

Discussion Paper Series

IZA DP No. 18491

March 2026

Market Size Disparity Moderates Competitive Balance Interventions in US Sports Leagues

Kyle J. Davis

University of Oklahoma

Tyler Ransom

University of Oklahoma, IZA@LISER
and GLO

Christopher Black

University of Oklahoma

Daniel J. Larson

University of Oklahoma

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.



Market Size Disparity Moderates Competitive Balance Interventions in US Sports Leagues

Abstract

This study examines the extent to which market size disparity across franchises — measured as the coefficient of variation of MSA populations — moderates the effectiveness of competitive balance interventions (CBIs) in Major League Baseball (MLB), the National Football League (NFL), the National Basketball Association (NBA), and the National Hockey League (NHL) from 1967 to 2023. Using two-way fixed effects models with multiple balance measures, we find that CBI effectiveness depends on the distribution of market sizes across league members. Jointly adopting a salary cap and floor improves competitive balance at low levels of market size disparity but is ineffective at high levels. Revenue sharing shows limited effects. Luxury taxes are associated with worsened competitive balance in high-disparity leagues. Our findings demonstrate that market size disparity not only affects competitive balance directly but also determines which interventions succeed. These results have direct relevance to recent discourse on competitive balance in MLB.

JEL classification

L83, L51, C23, D63, L11

Keywords

competitive balance, salary floor, salary cap, market size disparity, Major League Baseball

Corresponding author

Kyle J. Davis

kylejdavis@ou.edu

1 Introduction

Professional sports leagues in the U.S. have implemented competitive balance interventions (CBIs) for over 50 years to promote parity and maintain competitive uncertainty (Rottenberg, 1956). Sports leagues face a unique economic challenge: teams must cooperate to produce competitive games (Neale, 1964), yet profit-maximizing behavior in leagues with differing market sizes naturally produces competitive imbalance (El-Hodiri and Quirk, 1971). To address this tension, leagues have adopted various CBIs, including revenue sharing that redistributes income from high-revenue to low-revenue teams, salary caps that constrain maximum spending, salary floors that mandate minimum investment, and luxury taxes that penalize excessive payrolls.

The four major North American sports leagues have adopted markedly different intervention portfolios: the NFL and NHL utilize hard salary caps with floors and revenue sharing; the NBA uses a soft cap system with a floor in addition to a luxury tax and revenue sharing; while MLB uses the simplest structure with a luxury tax and revenue sharing (Rockerbie, 2024). Yet these structural differences alone do not seem to explain observed variation in competitive balance outcomes. Theoretical work predicts that market size disparities limit intervention effectiveness (Fort and Quirk, 1995), but empirical research has largely examined CBIs in isolation or within single leagues without examining how market context might moderate intervention effectiveness (Dietl, Lang, and Werner, 2010; Késenne, 2006; Lewis, 2008; Totty and Owens, 2011).

In this study, we fill this gap by empirically examining whether and to what extent market size disparity across teams moderates the effectiveness of CBIs. El-Hodiri and Quirk (1971) demonstrate mathematically that market size differences create corresponding differences in the marginal revenue product (MRP) of talent—a star player generates more revenue in New York than Kansas City—which drives large-market teams to rationally acquire talent from small markets. Empirical work confirms that revenue disparities depend on population and income (Borland and Macdonald, 2003). Fort and Quirk (1995) extend this by showing that intervention effectiveness depends on league-specific revenue structures and market characteristics. Our contribution is to test whether this theoretical mechanism operates

in practice by conducting a cross-league analysis that directly estimates the extent to which market size disparity moderates CBI effectiveness.

Using data from Major League Baseball, the National Football League, the National Basketball Association, and the National Hockey League from 1967–2023, we estimate two-way fixed effects models with interaction terms between four types of CBIs and market size disparity. We utilize three measures of competitive balance commonly used in the literature: the Noll-Scully Ratio (NSR; [Noll 1988](#); [Scully 1989](#)), the Herfindahl-Hirschman Index (HHI; [Owen, Ryan, and Weatherspoon 2007](#)), and the Gini coefficient ([Schmidt and Berri 2001](#)).¹ Each measure characterizes the distribution of winning across teams in a league within a single season, with lower values indicating greater parity. These measures focus on regular-season win dispersion rather than other dimensions of competitive balance (CB) such as championship concentration or between-season turnover. Our multi-measure approach allows us to assess whether intervention effects are consistent across different ways of measuring within-season CB.

The period from 1967 to 2023 presents an ideal window for evaluating intervention effectiveness. This timeframe captures the beginning of the Super Bowl era and extends through recent seasons, with substantial variation in CBI implementation providing identifying variation to estimate our parameters of interest. Importantly, this period includes substantial pre-intervention baseline data for most leagues. This allows us to isolate within-league variation attributable to policy changes rather than secular trends in competitive dynamics. Our period includes approximately 27 years before the NFL’s salary cap implementation (1994), 17 years before the NBA’s salary cap (1984–85), and 31 years before MLB’s pooled revenue sharing system (1998). The NHL’s modern CBIs began with the 2005–06 season following the league’s restructuring after the 2004–05 lockout.

Our analysis shows that CBI effectiveness depends critically on a league’s market structure. This is especially true for salary caps and floors. These salary restrictions improve

¹We also examine the Competitive Balance Ratio (CBR; [Humphreys 2002](#)) as a robustness check. The CBR captures between-season team mobility rather than within-season dispersion. Descriptive patterns across leagues are broadly similar (see Appendix), though the reader should understand that higher values of the CBR imply greater competitive balance, which is opposite the three measures we examine in the body of the paper. We do not apply our full analysis to the CBR due to its limited within-league temporal variation.

competitive balance in leagues where there is less disparity in market franchise sizes. On the other hand, we find limited impact for revenue sharing. For luxury taxes, we find that they worsen competitive balance in high-disparity leagues, but have no significant effect in MLB, suggesting this, too, is moderated by market structure. Overall, these patterns are consistent with [El-Hodiri and Quirk's \(1971\)](#) theoretical prediction that market size disparities fundamentally shape competitive outcomes, though we cannot fully separate market effects from sport-specific factors. Our results demonstrate that market size disparity not only affects competitive balance directly but also moderates which interventions succeed. Randomization inference tests, where we randomly implement the timing CBIs, corroborate these patterns. The interaction between market size disparity and salary floor is statistically distinct from random CBI re-timing across all three competitive balance measures ($p < 0.04$), while the salary cap and luxury tax interactions are not. This further suggests that salary floors are the active ingredient driving our core finding.

Our findings have important policy implications that challenge the idea that any single intervention portfolio can optimally serve all leagues. Ongoing MLB collective bargaining negotiations ([Gonzalez, Passan, and Rogers, 2025](#)) provide a salient example. Our cross-league evidence suggests that simultaneously imposing a salary cap and salary floor would improve competitive balance in MLB's relatively uniform market environment. Though we cannot fully separate market structure effects from MLB's unique competitive dynamics, the evidence points towards salary restrictions as more promising for low-disparity leagues like MLB. Policymakers should calibrate CBIs to their league's market structure. The results suggest that strategically choosing the location of franchises to manage market size disparity could prove as important as financial interventions in maintaining long-term competitive balance.

The remainder of this paper proceeds as follows. Section 2 details our data sources and measurement of our key variables (competitive balance, CBIs, and market size disparity). Section 3 presents a descriptive analysis of CBI implementation and market size disparity. In Section 4, we introduce our econometric framework and present its results. Section 5 discusses implications for theory and policy. Section 6 concludes.

2 Data and Measurement

2.1 Data Sources

This study employs a quantitative approach examining CBI effectiveness in the four major North American sports leagues from 1967 to 2023. We collect team performance data, including win-loss records, from Sports Reference websites. Market size data comes from U.S. Census Bureau Metropolitan Statistical Area (MSA) population estimates and Statistics Canada Census Metropolitan Area (CMA) estimates for Canadian teams, with intercensal estimates from MacroTrends. We manually compile league financial and intervention data from collective bargaining agreements (CBAs) and verify that the data matches related studies in the sports economics literature (Fort and Quirk, 1995; Rockerbie, 2024).

2.2 Measuring Competitive Balance

Competitive balance is the dependent variable in our analysis. The simplest measure of competitive balance is the standard deviation of win percentages within a season. However, this measure fails to account for differences in the number of games played, rendering it ineffective for cross-league comparisons. This shortcoming led to the development of the Noll-Scully Ratio, which compares the actual standard deviation of win percentages to the idealized standard deviation that would exist if all teams were equal (Noll, 1988; Scully, 1989). The idealized standard deviation equals $0.5/\sqrt{N_{jt}}$, where N_{jt} is the number of games played per team in league j and season t . NSR takes values in $[1, \infty)$, with 1 being perfect balance. The NSR is calculated as

$$\text{NSR}_{jt} = \frac{\sigma(w_{i,jt})}{\frac{0.5}{\sqrt{N_{jt}}}}, \quad (1)$$

where $w_{i,jt}$ is team i 's winning percentage and $\sigma(\cdot)$ is the standard deviation function.

Economic concentration measures have also been adapted to sports in the form of the Herfindahl-Hirschman Index (HHI) and the Gini coefficient. The HHI treats team wins as a market share of the total wins in each season (Owen, Ryan, and Weatherspoon, 2007).

HHI is calculated as the sum of squared market shares, $\text{HHI} = \sum_i s_i^2$, where lower values correspond to greater parity. The analysis of MLB by [Owen, Ryan, and Weatherspoon \(2007\)](#) shows that when measures are not adjusted for league size, this can falsely suggest an improvement in competitive balance during periods of expansion. In order to adjust for league size, we use a normalized HHI, calculated as

$$\text{HHI}_{jt}^* = \frac{\sum_{i=1}^{n_{jt}} s_{i,jt}^2 - \frac{1}{n_{jt}}}{\frac{2(2n_{jt} - 1)}{3n_{jt}(n_{jt} - 1)} - \frac{1}{n_{jt}}}, \quad (2)$$

where $s_{i,jt} = W_{i,jt} / \sum_{i=1}^{n_{jt}} W_{i,jt}$ is team i 's share of total wins in league j and season t , and n_{jt} is the number of teams in the league for that season. Throughout the rest of this paper, we use HHI to denote the normalized HHI (HHI*) listed above.

The Gini coefficient offers an alternative perspective on competitive balance by measuring the dispersion of wins among teams ([Lambert, 1993](#); [Schmidt and Berri, 2001](#)). The Gini coefficient has a defined range between 0 and 1, where 0 indicates perfect equality (all teams win 50% of games) and 1 indicates perfect inequality. The specific measure is calculated as

$$G_{jt} = \left(1 + \frac{1}{n_{jt}}\right) - \frac{2}{n_{jt}^2 \mu_{x,jt}} (x_{(1),jt} + 2x_{(2),jt} + 3x_{(3),jt} + \dots + n_{jt} x_{(n_{jt}),jt}), \quad (3)$$

where n_{jt} is as defined previously and $\mu_{x,jt}$ is the mean winning percentage in league j and season t , and $x_{(k),jt}$ represents the winning percentage for the k th-ranked team (ranked from highest to lowest) in that league and season. Lower values indicate greater competitive balance.

Our three competitive balance measures (NSR, HHI, and Gini) each measure balance in the regular season standings. Each measure will always yield a slightly different perspective, leading to varying results. Common patterns among all three measures pointing in the same direction are most compelling. While this is an important aspect of competitive balance, it is certainly not the only one. One might be concerned with whether end-of-season championships concentrate among a small number of franchises across seasons. Alternatively, one might measure competitive intensity through the distribution of margins of victory in games,

or even the distribution of within-game win probability ranges. We recognize that there are other ways of measuring competitive balance and leave such alternative analyses to future work.

2.3 Measuring Competitive Balance Interventions

CBIs are the treatment variables in our analysis. As mentioned in the Introduction, we analyze the following CBIs: revenue sharing, salary caps, salary floors, and luxury taxes. It is important to point out that there are other forms of CBIs that are non-monetary; for example, entry draft rules, roster and contract-structure rules, and scheduling or playoff-format adjustments that redistribute competitive advantages without directly affecting payroll. We focus on monetary interventions as they directly address the revenue disparities across market sizes that [El-Hodiri and Quirk \(1971\)](#) identified as the primary driver of competitive imbalance.

Our four intervention variables capture the primary financial mechanisms through which leagues have attempted to reduce the competitive disadvantages that arise from market size disparities. Revenue Sharing (RS) is a continuous variable measured as the proportion of local revenues each team contributes to a pooled redistribution system, taking on values in $[0, 1]$ in league-seasons where revenue sharing is implemented, and 0 otherwise. Salary Cap (SC) is a binary indicator equal to 1 in league-seasons where a maximum team payroll is in force. Salary Floor (SF) is a continuous variable measured as the floor-to-cap ratio, taking on values in $[0, 1]$ in league-seasons where a cap exists, and 0 otherwise. Luxury Tax (LT) is a binary indicator equal to 1 in league-seasons where teams face financial penalties for exceeding a specified payroll threshold.

Revenue-sharing systems reduce financial disparities by redistributing income from high-revenue teams to low-revenue teams. North American professional sports leagues have historically utilized two approaches: gate revenue sharing and pooled revenue sharing ([Rockerbie, 2024](#)). Gate revenue sharing, where visiting clubs received a fixed share of home gate receipts, created “revenue reversals” where teams’ revenue rankings could shift based on scheduling rather than market fundamentals. This led leagues to adopt pooled systems, where each club contributes a fixed percentage of local revenue to a central pool, then receives an equal

or formula-based share. As of 2023, contribution rates were 48% in MLB, 34% in the NFL, and 50% in the NBA (Rockerbie, 2024). Fort and Quirk (1995) demonstrate that while gate revenue sharing reduces talent costs without affecting competitive balance, local television revenue sharing is vital for achieving revenue-maximizing talent distributions, as TV revenues are substantially more responsive to market size than are gate revenues. Rascher et al. (2011) argue that revenue sharing created adverse CB incentives in the NFL.

Salary caps establish maximum team payrolls. Fort and Quirk (1995) argue that salary caps represent the most promising mechanism for achieving competitive balance by limiting the ability of large-market teams to capitalize on their revenue advantages. Salary caps improve competitive balance by constraining large-market clubs' talent accumulation while leaving small-market clubs unaffected, as caps remain non-binding for teams with lower revenues (Késenne, 2000). The NFL and NHL employ "hard" caps, accompanied by severe penalties for exceeding limits, whereas the NBA utilizes a "soft" cap, allowing teams to exceed thresholds through various exceptions. MLB has no salary cap. The NBA implemented its cap in 1984–85, the NFL in 1994, and the NHL in 2005–06 following the lockout. Despite theoretical predictions, empirical evidence shows mixed results, with effectiveness depending on enforcement mechanisms and whether accompanying salary floors prevent teams from spending substantially below the cap (Rascher, Maxcy, and Schwarz, 2021).

Salary floors address concerns that teams receiving substantial revenue-sharing payments might minimize payroll costs to maximize profits rather than investing in competitive rosters. Thus, we expect salary floors to improve competitive balance by raising the competitive floor, preventing weak teams from fielding uncompetitive rosters and compressing the distribution of wins across the league. Fort and Quirk (1995) suggest that floors are particularly effective when combined with revenue sharing, as they counteract moral hazard by requiring a minimum investment in player talent. In practice, the NBA, NFL, and NHL have all adopted systems that combine salary caps with floors, where the floor typically ranges from 74% to 90% of cap levels in recent years. Empirical evidence on the independent effectiveness of salary floors remains limited, as floors have generally been implemented simultaneously with caps (in NFL and NHL). The NBA provides useful variation due to its 15-year gap between cap and floor implementation. Additionally, we are able to leverage variation in

floor percentages over time in NFL and NHL to isolate the independent effects of caps and floors.

Luxury tax systems impose financial penalties on teams that exceed specified payroll thresholds, while still allowing such spending, using economic incentives rather than absolute prohibitions. Theoretically, luxury taxes with endogenous thresholds based on league average payrolls improve both competitive balance and social welfare (Dietl, Lang, and Werner, 2010). The mechanism operates through small-market clubs responding to subsidies from luxury tax revenues with investment increases exceeding any decreases by large-market clubs. MLB introduced its competitive balance tax in 1997, suspended it from 2000 to 2002, and then reintroduced it in 2003. The NBA’s luxury tax operates in conjunction with its soft salary cap. Empirical evidence shows mixed results, with effectiveness dependent on tax rate escalation for repeat offenders and whether thresholds keep pace with revenue growth (Rascher, Maxcy, and Schwarz, 2021).

2.4 Measuring Market Size Disparity

Market size disparity is the moderating variable in our analysis. Denoting it by MKT, we measure it as the coefficient of variation of MSA populations in cities that have franchises in the league, as follows:

$$\text{MKT}_{jt} = \frac{\sigma(\text{Pop}_{i \in j, t})}{\mu(\text{Pop}_{i \in j, t})}, \quad (4)$$

where i indexes member franchises in league j and season t , $\sigma(\cdot)$ is the standard deviation function, and $\mu(\cdot)$ is the sample average function. Higher values indicate greater market size disparity. Measuring market size disparity with CV allows us to use a unitless measure. This provides a simpler interpretation than other alternatives like the HHI or Gini coefficient.

3 Descriptive Analysis

We begin by describing patterns of competitive balance and CBI implementation across leagues to establish baseline trends and demonstrate the variation in market structure and intervention portfolios that drives our empirical analysis. This descriptive foundation moti-

vates the econometric approach presented in Section 4.

Table 1 summarizes the implementation of interventions across leagues. We code only pooled revenue-sharing systems, excluding earlier gate-sharing systems due to different incentive structures (Rockerbie, 2024). Revenue sharing and salary floors are coded as percentages, with 0 in years when they are not present. Both luxury tax and salary caps are coded as dummy variables. It is of note that, aside from temporary breaks in the luxury tax in MLB and salary cap and floor in the NFL, no CBIs in this study have been revoked once implemented.

Table 2 presents descriptive statistics for the three competitive balance measures across the four major North American professional sports leagues from 1967 to 2023. The data shows variation in competitive balance both across and within leagues over time.

The NBA has the highest average NSR (2.72), indicating the greatest competitive imbalance among the four leagues, with values ranging from 1.78 to 3.59. The NFL demonstrates the strongest competitive balance with the lowest average NSR (1.55) ranging from 1.22 to 1.91. MLB (1.81) and the NHL (1.98) fall between these extremes, though the NHL shows considerably more variation over time than MLB. The HHI diverges at the top where the NFL shows the highest average HHI (0.435), followed by the NBA (0.249), NHL (0.172), and MLB (0.057), with lower values indicating greater competitive balance. The Gini coefficient patterns reflect greater competitive balance in MLB (0.083) and NHL (0.143), with the NBA (0.175) and NFL (0.229) showing less parity. Market size disparity, measured as the coefficient of variation of MSA populations, differs substantially across leagues. The NFL exhibits the greatest market heterogeneity (68.4), followed by the NBA (66.9), NHL (66.1), and MLB (48.5).

Figure 1 depicts average competitive balance measures and market size disparity by league and year. The decadal trends reveal suggestive relationships between CBI adoption and competitive balance that motivate our econometric analysis. MLB’s competitive balance slightly worsened after the 1980s across multiple measures (NSR 1.66 to 2.02; Gini 0.077 to 0.102) despite implementing revenue sharing in 1996 and a more advanced luxury tax in 2003. The decline coincided with increasing market size disparity (46.4 in 1990s to 49.4 in 2020s). The NBA demonstrates considerable volatility, with NSR peaking at 2.99 in the

1990s shortly after the adoption of the salary cap, and HHI fluctuating between 0.197–0.308 across decades. The NFL maintained consistently strong balance throughout the period (NSR between 1.46–1.83), with HHI declining sharply from 0.654 in the 1960s to 0.376–0.421 in recent decades. This possibly reflects the effectiveness of their comprehensive CBI package. The NHL shows some improvement between 2000 and 2010 following the 2005 restructuring. NSR declined from 2.58 in the 1970s to 1.62 in the 2010s with HHI following suit, dropping from 0.303 to 0.094, though recent deterioration (NSR = 2.12, HHI = 0.165 in the 2020s) suggests that these gains may not be permanent. These patterns hint at differential CBI effectiveness across market contexts.

4 Identification, Estimation, and Results

The descriptive patterns in the previous section show substantial variation in competitive balance in all leagues, with some leagues improving over time and others staying flat or even worsening. At the same time, leagues have implemented greater numbers of CBIs over time, and market size disparity has declined in all four leagues over the span of the study. It is difficult to infer anything in these raw trends from either the presence or absence of correlations between CBIs and overall competitive balance or market size disparity. For example, market size disparity changes as leagues expand or relocate franchises, and macroeconomic conditions and labor market regulations affect all leagues at once. To isolate the relationship between specific CBIs and competitive balance—and to test whether market structure moderates these relationships—we turn to a two-way fixed effects regression framework that can account for these overlapping factors.

4.1 Identification and Measurement Challenges

Our central research question—whether market size disparity moderates CBI effectiveness—poses fundamental identification challenges. Ideally, we would run an experiment that would randomly assign CBI portfolios to leagues with varying market structures and then observe the resulting degree of competitive balance. Instead, we only observe four leagues that endogenously adopt different CBIs at different times, each with distinct market compositions,

competitive traditions, and sport-specific factors. With only four leagues and relatively few, selectively implemented CBI events, we cannot definitively establish causality. Leagues may implement or strengthen interventions in response to competitive imbalance. This potentially creates reverse causation, although lags between outcomes and CBA negotiations mitigate this concern. Our findings should therefore be taken as associations that illuminate patterns consistent with, but not definitive proof of, theoretical predictions about how market size moderates intervention effectiveness.

Three features of our setting limit clean identification. First, not all leagues implement all interventions. MLB has never adopted a salary cap or floor; the NFL and NHL have never implemented a luxury tax. This means CBI-specific coefficients are identified only from leagues exhibiting within-league variation in that intervention. The salary cap effect, for example, comes entirely from NBA, NFL, and NHL pre-post comparisons. MLB contributes no identifying variation. Thus, estimating how a salary cap would affect MLB requires extrapolating outside the observed data.

Second, our moderating variable—market size disparity—varies primarily across leagues rather than within leagues over time. Within-league changes in market composition occur mainly through rare expansion or relocation events, or gradual changes in the nationwide population distribution. Consequently, the interaction effects between CBIs and market size disparity are identified from very few data points: essentially, whether the pre-post effect of a given CBI differs across the small number of leagues that adopted it, each observed at a single (or slowly-evolving) level of market size disparity. More direct identification would leverage time series variation within leagues, which would avoid cross-league extrapolation but at the cost of being unable to estimate the moderating effect of market size disparity. We therefore choose cross-league models so that we can provide empirical evidence on our main hypothesis.

Third, identification of the independent effects of salary caps and salary floors poses challenges. The NFL and NHL simultaneously implemented both mechanisms, and a salary floor has never existed in those leagues independent of a salary cap, limiting our ability to cleanly separate their effects. The NBA provides some variation, with the salary cap implemented 15 years prior to the adoption of the salary floor. For all three leagues, some identifying

variation exists at the intensive margin because the percentage limits have changed over time. However, this variation is modest and potentially correlated with other CBI changes arising from adjustments to CBAs.

Beyond identification, several measurement limitations apply. Our binary and percentage measures cannot capture all implementation nuances (e.g., luxury tax penalty structures), and revenue sharing systems vary in their formulas, thresholds, and enforcement mechanisms across leagues over time. MSA population may not fully reflect revenue potential in the modern media era, though it remains the most consistently available measure across markets and time periods. Some CBIs have been implemented for shorter time periods (NHL cap since 2005–06; NBA revenue sharing since 2011), and non-monetary CBIs that we omit (draft systems, playoff formats, and free agency rules) may confound our estimates if their adoption correlates with the financial CBIs we study.

4.2 Estimation Strategy

Given the above constraints, we employ two-way fixed effects (TWFE) as a way to balance internal and external validity concerns. League fixed effects account for time-invariant league heterogeneity (e.g. sport-specific rules, competitive culture). Year fixed effects absorb common temporal shocks (macroeconomic conditions, population movement, and evolving media landscapes). Identification relies on comparing leagues to themselves before and after CBI adoption, after removing common temporal shocks.

$$\begin{aligned}
 CB_{jt} = & \beta_0 + \beta_1 RS_{jt} + \beta_2 SC_{jt} + \beta_3 SF_{jt} + \beta_4 LT_{jt} + \beta_5 MKT_{jt} + \\
 & \beta_6 (RS_{jt} \times MKT_{jt}) + \beta_7 (SC_{jt} \times MKT_{jt}) + \beta_8 (SF_{jt} \times MKT_{jt}) + \beta_9 (LT_{jt} \times MKT_{jt}) + \\
 & \alpha_j + \gamma_t + \varepsilon_{jt},
 \end{aligned} \tag{5}$$

where CB_{jt} represents competitive balance for league j in year t (as measured by NSR, HHI, or Gini); α_j captures league-specific fixed effects; γ_t captures time fixed effects; and ε_{jt} is the error term. We include interaction terms between interventions and market size disparity so that we can test our primary hypothesis regarding whether market structure moderates CBI effectiveness.

TWFE controls for time-invariant league heterogeneity and common temporal shocks, so identification rests on within-league, within-year variation in CBI adoption. The primary remaining threat is that our results reflect coincidental timing: leagues happening to adopt CBIs during periods when competitive balance was already trending in a particular direction. To address this, we conduct randomization inference tests that re-estimate equation (5) under 1,000 alternative CBI timing assignments drawn at random. These tests assess whether our results could arise from random variation in CBI adoption timing rather than true causal effects. For each league-CBI pair that has historically been implemented, we draw a random alternative adoption year and assign the intervention to be zero before that year and its observed trajectory in every year thereafter, then compare the actual coefficient estimate to the resulting null distribution. The p -value is the share of permutations that yield a coefficient at least as large in magnitude as the one we estimate. We report these results after presenting our preferred estimates and marginal effects in Section 4.3. We interpret our findings as informative associations rather than definitive causal estimates, but the design and robustness tests together provide meaningful empirical grounding for prior theoretical predictions.

4.3 Results

We produce three main sets of results. First are estimated regression coefficients, goodness of fit, and randomization inference tests from equation (5). Second are marginal effects evaluated at league-specific values of MKT. Finally, because salary floors have never existed independently of a salary cap outside of the NBA, we present results on the effectiveness of their joint adoption.

4.3.1 Estimated Coefficients and Goodness of Fit

Table 3 presents parameter estimates from equation (5) for our three competitive balance measures (NSR, HHI, and Gini). We use Arellano (1987) standard errors that correct for both heteroskedasticity and serial correlation. Standard cluster-robust approaches require many clusters for reliable inference; our setting with only four leagues means that serial

correlation is a much greater threat to valid statistical inference.

The pattern of results is broadly consistent across all three CB measures. Salary caps and salary floors each show large and significant main effects along with significant interaction terms of the opposite sign. Moreover, the estimated effects are mirror images of one another: whenever floors are positive, caps are negative, and vice versa. These patterns indicate that these two CBIs affect competitive balance in opposing ways, conditional on market size disparity. Revenue sharing shows a significant main effect and interaction term for NSR, though effects are not significant for HHI or Gini. Luxury taxes show a significant main effect for HHI and significant interaction terms across all three CB measures, indicating moderation through market size disparity. Market size disparity itself has no significant direct effect in HHI and Gini; in these models, it only impacts competitive balance through the interaction terms. For NSR, there is a small but statistically significant direct effect of MKT (-0.008 ; $p = 0.008$). The models explain between 18.5% (HHI) and 25.6% (Gini) of within-league, within-year variation in competitive balance.

Because our model includes interactions between CBIs and MKT, the main effect coefficients on each CBI represent the estimated effect when MKT is zero. This value is far outside the range of MKT values reported in Table 2. As such, assigning any interpretation to the main effect coefficients is unhelpful. Rather, we present marginal effects evaluated at league-specific values of MKT to assess whether CBIs improve competitive balance in the market environments that leagues actually face.

4.3.2 League-Specific Marginal Effects

The estimated marginal effects of market size disparity levels on competitive balance outcomes by league are shown in Table 4, again using [Arellano \(1987\)](#) standard errors that correct for both heteroskedasticity and serial correlation.

Table 4 presents the marginal effects of each CBI within the market context of each league. These marginal effects translate the interaction terms in Table 3 into interpretable effects, showing how each CBI performs in its respective market context. We report, but do not discuss, the marginal effects for salary caps or floors due to their high degree of collinearity resulting from their joint implementation in NFL and NHL. In the next subsection, we

present detailed results on how jointly implementing salary caps and floors would impact competitive balance.

Our attention in Table 4 focuses on luxury taxes and revenue sharing. In Table 3, revenue sharing shows a significant main effect and interaction term for NSR, while the luxury tax shows a significant main effect for HHI and significant interaction terms across all models. However, because of our interaction terms, these interventions' marginal effects need not be zero, nor statistically insignificant, because the delta method evaluates the precision of the combined linear function of both coefficients at specific values of MKT.

Table 4 reveals revenue sharing's marginal effect is significant only for MLB on NSR, where it is associated with worse balance. All other league-specific revenue-sharing marginal effects are insignificant, which suggests revenue sharing has a limited impact in higher market disparity contexts. As for luxury taxes, Panel D shows significant balance-worsening effects for NBA, NFL, and NHL across all three measures. MLB's luxury tax marginal effect is small and insignificant across all three measures, distinguishing it from the high-disparity leagues where the balance-worsening effect is concentrated. The marginal effect of luxury taxes on NSR ranges from 0.311 to 0.341 across the NBA, NFL, and NHL. This is equivalent to 0.8 to 2.0 standard deviations of within-league NSR variation and implies an economically meaningful deterioration in competitive balance.

4.3.3 Effects of Jointly Adopting Salary Caps and Salary Floors

Our results in Table 3 show that salary caps and salary floors show the largest and most systematic independent effects. However, as discussed earlier, their estimated impact on CBI effectiveness is limited by the fact that a salary floor has never existed independently of a salary cap, and the NFL and NHL implemented both simultaneously. Moreover, for MLB the results in Panels B and C of Table 4 are of opposing signs. This makes it difficult to interpret and is further indication of the multicollinearity problem. Further complicating interpretation is that the implementation effects of interest depend on both the level of market size disparity and the level of salary floor.

To provide more complete evidence on the combined efficacy of implementing salary caps and floors, especially in the context of MLB, we calculate the joint effect of adopting both

a salary cap and salary floor and separately vary the points of evaluation in terms of both market size disparity and salary floor level.

We report the results in separate panels of Table 5. Panel A evaluates the joint effect at the values of SF and MKT that are observed in the data for each league. For MLB (which has never imposed a salary floor), we assume it would adopt a salary floor of 75%. Panel A shows that SC and SF would improve MLB’s competitive balance across all three CB measures. The NBA, NFL, and NHL all show significant adverse effects on NSR, and the NFL shows additional significant adverse effects for Gini.

Panels B through E repeat the exercise, but at different values of SF and MKT. Panel B repeats Panel A but sets all leagues to a salary floor of 75%. Panel C repeats Panel B but at a salary floor of 90% for all leagues. Panel D imposes the league-specific salary floors but at the degree of market size disparity facing MLB (i.e., lowest value of MKT). Panel E repeats Panel D, but at the highest value of MKT (i.e., the NFL).

It is important to assess the magnitudes of Table 5’s results, while also keeping in mind that the interpretation should not be taken too literally due to the limitations of identification and measurement discussed above. Panel A’s result for MLB implies that adopting a salary cap and 75% salary floor would reduce NSR by about 1.95 standard deviations. The HHI and Gini effects are even larger, at over 4 standard deviations each. Panel C implies even larger effects for a 90% salary floor.

The results of Table 5 show our main point: that market size disparity is a key moderating variable that governs whether CBIs ultimately improve competitive balance. The key results are in Panels D and E. Panel D shows that, if all leagues faced the values of MKT that MLB faces, then SC and SF would improve competitive balance across the board for all of our CB measures. Likewise, if MKT were instead at the NFL’s level, SC and SF would have no effect or potentially be detrimental, particularly for the case of NSR and Gini in the NBA and NFL.

4.3.4 Randomization Inference Tests for Robustness

As discussed in Section 4.2, we further assess the plausibility and strength of our identification by conducting placebo randomization inference tests that randomly re-assign CBI timing.

Figure 2 shows the null distributions for each CBI×MKT interaction coefficient across 1,000 re-timed adoption assignments, with vertical lines marking our actual estimates. The salary floor interaction (SF×MKT) falls outside the null distribution in all three models (NSR $p = 0.038$; HHI $p = 0.028$; Gini $p = 0.031$). This confirms that the estimated effect for SF is unlikely to arise from coincidental timing. The salary cap interaction (SC×MKT) does not clear the conventional threshold in any model ($p = 0.059$ – 0.168), which is consistent with the interpretation that SC alone is not the active ingredient once SF is separately identified. The luxury tax interaction (LT×MKT) similarly does not pass ($p = 0.198$ – 0.539). Figure A.2 repeats the exercise for the main-effect coefficients; SF again passes in all three models ($p = 0.023$ – 0.043) while SC does not. This reinforces our conclusions.

5 Discussion and Policy Implications

Our findings that market structure shapes the effectiveness of CBIs provide empirical support for the theoretical framework established by [El-Hodiri and Quirk \(1971\)](#). The persistent disparity in competitive balance across leagues reflects both intervention portfolios and underlying market structures. The NFL manages to keep a strong competitive balance despite having the highest market size disparity. MLB achieves a strong balance with softer interventions operating in a favorable, low-disparity structure. The NBA’s persistent imbalance, despite comprehensive interventions, suggests that league-specific factors may limit the effectiveness of financial mechanisms, especially in high-disparity contexts.

These results challenge the idea that any single intervention portfolio can optimally serve all leagues. Our results in [Table 5](#) show that jointly adopting a salary cap and floor improves competitive balance when market size disparity is low, but has no effect, or is even detrimental, when market size disparity is high. [Table 4](#) shows that revenue sharing, at least in the way we have measured it, has limited bearing on competitive balance. Luxury taxes reduce competitive balance in high-disparity leagues, with no effect in MLB, indicating that luxury taxes are moderated by market structure.

For MLB, which has the lowest market size disparity and has never imposed a salary cap or floor, our estimates suggest that joint adoption would improve balance across all three CB

measures. In high-disparity leagues like the NFL, NBA, and NHL, the same interventions provide little benefit and may carry adverse effects. Mechanisms that improve balance in leagues with uniform market sizes may prove insufficient in high-disparity contexts.

More broadly, leagues must recognize that CBIs operate within constraints established by fundamental market structures. Our interaction effects imply that leagues should remain vigilant not only about whether their intervention portfolios are continuing to help them achieve their competitive balance goals, but also about keeping market structure in mind when evaluating decisions on expansion, relocation, and realignment. Strategic franchise location may be as important as financial interventions in maintaining long-term competitive balance.

6 Conclusion

This study presents a cross-league comparison of how market size disparity moderates the effectiveness of CBIs across North American professional sports leagues. Using 57 years of data spanning MLB, NFL, NBA, and NHL, we demonstrate that intervention impacts depend critically on league-specific market structures, consistent with the theoretical framework established by [El-Hodiri and Quirk \(1971\)](#). Jointly adopting a salary cap and floor improves competitive balance when market size disparity is low, but may carry detrimental effects when high. Revenue sharing has limited impact, and luxury taxes are associated with worse competitive balance in high-disparity leagues, with no significant effect in MLB. Though we cannot fully separate market structure effects from sport-specific competitive dynamics, our findings show that CBI effectiveness depends fundamentally on calibration to league-specific conditions rather than universal application.

For sports economics theory, our findings provide evidence for well-documented theoretical mechanisms, while also providing important caveats. Market size remains the dominant structural force shaping competitive balance, as demonstrated mathematically by [El-Hodiri and Quirk \(1971\)](#). However, modern interventions (e.g. salary floors) can partially counteract market forces in ways that depend on market context and the specific competitive balance dimension targeted. MLB's ongoing collective bargaining negotiations ([Gonzalez, Passan, and](#)

Rogers, 2025) illustrate this theoretical tension in practice: as the league confronts proposals to import salary cap structures from other sports, our findings suggest that jointly adopting a salary cap and floor could improve competitive balance in MLB's relatively uniform market environment, with the salary floor serving as the operative mechanism. More broadly, as professional sports leagues continue evolving with expanding media revenues and international markets, successfully managing the tension between market forces and competitive balance requires recognizing market size disparity as an active force moderating CBI effectiveness rather than as a background condition. Leagues that align their intervention portfolios with their market structures position themselves to maintain the competitive uncertainty upon which long-term success depends.

Future research should extend this analysis to other dimensions of competitive balance such as championship concentration or margin-of-victory balance. These alternative measures may respond differently to CBIs than regular season parity measures do. Additionally, our market size measure relies on MSA populations, but this may inadequately capture modern revenue potential. Professional sports teams are increasingly deriving revenues from national and international sources, such as media deals, digital platforms, and merchandising, that extend beyond local markets. More sophisticated measures incorporating media market reach, fan base geography, and digital engagement could improve future analyses. Studies should also examine the specific mechanisms through which interventions operate; e.g., whether salary floors primarily prevent tanking or force poorly managed franchises to invest. Within-league quasi-experimental designs could complement our cross-league comparisons by leveraging discrete policy changes or expansion/relocation events. International comparisons with European soccer leagues operating under different structures could test the generalizability of theories developed in the North American context.

References

- Arellano, Manuel. 1987. "Computing Robust Standard Errors for Within-groups Estimators." *Oxford Bulletin of Economics and Statistics* 49 (4):431–434.
- Borland, Jeffery and Robert Macdonald. 2003. "Demand for Sport." *Oxford Review of Economic Policy* 19 (4):478–502.
- Dietl, Helmut M., Markus Lang, and Stephan Werner. 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.
- El-Hodiri, Mohamed and James Quirk. 1971. "An economic model of a professional sports league." *Journal of Political Economy* 79 (6):1302–1319.
- Fort, Rodney and James Quirk. 1995. "Cross-subsidization, Incentives, and Outcomes in Professional Team Sports Leagues." *Journal of Economic Literature* 33 (3):1265–1299.
- Gonzalez, Alden, Jeff Passan, and Jesse Rogers. 2025. "What you need to know about MLB's looming labor battle." ESPN. URL https://www.espn.com/mlb/story/_/id/47108752/mlb-labor-battle-cba-salary-cap-owners-players-union-mlbpa-faq-manfred.
- Humphreys, Brad R. 2002. "Alternative Measures of Competitive Balance in Sports Leagues." *Journal of Sports Economics* 3 (2):133–148.
- Késenne, Stefan. 2000. "The Impact of Salary Caps in Professional Team Sports." *Scottish Journal of Political Economy* 47 (4):422–430.
- . 2006. "Competitive Balance in Team Sports and the Impact of Revenue Sharing." *Journal of Sport Management* 20 (1):39–51.
- Lambert, Peter J. 1993. *The Distribution and Redistribution of Income: A Mathematical Analysis*. Manchester, UK: Manchester University Press.
- Lewis, Michael. 2008. "Individual Team Incentives and Managing Competitive Balance in Sports Leagues: An Empirical Analysis of Major League Baseball." *Journal of Marketing Research* 45 (5):535–549.
- Neale, Walter C. 1964. "The Peculiar Economics of Professional Sports." *The Quarterly Journal of Economics* 78 (1):1–14.
- Noll, Roger G. 1988. "Professional Basketball." Studies in Industrial Economics Working Paper 144, Stanford University, Stanford, CA.
- Owen, P. Dorian, Michael Ryan, and Clayton R. Weatherspoon. 2007. "Measuring Competitive Balance in Professional Team Sports Using the Herfindahl-Hirschman Index." *Review of Industrial Organization* 31 (4):289–302.
- Rascher, Daniel A., Joel G. Maxcy, and Andy Schwarz. 2021. "The Unique Economic Aspects of Sports." *Journal of Global Sport Management* 6 (1):111–138.

- Rascher, Daniel A., Mark S. Nagel, Matthew T. Brown, and Chad D. McEvoy. 2011. "Free Ride, Take It Easy: An Empirical Analysis of Adverse Incentives Caused by Revenue Sharing." *Journal of Sport Management* 25 (5):373–390.
- Rockerbie, Duane. 2024. "Revenue Sharing in Professional Sports Leagues." *Encyclopedia* 4 (3):1173–1187.
- Rottenberg, Simon. 1956. "The Baseball Players' Labor Market." *Journal of Political Economy* 64 (3):242–258.
- Schmidt, Martin B. and David J. Berri. 2001. "Competitive Balance and Attendance: The Case of Major League Baseball." *Journal of Sport Economics* 2 (2):145–167.
- Scully, Gerald W. 1989. *The Business of Major League Baseball*. Chicago: University of Chicago Press.
- Totty, Evan S. and Mark F. Owens. 2011. "Salary Caps and Competitive Balance in Professional Sports Leagues." *Journal for Economic Educators* 11 (2):45–56.

Table 1: Adoption History of Competitive Balance Interventions by League

League	Revenue Sharing (Pooled)	Salary Cap	Salary Floor	Luxury Tax
NFL	1994 (34%)	1994	1994 ($\approx 58 \rightarrow 90\%$ of cap)	None
NBA	2011 (50%)	1984–85	1999–2000 ($\approx 75 \rightarrow 90\%$ of cap)	2001+
NHL	2005–06 (complex)	2005–06	2005–06 ($\approx 58 \rightarrow 74\%$ of cap)	None
MLB	1996 (20 \rightarrow 48%)	None	None	1997–1999; 2003+

NOTES.—Revenue Sharing percentages indicate the proportion of local revenues contributed to pooled systems. The NFL had revenue sharing arrangements prior to 1994 but these were not formalized in a CBA until 1994, which is the coding start date used in this study. MLB implemented revenue sharing in 1996; however, coding begins in 1998 as the 1996–1997 seasons operated under a hybrid transitional plan that differs from the pooled system measured here. NHL revenue sharing reflects a complex system in which only top-ten revenue teams contribute. Salary Floor percentages show floor as percentage of salary cap, with arrows indicating evolution over time. The NBA employs a soft cap, distinguishing it structurally from the hard caps used by the NFL and NHL; teams may exceed the NBA cap threshold through various exceptions. The NHL operates a hard cap system; in 2012–13, ongoing labor negotiations created a one-year exception permitting teams to exceed the cap up to a hard ceiling, effectively functioning as a soft cap with an absolute maximum for that season. For this study, the hard cap figure is used as the salary floor basis for the 2012–13 NHL season. Years indicate initial implementation. NBA luxury tax was first applicable in 2001–2002 but was not triggered in 2001–02 or 2004–05; every year thereafter was activated. MLB luxury tax was not active in 2000–2002, then permanently applied from 2003 on.

Table 2: Average Competitive Balance and Market Size Disparity by League (1967–2023)

League	<i>N</i>	Statistic	Competitive Balance Measure			Market CV
			NSR	HHI*	Gini	
MLB	57	Mean	1.81	0.057	0.083	48.5
		SD	(0.28)	(0.018)	(0.013)	(2.8)
		Range	[1.40, 2.49]	[0.033, 0.105]	[0.063, 0.114]	[44.3, 55.6]
NBA	57	Mean	2.72	0.249	0.175	66.9
		SD	(0.40)	(0.068)	(0.026)	(10.5)
		Range	[1.78, 3.59]	[0.101, 0.394]	[0.112, 0.233]	[49.4, 91.2]
NFL	57	Mean	1.55	0.435	0.229	68.4
		SD	(0.16)	(0.106)	(0.029)	(7.7)
		Range	[1.22, 1.91]	[0.254, 0.738]	[0.176, 0.307]	[58.0, 83.3]
NHL	56	Mean	1.98	0.172	0.143	66.1
		SD	(0.42)	(0.083)	(0.038)	(11.6)
		Range	[1.19, 2.85]	[0.047, 0.387]	[0.075, 0.226]	[54.0, 85.5]

NOTES.—This table reports statistics about competitive balance and market size disparity measures that we use in our analysis. Statistics are computed within-league across time.

Table 3: Two-Way Fixed Effects Panel Regression Results

Variable	NSR			HHI			Gini		
	$\hat{\beta}$	SE	p	$\hat{\beta}$	SE	p	$\hat{\beta}$	SE	p
Revenue Sharing (RS)	5.880*	2.642	0.027	0.424	0.525	0.420	0.274	0.209	0.192
Salary Cap (SC)	3.433***	0.660	0.000	0.995**	0.341	0.004	0.358***	0.098	0.000
Salary Floor (SF)	-8.926***	0.909	0.000	-1.924***	0.467	0.000	-0.825***	0.153	0.000
Luxury Tax (LT)	-0.548	0.515	0.289	-0.266*	0.112	0.019	-0.063	0.048	0.191
Market Size Disparity (MKT)	-0.008**	0.003	0.008	-0.001	0.001	0.594	-0.001	0.000	0.096
RS \times MKT	-0.091*	0.043	0.034	-0.007	0.008	0.413	-0.004	0.003	0.214
SC \times MKT	-0.050***	0.012	0.000	-0.015**	0.005	0.006	-0.005***	0.002	0.001
SF \times MKT	0.141***	0.013	0.000	0.029***	0.007	0.000	0.013***	0.002	0.000
LT \times MKT	0.013*	0.007	0.048	0.006***	0.002	0.000	0.002**	0.001	0.007
R^2	0.221			0.185			0.256		
Overall F-test	$F(9, 158) = 4.99$ $p < .001$			$F(9, 158) = 3.98$ $p < .001$			$F(9, 158) = 6.04$ $p < .001$		
F-test on MKT interactions (SC, SF, LT)	$F(3, 158) = 1918.36$; $p < .001$			$F(3, 158) = 1610.81$; $p < .001$			$F(3, 158) = 1439.51$; $p < .001$		

NOTES.— $N = 227$ league-years. All models include two-way fixed effects (league and year). Standard errors are corrected for heteroskedasticity and serial correlation according to [Arellano \(1987\)](#). The F -test on MKT interactions tests the joint significance of MKT interactions with SC, SF, and LT. F -tests on all four MKT interactions are not well-defined due to extreme collinearity. * $p < .05$, ** $p < .01$, *** $p < .001$. Figures listed as 0.000 are not truly zero but reflect extremely small quantities.

Table 4: Marginal Effects of Competitive Balance Interventions Evaluated at League-Specific Market Size Disparity Levels

League	NSR			HHI			Gini		
	MFX	SE	<i>p</i>	MFX	SE	<i>p</i>	MFX	SE	<i>p</i>
<i>Panel A. Revenue Sharing</i>									
MLB	1.442*	0.575	0.013	0.101	0.133	0.450	0.071	0.047	0.132
NBA	-0.238	0.240	0.323	-0.022	0.029	0.452	-0.005	0.017	0.751
NFL	-0.373	0.298	0.213	-0.031	0.037	0.401	-0.011	0.021	0.589
NHL	-0.163	0.209	0.435	-0.016	0.025	0.519	-0.002	0.014	0.895
<i>Panel B. Salary Cap</i>									
MLB	1.014***	0.092	0.000	0.274**	0.085	0.002	0.096***	0.024	0.000
NBA	0.099	0.167	0.553	0.001	0.027	0.979	-0.003	0.011	0.758
NFL	0.026	0.183	0.889	-0.021	0.032	0.514	-0.011	0.013	0.366
NHL	0.140	0.158	0.377	0.013	0.025	0.613	0.001	0.011	0.933
<i>Panel C. Salary Floor</i>									
MLB	-2.071***	0.295	0.000	-0.509***	0.137	0.000	-0.204***	0.044	0.000
NBA	0.523***	0.121	0.000	0.027	0.015	0.081	0.031***	0.005	0.000
NFL	0.732***	0.117	0.000	0.070***	0.011	0.000	0.050***	0.005	0.000
NHL	0.408**	0.124	0.001	0.003	0.020	0.871	0.021***	0.006	0.001
<i>Panel D. Luxury Tax</i>									
MLB	0.083	0.202	0.681	0.035	0.043	0.417	0.016	0.020	0.408
NBA	0.322***	0.089	0.000	0.149***	0.029	0.000	0.046***	0.010	0.000
NFL	0.341***	0.081	0.000	0.158***	0.029	0.000	0.049***	0.009	0.000
NHL	0.311**	0.093	0.001	0.144***	0.029	0.000	0.045***	0.010	0.000


NOTES.—MFX = Marginal effect; SE = Standard error (Arellano (1987) corrected). Marginal effects represent the impact of each intervention at the specific market size disparity level of each league, measured by the coefficient of variation of MSA population across markets in the league. **p* < .05, ***p* < .01, ****p* < .001. Figures listed as 0.000 are not truly zero but reflect extremely small quantities.

Table 5: Joint Effects of Salary Cap and Salary Floor Adoption on Competitive Balance Evaluated at Various Market Size Disparity and Salary Floor Levels

League	NSR			HHI			Gini		
	MFX	SE	<i>p</i>	MFX	SE	<i>p</i>	MFX	SE	<i>p</i>
<i>Panel A: Overall (League-specific SF and MKT values)</i>									
MLB	-0.539***	0.133	0.000	-0.108**	0.036	0.003	-0.057***	0.012	0.000
NBA	0.570**	0.176	0.001	0.025	0.036	0.490	0.025	0.014	0.076
NFL	0.684***	0.179	0.000	0.042	0.037	0.264	0.034*	0.014	0.020
NHL	0.446**	0.167	0.008	0.015	0.033	0.649	0.017	0.013	0.204
<i>Panel B: 75% Salary Floor and League-specific MKT</i>									
MLB	-0.539***	0.133	0.000	-0.108**	0.036	0.003	-0.057***	0.012	0.000
NBA	0.492**	0.170	0.004	0.021	0.034	0.544	0.020	0.013	0.136
NFL	0.574**	0.175	0.001	0.031	0.036	0.391	0.026	0.014	0.063
NHL	0.446**	0.167	0.008	0.015	0.033	0.649	0.017	0.013	0.204
<i>Panel C: 90% Salary Floor and League-specific MKT</i>									
MLB	-0.849***	0.176	0.000	-0.184***	0.051	0.000	-0.088***	0.017	0.000
NBA	0.570**	0.176	0.001	0.025	0.036	0.490	0.025	0.014	0.076
NFL	0.684***	0.179	0.000	0.042	0.037	0.264	0.034*	0.014	0.020
NHL	0.507**	0.174	0.004	0.016	0.036	0.659	0.020	0.014	0.149
<i>Panel D: League-specific SF and MLB MKT (lowest disparity)</i>									
MLB	-0.539***	0.133	0.000	-0.108**	0.036	0.003	-0.057***	0.012	0.000
NBA	-0.849***	0.176	0.000	-0.184***	0.051	0.000	-0.088***	0.017	0.000
NFL	-0.849***	0.176	0.000	-0.184***	0.051	0.000	-0.088***	0.017	0.000
NHL	-0.539***	0.133	0.000	-0.108**	0.036	0.003	-0.057***	0.012	0.000
<i>Panel E: League-specific SF and NFL MKT (highest disparity)</i>									
MLB	0.574**	0.175	0.001	0.031	0.036	0.391	0.026	0.014	0.063
NBA	0.684***	0.179	0.000	0.042	0.037	0.264	0.034*	0.014	0.020
NFL	0.684***	0.179	0.000	0.042	0.037	0.264	0.034*	0.014	0.020
NHL	0.574**	0.175	0.001	0.031	0.036	0.391	0.026	0.014	0.063


NOTES.—SE = Standard error (Arellano (1987) corrected). Effects represent the impact of jointly adopting salary cap and salary floor at the specific market size disparity level and salary floor level indicated. We use salary floor values of 90% for the NBA and NFL and values of 75% for the NHL and MLB. **p* < .05, ***p* < .01, ****p* < .001. Figures listed as 0.000 are not truly zero but reflect extremely small quantities.

Figure 1: Trends Over Time in Competitive Balance Measures and Market Size Disparity by League



NOTES.—For each competitive balance measure, lower values imply more competitive balance. Market size disparity is defined as the coefficient of variation of MSA populations across league member franchises.

Figure 2: Randomization Inference: CBI Implementation Re-timing and Market Size Disparity Interactions



NOTES.—Each histogram shows the distribution of interaction coefficients from 1,000 placebo simulations. In each simulation, a random adoption year is drawn for each CBI; SC and SF share one draw for NFL and NHL (reflecting simultaneous CBA introduction) but are drawn independently for NBA (SC 1984, SF 1999). Binary CBIs (SC, LT) are set persistently to 1 from the random start year. Continuous CBIs (RS, SF) use the actual post-adoption trajectory rescaled via linear interpolation; longer placebo periods hold the final value constant. MLB is excluded from randomization of SC and SF and set to SC = SF = 0 in all simulations, as it never adopted these CBIs. The dashed red line marks the actual estimated coefficient. Placebo p -values report the share of simulations in which the magnitude of the placebo coefficient is greater than or equal to the magnitude of the actual coefficient estimate.

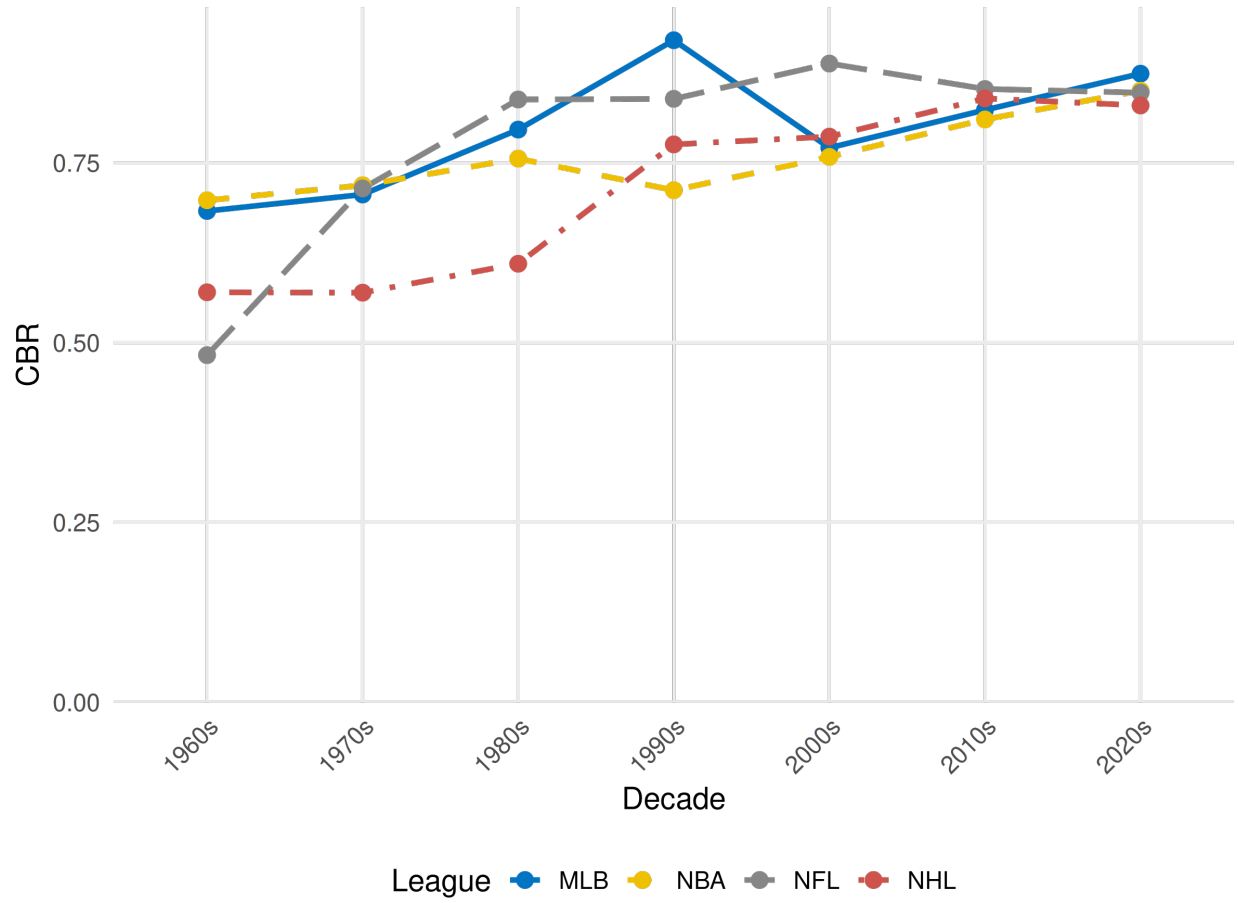
Appendix

Table A.1: Average Competitive Balance Ratio by League (1967–2023)

League	N	Statistic	CBR
MLB	55	Mean	0.81
		SD	(0.09)
		Range	[0.65, 0.95]
NBA	55	Mean	0.76
		SD	(0.06)
		Range	[0.62, 0.88]
NFL	55	Mean	0.82
		SD	(0.08)
		Range	[0.48, 0.91]
NHL	54	Mean	0.72
		SD	(0.12)
		Range	[0.50, 0.90]


NOTES.—This table reports summary statistics comparable to Table 2 but for the Competitive Balance Ratio (Humphreys, 2002). Note that CBR is oppositely defined from the other three measures: higher values of CBR imply greater levels of competitive balance.

Figure A.1: Trends Over Time in Competitive Balance Ratio by League



NOTES.—For the CBR, higher values imply more competitive balance. This is oppositely signed compared to the three measures we examine in the body of the paper.

Figure A.2: Randomization Inference: CBI Implementation Re-timing and CBI Main Effects



NOTES.—Each histogram shows the distribution of main effect coefficients from 1,000 placebo simulations. In each simulation, a random adoption year is drawn for each CBI; SC and SF share one draw for NFL and NHL (reflecting simultaneous CBA introduction) but are drawn independently for NBA (SC 1984, SF 1999). Binary CBIs (SC, LT) are set persistently to 1 from the random start year. Continuous CBIs (RS, SF) use the actual post-adoption trajectory rescaled via linear interpolation; longer placebo periods hold the final value constant. MLB is excluded from randomization of SC and SF and set to SC = SF = 0 in all simulations, as it never adopted these CBIs. The dashed red line marks the actual estimated coefficient. Placebo p -values report the share of simulations in which the magnitude of the placebo coefficient is greater than or equal to the magnitude of the actual coefficient estimate.