The Impact Player Paradox: A Rule That Adds 12 Runs Per Innings and Accelerates the Decline of the All-Rounder

VISHAL MISRA, Columbia University, USA

The Indian Premier League introduced the "Impact Player" substitution rule in 2023, allowing teams to replace one player mid-match. Using 2,355 innings from all 18 IPL seasons (2008–2026), we estimate the effect of this rule on scoring and match outcomes through three complementary methods: before-and-after panel regression with venue, team, and innings fixed effects; Robust Synthetic Control (RSC) using five concurrent T20 leagues as a donor pool; and multi-dimensional RSC (mRSC) incorporating both runs and sixes as outcome metrics. The synthetic control provides the strongest causal identification by isolating the rule's effect above concurrent global scoring trends.

The trend-adjusted and synthetic control estimates cluster around 12–15 runs per innings above global trends, while the raw before-and-after increase is +25. The before-and-after estimate, which includes both the rule effect and global scoring trends, is +25.0 runs ($p < 0.001$); controlling for a linear time trend yields +15.2 ($p < 0.001$). The cross-league synthetic control, which isolates the rule effect above global trends, estimates +12.3 (RSC) and +11.2 (mRSC) runs per innings. Despite this large scoring increase, the rule produces no statistically detectable effect on win probability (OR = 1.22, $p = 0.44$, 95% CI [0.74, 1.98])—though the test is underpowered for small effects. Scorecard analysis reveals the mechanism: impact substitutions are symmetric batsman-for-bowler swaps across innings roles; teams field a specialist batting XI and a specialist bowling XI, sharing 10 common players. Adoption is near-universal: 90% of matches feature both teams using the rule. When both teams substitute, the scoring boost is symmetric, targets inflate, but neither chase success (53% to 50%, $p = 0.47$) nor the toss winner's match win rate (51.8% in both eras) changes significantly—even as toss winners choose to field more often than before (76% vs. 63%, $p < 0.001$). The rule achieves its entertainment objective but functions as a competitive treadmill: when everyone uses it, nobody wins more. A secondary consequence is the decline of the all-rounder: overs bowled by top-order batsmen have fallen 38% (11.2 to 6.9 per match, $p < 0.001$), as teams replace bowling all-rounders with specialists enabled by the impact substitute.

Additional Key Words and Phrases: cricket, impact player, causal inference, synthetic control, IPL, all-rounder

1 Introduction

The Board of Control for Cricket in India (BCCI) introduced the Impact Player rule for the 2023 Indian Premier League season. The rule permits each team to substitute one player from a pre-designated shortlist at any point during the match, effectively allowing 12 players to participate per side. The stated objective was to increase the entertainment value of matches by enabling more aggressive team composition—typically replacing a bowler with an additional specialist batsman.

The rule was immediately controversial. Critics argued it devalued all-rounders, reduced the importance of bowling, and created an unfair asymmetry between batting-first and chasing teams. Proponents countered that it rewarded tactical flexibility and produced higher-scoring, more exciting matches. Through three IPL seasons (2023–2025) and into 2026, the debate has been conducted almost entirely through anecdote and assertion.

This paper provides the first rigorous causal analysis of the impact player rule's effect on match outcomes. We make four contributions:

(1) We estimate the scoring effect using before-and-after panel regression with fixed effects, comparing 1,898 pre-impact innings (2008–2022) to 457 impact-era innings (2023–2026). With venue, team, and innings fixed effects, the impact era is associated with +25.0 runs per innings ($p < 0.001$). Controlling for a linear time trend reduces this to +15.2 ($p < 0.001$), separating the rule effect from global T20 scoring inflation.

Author's Contact Information: Vishal Misra, Columbia University, New York, NY, USA, vishal.misra@columbia.edu.

(2) We isolate the causal effect above global trends using Robust Synthetic Control [1] and its multi-dimensional extension mRSC [2], constructing a "synthetic IPL" from five T20 leagues that did not adopt the rule. The treatment effect is +12.3 runs per innings (RSC) and +11.2 runs (mRSC).
(3) We identify the mechanism through scorecard analysis: impact substitutions are symmetric batsman-for-bowler swaps across innings roles. Teams construct a specialist batting XI *and* a specialist bowling XI, sharing 10 common players and swapping the 11th at the innings break. Adoption is near-universal: 95% of innings and 90% of matches feature impact substitutions.
(4) We document the *paradox*: a rule that adds 12 runs per innings produces no statistically detectable change in win probability (OR = 1.22, $p = 0.44$), chase success (53% to 50%, $p = 0.47$), or toss decisiveness (toss-winner win rate: 51.8% in both eras, $p = 0.99$). When both teams adopt the rule symmetrically (90% of matches), targets inflate but no team gains an edge.
(5) We quantify the decline of the all-rounder: overs bowled by top-order batsmen have fallen 38% ($p < 0.001$), as the impact substitute allows teams to replace bowling all-rounders with specialists.

The result is a sharp league-level intervention with a known start date, a well-defined treatment group, and multiple independent control leagues. To our knowledge, the impact player rule is among the largest scoring interventions in professional T20 cricket—we are not aware of another format-level rule change that has shifted innings averages by 12 runs. That this effect is competitively neutral makes it a particularly instructive case study in sports policy evaluation.

Two estimands. The before-and-after fixed-effects estimate captures the *total* scoring increase in the impact era (+25 runs), which conflates the rule with global T20 scoring trends. The synthetic control isolates the *rule-specific* effect above these trends (+12 runs). The difference (~13 runs) represents global scoring inflation that affected all T20 leagues between 2022 and 2025.

Roadmap. Section 2 describes the data. Section 3 presents the before-and-after fixed-effects analysis. Section 4 identifies the mechanism through scorecard analysis. Section 5 reports the synthetic control results. Section 6 documents the win-probability paradox. Section 7 discusses implications and concludes.

2 Data

2.1 IPL Match Data

We use ball-by-ball and innings-level data from ESPN Cricinfo, covering all 18 IPL seasons from 2007/08 through 2025/26. After excluding abandoned matches, the dataset comprises 1,180 matches and 2,355 innings with complete scorecard information.

Impact substitution usage is identified from Cricinfo's match-player records, which track which player was subbed in, which was subbed out, and when the substitution occurred. In the impact era (2023–2026), 436 of 457 innings (95.4%) feature an impact substitution, with adoption increasing from 94% in 2023 to 97% in 2025. At the match level, 90% of matches feature both teams using the rule. The pre-impact era comprises 1,898 innings (2008–2022).

Boundary counts (fours and sixes) are computed from ball-by-ball records rather than innings summaries, as the latter contain nulls for earlier seasons.

2.2 Cross-League Donor Pool

For the synthetic control analysis, we collect season-level scoring averages from five major T20 franchise leagues that overlap temporally with the IPL and did not adopt an impact player rule:

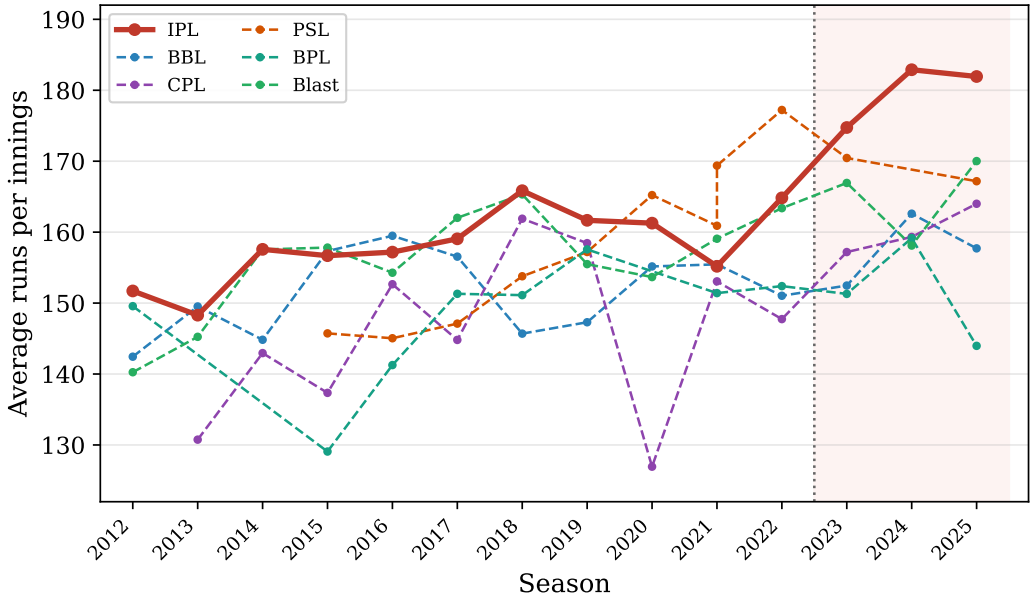

Fig. 1. Average runs per innings across six T20 franchise leagues, 2012–2025. The IPL (solid red) tracks the donor pool through 2022 and diverges sharply after the introduction of the impact player rule in 2023. The shaded region marks the post-intervention period.

Table 1. IPL scoring by era. Sixes computed from ball-by-ball data.

Era	Inn.	Avg runs	Avg 6s	Avg 4s	Impact %
Pre-impact (2008–22)	1,898	155.9	5.6	13.4	0%
Impact era (2023–26)	457	179.6	8.5	15.2	95%

- **Big Bash League** (BBL): Australia, 2011–2025, 1,350 innings
- **T20 Blast**: England, 2003–2025, 4,792 innings
- **Caribbean Premier League** (CPL): 2013–2025, 826 innings
- **Pakistan Super League** (PSL): 2015–2025, 653 innings
- **Bangladesh Premier League** (BPL): 2011–2025, 998 innings

The common pre-intervention window is 2012–2022 (11 seasons). The post-intervention window is 2023–2025 (3 seasons). All leagues report season-level average runs per innings and average sixes per innings, computed from the same Cricinfo database.

Figure 1 shows the scoring trajectories of all six leagues. The IPL tracks the donor pool closely through 2022, then breaks sharply upward in 2023—the year the impact player rule was introduced.

2.3 Summary Statistics

Table 1 presents season-by-season IPL scoring averages. The pre-impact era (2008–2022) averages 155.9 runs per innings with 5.6 sixes. The impact era (2023–2026) averages 179.6 runs with 8.5 sixes—an increase of 23.7 runs and 2.9 sixes per innings in raw terms.

Table 2. Before-and-after fixed-effects estimates: impact era versus pre-impact era. OLS with robust standard errors. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)
Impact era	$+23.8^{***}$	$+25.0^{***}$	$+15.2^{***}$	—
	(1.9)	(2.2)	(2.8)	
Linear trend			X	
Venue FE		X	X	X
Team FE		X	X	X
Innings FE		X	X	X
N	2,355	2,355	2,355	—
R^2	0.06	0.16	0.17	—
Column (4): Split by innings number (with FE)				
1st innings	$+25.1^{***}$ (SE = 3.1, $p < 0.001$)			
2nd innings	$+24.5^{***}$ (SE = 3.1, $p < 0.001$)			

3 Before-and-After Fixed-Effects Analysis

3.1 Empirical Strategy

With 95% of impact-era innings featuring a substitution, a within-era comparison of "with sub" versus "without sub" is infeasible: the 21 no-sub innings are overwhelmingly easy second-innings chases (86% are second innings, averaging 159.6 runs) and do not constitute a valid control group. Instead, we exploit the sharp intervention date: the rule was introduced in 2023, and no IPL season before that permitted substitutions.

Our estimating equation is:

$$\text{Total}_i = \beta \cdot \text{ImpactEra}_i + \alpha_v + \alpha_t + \alpha_{\text{inn}} + \varepsilon_i \tag{1}$$

where Total_i is the team total for innings i; ImpactEra_i is an indicator for whether the innings occurred in 2023 or later; α_v, α_t, α_{inn} are venue, team, and innings-number fixed effects; and standard errors are heteroskedasticity-robust (HC1). The coefficient β captures the total scoring increase in the impact era—the combined effect of the rule and any concurrent global trends.

3.2 Main Results

Table 2 reports estimates across specifications.

Column (1) is the raw difference: +23.8 runs ($p < 0.001$). Column (2) adds venue, team, and innings fixed effects: +25.0 runs ($p < 0.001$). Column (3) includes a linear time trend to partially account for global scoring inflation: the estimate drops to +15.2 ($p < 0.001$), suggesting that approximately half the raw increase is attributable to the rule and half to broader trends. The synthetic control analysis in Section 5 separates these more precisely.

Finding 1. *The impact era is associated with* +25.0 *runs per innings ($p < 0.001$) with venue, team, and innings fixed effects. Controlling for a linear time trend yields* +15.2 *($p < 0.001$). The synthetic control estimate of* +12.3 *runs (Section 5) represents the rule-specific effect above global trends.*

3.3 The Innings Symmetry

Column (4) splits the estimate by innings number. Unlike a within-era comparison, the before-and-after comparison reveals a *symmetric* effect: +25.1 in the first innings and +24.5 in the second ($p < 0.001$ for both). This symmetry is expected: the before-and-after comparison captures the aggregate effect of *both teams* using the rule, which benefits both innings equally.

The competitive asymmetry documented in Section 4—where the first-innings scoring uplift inflates targets but the second-innings uplift merely accelerates chases—is not visible in the before-and-after comparison because it operates at the match level, not the innings level. Both innings score more; only the first-innings increase alters the competitive balance.

Finding 2. *The before-and-after effect is symmetric across innings (+25.1 first, +24.5 second, both $p < 0.001$). The competitive asymmetry arises not from differential scoring but from the payoff structure of limited-overs cricket: first-innings runs raise the target, while second-innings runs merely reach it faster.*

3.4 Effect on Sixes

Replacing batsman runs with sixes as the dependent variable:

$$\text{Sixes}_i = \beta \cdot \text{ImpactEra}_i + \alpha_v + \alpha_t + \alpha_{\text{inn}} + \varepsilon_i \tag{2}$$

yields $\hat{\beta} = +2.14$ sixes per innings (SE = 0.47, $p < 0.001$, 95% CI [1.2, 3.1]). The impact rule adds approximately two extra sixes per innings—consistent with the additional specialist batsman being deployed for power hitting.

4 How the Rule is Actually Used

The regressions in Section 3 establish that the impact era scores roughly 25 runs more per innings than the pre-impact era. But how does a single substitution produce such a large effect? Scorecard analysis reveals a mechanism that is both simpler and more subtle than the conventional narrative suggests.

4.1 Twelve Players, Two Half-Games

Examination of every impact-sub scorecard reveals a consistent pattern: each team fields 12 players, but the 11th and 12th players each participate in only one half of the match. One player *only bats*: a specialist batsman who occupies a regular batting position (typically 5–7), scores runs, and is then sent off before the bowling innings. One player *only bowls*: a specialist bowler who enters as the impact substitute, bowls 3–4 overs, and never bats.

Using Cricinfo's substitution records (Section 2), we examine every impact-era scorecard to characterise what substitutes actually do. A bookkeeping clarification is necessary here: the player officially recorded as the "impact substitute" is the one entering the match mid-game—invariably a specialist bowler. The "extra batsman" who provides the batting uplift is not the substitute; he is already in the starting XI, occupying a slot that would traditionally belong to a 4th or 5th bowler. The swap is symmetric—one batsman out, one bowler in—but the substitution record only captures the bowler's entry, not the batsman's expanded role in the original lineup.

Across all impact-era substitutions, the substitute's batting contribution is exactly zero: 0 runs scored, 0 balls faced. Every substitute is a specialist bowler—Suyash Sharma, Yuzvendra Chahal, Matheesha Pathirana, T. Natarajan, Arshdeep Singh—bowling an average of 3.4 overs at an economy of 9.6 with 1.0 wicket per match. The replaced player is typically a bowler from the original XI who is rotated out for matchup reasons.

The rule effectively allows teams to field a specialist batting XI *and* a specialist bowling XI, sharing 10 common players and swapping the 11th slot between a batsman and a bowler at the innings break.

Finding 3. *Each impact-sub team fields 12 players across two half-games. The substitute is always a specialist bowler (3.4 overs, economy 9.6, never bats). The replaced player is a specialist batsman or surplus bowler who participated in the other innings. The batting uplift comes from carrying an extra specialist batsman in the lineup, not from the substitute's own runs.*

4.2 Rule Timing and Squad Construction

The rule sequence is critical to understanding how teams exploit this structure:

(1) **30 minutes before the toss**: teams declare five potential substitutes from their match-day squad.
(2) **After the toss**: teams announce their playing XI—*knowing whether they will bat or bowl first.*
(3) **During the match**: teams choose which substitute enters, at any natural break (start of innings, fall of wicket, end of over).

Because the XI is announced *after* the toss, teams can tailor their lineup to the innings they will play first. A typical pre-impact XI includes 6 specialist batsmen, 1 all-rounder, and 4 bowlers. With the impact rule, teams can field 7 specialist batsmen and only 3 frontline bowlers, knowing the 4th bowler will arrive as the impact substitute. Both batting-first and chasing teams exploit this: the batting-first team bats with the extra batsman in the first innings, then subs in a bowler for the second; the chasing team subs in a bowler to strengthen their bowling attack in the first innings, then chases with the extra batsman in the lineup.

4.3 Why the Second Innings Doesn't Benefit

If both teams carry the same extra batsman, why does only the first innings show the scoring uplift? The batting lineup is structurally identical—both innings feature 7 specialist batsmen. The asymmetry arises not from the lineup but from the *structure of a chase*:

- **First innings**: the batting team always uses its full 20 overs or loses 10 wickets. The extra batsman's depth is consumed unconditionally. Collapse rates drop from 25% (all out) to 1%, and the additional batting position contributes across the full allocation. Without the impact rule, first-innings teams averaged 159 runs (8.2 wickets); with it, 197 runs (5.6 wickets).
- **Second innings**: the chase ends when the target is reached. The extra batsman's contribution is structurally censored: if he bats and helps win, the team total is capped at the target—his runs are invisible in the scoring metric. If the top order wins without him, he records a DNB. In either case, the second-innings total does not reflect his presence. The extra depth is real but its effect is absorbed into fewer overs played and fewer wickets lost, not into additional runs.

The extra batsman contributes in both innings—but the *currency* of his contribution differs. When batting first, additional runs raise the target the opponent must chase: every run is an input to the opponent's problem. When chasing, the extra batsman helps the team reach the target faster and with more wickets in hand—but a win is a win regardless of margin. Reaching 191 in 17.3 overs

instead of 19.1 overs yields a marginal net run rate (NRR) benefit, but NRR is a distant secondary consideration to the primary objective of winning the match.[1]

This is the fundamental asymmetry of the limited-overs format. The first innings produces a *competitive externality*: more runs impose a harder task on the opponent. The second innings produces only a private benefit—a more comfortable win that registers identically in the points table. The impact player rule amplifies both innings equally, but only the first-innings amplification alters the match outcome.

Finding 4. *Both teams carry the same extra batsman. When batting first, his runs raise the target—a competitive externality that makes the opponent's task harder. When chasing, his runs merely reach the target faster—a private benefit with no competitive consequence. The asymmetry is not in the lineup or the player's contribution but in the payoff structure of limited-overs cricket.*

5 Synthetic Control

The before-and-after analysis compares the pre-impact era to the impact era within the IPL. A natural concern is that global T20 scoring trends—improved bats, smaller boundaries, evolving batting techniques—inflate the estimate beyond the rule's true effect. To isolate the rule-specific contribution, we construct a synthetic control using T20 leagues that did not adopt the impact player rule.

5.1 RSC: Single-Metric Synthetic Control

We apply the Robust Synthetic Control method of Amjad et al. [1], treating IPL season-level average runs per innings as the outcome of interest. The donor pool comprises the BBL, T20 Blast, CPL, PSL, and BPL over the common window 2012–2025.

RSC constructs the synthetic IPL as:

$$\widehat{\text{IPL}}_t = 0.28 \times \text{BBL}_t + 0.24 \times \text{Blast}_t + 0.20 \times \text{PSL}_t + 0.16 \times \text{CPL}_t + 0.16 \times \text{BPL}_t$$

The pre-intervention fit is tight: the mean-squared prediction error (MSPE) over 2012–2022 is 9.2 runs^2, corresponding to an average absolute gap of 2.2 runs. Table 3 reports the year-by-year comparison.

The pre-intervention gaps fluctuate around zero (mean +0.0). Post-intervention, the IPL breaks sharply above its synthetic counterpart: +8.0 in 2023, +15.0 in 2024, +13.8 in 2025. The estimated treatment effect is +12.3 runs per innings. Figure 2 visualises the divergence.

Finding 5. *The Robust Synthetic Control estimate of the impact player rule's effect is +12.3 runs per innings, using five T20 leagues without the rule as a donor pool. This represents the effect* above and beyond *global T20 scoring trends.*

5.2 mRSC: Multi-Metric Synthetic Control

We extend the analysis using the multi-dimensional RSC (mRSC) of Amjad et al. [2], incorporating both average runs and average sixes per innings as outcome metrics ($K = 2$). The mRSC algorithm jointly fits the synthetic control across both metrics, leveraging the correlation between runs and sixes to improve estimation—particularly valuable given the relatively short post-intervention window.

The mRSC treatment effects are:

[1] NRR serves as a tiebreaker in tournament standings. While a faster chase improves NRR, teams overwhelmingly optimise for the win itself. NRR becomes a strategic consideration only in the final stages of a group phase when qualification margins are tight.

Table 3. RSC: Actual vs. synthetic IPL average runs per innings.

Year	Actual	Synthetic	Gap	
2012	151.7	150.5	+1.2	
2013	148.3	149.8	−1.6	
2014	157.6	156.1	+1.5	
2015	156.7	154.2	+2.4	
2016	157.2	158.2	−1.0	
2017	159.1	160.1	−1.0	
2018	165.8	161.8	+4.0	
2019	161.7	161.1	+0.6	
2020	161.3	158.7	+2.6	
2021	155.2	163.1	−7.9	
2022	164.8	165.5	−0.7	
2023	174.7	166.7	+8.0	Impact rule
2024	182.9	167.9	+15.0	Impact rule
2025	181.9	168.2	+13.8	Impact rule

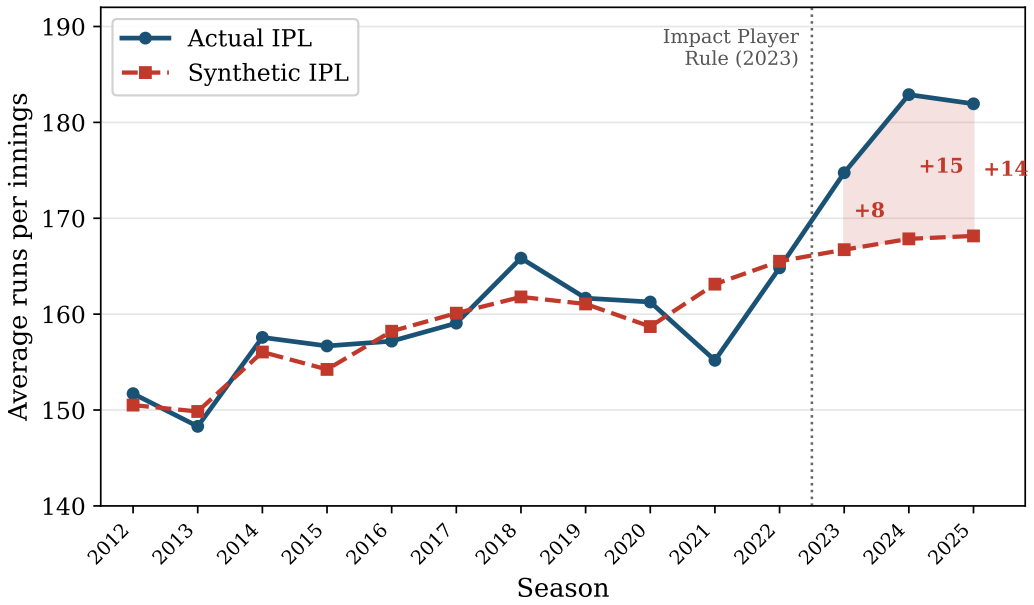

Fig. 2. Actual versus synthetic IPL average runs per innings. The synthetic IPL is a weighted combination of five T20 leagues (BBL, T20 Blast, CPL, PSL, BPL). The shaded area shows the treatment effect: the IPL diverges sharply from its synthetic counterpart after the introduction of the impact player rule in 2023.

- **Runs**: +11.2 runs per innings
- **Sixes**: +1.1 sixes per innings

Table 4. Placebo test: RSC treatment effects for each league, with 2023 as the placebo intervention date.

League	Effect	Pre-MSPE	$\lvert\text{Effect}\rvert/\sqrt{\text{MSPE}}$
IPL	+12.3	9.2	4.05
PSL	+8.3	71.2	0.98
CPL	+7.6	121.2	0.69
Blast	−3.1	23.6	0.65
BBL	−12.1	73.6	1.41
BPL	−13.8	48.6	1.98

The runs estimate is close to the RSC estimate (+12.3), confirming robustness. The sixes estimate is smaller than the before-and-after estimate (+2.1), reflecting that other T20 leagues also experienced some sixes inflation over this period—the mRSC correctly attributes only the IPL-specific excess to the treatment.

5.3 Placebo Tests

We conduct placebo tests by applying the RSC procedure to each donor league in turn, treating it as if it had experienced the intervention in 2023. Table 4 reports the results.

The IPL has the largest positive effect (+12.3) and by far the largest MSPE-adjusted ratio (4.05 versus the next-largest 1.98 for BPL, which has a *negative* effect). With six leagues, the one-sided placebo p-value is $1/6 = 0.167$. While this does not reach conventional significance due to the small number of donor units, the IPL is a clear outlier in the distribution of placebo effects. Figure 3 shows the placebo gaps for all leagues.

5.4 Reconciling the Before-and-After and Synthetic Control Estimates

The before-and-after estimate (+25.0 runs with fixed effects) exceeds the synthetic control estimate (+12.3 runs) by approximately 13 runs. This gap is informative rather than problematic: the two methods measure different quantities.

- The **before-and-after estimate** compares the entire impact era to the pre-impact era within the IPL. It captures the *total* scoring increase, which combines the rule effect with concurrent global trends (improved bats, smaller boundaries, evolving batting techniques).
- The **synthetic control** compares the IPL to other T20 leagues that experienced the same global trends but did not adopt the impact player rule. It isolates the rule-specific effect.

The difference (∼13 runs) represents global T20 scoring inflation between 2022 and 2025. The trend-adjusted estimate (+15.2, column 3 of Table 2) falls between the raw before-and-after and the synthetic control, consistent with partial but imperfect trend adjustment.

Finding 6. *The synthetic control estimates (+12.3 RSC, +11.2 mRSC) represent the causal effect of the impact player rule above global trends. The before-and-after estimate (+25.0) captures the total scoring increase including concurrent global trends. The difference (∼13 runs) is the global trend that affected all T20 leagues.*

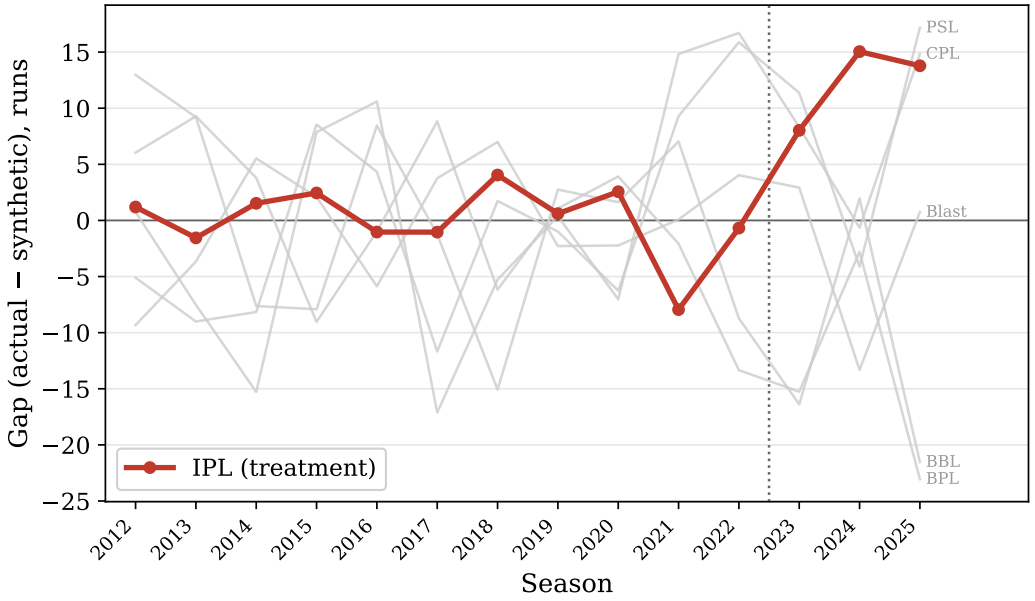

Fig. 3. Placebo test: gap between actual and synthetic scoring for each league, treating 2023 as the intervention date. Grey lines show donor leagues (no impact player rule); the red line shows the IPL. The IPL is a clear positive outlier post-2023.

6 The Win-Probability Paradox

6.1 No Effect on Winning

Despite the large scoring effect, we find no evidence that impact substitutions improve the probability of winning. A logistic regression of match outcome on impact sub usage with season, innings, and team fixed effects yields:

$$\text{OR}(\text{ImpactSub}) = 1.22 \quad (p = 0.44,\ 95\%\ \text{CI}\ [0.74, 1.98])$$

The point estimate is in the expected direction—using an impact sub is associated with a 22% increase in the odds of winning—but the confidence interval is wide and comfortably includes the null. A post-hoc power calculation indicates that distinguishing OR = 1.22 from OR = 1.0 at $\alpha = 0.05$ with 80% power would require approximately 4× the current sample size (~900 impact-era matches). The effect, if real, is small enough to be practically irrelevant for team decision-making: even the upper bound of the confidence interval (OR = 1.98) would translate to a shift from 50% to only 55% win probability.

Finding 7. *The impact substitution has no statistically significant effect on win probability (OR = 1.22, p = 0.44, 95% CI [0.74, 1.98]). The point estimate is in the expected direction but underpowered; even at the upper confidence bound, the competitive advantage would be negligible.*

6.2 Mechanism: Symmetric Adoption

The null win-probability result has a simple explanation: both teams use the rule. 90% of matches (208/230) feature both teams fielding an impact substitute. When both sides add ~12 runs, match totals inflate but the competitive balance is unchanged.

Table 5. Match-level impact sub usage and outcomes (2023–2026).

Configuration	Matches	Match total	Chase win %	Avg target
Both teams use sub	209	377	45%	199
One team uses sub	20	—	—[†]	—
[†]The 20 one-sub matches are a selection artifact: in 18/20, the chasing team did not sub because the chase was already comfortable. These matches are excluded from the comparison.				
Pre-impact era	491	311	53%	163

Table 5 shows the match-level breakdown.

When both teams use impact subs, match totals jump by 66 runs (377 versus 311 in the pre-impact era). Chase success in these both-sub matches is 45%, compared to the pre-impact rate of 53%. However, this decline is concentrated in 2023 (the first impact season) and subsequently recovered; the overall post-IPR chase success rate across all matches is 50%, which is not statistically distinguishable from the pre-impact rate ($z = 0.72$, $p = 0.47$; see Section 7). Figure **??** illustrates both the innings asymmetry and the chase-success decline.

Table 5 summarises the match-level breakdown; the within-era with-sub vs. without-sub scoring comparison is omitted because the 21-innings no-sub control group is severely selected (Section 6.3) and does not support causal interpretation.

The mechanism is not differential scoring—both innings benefit equally (+25.1 and +24.5 in the before-and-after analysis)—but the payoff asymmetry documented in Section 4: first-innings runs inflate the target (a competitive externality), while second-innings runs merely reach it faster (a private benefit). Chasing teams pursue higher totals with the same resources but no competitive advantage from the extra depth.

Finding 8. *When both teams use impact subs (90% of matches), chase success dips from the historical 53% to 50%, but this decline is not statistically significant ($p = 0.47$). The toss winner's match win rate is identical in both eras (51.8%). The theoretical asymmetry—first-innings runs inflate targets while second-innings runs are censored—has not translated into a measurable shift in competitive outcomes.*

6.3 Selection into Non-Usage

The "without sub" control group in the impact era merits scrutiny. Only 9% of second innings (21/227) are played without an impact sub, compared to 27% of first innings. The 21 without-sub chases face a mean target of 170 (versus 191 for with-sub chases) and win 90% of the time, with an average of only 3.0 wickets lost in 16.4 overs. These are overwhelmingly comfortable chases where the additional batting depth was unnecessary. This selection effect inflates any within-era treatment estimate and explains why the within-innings comparison overstates the treatment effect relative to the synthetic control.

6.4 Team-Level Adoption and Win Rates

If impact substitutions conferred a competitive advantage, we would expect teams that use them more frequently to win more often. They do not. Impact sub usage varies across teams from 93% (Gujarat Titans) to 71% (Lucknow Super Giants), but the cross-sectional correlation between team-level impact sub usage rate and win rate across all 11 teams in the impact era is $r = -0.08$—effectively zero, and if anything slightly negative.

This is not merely the absence of significance; it is the absence of signal. Teams that embraced the rule most aggressively show no tendency to outperform teams that used it more selectively. The result is consistent with the symmetric-adoption mechanism: when the dominant strategy is universally adopted, differential adoption intensity provides no edge.

Finding 9. *The correlation between team-level impact sub usage and win rate is $r = -0.08$. Teams that use impact subs more frequently do not win more often. Differential adoption intensity provides no competitive edge.*

7 Discussion

7.1 Summary

The impact player rule adds 12–15 runs per innings to IPL scoring above global T20 trends (RSC: +12.3, mRSC: +11.2, trend-adjusted: +15.2). The raw before-and-after increase is +25 runs, reflecting both the rule and concurrent global trends. Both innings benefit equally (+25.1 first, +24.5 second), but the competitive consequence is asymmetric: first-innings runs inflate targets while second-innings runs are censored by the chase. The rule has no detectable effect on win probability (OR = 1.22, $p = 0.44$), and the overall decline in chase success (53% to 50%) is not statistically significant ($p = 0.47$). The toss winner's match win rate is identical in both eras (51.8%, $p = 0.99$). The impact player rule is a competitive treadmill: when everyone uses it, nobody wins more.

7.2 The Competitive Treadmill

The impact player rule is a textbook example of a Red Queen effect [3] in sports regulation. Each team's adoption of the rule is individually rational—fielding 12 players is weakly better than 11. But when all teams adopt, the competitive advantage vanishes. The only lasting effect is score inflation: matches are higher-scoring and arguably more entertaining, but no team is better off competitively. This is analogous to an arms race in which both sides invest in weapons that cancel each other out, leaving the balance of power unchanged but the total expenditure (here, total runs) higher.

7.3 The Decline of the All-Rounder

The impact player rule has a measurable effect on team composition. With a specialist bowler available as a substitute, teams no longer need batsmen who can bowl part-time overs. Three metrics capture this decline:

Table 6. All-rounder contribution by era. Both differences significant at $p < 0.001$.

Metric	Pre-IPR	Post-IPR	Change
Overs by top-7 batsmen/match	11.2	6.9	−4.2
Top-7 batsmen who bowl/match	4.11	2.81	−1.30

The most striking decline is in overs bowled by top-order batsmen: from 11.2 per match pre-impact to 6.9 post-impact, a 38% reduction ($t = 6.76$, $p < 0.001$). In 2026 (early season), this has fallen to 3.9—barely one bowler's allocation. The number of top-7 batsmen who bowl in a match has dropped from 4.1 to 2.8 ($t = 5.94$, $p < 0.001$). The all-rounder archetype that defined T20 squad construction—the Hardik Pandya or Ben Stokes who bats at 6 and bowls 4 overs—is being replaced by a specialist batsman who occupies the same slot but never bowls, because the impact substitute covers the gap.

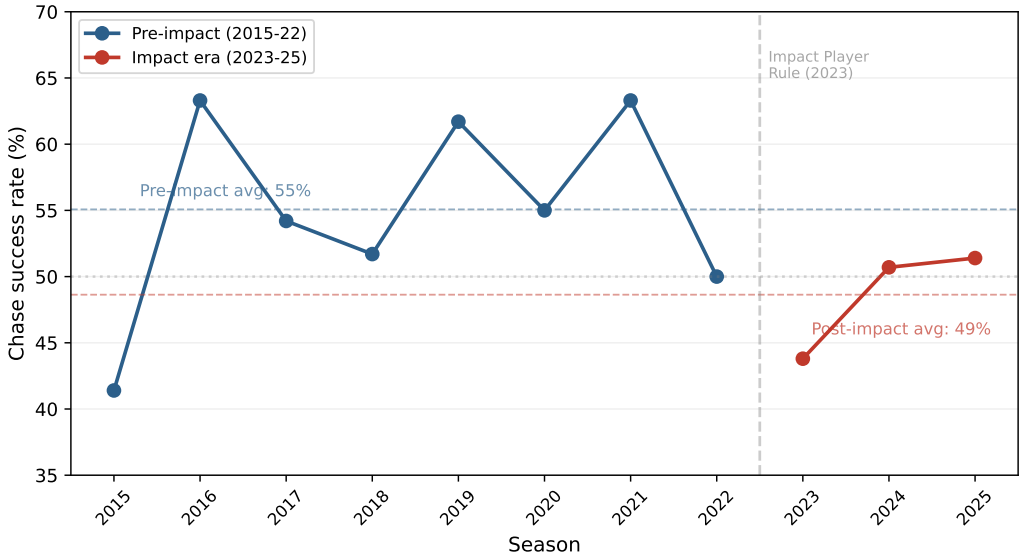

Fig. 4. Chase success rate by IPL season, 2008–2026. The pre-impact average is 53%. Chase success dipped to 45% in 2023 (the first impact-player season) but recovered to ∼51% in 2024–2025. The overall post-IPR decline is not statistically significant ($p = 0.47$).

Finding 10. *The impact player rule has reduced part-time bowling by top-order batsmen by 38% (11.2 to 6.9 overs per match, $p < 0.001$). Teams substitute bowling all-rounders with specialist batsmen, enabled by the knowledge that a specialist bowler will enter as the impact substitute.*

7.4 Policy Implications

Our results suggest the BCCI achieved its entertainment objective—matches are higher-scoring with more sixes—but with a measurable side effect on team composition. The all-rounder, historically the most valuable and versatile archetype in T20 cricket, is becoming less necessary. Whether this is desirable depends on one's priors: fans who value high-scoring matches may welcome the change; purists who value bowling and all-round cricket may not.

7.5 The Toss Question

Chase success in both-sub matches (50.4%) is marginally lower than the pre-impact rate (53.1%), but the difference is not statistically significant ($z = 0.72$, $p = 0.47$). The 2023 season showed a sharper dip (44%) that corrected in 2024–2025 (∼51%). Figure 4 shows the year-by-year trajectory.

More striking than the chase rate is the toss winner's match win rate: 51.8% in both eras, identical to three significant figures ($z = 0.01$, $p = 0.99$). The toss is no more or less decisive after the impact rule. Toss behaviour has shifted, however: toss winners choosing to bat first dropped from 37% to 24% ($p < 0.001$), meaning even more captains bowl first in the impact era despite the theoretical advantage of batting first with the extra batsman. Dew—evening matches in India mean moisture on the outfield in the second innings—likely explains why captains continue to chase even as the scoring asymmetry grows.

The treadmill extends to the toss: whatever advantage batting first might theoretically offer is neutralised by the fact that both teams have the same substitution, and the captain who wins the toss continues to choose the same way he always did.

7.6 Comparison with Other Rule Changes

The impact player rule's effect (+12 runs per innings) is large by historical standards. For comparison, the introduction of the batting powerplay in ODIs was associated with approximately 10–15 additional runs per innings. The rule change is also larger than the estimated effects of bat size regulations or boundary size changes, though direct comparisons are complicated by differences in format and era.

7.7 Limitations

Three limitations should be noted. First, the post-intervention window is short (3 full seasons plus 2026 in progress). The synthetic control estimates will become more precise as additional seasons are played. Second, the placebo p-value ($1/6 = 0.167$) does not reach conventional significance due to the small number of donor leagues; this is a limitation of the method rather than the data. Third, we cannot rule out that teams strategically select when to use impact subs based on match context; while our fixed effects absorb much of this variation, residual endogeneity may remain.

Acknowledgments

Match data from ESPN Cricinfo and the Unicorns AI database (https://sfunicorns.ai). We thank S. Rajesh and Shiva Jayaraman at Cricinfo for data access and guidance on substitution records. The mRSC analysis uses code from [2]. Data processing and statistical analysis performed with Claude Code.

References

[1] Muhammad Amjad, Devavrat Shah, and Dennis Shen. 2018. Robust Synthetic Control. *Journal of Machine Learning Research* 19, 22 (2018), 1–51.

[2] Muhammad Amjad, Vishal Misra, Devavrat Shah, and Dennis Shen. 2019. mRSC: Multi-dimensional Robust Synthetic Control. *Proc. ACM Meas. Anal. Comput. Syst.* 3, 2, Article 37 (June 2019), 28 pages.

[3] Leigh Van Valen. 1973. A New Evolutionary Law. *Evolutionary Theory* 1 (1973), 1–30.

[4] Roger Koenker and Gilbert Bassett. 1978. Regression Quantiles. *Econometrica* 46, 1 (1978), 33–50.