_1^* and S_2^* are increasing in α.

Proposition 3 highlights the first insight of the model: under reasonable assumptions (i.e. when $\alpha > \underline{\alpha}(F_1, F_2)$), teams' hiring incentives are increasing in the productivity of skills, α. The intuition behind the result is simple. When α is high, the effect of an additional unit of skills on the probability of winning, and thus on revenue, is stronger. The reason for the restriction $\alpha > \underline{\alpha}(F_1, F_2)$ has to do with the tail behavior of the win probability function when one of the skill levels is close to zero (see appendix 6.3).

Figure 1 provides numerical results that illustrate the main insights of Proposition 3. Using the functional form from assumption 2 and a given set of other model parameters, it shows the teams' best response functions, the associated Nash equilibrium as well as the efficient allocation when $\alpha = 0.6$ and $\alpha = 1$. From the figure, we see that at these parameters, $\alpha = 0.6$ is high enough that the "zero sum" externality dominates and the Nash equilibrium involves both teams choosing higher levels of skills than the efficient outcome. As discussed above, a higher α strengthens the "zero sum" externality, which causes the best response functions to shift outward and upward, thereby increasing the equilibrium level of skills for the two teams. In contrast, the efficient level of total talent is independent of α so the efficient allocation only moves along the negatively sloped 45-degree line.

Figure 1: Decentralized Nash equilibria and efficient outcomes

Model outcomes for $\alpha = 0.6$ and $\alpha = 1$:

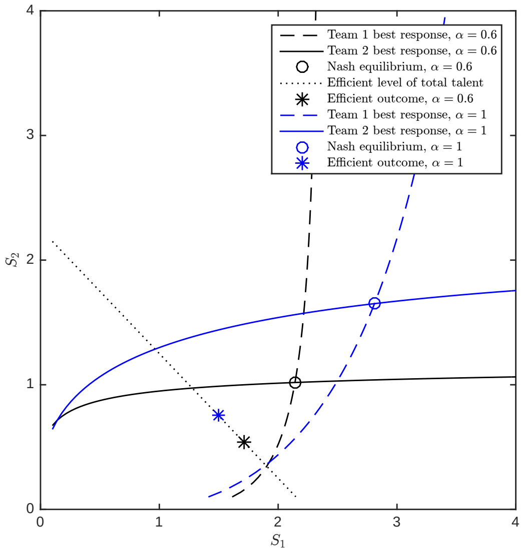

The graph shows numerical results from the model under assumption 2. Solid and long-dashed lines shows teams' best response functions and circles show Nash equilibria. Stars show efficient allocations, while the short-dashed lines shows allocations involving the efficient level of total talent. Note that the efficient level of total talent is independent of α. The additional parameter values used are $F_1 = 2$, $F_2 = 1$, $\bar{S}^* = 2.2$, and $c = 1$.

3.8 The role of regulation and redistributive taxation

In the previous section, we saw that the decentralized Nash equilibrium in the league generally does not maximize total league profits as teams hire too much (or too little) skill. This suggests that the league may centrally want to introduce regulatory institutions that affect hiring incentives.

We will focus on revenue sharing since it is the canonical example of a redistributive, regulatory institution in team sports. In addition, we assume the league can introduce a salary tax that only affects one team. The motivation for this is that league revenue depends on choices by both teams, thus it is generically not possible to implement the profit maximizing allocation using only a single policy instrument. Introducing an asymmetric hiring tax solves this issue allowing us to characterize first-best institutions.[33]

Following the literature we introduce revenue sharing by assuming that each team keeps $1-\tau$ of its own revenues and receives $\frac{\tau}{N}$ of aggregate team's revenues (including own revenue). Since sports leagues always have a relatively modest amount of teams we also follow the sports literature in assuming that teams are not myopic but take into account that their decisions affect aggregate shared league revenues.[34] For the strong team we introduce a hiring tax at rate s whose proceeds are rebated lump sum to both teams.In this case, team i's problem becomes:

$$\max_{S_i} \pi_i\left(S_i\right)=\left(1-\frac{\tau}{2}\right) F_i R\left(S_i+S_j, w_i\right)+\frac{\tau}{2} F_j R\left(S_i+S_j, w_j\right)-c(1+s \mathbb{1}_{i=1}) S_i+T_i,$$

where the indicator function tells us that for team 1 the per unit cost of skills is now $c(1+s)$, and T_i are the lump sum transfers, $T_1+T_2=csS_1^*$. We are interested in examining how regulations affect total league revenue. We begin by establishing that the model has a unique Nash equilibrium when regulatory institutions are introduced.

Proposition 4. Under assumption 2, the model has a unique Nash equilibrium.

Next we analyze the effect of regulations on team behavior and their role in increasing

[33]For the same reason, in a two-team league other institutions that directly or indirectly affect teams' hiring of talent are likely to implement the profit maximizing allocation.

[34]It is worth noting that under this non-myopic assumption, an easy way for the planner to always implement the efficient outcome is to implement full profit sharing. With non-myopic teams and full profit sharing each team will choose their own skill so as to maximize league profits. We view full profit sharing as infeasible in practice, however, because teams in the real world can spend money on things other than talent. Under full profit sharing, teams would therefore have the incentive to incur unnecessary costs to lower their own profits. For example, the owner may hire a relative into an overpaid consultant position.

total league profits. The first-order condition for team i is:

$$\left(1 - \frac{\tau}{2}\right) F_i \left(\frac{\partial R}{\partial S_i} + \frac{\partial R}{\partial w_i}\frac{dw_i}{dS_i}\right) + \frac{\tau}{2} F_j \left(\frac{\partial R}{\partial S_i} + \frac{\partial R}{\partial w_j}\frac{dw_j}{dS_i}\right) - c(1 + s\mathbb{1}_{i=1}) \leq 0. \tag{2}$$

Comparing this new first-order condition with the one without regulations, we see that the terms corresponding to the effect of additional skill on the other team now appear, although only with a weight of $\frac{\tau}{2}$. In addition, the terms corresponding to the effect of additional skill on team i's own revenue now only enter with a weight of $1 - \frac{\tau}{2}$. Intuitively this shows that revenue sharing has two effects: a) it gets teams to partially internalize the externalities from before, and b) it tends to lower the incentive to invest in skill because teams now only get to keep $1 - \frac{\tau}{2}$ of own revenue. The hiring tax reduces team 1's incentive to hire skills (or increases it if the tax is negative), while having no effect on team 2.

When the decentralized equilibrium involves teams choosing too high skill levels, this suggests that (higher levels of) revenue sharing can improve league profits by causing teams to lower the amount of skill they hire in equilibrium. In particular we have the following proposition.

Proposition 5. When the decentralized equilibrium involves teams choosing too high skill levels, there exists a policy $(\tau^*(\alpha), s^*(\alpha))$ that leads to the profit maximizing allocation, and $\frac{d\tau^*(\alpha)}{d\alpha} \geq 0$. An increase in α reduces the dispersion of talent across teams. An increase in the dispersion of fan bases, $\frac{F_1}{F_2}$, increases revenue sharing. Alternatively, the profit maximizing allocation can be achieved with revenue sharing $\tau^*(\alpha)$ and a binding salary cap on team 1.

Proposition 5 shows that regulatory institutions can achieve the profit-maximizing outcome if sports have a sufficiently high α. Moreover, leagues for sports that involve higher levels of α will optimally want to introduce higher levels of revenue sharing to counter the externalities discussed above. This is the second main insight of our model. In other words, in sports where teams' win probabilities are more sensitive to relative skill, teams have inefficiently strong incentives to hire talent. Leagues authorities in these sports therefore choose to introduce higher levels of revenue sharing to dampen teams' individual incentives. While proposition 5 is silent about the case when the decentralized equilibrium features too low skill levels, in Gonzalez-Eiras et al. (2018) we show that in this case leagues have an incentive to increase season length.[35]

[35] Basically, increasing the number of times that two teams play against each other increases the "effective" productivity of skills, the elasticity of winning at least half the matches with respect to skills.

Proposition 1 showed that, conditional on fan bases, all sports have the same profit-maximizing win probability, w_1^P. Since this allocation is attainable to all sports for which the decentralized equilibrium involves teams choosing too high skill levels, and given that $\frac{w_1}{w_2} = \frac{S_1^\alpha}{S_2^\alpha}$, sports with a higher α tend to produce more dispersed win probabilities for a given distribution of talent. Thus, as the productivity of skills is increased, to attain w_1^P talent has to be more evenly distributed across teams. If we increase the dispersion of fan bases, which would happen when the league expands incorporating cities with lower fan bases, the incentives to redistribute revenue increase.

Figure 2 illustrates these points. For a given set of other model parameters and $\alpha = 0.6$, the black lines and circle show the best response functions and the Nash equilibrium without revenue sharing. The black asterisk shows the efficient outcome. The Nash equilibrium has the two teams hiring more talent than what is efficient. There is thus scope for improving efficiency by introducing revenue sharing. The blue lines and markers therefore show the effect of introducing a revenue sharing of 40% ($\tau = 0.40$). As discussed above, revenue sharing reduces team's individual incentives to hire skills and, as a result, shifts the best response functions and the equilibrium level of skills for both teams inwards so that they are much closer to the efficient levels.

Note that in figure 2 introduction of revenue sharing alone is not enough for the league to implement the efficient allocation. To do this, the league must also introduce an additional instrument that affects hiring decisions. Figure 3 shows how the optimal level of revenue sharing, τ, varies with the level of α assuming that the league also introduces an optimally set hiring tax, or a binding salary cap, for the strong team. As the figure illustrates revenue sharing is increasing in the productivity of skills when α is high, and there is no revenue sharing when α is low.

4 Empirical results

In this section we use data on match outcomes and players' payroll to estimate how the productivity of skills differs across the sports played in the Big Four US leagues. We then relate this to the use of regulatory institutions in the different leagues.

4.1 Data

The first piece of data we need is information about team skills. As discussed in detail in the next section, this data will be based on team's total payroll (skill expenditure). Since our objective is to make comparisons across the different sports' leagues, it is important

Figure 2: Effects of introducing revenue sharing

Improving efficiency with revenue sharing, $\alpha = 0.6$:

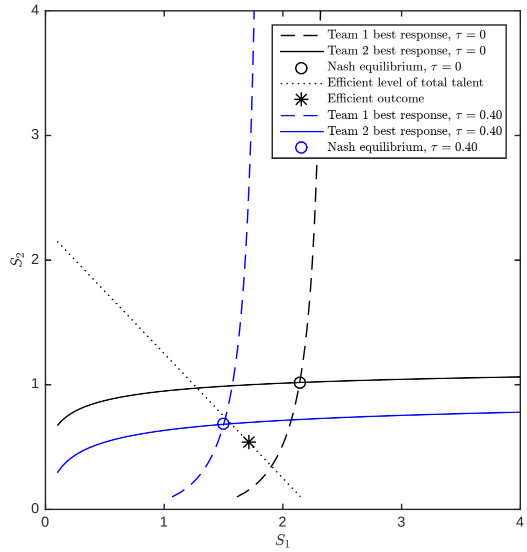

The graph shows numerical results from the model. Solid and long-dashed lines shows teams' best response functions and a circles show Nash equilibria. A star shows the efficient allocation. Black lines correspond to outcomes without revenue sharing, while blue lines correspond to outcomes with 40% revenue sharing. The additional parameter values used are $F_1 = 2$, $F_2 = 1$, $\bar{S}^*$, $c = 1$ and $s = 0$.

Figure 3: Optimal revenue sharing

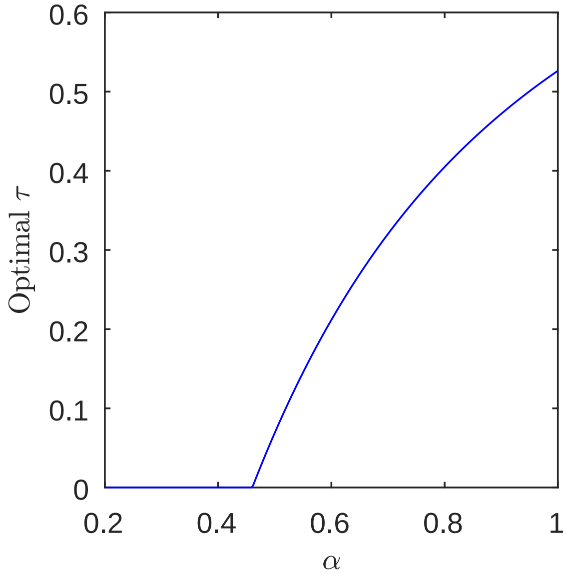

The graph shows numerical results from the model. The line shows the optimal level of revenue sharing, τ, as a function of α, when the league simultaneously introduces an optimally set wage subside s. The additional parameter values used are $F_1 = 2$, $F_2 = 1$, $\sigma = 0.5$, $n = 1$ and $K = 2$, $c = 1$.

that the payroll data we use is comparable across the different leagues. For this reason we use salary data from the yearly Global Sports' Salaries Survey (GSSS). The GSSS reports team salaries for a range of different sports leagues, including the NFL, NBA, NHL and MLB. Importantly, the GSSS data is constructed explicitly with the aim of making comparisons across different sports' leagues.

We base our measure of total team skill expenditure on the GSSS definition of player salaries, which includes both players' base salaries and any performance bonuses that have been paid out.[36] From the 2012-2015 GSSS data, we construct total yearly team payrolls for each team for each season.[37]

From these base payroll data we make one correction that aims to deal with a phenomenon paticular to the NFL, namely quarterback injuries. Relative to the other sports, the quarterback on an NFL team makes up a very large fraction of the yearly payroll. Combined with the high incidence of injuries in the NFL, (Hootman et al., 2007) this creates a particularly strong measurement error issue when the yearly payroll is used to measure skills in individual games: In games where the starting quarterback is injured, the yearly payroll figure will vastly overstate the team's actual skill expenditure. Accordingly, we use game starter rosters from Pro Football Reference and quarterback salaries from Over the Cap to subtract out the salary of the starter quarterback in games where he is injured.[38] The validity of this adjustment is strongly supported in our data: If we do not adjust the NFL data for quarterback injuries the empirical results in the next section show *no* significant effect of skills on win probability for the NFL.[39]

The second piece of data we use is data on game outcomes in each of the sports leagues. For each league we collected data on all regular season games during the four seasons that correspond to our salary data.[40]

[36]We include bonus payments in our measure of skill expenditures for two reasons: First, the key choice variable in our model is how much teams choose to spend in total on talent, regardless of whether this is paid out as base salary or as performance related bonuses. Second, as a purely practical consideration, we are unaware of any data source on team payrolls which makes it possible to separate out bonus and performance pay in a consistent way across multiple sports. Note that GSSS reports salaries before the end of the playing season so some performance bonuses, those of the most succesful teams, are not included.

[37]The GSSS reports data on average player salary. We convert this to total team payrolls by multiplying by the typical number of players on a team in each of the four leagues (53 in the NFL, 29 in MLB, 25 in the NHL and 15 in the NBA). To the extent that not all teams within a league carry the same number players on the roster for the full year, this will introduce some measurement error.

[38]Data taken from https://www.pro-football-reference.com and https://overthecap.com.

[39]Without correcting for quarterback injures, our estimate of α_{NFL} in Table 3 is never signficantly different from zero.

[40]For the NFL, NBA and NHL, where seasons run from the Fall of year t to the Spring of year $t+1$, these are: 2011-2012, 2012-2013, 2013-2014, 2014-2015. For the MLB where seasons run from the Spring of year t to the Fall of year t, these seasons are: 2012, 2013, 2014, 2015.

Based on these sources, we construct a data set where each observation corresponds to a particular game. For game g taking place in league l in season t between the home team i and the away team j, we define the $SkillRatio_{tlgij}$ as the ratio of the skill expenditure (payroll) of the home team to the away team. We define the dummy variables NFL_l, NBA_l, NHL_l and MLB_l as indicators for the different leagues. We define the variable $RegulatoryRank_l$ as the ranking of league l in terms of the extent of regulatory institutions such as revenue sharing (where NFL has rank 1, NBA rank 2, NHL rank 3, and MLB rank 4, cf. Table 1). We define the variable $HometeamWin_{tlgij}$ as an indicator for whether the home team i won against the away team j in game g in league l in season t.[41]

Table 2 presents summary statistics of our data.

4.2 Empirical specification and estimation

We now discuss how to take the theoretical model from Section 3 to the data and estimate the relationship between skill input and win probabilities for each sport. To accommodate the fact that we now consider four different leagues and use data from multiple seasons and games, we adapt the model setup and notation from Section 3 as follows: We let c_{tl} denote the cost of talent in league l during season t. We let S_{tli} denote the skill employed by team i from league l during season t. We let w_{tlgij} denotes the probability that team i beats team j in game g during season t in league l. Finally, we let α_l denote the returns to skill in league l.

With this notation, the win probability in a given game depends on skill inputs as follows:

$$w_{tlgij} = \frac{S_{tli}^{\alpha_l}}{S_{tli}^{\alpha_l} + S_{tlj}^{\alpha_l}}. \tag{3}$$

Conveniently for the empirical implementation, simple algebra shows that this formulation is equivalent to a standard logit model:

$$w_{tlgij} = \text{logistic}\left(\alpha_l \log\left(\frac{S_{tli}}{S_{tlj}}\right)\right). \tag{4}$$

We use this logit formulation to estimate the different α_l's from the data described in

[41] A minor complication arises in defining this variable and taking our theoretical model to the data, due to the possibility of tie games. While our model does not allow for ties, ties are possible in some of the sports we analyze, although they are extremely rare. We present results here where tie games are normalized as wins for the home team. Due to the very few ties in our data, however, dropping ties or instead treating them as losses for the home team leads to virtually identical estimates.

Table 2: Summary statistics

NFL

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Hometeam skill expenditure	1,024	110,710	13,389	59,519	155,820
Skill ratio	1,024	1.016	0.189	0.488	2.051
Log skill ratio	1,024	0.000	0.179	-0.718	0.718
Hometeam wins	1,024	0.578	0.494	0	1

NBA

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Hometeam skill expenditure	4,679	67,473	11,600	33,150	102,150
Skill ratio	4,679	1.034	0.276	0.354	2.828
Log skill ratio	4,679	0.000	0.258	-1.040	1.040
Hometeam wins	4,679	0.588	0.492	0	1

NHL

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Hometeam skill expenditure	4,116	61,353	8,384	38,000	83,500
Skill ratio	4,116	1.018	0.196	0.502	1.993
Log skill ratio	4,116	0.000	0.190	-0.690	0.690
Hometeam wins	4,116	0.548	0.498	0	1

MLB

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Hometeam skill expenditure	9,720	109,893	41,870	23,780	232,870
Skill ratio	9,720	1.167	0.726	0.115	8.720
Log skill ratio	9,720	0.000	0.554	-2.166	2.166
Hometeam wins	9,720	0.536	0.499	0	1

The table shows summary statistics for key variables for each of the sports leagues: NFL, NBA, MLB, and NHL. The unit of observation is a game. Hometeam skill expenditure is measured in 1,000s of dollars. Differences in the number of observations across the different leagues reflect differences in typical number of games per season as well as idiosyncratic differences in the number of games played across seasons for example due to player strikes.

Section 3. We can relate w_{tlgij} to the variables in our data as follows:

$$P\left(HometeamWin_{tlgij}\right) = w_{tlgij}. \tag{5}$$

For the ratio $\frac{S_{tli}}{S_{tlj}}$, we note that under the assumption that the cost per unit of skill, c_{lt}, is constant across teams in a given league and season, the ratio of employed skills is simply equal to the ratio of skill expenditures from our dataÆ

$$\frac{S_{tli}}{S_{tlj}} = \frac{c_{tl}S_{tli}}{c_{tl}S_{tlj}} = SkillRatio_{tlij}. \tag{6}$$

Accordingly, we use the ratio of skill expenditures (payrolls) as a measure of skill ratio between teams throughout the empirical analysis.[42] Combining (4), (5) and (6), we see that we can estimate α_l from our data using the following logit model:

$$P\left(HometeamWin_{tlgij}\right) = \text{logistic}\Big(\alpha_l \log\left(SkillRatio_{tlij}\right)\Big) \tag{7}$$

Rather than estimate this equation individually for each of the four leagues in our data, however, it will be convenient to use dummy variables for the different leagues and formulate one model:

$$\begin{aligned} P\left(HometeamWin_{tlgij}\right) = \text{logistic}\Big(&\alpha_{NFL} \log\left(SkillRatio_{tlij}\right) \times NFL_l \\ &+\alpha_{NBA} \log\left(SkillRatio_{tlij}\right) \times NBA_l \\ &+\alpha_{NHL} \log\left(SkillRatio_{tlij}\right) \times NHL_l \\ &+\alpha_{MLB} \log\left(SkillRatio_{tlij}\right) \times MLB_l\Big) \end{aligned} \tag{8}$$

Estimation of (8) will be our main approach to estimating α_l for the individual leagues. The conclusion from the theoretical model from Section 3 is that sports with inherently higher levels of α_l tend to adopt more regulatory institutions. A crude way of capturing this in the empirical model is to substitute in $\alpha_l = \bar{\alpha} + \rho \cdot RegulatoryRank_l$ to get:

$$\begin{aligned} P\left(HometeamWin_{tlgij}\right) =\text{logistic}\Big(&\bar{\alpha} \log\left(SkillRatio_{tlij}\right) \\ &+ \rho \log\left(SkillRatio_{tlij}\right) \times RegulatoryRank_l\Big) \end{aligned} \tag{9}$$

[42]There is strong empirical evidence that payrolls are significant predictors of playing success in different sports. See Simmons and Forrest (2004) and Szymanski and Kuypers (1999). Section 3.3 and footnote 28 provides additional discussion regarding the assumption that total skill expenditures are proportional to the employed level of skill within each league

In this model, ρ measures the relationship between α_l and the degree of regulation in the league. If leagues with higher α_l optimally adopt more regulatory institutions, we should have $\rho < 0$.

We note that the estimating equations above treat the home and away team symmetrically. In practice, however, there is evidence that in many sports home teams are more likely to win. While our theoretical model does not include this, for robustness we also estimate versions of the logit models above that allow for a league-specific homefield advantage, ψ_l. This simply corresponds to include the un-interacted indicator variables NFL_l, NBA_l, NHL_l and MLB_l as regressors:

$$\begin{aligned} P\left(HometeamWin_{tlgij}\right) = \text{logistic}\Big(&\alpha_{NFL} \log\left(SkillRatio_{tlij}\right) \times NFL_l \\ &+\alpha_{NBA} \log\left(SkillRatio_{tlij}\right) \times NBA_l \\ &+\alpha_{NHL} \log\left(SkillRatio_{tlij}\right) \times NHL_l \\ &+\alpha_{MLB} \log\left(SkillRatio_{tlij}\right) \times MLB_l \\ &+\psi_{NFL} NFL_l + \psi_{NBA} NBA_l + \psi_{NHL} NHL_l + \psi_{MLB} MLB_l\Big) \end{aligned} \tag{10}$$

$$\begin{aligned} P\left(HometeamWin_{tlgij}\right) = \text{logistic}\Big(&\bar{\alpha} \log\left(SkillRatio_{tlij}\right) \\ &+\rho \log\left(SkillRatio_{tlij}\right) \times RegulatoryRank_l \\ &+\psi_{NFL} NFL_l + \psi_{NBA} NBA_l + \psi_{NHL} NHL_l + \psi_{MLB} MLB_l\Big) \end{aligned} \tag{11}$$

4.3 Marginal productivity of skills and regulatory institutions

We now turn to estimating α_l for the different sports using the logit models from above. Table 3 presents estimates of the logit models in equations 8-11. Because each observation in our data pertains to a pair (dyad) of teams playing against each other, and because team outcomes may be correlated across seasons, we use dyadic cluster robust standard errors throughout and cluster at the team-level.[43] Column (1) presents estimates of a logit model where the explanatory variables are interactions between the log salary ratio and dummy variables for the four different sports. This corresponds to equation (8) from above and accordingly the estimated coefficients on each of the interaction terms corresponds to estimates of α_l for each of the sports. We see substantial differences across the four sports as $\hat{\alpha}_l$ ranges from 0.16 to 1.54. Moreover we can strongly reject the

[43] See Cameron and Miller (2014) and Aronow et al. (2015). There are 122 teams (clusters) in our data in total.

Table 3: Differences in productivity of skills, α_l, across sports

VARIABLES	(1) Logit: Hometeam win	(2) Logit: Hometeam win	(3) Logit: Hometeam win	(4) Logit: Hometeam win
Log skill ratio × NFL (α_{NFL})	0.983**		1.005**	
	(0.464)		(0.476)	
Log skill ratio × NBA (α_{NBA})	1.538***		1.590***	
	(0.361)		(0.374)	
Log skill ratio × NHL (α_{NHL})	0.739***		0.745***	
	(0.201)		(0.200)	
Log skill ratio × MLB (α_{MLB})	0.161***		0.162***	
	(0.0590)		(0.0603)	
Log skill ratio		2.559***		2.639***
		(0.548)		(0.565)
Log skill ratio × Regul. rank (ρ)		-0.598***		-0.618***
		(0.140)		(0.144)
NFL (ψ_{NFL})			0.320***	0.327***
			(0.0533)	(0.0550)
NBA (ψ_{NBA})			0.371***	0.368***
			(0.0286)	(0.0285)
NHL (ψ_{NHL})			0.194***	0.194***
			(0.0297)	(0.0300)
MLB (ψ_{MLB})			0.143***	0.143***
			(0.0230)	(0.0230)
Observations	19,538	19,538	19,538	19,538
p-value, α equal across sports	0.000		0.000	

The table shows estimates from Logit models with home team victory as the outcome. The explanatory variables used are the log skill expenditure ratio of the home team to the away team, indicators for the different leagues, the regulatory rank of each league, as well as interaction terms between these regressors as shown. Dyadic cluster robust standard errors clustered at the team level are in parenthesis. $* : p < 0.10$, $** : p < 0.05$, $*** : p < 0.01$

Table 4: Pairwise tests of differences across sports

Tests of $H_A : \alpha_l \neq \alpha_k$ vs. $H_0 : \alpha_l = \alpha_k$, p-values:

k: :	NFL	NBA	NHL	MLB
NFL		0.347	0.630	0.082*
NBA			0.055*	< 0.001***
NHL				0.007***
MLB				

The matrix shows p-values from pairwise t-tests of differences in α_l across sports based on the logit model in column (1) of Table 3. The test are based on dyadic cluster robust standard errors clustered at the team level. $* : p < 0.10$, $** : p < 0.05$, $*** : p < 0.01$

hypothesis of a constant α ($p < 0.001$).

Looking at the magnitudes of the estimated α_l's in Column (1), we also note that they accord very well with the prediction of our theoretical model regarding league regulation and the relative regulatory ranking of the four leagues described in Section 3. For three of the four leagues, NBA, NHL and MLB, the ranking in terms of $\hat{\alpha}$ exactly mirrors the ranking in terms of regulation. Among these leagues, the NBA is the most heavily regulated and also has the highest estimated productivity of skills, 1.54, while MLB is the least regulated and has the lowest estimated productivity of skills, 0.16. Moreover, as shown in Table 4, the estimated differences in productivity across these three leagues are statistically significant at least at the 10 percent level.

Turning to the fourth league, the NFL, the point estimate $\hat{\alpha}$ does not fit with the predictions of the model given the NFL's regulatory ranking. The NFL is the most heavily regulated league, and while its estimated productivity of skills of 0.98 is high, it is still noticably lower than the estimated productivity of skills for the NBA . At the same time, however, because the NFL has so few games per season and also has the smallest variance in skill expenditures across teams, the standard error on the estimated productivity of skills in the NFL is relatively large. As a result, we cannot reject that the productivity of skills in the NFL is in fact larger than in each of the other sports (see Table 4).[44]

As a succinct way of capturing the relationship between productivity of skills, $\hat{\alpha}_l$ and the amount of regulation, column (2) of Table 3 presents estimates a model that includes the log salary ratio and an interaction term between the log salary ratio and the leagues ranking in terms of regulation, corresponding to equation (9) from above. We find a highly significant negative coefficient on the interaction term ($\rho < 0$); leagues with higher levels of alpha have systematically more regulation (a lower rank).

Finally, columns (3) and (4) of Table 3, check how estimates are affected by allowing for a league-specific home field advantage, as in equations (10) and (11). We see that our estimates are virtually unaffected by these alternative specifications.

Overall, we conclude that the predictions from our theoretical model fit very well with the observed relationship between productivity of skills and amount of regulatory

[44] Note also that this problem is not one that can readily be solved by additional data collection. Given our current estimates and assuming that standard errors decrease with the square root of the sample size (which due to clustering at the team level is likely an optimistic assumption), we would need to more than double the number of seasons in our data set to be able to reject that α_{NFL} is higher than α_l for the other three leagues at the 5 percent level of significance. Even ignoring the issue of comparability across leagues, we are unaware of any data source that contains consistent payroll data for nine NFL seasons.

Table 5: Productivity of skills and wage dispersion

League	estimated α	wage dispersion
NFL	0.98	0.13
NBA	1.54	0.19
NHL	0.74	0.15
MLB	0.16	0.75

institutions across the four leagues.[45]

4.4 Wage dispersion

The main focus of this paper is the relationship between marginal productivity of skills and regulatory institutions. Our model however also yields an prediction about skill dispersion that we can also confront with data.

According to proposition 5, sports with higher skill productivities will have more skill dispersion. In table 5 we report the estimated marginal productivity of skills together with wage dispersion for the four sports leagues.[46] As expected, the fact that NFL and NHL have salary caps results in these sports having the lowest wage dispersions. But wage dispersion is lower in the NFL than the NHL as would be expected if the marginal productivity of skills were higher in the former. More importantly, wage dispersion is lower in NBA than MLB, which is consistent with our finding that the marginal productivity of skills is significantly higher in the former.

4.5 Alternative explanations for observed regulations

The results above show that our theoretical predictions fit the data well and that differences in the productivity of skills can explain the observed differences in regulatory institutions across league. In this section, we briefly consider whether other, alternative explanations might explain the observed differences as well.

In light of occasional league shutdowns, one can argue that low revenue teams can threaten not to participate unless they get compensation from high revenue teams. Ac-

[45] Szymanski (2003) presents estimates of the pay-performance sensitivity of these, and other, sports using season winning percentage and relative wage bills. He finds that the pay performance sensitivity is highest in NFL, followed by NBA, NHL and lowest in MLB. While Szymanski's results do not capture a fundamental characteristic of the corresponding sports, it is reassuring that we find the same relative standing for three of the four leagues.

[46] Wage dispersion is measured as the average of standard deviations of teams' yearly payrolls.

cordingly, differences in the level of revenue sharing could merely reflect differences in the strength of this threat across leagues. Under this alternative explanation, however, we would expect that revenue sharing is systematically related to the dispersion of fan bases. This is not what we observe in the data.[47]

Another alternative explanation is that leagues that depend more on national broadcasting, and thus depend more on having "neutral" fans, would have more revenue sharing in an attempt to promote competitive balance. While we can not rule it out, we view this as an unsatisfactory explanation for a couple of reasons. First, we note that while revenues from national broadcasting are in fact subject to 100% revenue sharing in all four sports, leagues chose differently how to share other revenue sources. Second, this alternative explanation is incomplete in that it does not explain why the importance of national broadcasting would be correlated with the marginal productivity of skills.

Peeters (2015) studies how leagues use revenue sharing to coordinate skill demand. Leagues have the same productivity of skills, but differ in how homogeneous their fan bases are. Revenue sharing should be higher if teams in the league serve more homogeneous local markets. In the proof of proposition 5 we show that an increase in the dispersion of fan bases increases revenue sharing. Thus, our model, contrary to Peeters (2015), predicts that more heterogeneity would lead to more revenue sharing. We believe this can be a possible explanation for the secular increase in revenue sharing that has accompanied league expansions in recent decades. It is also worth noting that heterogeneity in revenue bases cannot account for the observed discrepancy in the relative ranking of the NFL. While the proportional effect of α and fan base dispersion on the marginal incentive to hire skills is roughly the same, the difference in α across sports is tenfold, while the difference in fan base dispersion is much lower.[48]

In reality different leagues do differ in the heterogeneity of their revenue sources, on the importance of neutral fans, and on a number of dimensions not covered by our analysis. We thus view our explanation for observed institutions as complementary to those provided in the literature.

[47] The ratio of fan bases of teams ranked 8^{th} and 23^{rd} is a convenient measure of dispersion of fan base for a league with 30 teams. Using actual population data for the different US sports we find that this ratio is 2.76 for NFL, 2.56 for NBA, 3.08 for NHL, and 2.36 for MLB. Thus, dispersion of fan bases is only 30% higher in NHL relative to MLB.

[48] To make this comparison we evaluate the marginal productivity of skills at the social planner allocation. Note from footnote 47 that the dispersion of fan bases is only 30% higher in NHL relative to MLB.

5 Conclusions

In this paper, we examined institutional choice across the Big Four US sports leagues. Despite having very similar business models and facing the same economic and legal environment, these leagues exhibit large differences in their use of regulatory institutions such as revenue sharing, salary caps or luxury taxes. Since the four leagues differ in that they play sports with very different rules, it seems natural to associate institutional choice as the optimal response to sports' fundamentals.

Building on a standard model of sports leagues, we showed theoretically that heterogeneity in the characteristics of the underlying sports regarding how skills translate into win probabilities may make it optimal for some leagues to adopt systematically more regulatory institutions. Because they play games against each other, teams' hiring decisions are subject to externalities that may lead to inefficiently high levels of talent. The strength of the externalities depend on the marginal productivity of skills. We find that for sports with a higher marginal productivity of skills, stronger hiring externalities make it optimal for leagues to introduce regulatory institutions that constrain teams' hiring incentives, such as salary caps, revenue sharing or payroll taxes.

Using data on game outcomes and team payrolls we then estimated the marginal productivity of skills for the four US sports leagues and found them to be significantly different, ranging from 0.16 for MLB, to 1.54 for the NBA. Comparing the estimated productivity of skills with the institutions actually observed in the different leagues, we find that the theoretical predictions from the model fit for three of the four leagues. Overall, the observed differences in adopted institutions is thus in part explained as optimal responses to these differences in the marginal productivity of skills.

Our work highlights some pathways for future research in the field of sports economics. We have here focused on one fundamental parameter of a sport, the productivity of skills, and shown how it relates to the optimal choice of regulatory institutions that constrain teams' hiring decisions. It is natural to try and extend this approach to focus on other fundamental sports' parameters (e.g. sports differ in the way in which suspense and surprise develop over time), as well as possibly other, more complex institutional differences across leagues and over time.

More broadly, we view the results of our case study of US sports leagues as relevant for the economics literature on institutions. A central tenet of this literature, and the related policy debate, is that economic outcomes can be improved by transplanting successful institutions from one context to another. In the specific case of the sports leagues we consider, for example, commentators periodically suggest that one or more of leagues adopt regulatory institutions from the other leagues. Our work suggests that such policy prescriptions may be misguided. When observed differences in institutions reflect optimal responses to differences in underlying fundamentals, then transplanting institutions across settings may in fact be harmful.

References

Acemoglu, D., Robinson, J.A. and Verdier, T. (2017), 'Asymmetric Growth and Institutions in an Interdependent World', *Journal of Political Economy* **125**(5), 1245-1305.

Andreff, W. and Szymanski, S. (2006), Handbook on the Economics of Sport, *Edward Elgar Publishing.*

Aronow, P.M., Samii, C. and Assenova, V.A. (2015), 'Cluster-Robust Variance Estimation for Dyadic Data', *Political Analysis* **23**, 564-577.

Bizzozero, P., Flepp, R. and Franck, E. (2016), 'The importance of suspense and surprise in entertainment demand: Evidence from Wimbledon', *Journal of Economic Behavior & Organization* **130**, 47–63.

Cameron, A.C. and Miller, D.L. (2014), 'Robust Inference for Dyadic Data', working paper, University of California - Davis.

Christenfeld, N. (1996), 'What makes a good sport', *Nature* **383**(24), 662.

Dietl, H., Lang, M. and Werner, S. (2010), 'The Effect of Luxury Taxes on Competitive Balance, Club Profits, and Social Welfare in Sports Leagues', *International Journal of Sport Finance* **5**(1), 41–51.

Downward, P. and Dawson, A. (2000), The Economics of Professional Team Sports, Psychology Press.

Downward, P., Dawson, A. and Dejonghe, T. (2009), 'Sports Economics: Theory, Evidence and Policy', Butterworth-Heinemann.

El-Hodiri, M. and Quirk, J. (1971), 'An Economic Model of a Professional Sports League', *Journal of Political Economics* **79**(6), 1302–1319.

Falconieri, S., Palomino, F. and Sákovics J. (2004), 'Collective versus individual sale of television rights in league sports', *Journal of the European Economic Association* **2**(5), 833-862.

Fort, R. and Quirk, J. (1995), 'Cross-subsidization, incentives, and outcomes in professional team sports leagues' *Journal of Economic Literature* **33**, 1265–1299.

Gonzalez-Eiras, M., Harmon, N. and Rossi M. (2018), 'Fundamentals and Optimal Institutions: The case of US Xports Leagues', Working paper 128, Universidad de San Andrés.

Gordon, S. (2012), NHL should look to the success of revenue sharing in other pro leagues, The Globe and Mail, http://www.theglobeandmail.com/, Toronto.

Heitner, D. (2015), Sports Industry To Reach $73.5 Billion By 2019, Forbes, http://www.forbes.com, Jersey City.

Hootman, J.M., Dick, R. and Agel, J. (2007), 'Epidemiology of Collegiate Injuries for 15 Sports: Summary and Recommendations for Injury Prevention Initiatives', *Journal of Athletic Training* **42**(2), 311–319.

Levin, R., Mitchell, G., Volcker, P. and Will G. (2000), The Report of the Independent Members of the Commissioner's Blue Ribbon Panel on Baseball Economics. New York: Major League Baseball.

Massey, C. B. and Thaler R. H. (2013), 'The Loser's Curse: Decision Making and Market Efficiency in the National Football League Draft', *Management Science* **59**(7), 1479–1495.

Peeters, T. (2015), 'Profit-maximizing gate revenue sharing in sports leagues', *Economic Inquiry* **53**(2), 1275–1291.

Rogers, J. (2015), Communist sports are better, Forbes, http://www.forbes.com, Jersey City.

Rosen, S. and Sanderson, A. (2001), 'Labour markets in professional sports', *Economic Journal* **111**(469), 47–68.

Rossi, M. and Ruzzier, C. (2018), 'Career Choices and the Evolution of the College Gender Gap', *The World Bank Economic Review* **32**(2), 307–333.

Salaga, S., Ostfeld, A., and Winfree, J. (2014), 'Revenue Sharing with Heterogeneous Investments in Sports Leagues: Share Media, Not Stadiums', *Review of Industrial Organization* **45**(1), 1–19.

Scully, G. W. (1974), 'Pay and Performance in Major League Baseball', *American Economic Review* **64**(6), 915–930.

Simmons, R. and Forrest, D. (2004), 'Buying success: Team Performance and wage bills in U.S. and European sports leagues', in Fort, R. and Fizel, J. (eds) '*International Sports Economics Comparisons*', New York: Praeger.

Skaperdas, S. (1996), 'Contest success functions', *Economic Theory* **7**(2), 283–290.

Stiglitz, J. (2008), 'The future of global governance', in Serra, N. and Stiglitz, J. (eds) '*The Washington consensus reconsidered: Towards a new global governance*', Oxford University Press, Oxford.

Szymanski, S. (2003), 'The economic design of sporting contests' *Journal of Economic Literature* **41**, 1137–1187.

Szymanski, S. and Kuypers T. (1999), Winners and Losers: The Business Strategy of Football, Viking Press, London.

Vrooman, J. (1995), 'A general theory of professional sports leagues' *Southern Economic Journal* **61**, 971–990.

6 Appendix

6.1 Proof of proposition 1

In solving the planner's problem, it will be convenient to reparameterize it so that the planner directly chooses total talent, $\bar{S}$, and win probabilities, w_i $(i = 1, .., N-1)$, instead of choosing team skills S_i. The reparameterized problem is equivalent to the original one because, for any $\vec{S}$ there is a one-to-one correspondence to a set of total skills $\bar{S}$ and win probabilities, w_i, and vice versa.

The reparameterized maximization problem is:

$$\max_{\bar{S},\vec{w}} = Q(\bar{S}, \vec{w}; \vec{F}) = (N-1)\sum_{i=1}^{N} F_i R(\bar{S}, w_i) - c(N-1)\bar{S} - X(n).$$

First-order conditions imply (using $w_N = 1 - \sum_{i=1}^{N-1} w_i$)

$$\frac{\partial Q}{\partial w_i} = F_i \frac{\partial R(\bar{S}, w_i)}{\partial w_i} - F_N \frac{\partial R(\bar{S}, w_N)}{\partial w_N} =0, \quad i = 1, .., N-1,$$

$$\frac{\partial Q}{\partial \bar{S}} = \sum_{i=1}^{N} F_i \frac{\partial R}{\partial \bar{S}} - c =0.$$

Note that $\bar{S}^P$ and w_i^P do not depend on α, and that assumption 3 implies $\frac{w_i^P}{w_j^P} = \frac{S_i^\alpha}{S_j^\alpha}$, thus an increase in α should be met with a more equal distribution of skills.

To guarantee uniqueness of the planner's allocation, the objective must be strictly concave which requires a negative definite Hessian matrix. A sufficient condition for this is $\frac{\partial^2 R}{\partial \bar{S} \partial w_i} = 0$, which is satisfied under assumption 2. Under this assumption we furthermore have:

$$w_i^P = \frac{F_i}{\sum_{i=1}^{N} F_i},$$

$$\bar{S}^P = \bar{S}^*.$$

It is worth noting that the result that $\bar{S}^P = \bar{S}^*$ is trivial in this setting as teams' revenues do not depend directly on aggregate skills. Thus, since skills are costly, a social planner would set S_i^P as low as possible. But the property that $\bar{S}^P$ does not depend on α is general.

6.2 Proof of proposition 2

We start by showing that there exists $\underline{\alpha}$ such that when $\alpha < \underline{\alpha}$, $S_1^* + S_2^* < \bar{S}$ and thus no revenue is generated. For this we start by noting that the ratio of skills, $\frac{S_1}{S_2}$ (and thus w_1), is independent of α. Combining first order conditions for teams 1 and 2, we get

$$\frac{F_1}{w_1} = \frac{F_2}{1 - w_1} \frac{S_1^*}{S_2^*}.$$

Next, we note that the hiring incentive is stronger the higher α. The first order condition of team i is

$$\frac{F_i}{w_i} \frac{\left(\frac{S_j^*}{S_i^*}\right)^\alpha}{1 + \left(\frac{S_j^*}{S_i^*}\right)^\alpha} \frac{\alpha}{S_i^*} = c.$$

Thus, given that $\frac{S_1^*}{S_2^*}$ and w_i are independent of α, it must be the case that $\frac{dS_i^*}{d\alpha} > 0$. This completes the proof that only sports for which $\alpha \geq \underline{\alpha}$ generate positive revenues.

Plotting the graphs of $S_1^*(S_1)$, $S_2^*(S_1)$ against each other in (S_1, S_2)-space, Nash-equilibria of the model occur at (and only at) intersections of the two graphs. Based on this, we prove equilibrium existence in two steps by showing that:

i. $\lim_{S_j \to 0} S_i^*(S_j) > S_j$.

ii. $\lim_{S_j \to \infty} S_i^*(S_j) < S_j$.

We now go through the two steps. Step 1: From the first order condition for team i, (1), when $S_j \to 0$, choice of S_i is given by

$$\alpha \frac{1 - w_i}{S_i} - \frac{c}{F_i} \leq 0.$$

Note that if $S_i = 0$ then $w_i = w_i = 0$, and the first term diverges in this case. Thus $S_i > 0$ which implies $w_i = w_i = 1$ and $S_i > S_j$.

Step 2: From the first order condition for team i, we note that $w_i = 0$, since if $w_i > 0$, the first order condition would imply $S_j < \infty$, which would be a contradiction. For $w_i = 0$, it has to be the case that $S_i < S_j$.

Since the best responses $S_i^*(S_j)$ are continuous, the above implies that $S_1^*(S_2)$ and $S_2^*(S_1)$ must cross at least once and thus that there is at least one Nash equilibrium in the model.

To prove uniqueness, we start by observing that if two equilibria have the same win probabilities, $w_i = w_i'$, then it must be the case that $S_i = S_i'$ and the equilibria are the same. Assume then that for the two equilibria it is the case that $w_2' > w_2$. Then from the first order conditions for team 1 it must be the case that $S_1' > S_1$ and from the first order conditions for team 2 it must be that $S_2' < S_2$. But if this were the case then $w_1' > w_1$, contradicting our assumption. We thus conclude that the equilibrium is unique.

6.3 Proof of proposition 3

From the first order conditions for hiring talent,

$$\frac{\partial R}{\partial S_i} + \frac{\partial R}{\partial w_i}\frac{dw_i}{dS_i} - \frac{c}{F_i} = 0.$$

An increase in α will affect the term $\frac{dw_i}{dS_i}$, with

$$\frac{dw_i}{dS_i} = \alpha \frac{w_i(1-w_i)}{S_i}.$$

Its derivative with respect to α is given by

$$\begin{aligned}
\frac{d\frac{dw_i}{dS_i}}{d\alpha} &= \frac{w_i(1-w_i)}{S_i} \\
&\quad + \alpha \frac{1}{S_i} w_i(1-w_i)\left[(1-w_i)\ln S_i + w_i \ln S_j\right], \\
&= \frac{dw_i}{dS_i}\left(\frac{1}{\alpha} + \left[(1-w_i)\ln S_i + w_i \ln S_j\right]\right).
\end{aligned}$$

When $\alpha > 0$, this derivative might be negative only when one of the teams' skill choice is close to zero. A sufficient condition for the derivative to be positive would be that F_1 and F_2 are large enough such that $S_2 = 1$ (since then $S_1 > 1$). A lower threshold for α, $\underline{\alpha}(F_1, F_2)$, such that this sufficient condition is satisfied, is characterized by the following system of equations from teams' first order conditions

$$\begin{aligned}
\frac{\partial R}{\partial \bar{S}}|_{\bar{S}=S_1+1} + \underline{\alpha}(F_1,F_2)\frac{w_1(1-w_1)}{S_1}\frac{\partial R}{\partial w}|_{w=w_1} &= \frac{c}{F_1}, \\
\frac{\partial R}{\partial \bar{S}}|_{\bar{S}=S_1+1} + \underline{\alpha}(F_1,F_2) w_1(1-w_1)\frac{\partial R}{\partial w}|_{w=1-w_1} &= \frac{c}{F_2}.
\end{aligned}$$

Since $\lim_{\bar{S}\to\infty} \frac{\partial R}{\partial \bar{S}} = 0$, it must be that $\lim_{F_2\to\infty} \underline{\alpha}(F_1, F_2) = 0$. Thus, increases in α when $\alpha > \underline{\alpha}(F_1, F_2)$ always increases the Nash equilibrium hiring levels, S_1^* and S_2^*.

Proof of Proposition 4

In parallel as we did in the proof of proposition 2, we prove equilibrium existence by studying properties of best response functions $S_1^*(S_1), S_2^*(S_1)$ and showing:

i. That $\lim_{S_j \to 0} S_i^*(S_j) > S_j$.

ii. That $\lim_{S_j \to \infty} S_i^*(S_j) < S_j$.

We now go through the two steps. Step 1: From the first order condition for team i, (2), when $S_j \to 0$, choice of S_i is given by

$$\frac{\alpha}{S_i}\Big((1-\frac{\tau}{2})F_i(1-w_i) - \frac{\tau}{2}F_j w_i\Big) - c(1 - s\mathbb{1}_{i=2}) \leq 0.$$

Note that if $S_i = 0$ then $w_i = 0$ and the first term diverges. Thus $S_i > 0$ which implies $w_i = 1$ and $S_i > S_j$.

Step 2: From the first order condition for team i, note that if $S_i \to \infty$ this would be negative. This implies $S_i < \infty$, $w_i = 0$, and thus $S_i < S_j$.

Since the best responses $S_i^*(S_j)$ are continuous, the above implies that $S_1^*(S_2)$ and $S_2^*(S_1)$ must cross at least once and thus that there is at least one Nash equilibrium in the model.

The proof of uniqueness parallels that of proposition 2.

6.4 Proof of proposition 5

First, we are going to show that when $S_1^* + S_2^* > \bar{S}^P$, a pair $0 < \tau < 1$, and s exists such that both teams first order conditions are satisfied when evaluated at the optimal allocation. Using the planner's profit maximizing first order conditions, we can rewrite the marginal effect of hiring for team 2, evaluated at the profit maximizing allocation (and thus satisfying $(F_1 + F_2)\frac{\partial R}{\partial \bar{S}} = c$), as

$$\begin{aligned}
& (1-\frac{\tau}{2})F_2\frac{dR}{dS_2} + \frac{\tau}{2}F_1\frac{dR}{dS_1} - (F_1+F_2)\frac{\partial R}{\partial \bar{S}} \\
= \;& -\frac{\tau}{2}F_2\frac{\partial R}{\partial S_2} - (1-\frac{\tau}{2})F_1\frac{\partial R}{\partial S_2} + (1-\frac{\tau}{2})F_2\frac{\partial R}{\partial w_2}\frac{\partial w_2}{\partial S_2} + \frac{\tau}{2}F_1\frac{\partial R}{\partial w_1}\frac{\partial w_1}{\partial S_2} \\
= \;& -\frac{\tau}{2}F_2\frac{\partial R}{\partial S_2} - (1-\frac{\tau}{2})F_1\frac{\partial R}{\partial S_2} + (1-\tau)F_2\frac{\partial R}{\partial w_2}\frac{\partial w_2}{\partial S_2}. \qquad (12)
\end{aligned}$$

Where in the last equality we have used $F_1\frac{\partial R}{\partial w_1} = F_2\frac{\partial R}{\partial w_2}$ when evaluated at the profit maximizing allocation, and $\frac{\partial w_1}{\partial S_2} = -\frac{\partial w_2}{\partial S_2}$.

When $\tau = 1$ the above equation is negative. And when $\tau = 0$ it is simply the first order condition for team 2 in the absence of regulation evaluated at $S_i = S_i^P$. Since $S_1^* + S_2^* > \bar{S}^P$, the first order condition of at least one team would be positive when evaluated at $S_i = S_i^P$. We proceed under the assumptions that the marginal benefit of hiring for team 1 when evaluated at $S_i = S_i^P$ is weakly larger than that of team 2 when $S_1^* + S_2^* > \bar{S}^P$, and that both are positive.[49] Therefore, by continuity, when $S_1^* + S_2^* > \bar{S}^P$ there exists a $0 < \tau < 1$ that would make team 2 to choose $S_2 = S_2^P$ when team 1 is choosing S_1^P.

When the marginal hiring benefit for team 1 when evaluated at $S_i = S_i^P$ is weakly larger than that of team 2, the first order condition for team 1 evaluated at S_1^P would be positive when the revenue sharing tax is set at the level that leads team 2 to choose S_2^P, and assuming there are no luxury taxes, i.e. $s = 0$. In this case a salary cap set at S_1^P would only be binding for team 1 and would lead to the profit maximizing allocation.

Now we consider the case of luxury taxes instead of salary cap, and evaluate team 1's first order condition at the profit-maximizing allocation and the tax rate that makes team 2 to choose S_2^P. This gives

$$\begin{aligned} & (1-\frac{\tau}{2})F_1\frac{dR}{dS_1} + \frac{\tau}{2}F_2\frac{dR}{dS_1} - (F_1+F_2)\frac{\partial R}{\partial \bar{S}} - cs, \\ = \quad & -\frac{\tau}{2}F_1\frac{\partial R}{\partial S_1} - (1-\frac{\tau}{2})F_2\frac{\partial R}{\partial S_1} + (1-\tau)F_1\frac{\partial R}{\partial w_1}\frac{\partial w_1}{\partial S_1} - cs. \end{aligned}$$

Clearly there always exists a hiring tax/subsidy s that would make this equation to equal zero when $S_1 = S_1^P$. Now we want to prove that $s > 0$, i.e. the league taxes team 1. To do this we substract team 2's first order condition from team 1's to get

$$(1-\tau)\left[(F_1-F_2)\frac{\partial R}{\partial \bar{S}} + F_1\frac{\partial R}{\partial w_1}\frac{\partial w_1}{\partial S_1} - F_2\frac{\partial R}{\partial w_2}\frac{\partial w_2}{\partial S_2}\right] - cs. \tag{13}$$

If $\tau = s = 0$, (13) is positive, under our assumption that the first order condition of team 1 is weakly higher than that of team 2 when evaluated at the planner's allocation. This implies that the term in brackets is positive, and thus that s must be positive if (13) is to be zero. Therefore, the strong team is taxed. We conclude then that the profit maximizing allocation can be decentralized through a combination of revenue sharing and a hiring tax for the strong team provided that the autarky allocation satisfies $S_1^* + S_2^* > \bar{S}^P$.

[49] Otherwise the proof would go along using a parallel reasoning. If the marginal benefit of the weak team is negative then we would need to have a hiring subsidy for this team instead of a hiring tax on the strong team. And if the marginal benefit of the strong team is lower than that of the weak team then the strong team should be given a hiring subsidy instead of a tax.

To prove that an increase in α leads to an increase in optimal level of revenue sharing we consider the effect that this has on (12). The only term that is directly affected is $\frac{\partial w_1}{\partial S_1}$, which we found increases with α when $\alpha > \underline{\alpha}(F_1, F_2)$. Thus to keep this first order condition at zero the revenue sharing tax, τ, has to increase as well: $\frac{d\tau}{d\alpha} > 0$.

The result that an increase in α reduces the dispersion of skills across teams follows from the fact that in the optimal allocation, w_1^P is independent of α but w_1 is decreasing in α for a given skills ratio, since $\frac{w_1}{w_2} = \frac{S_1^\alpha}{S_2^\alpha}$. Thus $\frac{S_1}{S_2}$ has to decrease with α.

Finally, the result that an increase in in the dispersion of F_i (which proxies for an expansion of the league's size, N) leads to an increase in revenue sharing comes from considering $\Delta F_1 = -\Delta F_2 > 0$ in (12). This results in $(1-\tau)\left[\frac{\partial R}{\partial S_1} + \frac{\partial R}{\partial w_1}\frac{\partial w_1}{\partial S_1}\right]\Delta F_1 > 0$. Thus this change increases incentives for revenue sharing resulting in a higher tax rate.