The Deterrence Trap:

Reclassifying Rim Protection by Career Success in the Modern NBA

Nathan Humphrey*

Independent Researcher; Annapolis, MD
ORCID: 0009-0004-8512-4105

June 2026

Abstract

A decade of interior-defense analytics, beginning with the Dwight Effect (Goldsberry and Weiss, 2013), established that a rim protector's value lies in suppressing close attempts rather than in blocking them, and the modern public-metric ecosystem now measures block-independent deterrence routinely. What that literature has not asked is longitudinal: which deterrence archetype predicts a successful career? We classify 46 modern NBA centers (419 center-seasons) on two axes – a box axis (BLK/36 × AST/36) and a tracking-deterrence axis (rim-suppression × AST/36) – and relate the resulting quadrants to career outcomes. Quadrant membership discriminates career production (PRA/36) at Kruskal-Wallis $p = 0.003$, surviving multiple-comparison correction; the same test on the block axis is weaker ($p = 0.008$), so sorting centers by measured deterrence rather than by blocks separates careers more sharply than the box score does. Reclassifying centers from the box axis to the deterrence axis moves 30% of the pool (14 of 46) across quadrant boundaries. A secondary, directional finding identifies a failure archetype: across the four deterrence quadrants the low-playmaking swiper posts the lowest mean career-success composite (−0.195) and the floor general the highest (+0.411). The omnibus test on that composite does not reach significance ($p = 0.16$), and the sharpest pairwise contrast – floor general over swiper – is a large effect (Cliff's $\delta = 0.474$) that does not survive family-wise correction; we report the ranking as directional, and note that it is stable across both robustness panels. We state the limits plainly: $n = 46$ careers, tracking-deterrence measured 2019–25 only, and one theory-motivated robustness trim. A forward replication is pre-registered for the 2026–27 season.

1 Lineage and the gap

Basketball's conventional statistics summarize offense far more richly than defense, and the metrics that do exist for defense (blocks, steals, defensive rebounds) persist because they capture salient,

*Contact: mr.nathanhumphrey@gmail.com

easily observed events rather than because they measure defensive value well (Franks et al., 2015). Interior defense is the sharpest case. Goldsberry and Weiss (2013) showed with optical tracking that the value of a rim protector comes from deterrence – opponents declining to attack the rim – not from the block itself; in their analysis the league's blocks leader was not the league's best interior defender, and because roughly 70% of close attempts produce points, free throws, or offensive rebounds, suppressing the frequency of those attempts outweighs marginal changes in their efficiency. The block, in their framing, is a side effect of deterrence, not its mechanism.

Franks et al. (2015) operationalized this with player-tracking matchup estimates, introducing a volume score (how often a defender's assignment even attempts a shot) and a disruption score (how much that defender lowers efficiency when shots occur), and found that centers as a class derive more value from volume reduction than from disruption. The intervening decade has commoditized the measurement of deterrence. Public and proprietary systems – Bruin Sports Analytics (2021), and the family of impact metrics that include Basketball Index's Help Defense and Defensive Playmaking measures – now routinely identify defenders who anchor the rim without accumulating blocks. The existence of the block-shy deterrer is, by 2026, well known.

Two adjacent literatures bound the question from the other side. Player-archetype clustering is by now a saturated genre: unsupervised methods reliably recover the rim-protecting big as a high-block, high-rebound, low-usage cluster, but these descriptions are cross-sectional and say nothing about career trajectory. The one prior link between a big-man archetype and career outcome is offensive – the observation that score-first centers who cannot protect the rim show depressed win-share trajectories (Cheng, 2017) – not a statement about deterrence.

This leaves a specific gap. The mechanism is understood (2013); the measurement is mature (2015–2025); the archetypes are mapped descriptively. No published work asks the longitudinal questions: within the set of rim-protecting centers, which deterrence archetype predicts a successful career, and is any archetype a measurable career trap? That is the question this paper answers.

2 The variance-coupling bridge

The study began with an observation from the author's ongoing work on the statistical structure of basketball projections – an observation that this paper treats as motivation rather than as an established result, and which will be developed rigorously elsewhere. In that work, the residual variance of blocks for centers appears to be systematically elevated: after conditioning on prior-season blocks, position, age, and team, centers carry markedly wider block-residual variance than the rest of the league. The pattern recurred across every league and projection method examined – the NBA, the WNBA, and NCAA Division I men's and women's, under both a career-mean and a hierarchical negative-binomial method – with variance ratios falling in a tight band of roughly 1.3 to 2.0. We report these figures here only as the spur for the present study, not as a vetted cross-league claim; the full analysis is the subject of separate work in progress.

The reason the observation pointed toward this paper connects directly to Goldsberry and Weiss

(2013). If a center's block count is a noisy proxy for a latent rim-deterrence skill, then center block residuals should inherit both the variance of the latent skill and the variance of the proxy noise – which is exactly the elevated, replicated ratio the projection work kept surfacing. The 2013 paper diagnosed that proxy relationship qualitatively; the variance pattern suggested its magnitude; and that, in turn, raised the question this paper actually answers – whether the deterrence skill the blocks imperfectly track predicts career outcomes. We draw one scope line before proceeding: the variance observation spans several leagues, but the archetype-to-success map developed below depends on NBA tracking data that exists only from 2019 onward and is therefore a modern-NBA result. The two should not be conflated.

3 Data and methods

3.1 Pool

The universe is NBA player-seasons from 1998–99 through 2025–26. Position is assigned by an inclusive listed-position classifier (any "Center" designation maps to center), matching the protocol used in the variance work so the two analyses share a definition. We require at least three career center-seasons, and at least 5,000 career minutes at center. The resulting pool is 46 centers and 419 center-seasons. We disclose the era seam honestly: the reachable data lake is effectively a post-2010-debut pool, so prime-era anchors such as Olajuwon, Mutombo, Robinson, and O'Neal are absent, and the era split below contrasts a 2010–13 transition window (83 seasons) with a 2014–25 pace-and-space window (336 seasons) rather than spanning the shot-blocking heyday.

3.2 Deterrence metric

Tracking deterrence is computed from NBA Stats closest-defender rim data for 2019–20 through 2024–25. For each center we take the field-goal percentage allowed within six feet when that player is the closest defender and subtract it from the league's normal-shooting baseline for those same attempts; the difference is a rim-suppression rate, positive when the defender holds opponents below baseline. Across the pool the suppression rate has mean 0.051 and median 0.055 (standard deviation 0.041), ranging from a small negative value for a late-career, low-sample season to +0.117 for the pool's strongest rim deterrent.

3.3 Two axes and four quadrants

We define a box axis from BLK/36 and AST/36, and a deterrence axis from rim-suppression and AST/36, each z-scored within the pool. Splitting each axis at zero yields four quadrants with hand-coded labels: high suppression/low assists is a **swiper**; high suppression/high assists is a **two-way unicorn**; low suppression/high assists is a **floor general**; and low suppression/low assists is a **dead zone**. The same labels apply on both axes, which is what makes the reclassification in Section 6 legible. The assist dimension is shared across both axes deliberately: it is the

playmaking-and-connective-passing signal that, as the results below show, does most of the work in separating successful careers from unsuccessful ones, while the vertical axis swaps blocks for measured deterrence.

3.4 Success composite

Career success is an equal-weight average of four z-scored components: career minutes, career playoff minutes, points-plus-rebounds-plus-assists per 36 minutes (PRA/36), and true-shooting percentage. The composite is deliberately simple; a robustness check on weighted composites is noted as a follow-up. Throughout, we keep two success criteria distinct: PRA/36 alone (the production rate, which carries the primary result) and the four-part composite (which carries the secondary archetype ranking). They behave differently under test, and we report each at its own strength rather than letting the stronger lend significance to the weaker.

3.5 Pre-specified analytic choices and multiple-comparison hygiene

Panel A (all 46 centers) is primary. Panel B is a theory-motivated robustness panel that removes four players – Howard, Jordan, Drummond, and Chandler – whose careers were built primarily before 2019 and whose tracking-window deterrence therefore reflects decline-phase defense rather than prime. The trim is justified by data-window-versus-career-window misalignment, not by effect maximization, and both panels are reported. We flag this explicitly as a researcher degree of freedom. Across the study we run 42 hypothesis tests in six phases; we correct by family (Bonferroni or equivalent), lead with results that survive correction, and demote the rest to directional. Table 1 records the accounting.

Table 1. Test accounting across the analysis phases. The single claim surviving family-wise correction is the PRA/36 deterrence-quadrant discrimination; the floor-general-over-swiper contrast is the strongest of the success-composite pairwise tests and is reported as directional.

Test family	n	Sig. @ .05	Bonferroni α	Survives
Success axis × scheme (KW)	14	2	0.0036	1 — PRA/36 × deterrence ($p = .003$)
Within-archetype Spearman	16	1	0.0031	0
Pairwise vs. dead zone (Panel A)	3	0	0.0167	0
Full pairwise (archetypes, Panel A)	6	1	0.0083	0 — floor general > swiper ($p = .041$, $\delta = .474$)
BLK–deterrence by year	6	6	0.0083	5 of 6

4 Primary results

P1. Career production is discriminated by deterrence quadrant. A Kruskal-Wallis test of PRA/36 across the four deterrence quadrants returns H $= 14.27, p = 0.003$, which survives Bonferroni correction within the 14-test family of success-axis-by-scheme comparisons ($\alpha = 0.0036$). The identical test run on the block-axis quadrants is weaker (H $= 11.93, p = 0.008$): it also clears correction but discriminates career production about 20% less sharply. Figure 1 shows the contrast. The substantive point is the axis swap itself – sorting centers by deterrence rather than by blocks produces tighter separation on the career-production criterion that the prior literature never tested. The deterrence axis does not merely relabel the same centers; it relabels them in a way that tracks who actually had the better career, which is the property a classification scheme should have and which the block axis has more weakly ($H = 14.27$ vs. block 11.93).

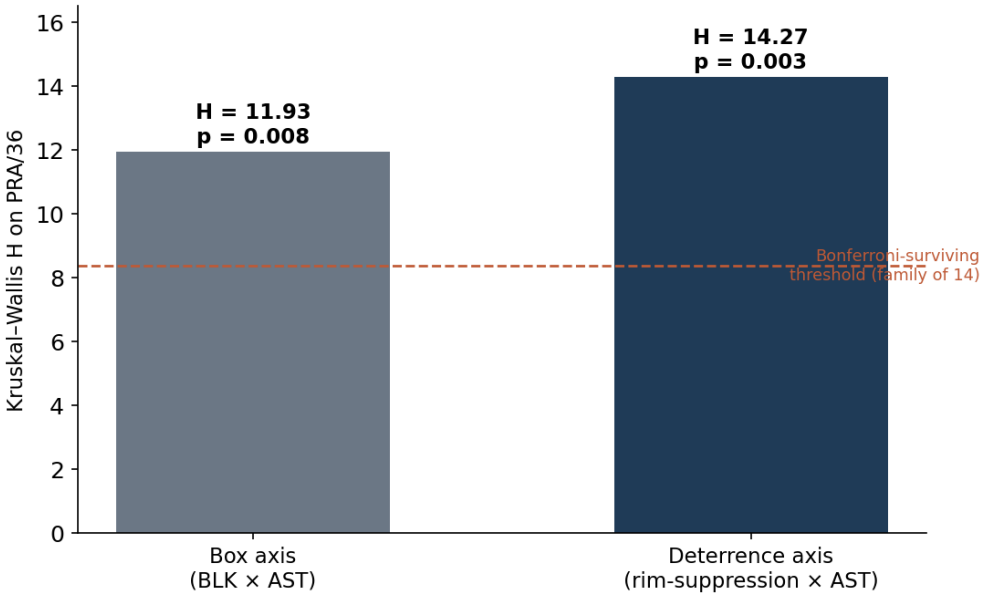

Figure 1. The axis swap. Kruskal–Wallis H for PRA/36 across the four quadrants, computed on the box axis (BLK × AST) versus the deterrence axis (rim-suppression × AST). Both clear family-wise correction, but the deterrence axis discriminates career production about 20% more sharply ($H = 14.27$ vs. 11.93). Sorting centers by deterrence rather than blocks is the move the prior literature never tested against a career criterion.

P2. The block-to-deterrence coupling strengthens over time. Within each tracking season we correlate BLK/36 with rim-suppression across the pool. The Spearman coefficient rises from 0.43 in 2019–20 to a peak of 0.64 in 2022-23 and holds near 0.56–0.63 thereafter; five of the six seasons survive Bonferroni correction ($\alpha = 0.0083$), the lone exception being the smallest-sample first season (Figure 2). By block quartile, mean rim-suppression climbs monotonically from 0.024 in the lowest quartile to 0.087 in the highest.

This inverts our pre-specified expectation. We expected the three-point-attempt explosion to

Table 2. Year-by-year rank correlation between BLK/36 and rim-suppression. The coupling strengthens across the window.

Season	Spearman r	p	Survives $\alpha = .0083$
2019-20	0.432	0.015	no
2020-21	0.554	0.003	yes
2021-22	0.532	0.004	yes
2022-23	0.644	0.0005	yes (peak)
2023-24	0.631	0.001	yes
2024-25	0.562	0.007	yes

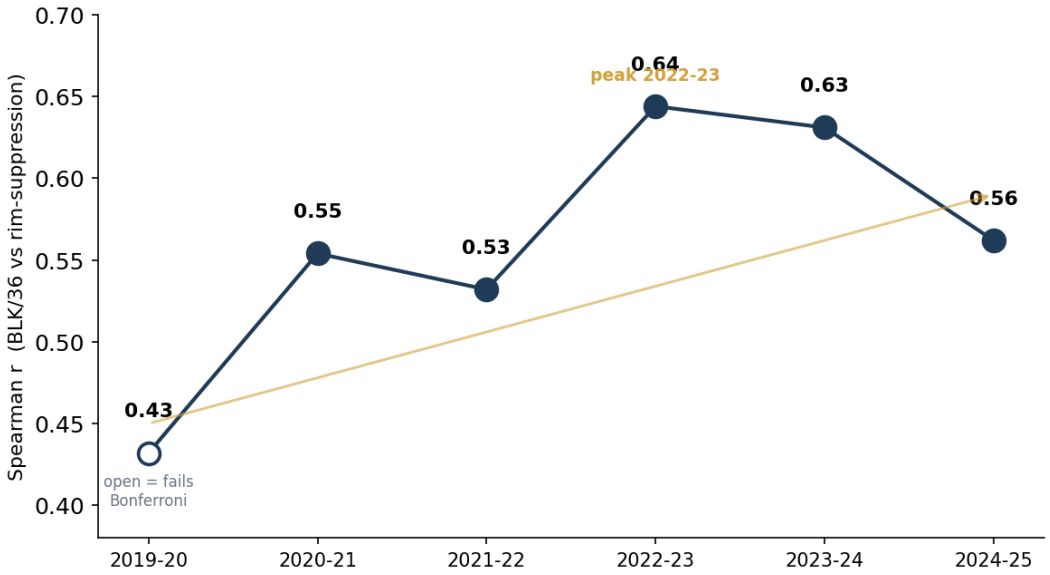

Figure 2. The block-to-deterrence coupling strengthens across the tracking window. Year-by-year Spearman correlation between BLK/36 and rim-suppression, rising from 0.43 (2019-20) to a 0.64 peak (2022-23) and holding near 0.56–0.63 thereafter. Filled points survive Bonferroni correction ($\alpha = 0.0083$); the open 2019-20 point is the smallest-sample season.

make blocks a noisier deterrence proxy over time; the data show the opposite, and we report the prediction as wrong. The likely mechanism is selection: as marginal rim attempts migrate to the perimeter, the close attempts that remain are more selective and elite-driven, which makes a block a cleaner signal of genuine rim deterrence in 2024 than it was in 2013. This is the era-honest correction to the Dwight Effect – the 2013 diagnosis was right for its time, and the proxy noise has since shrunk. It is worth being precise about what this does and does not soften: blocks are a better deterrence proxy now than they were, but `P1` still shows the deterrence axis out-discriminating the block axis on careers, so the proxy has improved without becoming a substitute.

5 The failure archetype

The primary result establishes that deterrence quadrants discriminate production. The natural follow-up is whether they discriminate the broader career-success composite, and in particular whether any single archetype is a measurable trap. Here the evidence is real but weaker than the production result — the omnibus does not reach significance, but the archetype ranking is stable across both robustness panels.

Across the four deterrence quadrants, the career-success composite ranks cleanly even though the omnibus test does not clear significance. The floor general is the highest-mean archetype ($+0.411, n = 10$), the dead zone sits near zero ($+0.025, n = 10$), the two-way unicorn is mildly negative ($-0.093, n = 7$), and the swiper is the lowest-mean archetype ($-0.195, n = 19$). The Kruskal-Wallis omnibus on the composite across quadrants returns $\mathrm{H} = 5.16, p = 0.16$ – not significant – so the ranking is descriptive, and we treat it as directional rather than confirmed. The single pairwise contrast that reaches single-test significance is floor general over swiper (Cliff's $\delta = 0.474, p = 0.041$), a large effect by the Romano-Olkin convention (whose threshold for "large" is 0.474) that does not survive the family-wise correction. Table 3 gives the quadrant means.

Table 3. Career-success composite by deterrence quadrant (Panel A, full pool). The swiper is the lowest-mean archetype and the floor general the highest; the omnibus Kruskal–Wallis on the composite is not significant ($p = 0.16$), and the floor-general-over-swiper pairwise ($\delta = 0.474$, $p = 0.041$) does not survive family correction. The ranking is reported as directional.

Deterrence quadrant	Mean success composite	n	Reading
Floor general	+0.411	10	Highest — the connective, high-assist big
Dead zone	+0.025	10	Middle of pack
Two-way unicorn	−0.093	7	Mildly negative
Swiper	**−0.195**	**19**	**Lowest — the failure archetype**

The basketball reading survives the statistical caution, and it is the inversion of the intuition that low-event centers are the weak ones. The worst career outcomes do not belong to the dead zone – the low-suppression, low-assist quadrant one might expect to be the bottom – but to the swiper, the center who suppresses or blocks without the connective playmaking that the floor general supplies. Read alongside the primary result, the discriminating dimension is less the rim-protection axis than the assist axis: floor generals, who pair adequate rim impact with playmaking, post the best careers, while swipers, whose game is rim impact without playmaking, post the worst. This is consistent with the Dwight Effect's qualitative warning at one remove – blocks alone, even backed by genuine suppression, do not make a career; the centers who last are the ones who also touch the offense.

A timing-confound robustness panel (Panel B) removes the four players whose production predates the tracking window. Under the trim the dead-zone mean falls to -0.296 ($n = 6$) and the floor-general-over-dead-zone contrast sharpens to $\delta = -0.567$ ($p = 0.073$), still short of significance at this sample size. The swiper and floor-general quadrants are unaffected by the trim, so the

central ranking – floor general highest, swiper lowest – is stable across both panels. We do not over-read Panel B; we report it because the trim is a researcher degree of freedom and both panels should be visible.

We are explicit about what this section does and does not claim. It does not claim a significant omnibus discrimination of the success composite; that test fails ($p = 0.16$). It does claim a stable, directional archetype ranking in which the swiper is the lowest-success quadrant and the floor general the highest, supported by the one large-effect pairwise contrast between them. The strength of the finding is the consistency of the ranking and the size of the floor-general-over-swiper effect, not a p-value, and we lean on the pre-registered forward replication in Section 9 to convert it from directional to confirmed or to fail it.

6 The reclassification

The most legible result is purely descriptive and therefore carries no inferential risk. Moving centers from the box axis to the deterrence axis reclassifies 14 of 46 (30%) across quadrant boundaries. Table 4 is the confusion matrix and Figure 3 renders it.

Table 4. Box-axis to deterrence-axis reclassification. Off-diagonal cells (14 centers) are the reclassified players.

Box \ deterrence	Dead zone	Floor general	Swiper	Two-way unicorn
Dead zone	7	0	8	0
Floor general	0	10	0	3
Swiper	3	0	11	0
Two-way unicorn	0	0	0	4

Three groups move. Eight box-dead-zone centers become deterrence-swipers – the block-shy deterrers the literature anticipated qualitatively (Tristan Thompson, Zubac, Jarrett Allen, Aldridge, and others): centers whose modest block totals understated genuine rim suppression. Three box-floor-generals become two-way unicorns (Wendell Carter Jr., Hartenstein, Naz Reid). Three box-swipers fall to the deterrence dead zone (Jordan, Drummond, Howard) – and these three are precisely the Panel B trim cases, whose prime deterrence predates the tracking window, so the fall is partly a measurement-window artifact and is flagged as such rather than presented as a verdict on their careers. The reclassification is where the deterrence axis earns its keep descriptively: nearly a third of the pool sits in a different archetype once blocks are replaced by measured suppression, and the largest single move is exactly the block-shy deterrer the 2013 framework predicted would exist.

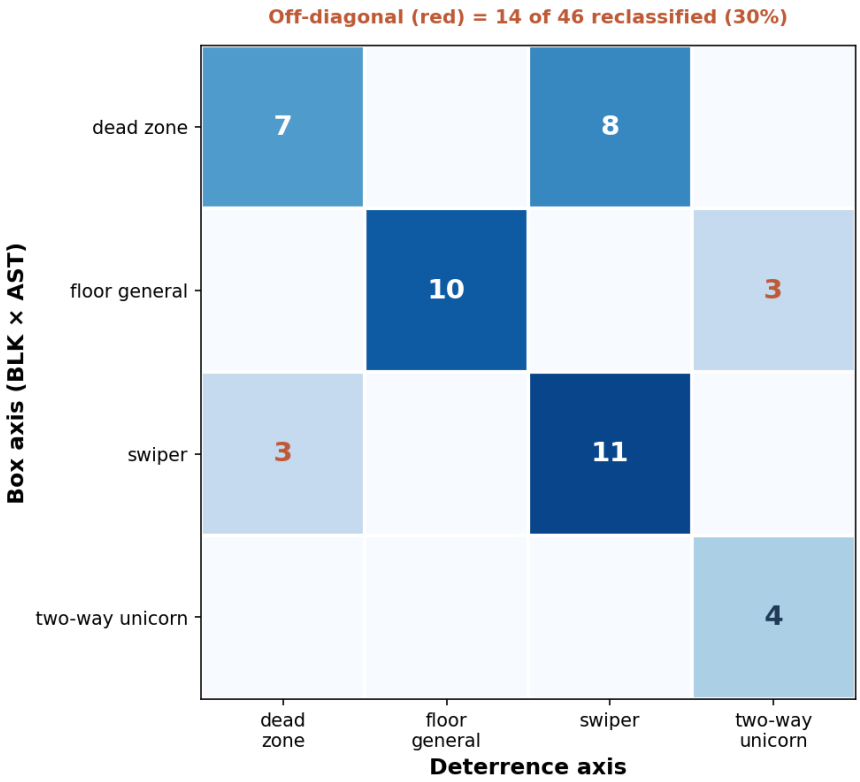

Figure 3. Reclassification from the box axis to the deterrence axis. Diagonal cells (blue) are centers whose quadrant is unchanged; off-diagonal cells (red) are the 14 of 46 (30%) reclassified. The two largest moves are eight box-dead-zone centers becoming deterrence-swipers (the block-shy deterrers the literature anticipated) and three box-swipers falling to the deterrence dead zone (the timing-confound cases handled in Panel B).

7 Era and scope

The floor-general center is the archetype that expanded most in the pace-and-space era: from 8 unique players (25 player-seasons) in the 2010–13 transition window to 25 unique players (98 player-seasons) in 2014–25. All four archetypes exist in both eras. The strengthening block-to-deterrence coupling documented in Section 4 lives here as the era story, which inverts the noisier-proxy prediction we pre-specified. We restate the scope limit: the archetype-to-success map is a modern-NBA result, dependent on 2019–25 tracking data, and does not generalize to the WNBA, the NCAA, or the pre-2019 NBA, even though the variance pattern that motivated the look replicates across those leagues; the career result here does not. The expansion of the floor general is also the era reading of Section 5: the archetype the league has been adding is the highest-success one, while the swiper it has been able to source cheaply off the bench is the lowest.

8 Limitations

We state the constraints rather than soften them. The pool is 46 careers, which is power-limited; the success-composite omnibus does not reach significance for exactly this reason, and we have leaned on the stability of the archetype ranking and on effect sizes accordingly. Tracking deterrence exists only for 2019–25, so for any career built largely before 2019 the metric captures decline-phase defense; Panel B handles the four most affected players explicitly, but the same confound is partially present for others with substantial pre-2019 minutes. The inclusive listed-position classifier misses players who logged heavy center minutes under a forward listing. The success composite is unweighted. The four-player Panel B trim is a researcher degree of freedom, defended on data-window logic and reported in both panels. Finally, the analysis sequence was specified in advance but not filed as a public pre-registration – which is the motivation for Section 9.

9 Pre-registered forward replication

To convert the absence of a prior pre-registration into a falsifiable forward commitment, we register the following before the 2026–27 NBA season generates any tracking data. Tier-1 hypothesis: in the 2026–27 center pool, PRA/36 will be discriminated by deterrence quadrant at $p < 0.05$ under the same Kruskal-Wallis test, pool-construction rules, and deterrence definition used here. Tier-2 hypothesis: the deterrence-quadrant success-composite ranking will again place the floor general above the swiper. Falsification: a null at $p \geq 0.05$ on Tier-1 in 2026–27 fails the primary replication and walks the production claim back to a single-window result; a reversal or collapse of the Tier-2 ranking fails the directional archetype claim. The registration is to be timestamped before the season opener.

References

Bruin Sports Analytics (2021). Defensive deterrence I: Quantifying defenders' off-ball impact at the rim and beyond. `https://www.bruinsportsanalytics.com/post/defensive_deterrence_i`. Public-data deterrence metric on a 63-player sample spanning strong and weak rim protectors; labeled Part I of a series.

Cheng, A. (2017). Using machine learning to find the 8 types of players in the NBA. Fastbreak Data, Medium. Representative of the player-archetype clustering genre; notes depressed win-share trajectories for score-first centers who cannot protect the rim.

Franks, A., Miller, A., Bornn, L., and Goldsberry, K. (2015). Characterizing the spatial structure of defensive skill in professional basketball. *The Annals of Applied Statistics*, 9(1):94–121. Peer-reviewed companion to the Counterpoints conference paper; Bayesian hierarchical / HMM / NMF treatment of defensive matchups.

Goldsberry, K. and Weiss, E. (2013). The Dwight effect: A new ensemble of interior defense analytics for the NBA. In *Proceedings of the 2013 MIT Sloan Sports Analytics Conference*, Boston, MA. Establishes that rim-protection value derives from deterrence (suppressed close attempts) rather than blocks; the block is a side effect of the mechanism, not the mechanism.