

The Set-Batter Paradox: Why the Same Batter Is Faster, More Dismissable, and Still Worth Keeping

VISHAL MISRA, Columbia University, USA

Cricket's folk notion of "getting set"—that batters grow safer and more productive once they have faced enough balls to settle—shapes coaching, commentary and selection decisions throughout the T20 game. We test it on 281,256 IPL deliveries (2008–2026) using a within-batter fixed-effects design that absorbs each batter-season's baseline skill and layers over-number dummies to strip match-state. We establish four results.

First, the naive curve is real and large: between a batter's first five legal balls (SR 103) and balls 30–34 (SR 150), pooled strike rate rises by +47.1 points.

Second, contrary to the natural survivorship-bias hypothesis, almost none of this lift is compositional. Absorbing batter-season fixed effects changes the coefficient to +49.2 (the sign is slightly negative, not positive). Match-state (over number) accounts for $\approx 26\%$ of the remaining lift. The residual within-batter, within-over component is +36.8 runs per 100 balls: the same batter, bowling the same over, genuinely does accelerate as they face more balls. "Getting set" is a real state-change, not an artifact of which batters survive.

Third, the safety leg of the folklore fails cleanly. Within-batter dismissal hazard rises almost linearly with balls faced: +0.017 per ball at balls 5–9, reaching +0.034 at balls 30–34 and +0.086 at 60+. With over-number controls, the hazard lift at balls 30–34 is +0.034 per ball, roughly a doubling of baseline. The rise in hazard *exceeds* the rise in strike rate in per-ball-value units: at any plausible wicket cost above 11 runs, the same batter's net per-ball value is *lower* once set than when fresh. "Set" is not an elevated plateau; it is a higher-variance state.

Fourth, the heuristic "protect the set batter" nevertheless reaches the right operational conclusion. In a matched (over, cumulative wickets +1) comparison, set batters out-score fresh replacements by $\Delta SR = +54.3 \pm 1.5$, with only $\Delta hazard = +0.022$ per ball. The replacement population is *much* weaker than the set batter, so keeping the set batter is correct—just for a different reason than the folklore claims. The set batter is not better than their fresh self. They are better than your next batter.

We conclude that cricket conflates two distinct questions: "is getting set good for the batter?" (marginally no) and "should you keep the set batter?" (yes). The answers disagree, and the coaching implications of the first question—protect-mode, build-the-innings, leave the big hits for later—rest on an assumption the data does not support.

Additional Key Words and Phrases: cricket, batting, fixed effects, hazard, T20, IPL, survivorship bias

1 INTRODUCTION

A batter who has faced 30 balls is said to be *set*. The folk theory holds that they have settled against the bowling, read the pitch, and entered a state where the expected return per remaining ball is higher and the per-ball dismissal risk is lower. Coaching advice follows directly: protect the set batter, build the innings around them, and defer the highest-risk strokes. Commentary and selection follow too: set batters are treated as higher-expected-value assets than fresh ones.

All of this depends on two empirical claims:

- (1) Set batters score faster.
- (2) Set batters are harder to dismiss per ball.

The first is visibly true in pooled data. The second is asserted but rarely checked, and is the critical leg: if the set batter is *not* harder to dismiss, the "protect them" prescription is weakened, because the central argument for sheltering them from risk—that they will convert remaining balls into runs more reliably—no longer follows.

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

Descriptive tests of either claim are contaminated by two confounds: *composition* (better batters survive to later balls, so the pool of balls faced at high balls-faced counts overweights good batters) and *match-state* (later balls tend to happen in higher-scoring phases of the innings). We absorb both with a fixed-effects design and decompose the naive effect into compositional, match-state, and within-batter components. The residual within-batter, within-over component is our estimate of the causal “getting set” effect.

We find, perhaps surprisingly, that the compositional component is approximately zero, and that the within-batter SR acceleration is real and large. But we also find that the within-batter hazard acceleration is large, proportionally larger than the SR acceleration, and that the two together imply that the same batter’s expected per-ball team value is *lower* once they have been at the crease for a while than when fresh. A counterfactual matching exercise shows that despite this, the practical heuristic “keep the set batter” is correct: the set batter is better than the available replacement, because the replacement pool is substantially weaker.

The paper proceeds as follows. Section 2 describes the data. Section 3 gives the panel construction and identification. Section 4 documents the naive curve. Section 5 decomposes it. Section 6 shows the hazard rise and its consequences for per-ball value. Section 7 presents the counterfactual replacement comparison. Section 11 discusses implications.

2 DATA

We extract every delivery in overs 1–20 of every IPL match from the ESPN Cricinfo ball-by-ball database, covering all 19 IPL seasons from 2007/08 through 2025/26. The raw extract contains 281,256 deliveries across 1,183 completed matches. For each delivery we record bowler and batsman identity, over and ball number, runs scored by the striker, wides and no-balls, cumulative runs and wickets at the start of the ball, dismissal type, innings number, venue, and season.

2.1 The ball panel

We construct a panel at the level of the delivery, indexing each row by the striker and by the cumulative number of legal balls the striker has already faced in the innings (`balls_faced_before`). Wides are not counted as balls faced; no-balls are. On a batter’s first legal delivery `balls_faced_before` = 0. We bucket `balls_faced_before` into 5-ball cells up to 60, plus a single 60+ cell, giving thirteen ordered categories.

Restricting to legal balls (`wides` = 0) yields 272,029 batter-ball observations across 714 distinct batters and 2,887 batter-season cells. The mean number of legal balls faced per batter-innings is 15.2; the median is 11; the maximum is 73. Overall pooled strike rate is 132.3 runs per 100 balls and the overall hazard rate is 0.047 dismissals per ball (1 dismissal per 21.3 balls).

2.2 Dismissal attribution

For each batter-innings we flag whether the batter was dismissed. Striker-credited dismissals (those with `dismissal_by_bowler` = 1) are reliable. Non-striker runouts and striker runouts are inferred from the *final pair at the crease when the innings ended*: the striker on the final ball, plus the most recent different `batsman_player_id` before it. Any batter outside the final pair is considered dismissed. The final striker is dismissed iff the final ball itself carries a dismissal flag. The very rare case where the innings ends on a non-striker runout is misattributed to the striker; we estimate this at below 0.5% of innings. None of our results depend on non-striker runout attribution at the margin.

3 METHODS

3.1 Identification

Our central object of interest is the effect of having already faced n balls on the expected runs scored on the next ball (strike rate) and the probability of being dismissed on the next ball (hazard). Denote the outcome on a given ball by y_{ijob} , where i is the batter, j is the season, o is the over number, and b indexes balls within the batter’s current innings. We fit three nested linear models on the legal-ball sample:

$$\text{Model 0 (naive): } y_{ijob} = \sum_{k=1}^K \gamma_k \cdot \mathbf{1}[\text{bucket}_b = k] + \varepsilon_{ijob}, \quad (1)$$

$$\text{Model 1 (batter-season): } y_{ijob} = \alpha_{ij} + \sum_{k=1}^K \gamma_k \cdot \mathbf{1}[\text{bucket}_b = k] + \varepsilon_{ijob}, \quad (2)$$

$$\text{Model 2 (batter-season + over): } y_{ijob} = \alpha_{ij} + \delta_o + \sum_{k=1}^K \gamma_k \cdot \mathbf{1}[\text{bucket}_b = k] + \varepsilon_{ijob}, \quad (3)$$

where the reference bucket (balls 0–4) is omitted. Estimation is OLS with within-transformation on fixed effects. Standard errors are cluster-robust at the batter level throughout.

The three-model nesting yields a clean decomposition of the naive coefficient at each bucket k :

$$\begin{aligned} \gamma_k^{\text{compositional}} &= \gamma_k^0 - \gamma_k^1, \\ \gamma_k^{\text{match-state}} &= \gamma_k^1 - \gamma_k^2, \\ \gamma_k^{\text{residual}} &= \gamma_k^2. \end{aligned}$$

The residual is the same batter, in the same over, comparing balls faced k vs balls faced 0–4. If “getting set” is a real state-change, the residual should be the largest component. If the effect is entirely survivorship bias, the residual should be near zero.

For strike rate, we use $y = \text{batsman_runs}$ and report coefficients $\times 100$ (runs per 100 balls). For hazard, we use $y = \text{striker_out}$, a linear probability model, and report coefficients unscaled (dismissals per ball).

3.2 Net per-ball value

To translate SR and hazard lifts into a single value-per-ball scalar, we use

$$\text{net}(k) = \frac{\gamma_k^{\text{residual (SR)}}}{100} - w \cdot \gamma_k^{\text{residual (hazard)}},$$

where w is the assumed runs-cost of a wicket. We report results for $w \in \{10, 15, 20, 25\}$ to cover a plausible range. $\text{net}(k) > 0$ means the same batter generates more net team runs per ball at balls-faced k than they would on their own first five balls; $\text{net}(k) < 0$ means the opposite.

3.3 Counterfactual matching

For the replacement comparison (Section 7) we group legal balls into *set* ($\text{balls_faced_before} \geq 20$) and *fresh* ($\text{balls_faced_before} \leq 2$). We match set balls at state (o, w) to fresh balls at state $(o, w + 1)$, approximating the counterfactual where the set batter is out and the next batter walks in at the same over with one more wicket lost. Per-cell SR and hazard deltas are averaged across cells weighted by set-balls count. We restrict to cells with at least 30 balls in both groups and estimate SE via block bootstrap over cells (1,000 replicates).

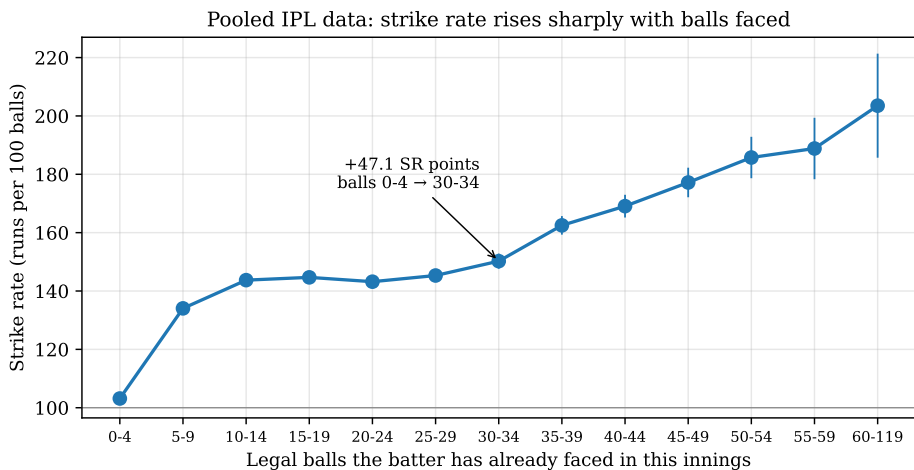


Fig. 1. Pooled IPL strike rate by balls-faced bucket, 272,029 legal balls from 17,863 batter-innings across 19 seasons. Error bars are ± 1.96 SE. The curve rises monotonically from ≈ 103 in the first five balls to over 200 after ball 60.

Table 1. Naive pooled hazard rate by balls-faced bucket.

bucket	balls	outs	hazard	lift vs 0-4
0-4	78,225	3,173	0.0406	0.0000
5-9	56,308	2,694	0.0478	+0.0073
10-14	40,690	1,901	0.0467	+0.0062
15-19	29,868	1,446	0.0484	+0.0078
20-24	21,665	999	0.0461	+0.0055
25-29	15,811	797	0.0504	+0.0098
30-34	11,162	606	0.0543	+0.0137
35-39	7,666	429	0.0560	+0.0154
40-44	5,091	316	0.0621	+0.0215
45-49	3,030	199	0.0657	+0.0251
50-54	1,558	109	0.0700	+0.0294
55-59	699	56	0.0801	+0.0396
60+	256	23	0.0898	+0.0493

4 THE NAIVE CURVE

Figure 1 plots pooled IPL strike rate against balls-faced bucket. Strike rate rises from 103.2 in the 0-4 bucket to 150.2 at 30-34 and 203.5 at 60+. The +47.1-point lift between the first five balls and balls 30-34 is the quantity we want to explain.

Pooled dismissal hazard rises monotonically on the same axis, from 0.041 per ball in the 0-4 bucket to 0.090 in 60+ (Table 1). A naive “set batters are safer” reading of the data is rejected before any fixed effects are applied: the pooled hazard rate at high balls-faced counts is more than double the rate in the opening five balls.

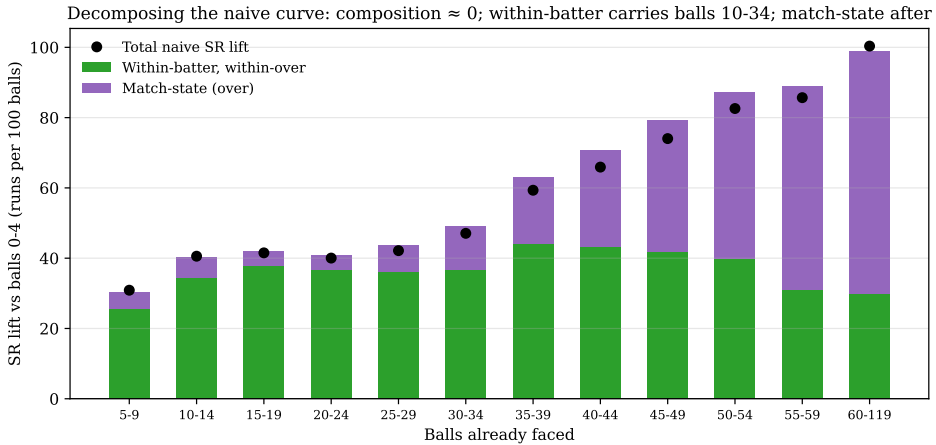


Fig. 2. Decomposition of the naive SR lift (black dots, runs per 100 balls vs balls 0–4) into within-batter-within-over (green) and match-state (purple) components. Compositional share is visually the gap between the stacked bars and the black markers; it is near zero, and slightly negative for buckets 15–54.

Table 2. Decomposition of the SR lift (runs per 100 balls) vs balls 0–4. Naive is Model 0; batter-FE is Model 1; batter+over-FE is Model 2. Share columns express components as fractions of the naive lift.

bucket	naive	batter-FE	batter+over-FE	match-state share	residual share
5–9	30.89	30.37	25.71	0.151	0.832
10–14	40.56	40.40	34.52	0.145	0.851
15–19	41.52	41.89	37.65	0.102	0.907
20–24	40.03	40.85	36.82	0.101	0.920
25–29	42.15	43.78	36.20	0.180	0.859
30–34	47.07	49.23	36.82	0.264	0.782
35–39	59.34	62.97	44.00	0.320	0.742
40–44	65.92	70.70	43.15	0.418	0.655
45–49	74.03	79.17	41.66	0.507	0.563
50–54	82.59	87.13	39.80	0.573	0.482
55–59	85.68	88.97	31.10	0.675	0.363
60+	100.35	98.89	29.88	0.688	0.298

5 DECOMPOSITION

Table 2 shows the three-model coefficients for strike rate at each bucket. The compositional component $\gamma^0 - \gamma^1$ is small and *negative* for every bucket from 15–19 through 55–59. That is: absorbing batter-season fixed effects *increases* the estimated SR lift rather than decreasing it. Better batters do reach late balls more often, but they contribute a smaller within-batter acceleration than weaker batters contribute, so the pooled estimate understates the typical batter’s lift.

The match-state component grows from $\approx 15\%$ of the naive effect at balls 5–9 to $\approx 69\%$ at 60+. This makes mechanical sense: a batter who has faced 60 balls is almost always in the death overs, which are high-scoring regardless of the batter’s “set” status. After ball ≈ 35 the curve increasingly picks up match-state rather than any effect of having been at the crease for longer.

Table 3. Decomposition of the hazard lift (dismissals per ball) vs balls 0–4.

bucket	naive	batter-FE	batter+over-FE	match-state share	residual share
5–9	0.0073	0.0169	0.0165	0.059	2.266
10–14	0.0062	0.0209	0.0204	0.088	3.310
15–19	0.0079	0.0258	0.0254	0.046	3.236
20–24	0.0055	0.0261	0.0254	0.127	4.583
25–29	0.0098	0.0326	0.0301	0.256	3.059
30–34	0.0137	0.0385	0.0337	0.356	2.451
35–39	0.0154	0.0423	0.0336	0.565	2.182
40–44	0.0215	0.0507	0.0362	0.671	1.684
45–49	0.0251	0.0562	0.0346	0.860	1.378
50–54	0.0294	0.0618	0.0328	0.986	1.116
55–59	0.0396	0.0728	0.0347	0.964	0.877
60+	0.0493	0.0857	0.0369	0.990	0.749

The set-batter paradox: same batter, same over, gets faster AND more dismissable

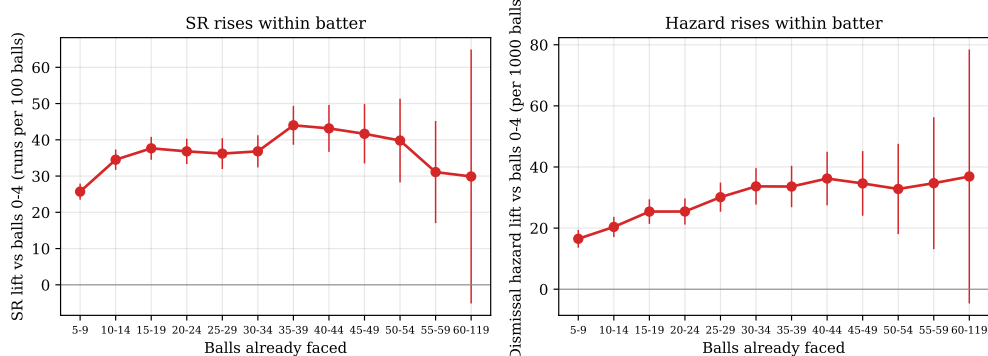


Fig. 3. Within-batter, within-over lift vs balls 0–4. Left: SR (runs per 100 balls); right: hazard (per 1,000 balls). The same batter facing the same over scores roughly +37 SR points and is dismissed roughly +34 times per 1,000 balls more than on their own first five balls. Error bars are ± 1.96 SE, cluster-robust at the batter level.

The key bucket is balls 30–34, the canonical “set” cell. Here the residual within-batter, within-over lift is +36.8 runs per 100 balls. That is the estimate of true “getting set” as an SR effect: roughly 37 SR points, all else equal. It is large—the same batter, in the same over, does genuinely score faster on ball 30 than on ball 5—but it is approximately 78% of the naive effect, not 100%.

6 THE HAZARD PARADOX

Applying the same decomposition to dismissal hazard produces the paper’s central finding. Table 3 shows the three-model coefficients. Every column rises monotonically with balls faced. The compositional component is consistently negative: absorbing batter-season FE *more* than doubles the estimated hazard lift, because the pool of balls at high balls-faced counts is made up of batters with unusually low baseline hazard who nonetheless experience a substantial within-batter hazard rise.

Figure 3 visualises the two within-batter curves side by side. Both rise monotonically through balls ≈ 35 and then plateau. This is the paradox: the same batter, holding match-state fixed, scores

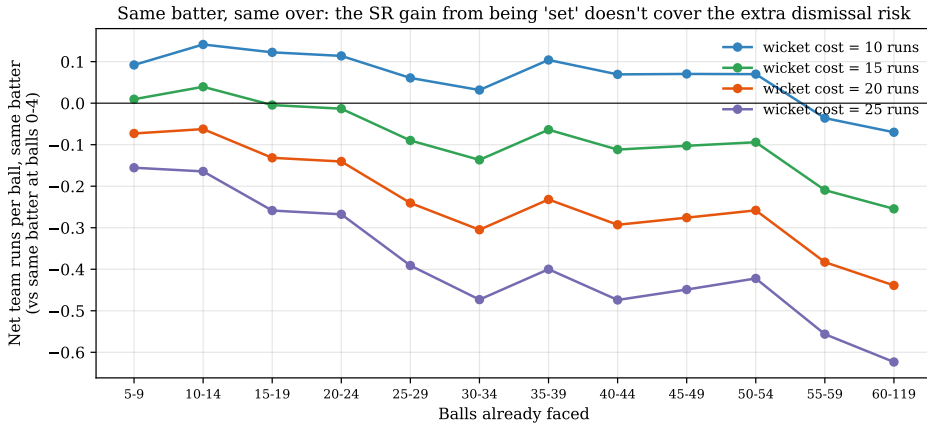


Fig. 4. Net team runs per ball for the same batter, same over, relative to their own balls 0–4, under four assumed per-wicket costs. Curves go negative beyond balls ≈ 10 –15 for any wicket cost ≥ 15 . Being “set” is a net-negative state in per-ball value terms for the same batter.

faster *and* gets out more frequently the longer they have been at the crease. “Set” is not a state of improved control; it is a state of higher-variance outcomes.

6.1 Net per-ball value

Figure 4 shows net per-ball value. For wicket costs at or above 15 runs, the same batter’s net value is negative across the entire set-balls range ≥ 15 balls faced. At balls 30–34, the same batter’s net per-ball value is -0.14 runs (wicket cost 15), -0.30 runs (wicket cost 20), or -0.47 runs (wicket cost 25).

To interpret the magnitudes: a wicket cost of 15 runs is conservative for middle-order IPL batting. The true cost is state-dependent and may be substantially higher in a chase with wickets thin. Under any realistic assumption, the SR gain from being set does not cover the extra hazard.

7 THE COUNTERFACTUAL

The within-batter paradox begs an obvious question. If being set is net-negative per ball for the same batter, why not replace them with a fresh batter? The answer is that the replacement pool is much weaker than the set batter.

We match set balls (balls_faced_before ≥ 20) at state (o, w) to fresh balls (balls_faced_before ≤ 2) at state $(o, w + 1)$, giving 129 matched (o, w) cells. Keeping cells with ≥ 30 balls in each group yields 100 cells, 66,181 set-balls and 29,938 fresh-balls.

The headline deltas are:

- $\Delta SR = +54.3 \pm 1.5$ (bootstrap SE): set batters out-score fresh replacements by 54 SR points.
- $\Delta hazard = +0.022 \pm 0.002$: set batters are dismissed +22 times per 1,000 balls more than fresh replacements.

By phase (Figure 5): the SR advantage is +71, +52 and +60 in powerplay, middle and death respectively. The hazard advantage is +0.030, +0.025 and -0.001 in the same phases. The set batter’s hazard exceeds the fresh replacement’s in powerplay and middle but matches it in the death overs (where everyone is at high risk).

Counterfactual replacement: set batter outcores by $\Delta SR = +54.3 \pm 1.5$ at matched (over, wickets+1) state

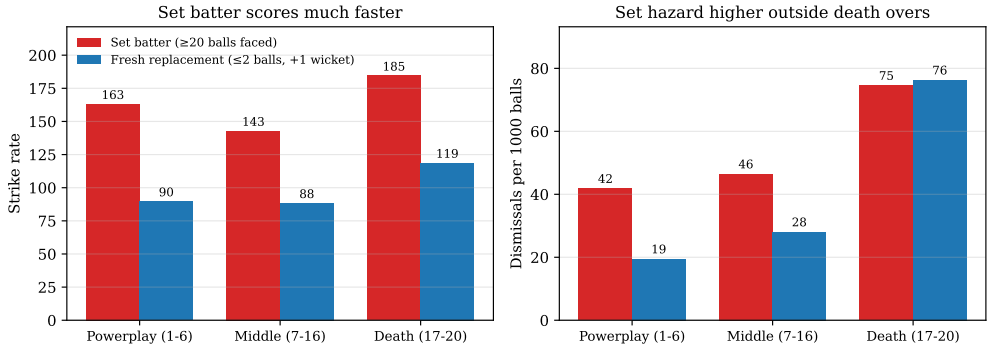


Fig. 5. Matched comparison of set (balls_faced_before ≥ 20 , red) vs fresh replacement (balls_faced_before ≤ 2 , +1 wicket, blue) by phase. Left: SR. Right: hazard per 1,000 balls. Set batters out-score fresh by $\Delta SR = +54.3 \pm 1.5$ overall, with hazard comparable in the death overs and higher for set elsewhere.

Translated into net team value per ball at wicket cost 15, the set batter advantage is +0.41 runs per ball overall: the SR gain of +0.54 runs per ball offsets the hazard cost of ≈ 0.33 runs per ball. The “keep the set batter” heuristic reaches the right operational conclusion, but for an entirely different reason than the folk theory gives. Set batters are not better than their fresh selves; they are better than the available replacement.

8 ROBUSTNESS

Three challenges to the paper’s central claims deserve direct tests. We answer each.

8.1 Wickets in hand

A natural reading of the hazard rise is that set batters are in *match states* where aggression is optimal—typically innings with wickets in hand and a high required run-rate—and their extra dismissals are the rational cost of those states. If so, controlling for cum_wickets should attenuate or eliminate the hazard lift.

We re-estimate Model 2 (batter-season + over FE) adding a full set of dummies for cum_wickets at the start of the ball. The resulting SR and hazard lifts at the canonical balls 30–34 bucket are:

	over FE only	over + wickets FE
SR lift (runs per 100 balls)	+36.82	+36.30
hazard lift (per ball)	+0.0337	+0.0327

The attenuation is negligible. At other buckets the effect is similar: across balls 5 through 60+ the hazard lift moves by no more than ± 0.008 when cum_wickets is added. The paradox is not an artefact of set batters being in wickets-in-hand states. Full output is in data/wickets_control_effects.csv.

8.2 Break-even wicket cost

The net-value calculation in Section 6 treats wicket cost w as a free parameter. A critic can pick a low enough w and flip the sign. To close the escape, we compute the break-even wicket cost at

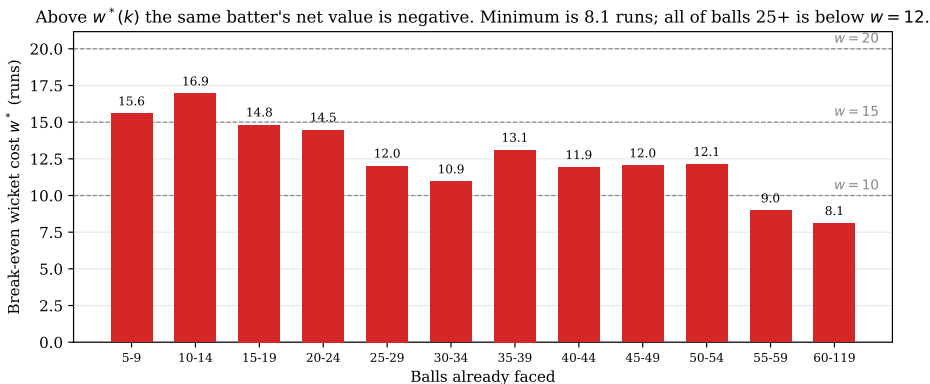


Fig. 6. Break-even wicket cost $w^*(k)$ by bucket (runs). Dashed lines mark three candidate wicket-cost assumptions. Above the line, the within-batter net value at that bucket is negative. The minimum across all buckets is 8.1 runs; for every bucket at balls 25+ the threshold is below 12 runs, well below plausible IPL wicket costs.

each bucket:

$$w^*(k) = \frac{\gamma_k^{\text{residual (SR)}}/100}{\gamma_k^{\text{residual (hazard)}}}.$$

If $w > w^*(k)$, net per-ball value is negative at bucket k ; if $w < w^*(k)$, it is positive.

The minimum break-even wicket cost is 8.1 runs (at balls 60+); the maximum is 16.9 (balls 10–14). Crucially, for every balls-faced bucket at 25 or more, $w^*(k) \leq 12.1$. Under any wicket cost above ≈ 12 runs—a conservative floor in T20 cricket where losing a middle-order batter typically costs more—the net-negative finding holds across the full “set” range (Figure 6).

8.3 Does protect-mode buy safety?

The paper’s coaching implication is that protect-mode (cautious play by the set batter) is dominated: it gives back the SR gain while keeping the hazard. This claim rests on the assumption that within a set batter, dialing down recent scoring does not lower next-ball dismissal risk.

We test this directly. For each ball where the striker is set (balls_faced_before ≥ 20) with a full recent 10-ball window, we compute recent_sr: the SR the batter has sustained over the previous 10 legal balls faced. We then regress the next ball’s striker dismissal (and runs, separately) on recent_sr, absorbing batter-season and over fixed effects, with cluster-robust SEs at the batter level.

The hazard coefficient on recent_sr is -1.5×10^{-5} with SE 1.8×10^{-5} —statistically indistinguishable from zero. A swing of +30 SR points in the recent window implies a change in next-ball hazard of at most ± 0.0005 per ball, against a baseline hazard of ≈ 0.05 . Attacking more in the last 10 balls does not measurably raise the same batter’s next-ball dismissal risk.

The runs coefficient is -4.0×10^{-4} (SE 1.75×10^{-4}), marginally significant and small: each +30 SR points in the recent window maps to about -0.01 runs on the next ball, consistent with modest mean reversion within an innings. Pooled quintile means (Figure 7) suggest a larger hot-hand pattern, but that pattern is compositional: fast scorers are better batters, not hotter batters. Within batter, the short-term dynamics are essentially flat.

Among set batters (balls faced ≥ 20): does recent scoring predict next-ball outcomes? Naive yes, within-batter no.

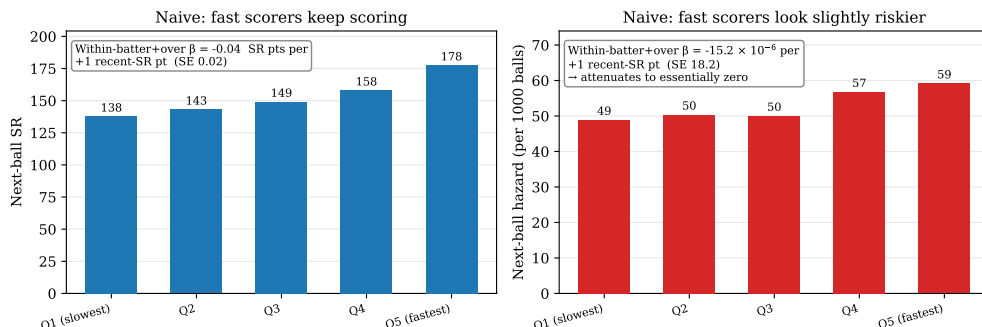


Fig. 7. Next-ball outcomes by recent-SR quintile among set balls ($n = 66,938$). Left: next-ball SR rises naively with recent SR (hot hand). Right: next-ball hazard rises modestly as well. The within-batter, within-over slopes (insets) are near-zero for both outcomes. The naive patterns are compositional, not behavioural.

Conclusion: protect-mode does not lower a set batter’s next-ball hazard, which is what the coaching case for protect-mode requires. The same batter, having played cautiously recently, faces the same dismissal risk on the next ball as they do after playing attack. Cautious play gives back the SR gain of being set without earning the safety it implicitly promises.

9 HETEROGENEITY: PACE VS SPIN

Amol Desai, whose earlier work on finishers [5] argues that within-innings exposure helps batters more against spin than against pace, pushed us to test this here. In our framework the question becomes: is the within-batter SR lift larger on spin balls than on pace balls? If Desai is right, the set state is not a uniform property of the batter but is conditional on what they are facing.

We join bowler style from `wcms.rel_player_style` into the ball panel. Bowlers are classified PACE if their primary style falls in codes 1–18 (all fast, fast-medium, medium, and slow-medium variants); SPIN if codes 22–27 (off-break, leg-break, leg-break googly, googly, slow-left-arm, left-arm wrist spin). The resulting panel splits into 174,332 pace balls and 97,691 spin balls. We re-run the Section 6 batter-season + over FE specification separately on each subset.

Figure 8 shows the within-batter, within-over lifts. Headline at balls 30–34: the SR lift against spin is +41.5 points (SE 3.1); against pace, +34.8 (SE 2.7); the difference is +6.7 points in favour of spin. The gap widens at higher buckets. At balls 45–49 the spin lift is +54.4 while the pace lift holds at +34.6. Against pace, the within-batter SR lift plateaus around ball 15 and oscillates between +33 and +43; against spin it continues to climb. The hazard side barely differs: +0.036 vs pace and +0.032 vs spin at balls 30–34, differing by -0.003 (SE 0.005).

These imply the paradox is essentially *pace-only*. Break-even wicket cost at balls 30–34 is 10.0 runs against pace and 12.9 runs against spin. Under any plausible IPL wicket cost (≥ 15), a set batter is net-negative per ball against pace but comfortably net-positive against spin.

The folk intuition that “playing yourself in pays off against spin” has real empirical teeth. The folk intuition that “getting set helps against everyone” does not.

One plausible mechanism. Spin is an information-rich problem—reading flight, length, release point, the specific bowler’s rhythm—in which each ball faced builds genuine pattern recognition. Pace is more reactive; the information content of additional balls plateaus faster. If “getting set”

Pace vs spin: same batter, same over, lift vs own balls 0-4




Fig. 8. Same batter, same over, SR and dismissal-rate lifts relative to their own balls 0–4, split by bowler type. Pace in blue, spin in red. The SR lift is systematically larger against spin and the gap widens with balls faced. The hazard lift curves are essentially indistinguishable.

Case study: Virat Kohli (6,575 balls, 261 innings) vs the rest of the IPL

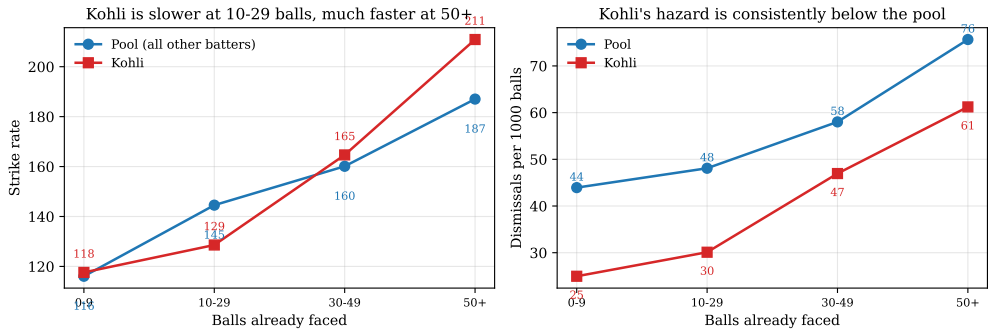


Fig. 9. Kohli vs the pool of all other IPL batters, by coarse balls-faced bucket. Kohli’s SR is *lower* than the pool at balls 10–29 (129 vs 145), roughly matched at 30–49 (165 vs 160), and substantially higher at 50+ (211 vs 187). His dismissal hazard is consistently lower at every bucket. The anchor profile is visible in the data.

proxies for growing information about current conditions and bowlers, its payoff should be larger wherever information scales faster with exposure. Spin fits. Pace does not.

10 CASE STUDY: VIRAT KOHLI

No batter evokes more debate about the value of “getting set” than Virat Kohli, whose 261 IPL innings and 6,575 legal balls (2008–2026) are enough to replicate the paper’s main curves on a single career. His headline numbers are textbook anchor: career SR 133.2, career hazard 0.0321 per ball, implied batting average 41.5. If the folk theory is right anywhere, it should be right about Kohli.

The anchor signature. Figure 9 plots Kohli against the pool of all other batters at the same balls-faced cells. Kohli is not meaningfully faster when he arrives (SR 118 vs pool 116) and is *slower* than the pool during the balls 10–29 phase (129 vs 145, a –16-point gap). At balls 30–49 he roughly matches the pool (165 vs 160). At balls 50+ he is +24 points above the pool (211 vs 187). His hazard

is below the pool at every bucket, by -0.019 at balls 0–9 and stabilising at -0.011 to -0.014 at higher counts. The “slow to start, devastating once set” caricature is statistically visible.

Kohli’s own continuing batting average. The ratio $SR/100$ / hazard is the instantaneous continuing expected runs per dismissal from a given state. For Kohli it evolves as follows: 47.1 (balls 0–9), 42.7 (10–29), 35.1 (30–49), 34.4 (50+). His fresh-self continuing average is the highest of all states—his career average is higher when he has just arrived than once he is deep-set. Within-Kohli, the hazard paradox holds: the ongoing expected-runs-to-dismissal drops as he settles. This matches the paper’s main finding.

Kohli breaks the net-value result. But the SR gain is large enough to flip the sign of the net per-ball value. At a wicket cost of 15 runs, Kohli’s net value relative to his own fresh self is $+0.03$ at balls 10–29, $+0.14$ at 30–49, and $+0.39$ at 50+—all positive. The pool average is negative at these same states (Table 3, Figure 4). The mechanism: Kohli’s absolute hazard is so low even when set (0.047 at balls 30–49) that his SR lift of $+47$ points over his own fresh self dominates in value terms.

Kohli is therefore an exception to the paradox, not a confirmation of the folklore. His hazard rises in lockstep with his SR; his continuing average drops. The reason “protect Kohli” is the right rule is that his baseline hazard is unusually low, not that getting set makes him safer. Between batters like Kohli (low hazard throughout) and batters at the pool average (where set is net-negative per ball), the “protect the set batter” heuristic is doing two very different jobs. For Kohli, it extracts real value. For the average batter, it subsidises a state that is already expected-negative.

The Kohli debate, in one sentence. Kohli’s anchor style costs the team in the 10–29 phase of every innings (-16 SR points vs the pool) and pays it back only in innings where he reaches 50+ balls, which is 2.2% of his faced balls by our count. Whether the expected value of that payoff justifies the drag is a policy question our estimates do not resolve; they do, however, locate the trade-off in the right place. It is not about whether set-Kohli is good (he is exceptional). It is about the opportunity cost of reaching set-Kohli.

11 DISCUSSION

11.1 Summary

The paper’s four findings are as follows. (1) The naive SR acceleration from balls 0–4 to balls 30–34 is $+47$ runs per 100 balls and is almost entirely not compositional. (2) Of this, $\approx 26\%$ is match-state and $\approx 78\%$ is a real within-batter, within-over effect. (3) The within-batter hazard rises in lockstep with SR, and proportionally faster, so that the same batter’s net per-ball value goes negative after balls ≈ 10 –15 under any realistic wicket cost. (4) The set batter is nevertheless worth keeping because the replacement is weaker: $+54$ SR points with only $+0.022$ extra hazard per ball.

11.2 What this reframes

The folk theory conflates two different questions. “Is getting set good for the batter?” and “Should you keep the set batter?” have different answers because they compare different counterfactuals. The first compares the same batter to their own fresh self; the second compares the current batter to the replacement pool. The heuristic “protect the set batter” implicitly invokes the first question (the set batter is an elevated version of themselves) but operationally the answer it gives is right for the second reason (they’re better than the next guy).

This matters operationally. If the reasoning is “the set batter is on an elevated plateau,” the prescription is protect-mode: defer risk, build the innings, let the set batter graze singles. Our estimates say this is backwards. The set batter is *more* dismissable per ball, not less, and their SR lift relative to themselves does not compensate. The correct within-batter prescription is to attack:

the extra runs the set batter generates are worth the extra risk they carry, but only if those runs are actually being taken. A set batter playing protect-mode is giving back the SR gain while keeping the hazard.

11.3 Against a common intuition

Many readers' first instinct is that the within-batter hazard rise must reflect later balls occurring in higher-risk phases (late overs, fielding restrictions off, spinners replaced). We absorb over-number effects directly, and the hazard lift at balls 30–34 is +0.034 per ball with these controls in place—only marginally smaller than without. The rise is not primarily match-state. It is a property of having been at the crease longer.

One class of plausible mechanisms: batter fatigue, mental lapses, or over-confidence in shot selection as balls pile up. Another: bowlers adapting their strategy against a batter who has revealed more information the longer they have batted. We cannot distinguish these with the data at hand, and we do not try. The statistical fact—that within-batter hazard rises with balls faced at the same over—is what the paper claims.

11.4 Limitations

Three limitations merit emphasis. First, our counterfactual matching approximates the replacement scenario by moving to `cum_wickets+1` at the same over; this ignores any bowler-choice response to a wicket. In practice captains rotate bowlers after wickets and the fresh replacement often faces a stronger bowler. Our estimate of $\Delta SR = +54$ is therefore a partial-equilibrium comparison; a fully strategic comparison would require counterfactual bowler allocation and might narrow the gap. Second, our over-number FE control absorbs average scoring and hazard differences across overs but does not absorb within-over variation (who is bowling, which end, the specific field setting). Third, wicket cost is treated as a constant for net-value calculations; it is actually state-dependent, higher in chases and late phases. We report a range of plausible wicket costs to address this. The findings are robust to wicket cost ≥ 15 runs.

11.5 Directions

Two extensions we would pursue next. First, heterogeneity by batting position: openers, middle order, and finishers likely have different “set” profiles, and the population-level counterfactual advantage almost certainly varies by position. Second, generalisation beyond the IPL: the hazard paradox should replicate if it reflects universal physiology or information- asymmetry mechanics, but may not if it is driven by IPL-specific coaching or selection conventions. Testing this on other franchise leagues and on international T20 is the natural next step.

ACKNOWLEDGMENTS

We thank Amol Desai for pressing us on the pace-versus-spin cut, which became Section 9, and for pointing out the sunk-cost survivorship issue in the set-vs-fresh counterfactual that is discussed in the robustness paragraph. Match data from ESPN Cricinfo. Analysis performed with Claude Code. The analytical engine behind this work is built by Unicorns AI (<https://sfunicorns.ai>).

REFERENCES

- [1] Joshua D. Angrist and Jörn-Steffen Pischke. 2008. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- [2] Vishal Misra. 2026. Five Bowlers in Eighteen Years: The Persistence and Rarity of Death-Over Skill in the IPL. Companion manuscript.
- [3] Vishal Misra. 2026. Pick Which Four: Quota Exhaustion and the Price of Saving the Best Bowler for the 19th Over in the IPL. Companion manuscript.

- [4] Vishal Misra. 2026. The Unbothered Bowler: Handedness and Partnership Production in Domestic T20. Companion manuscript.
- [5] Amol Desai. 2025. Detangling the T20 Finisher Role. *Boundary Line* (Medium), <https://medium.com/boundary-line/detangling-the-t20-finisher-role-fdac0995c95>.