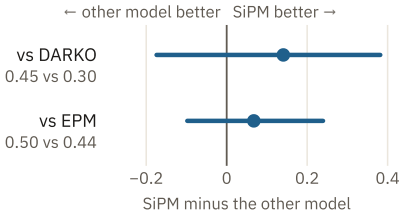
Running The Bulls, Ahead of the Herd

The NBA is driven by stars. How do the Chicago Bulls find the next star and put the best supporting cast around him?

Winning in the NBA is a constant battle to scrap for every edge you can possibly find. This is traditionally thought of as the domain of analytics, which is often seen as being in tension with the eye test. **In fact, analytics and the eye test are complementary.**

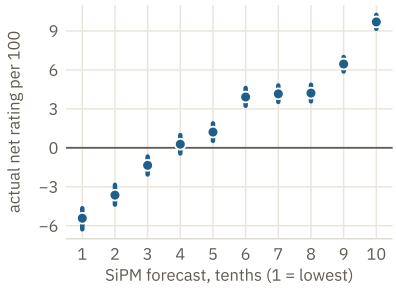
An NBA analytics department must have an impact model that blends lineup ratings and box-score data. **My model is on par with the leading public models.**

SiPM, my model, forecasts as accurately as DARKO and EPM



Both intervals cross zero: the models are level.

The higher SiPM rates a lineup, the better it plays



14,274 lineups from seasons the model never saw.

Analytics work is only as good as the underlying data.

Even the official sources contain errors.

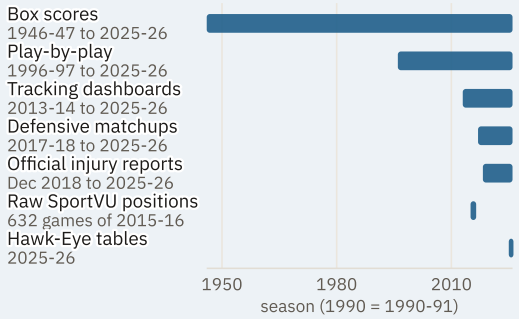
51 errors found in the providers' own data

37 corrected

6 kept out of the analysis

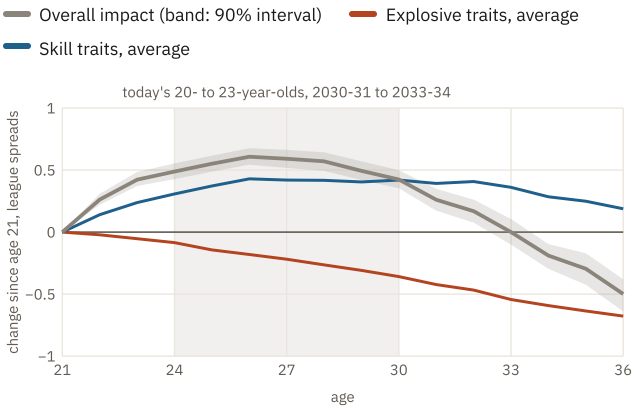
8 flagged where used

Seasons covered, by source

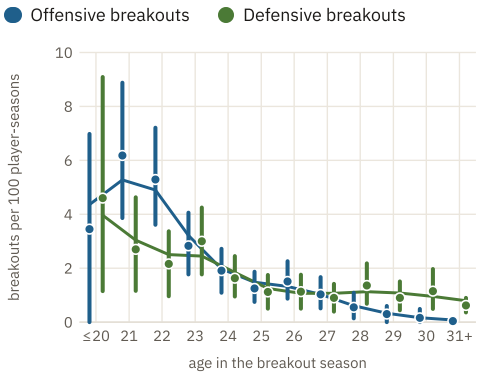


Every error, by source

The question is what you do from there. I want to model trajectories and test questions. **Everything works backwards from the end goal: winning NBA championships.**



Shaded: the ages today's 20- to 23-year-olds reach in 2030-34.

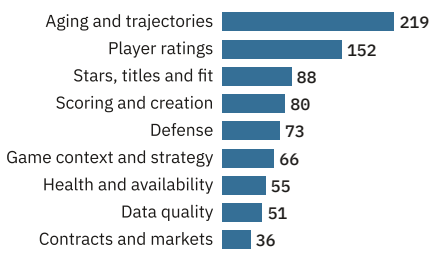


The NBA has singular star talents. Looking at the best players today, each has a unique game:

Wemby Jokic SGA Giannis Kawhi LeBron Steph KD Luka Ant

The name came from my hypothesis that was found to be unproven: that identifying the next singular star could be done by data alone. Data narrows the list, shows whose game is still growing, and finds who fits around him. **Singularity is identified by watching the game.**

820 questions tested, by topic



1 Trajectories: how each part of a player's game grows and fades

A rebuilding team invests in primes. This figure shows which traits a roster gains and loses as its players age.

FIGURE 1

Most explosive box-score traits peak at 21 or younger and fade first; passing peaks at 26 and ball security at 25, while three-point and mid-range shooting hold into the 30s



A roster's explosive strengths are the first to go. Rim shots, dunks, transition shots, blocks and offensive rebounds peak at 21 or younger and slide from there; passing peaks at 26, free-throw shooting at 28 and ball security at 25; three-point and mid-range shooting hold into the 30s; overall impact peaks at 26 and holds to 28 before it falls. Guards age best: their overall impact peaks at 28, against 26 for wings and bigs, while bigs lose shooting efficiency and ball security first.

Line: typical change since age 21, in league spreads (up = better, so ball security is turnovers turned upside down). Dot: peak age. Each player's change from one season to the next counts once; players likely to leave the league are weighted up, using only what was known that season. A line stops once fewer than 30 player-seasons of that position remain, which is why bigs' three-point line ends early. Late-career ages carry the widest uncertainty. Box-score defense is the weakest measure here: age is one of its model's inputs, so its curve partly echoes that age term.

2 Findings

Basketball minds are constantly asking how to play the game better. Those questions are testable, and testing them shows how well the eye test and analytics go together.

Findings at a glance

- **Trajectories.** **Explosive box-score traits peak at 21 or younger** and fade first. **Passing peaks at 26** and **ball security peaks at 25**. **Three-point and mid-range shooting hold into the 30s. Guards age best;** bigs lose shooting efficiency and ball security first. Defensive breakouts come about two years after offensive ones.
- **Usage and efficiency.** When a teammate's absence hands a player a bigger role, his true shooting falls about 0.5 points per usage point. **The drop comes mostly from losing the teammate who set him up, not from the extra shots.**
- **Fit.** Spacers and interior bigs add value next to a star creator. **Two star creators together are worth about 0.9 per 100 less than their ratings add up to.**
- **Threes versus twos.** In the half court, threes and twos are now worth the same once offensive rebounds and the opponent's next trip are counted: NBA teams have found the balance. Shots at the rim beat a three. Midrange twos do not. **What I want to track is how the balance between threes and twos shifts, so the Bulls can lead the pack on the next play-style revolution.**
- **What data cannot see.** **On defense, award votes add signal the numbers miss.** A model built on public data captures about half of the signal in a player's next three seasons.
- **Coaching decisions.** Leaving a starter in after an early second foul is worth about a third of a point of win probability. **Coaches give up about 0.3 wins a season, mostly by sitting starters in foul trouble.**
- **Strategy.** **Crashing the offensive glass costs almost nothing in transition.** An isolation-heavy offense shows no measurable playoff penalty, and playing faster with the ball does not buy offense.
- **How it was tested.** 820 questions tested as statistical hypotheses.

FIGURE 2

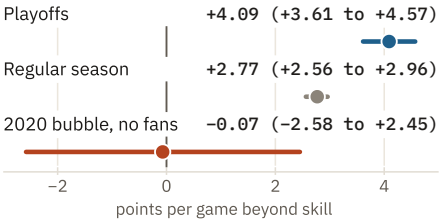
I thought to test it: fifteen common beliefs, measured

Many basketball truisms hold up; what's notable is their size. A coaching staff can prepare with those sizes.

Dot: estimate. Bar: 90% interval (95% for playing hurt). No bar: the test reported a single estimate or a spread.

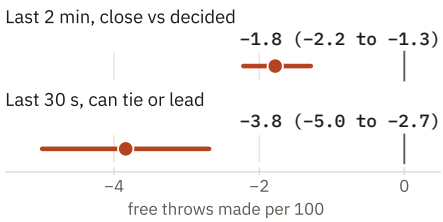
Playoff home court

Worth about 4 points a game beyond skill, 1.5 times the regular season's; it vanished in the 2020 bubble.



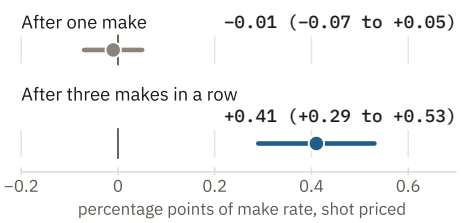
Free throws under pressure

Players make fewer free throws when the game is on the line.



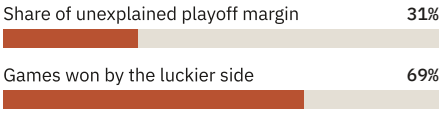
The hot hand

Real but small once shot difficulty is priced: hot players take harder shots.



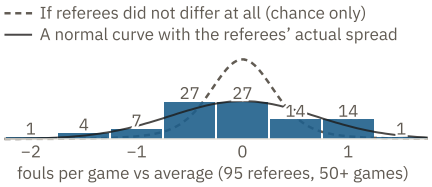
Luck in playoff surprises

Three-point luck explains about a third of what regular-season margins miss, though a real playoff shooting skill would count as luck here.



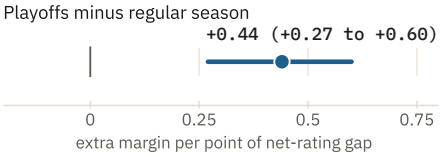
Referees

Referees differ in how many fouls they call, and consistently so. A crew of three differs by about 1.2 fouls a game (one SD). Each referee's effect is shrunk toward average by how noisy it is.



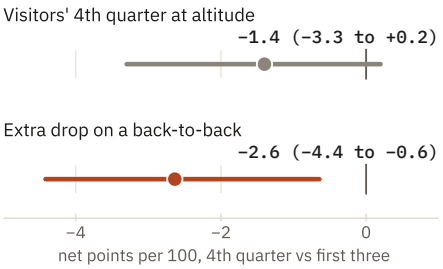
Playoffs widen quality gaps

Between playoff teams, a point of net-rating gap buys 1.19 points of playoff margin against 0.75 in the regular season.



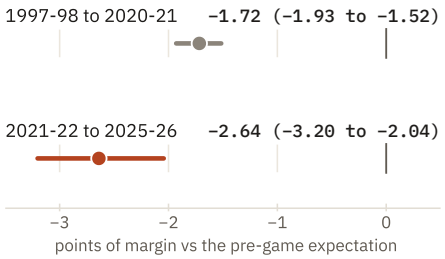
Altitude in Denver and Utah

Their home edges are the league's largest. Visitors' fourth-quarter fade is unclear on its own but shows on the second night of a back-to-back.



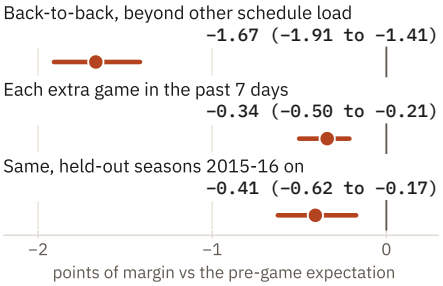
Back-to-backs cost more now

A back-to-back costs a team about 2.6 points of margin since 2021-22, up from 1.7.



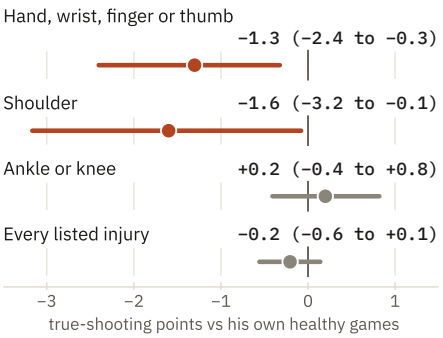
Schedule density

Tough stretches wear teams down beyond back-to-backs: each extra game in the past week costs about 0.34 points of margin, about a fifth of a back-to-back, and it shows up on defense.



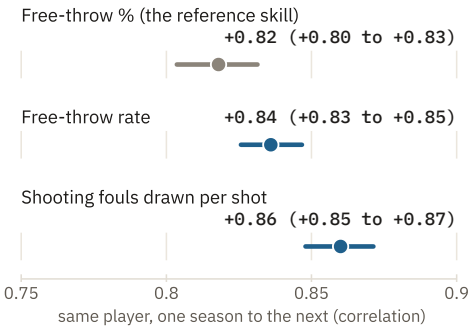
Playing hurt

Hand and shoulder injuries cut true shooting (about 1.3 and 1.6 points); ankle and knee injuries, and the average listing, do not move it.



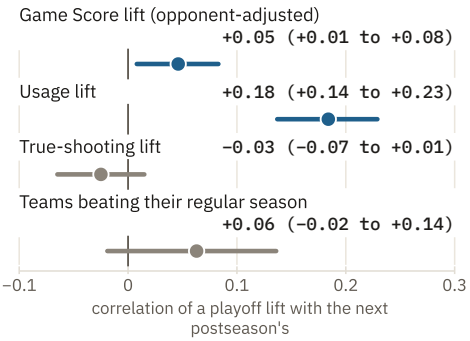
Drawing fouls is a skill

Season to season, how often a player draws fouls is as stable as his free-throw percentage.



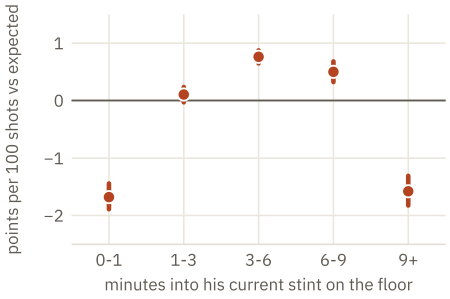
Playoff risers

Playoff stars are mostly regular-season stars (correlation 0.83). Some players repeatably take on more scoring in the playoffs and hold their production better than others, but the edge is small, it is not in shooting efficiency, and one postseason's lift is about 4% player and 96% noise.



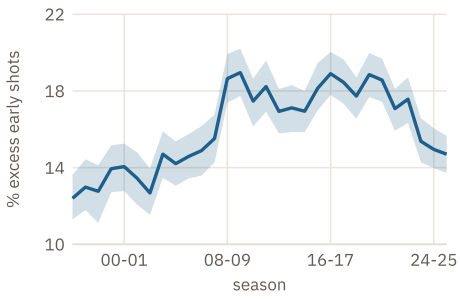
Starters fade after about nine minutes

Players who average 20+ minutes make shots worth 2.3 points per 100 fewer at nine or more straight minutes on the floor than at three to six (2.0 to 2.7).



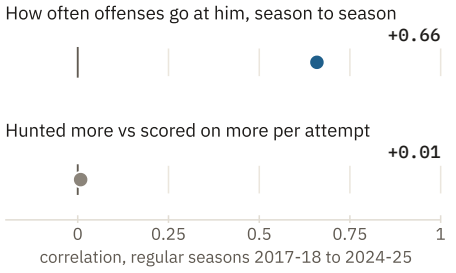
The two-for-one has faded

Quarter-end two-for-ones rose through the 2000s, jumped in 2008-09 and stayed high until 2022-23; since then teams go for them less. Why is not yet tested.



Hunted defenders

Offenses go after the same defenders season after season (correlation 0.66), but being hunted does not mean being scored on efficiently (correlation 0.01).



3 What I'm working on further

An analytics department can give the decision makers in the front office **invaluable tools**. Here are a few, each marked built or proposed.

PROPOSED	**Player trajectory profiles.** Each player's forecast for next season and beyond, split into shooting, creation, rim pressure, defense, minutes and value. Staff could set a custom path for one part (say a rookie's shooting improves faster than the model expects) and see what it does to the rest and to his long-term value.
BUILT	**Cap and rules engine.** The current CBA as code, with the league's announced 2026-27 cap, tax and apron figures and every team's cap position by season. It passes 103 of 113 cap-expert edge cases, from the stretch provision and apron trade matching to sign-and-trade hard caps. The other 10 are listed as known gaps.
PROPOSED	**Trade-deadline scenario simulator.** On that engine: forecast each team's likely moves and run thousands of league-wide deadline scenarios, updated as moves happen.
PROPOSED	**Trade fit simulator.** Score any trade by fit with the current core: switchable defenders, and spacers and interior bigs next to the star creator (Figures 14 and 15).
PROPOSED	**Organization forecasting app.** Anyone in the organization would rate players from a phone: underrated, overrated or fair. Each call would be scored as seasons play out, showing which signals and whose judgments predict best.
PROPOSED	**More qualitative signals.** Coaches' All-Star votes, Most Improved votes and scouting language, each held to the test the All-Defensive votes passed (Figure 16).
PROPOSED	**Free-agent supply forecast.** Who will reach free agency and when, from the rookie-extension and restricted-free-agency models already fitted.
PROPOSED	**Automated briefings.** Pre-game and event briefings for every part of the organization: schedule density, altitude back-to-backs, the referee crew's foul rate, and new data as it arrives.
PROPOSED	**Defensive matchup ratings, effort and gravity from tracking.** Score each defender against each attacker's own usual rates; measure effort as speed and distance on like-for-like plays; measure gravity directly from where defenders stand. Public data cannot separate these from role.
PROPOSED	**A day-one plan for the team's tracking feed.** A 15-step plan, already written: week one would check the feed against numbers reproduced from public data; month one would build what public data cannot.

The model, in depth

- **What it is.** SiPM (Singular Plus-Minus) rates how many points per 100 possessions a player adds, built from public box-score, play-by-play and lineup data.
- **How it is tested.** Its forecasts are scored on seasons it never saw.
- **Against the leading public models.** With ratings frozen on every game day of 2025-26 and every week of the two seasons before, it is level with two leading public impact models: DARKO in every replay, and EPM on every 2025-26 date with an archived copy.

0.81
correlation with next season's team net rating

0.98
slope: a point of forecast is a point of margin (90% interval 0.93 to 1.03)

0.60
the same correlation using last season's team net rating

Add up each player's prior-season SiPM, weighted by the minutes he went on to play, and you get the team's next-season net rating almost point for point. It uses the real minutes, so it tests the ratings, not a minutes forecast.

SiPM against two leading public models

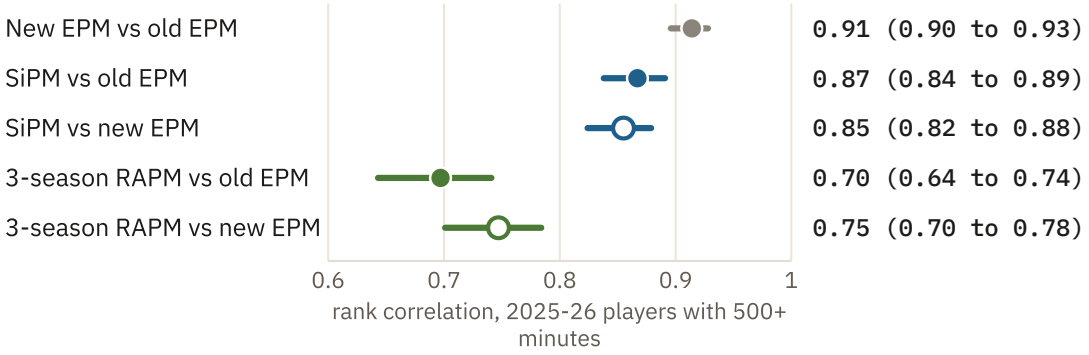
Comparison and replay	Dates	Difference (90% interval)	Reading
SiPM minus DARKO every game day, 2025-26	135	+0.14 (−0.17 to +0.38)	level
SiPM minus DARKO every week, 2025-26	21	+0.19 (−0.16 to +0.39)	level
SiPM minus DARKO every week, 2024-25	21	−0.25 (−0.56 to +0.12)	level
SiPM minus DARKO every week, 2023-24	21	−0.01 (−0.37 to +0.25)	level
SiPM minus EPM game days with an archived EPM, 2025-26	22	+0.07 (−0.10 to +0.24)	level
SiPM minus EPM weeks with an archived EPM, 2025-26	3	+0.12 (−0.09 to +0.30)	level

Score: how much of the error in forecasting the next 30 days of game margins each model's ratings remove, in percentage points; above 0 favors SiPM. "Level" means the interval includes 0.

Like other impact models, it overvalues non-creating bigs and undervalues primary creators. Finding the underlying cause of these residuals, and framing the data and the analysis accordingly, is a primary focus of this work.

FIGURE 3

EPM re-issued its 2025-26 ratings with a wider spread; the benchmark stands because it scored what EPM published at the time



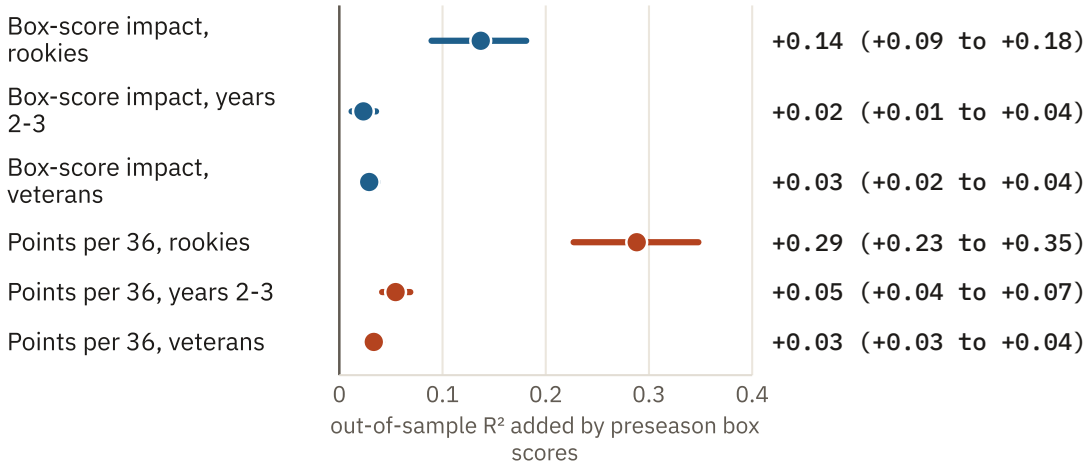
The benchmark against EPM still stands. Between October 7 and 20, 2026, EPM's publisher replaced its 2025-26 ratings with a version whose spread is 14% wider (10 to 17%), though the order barely moved. The replays in the table above scored saved copies of what EPM published on each date, so a later refit does not change them.

Dot: rank correlation. Bar: 90% interval, players resampled. Filled: old version; hollow: new version.

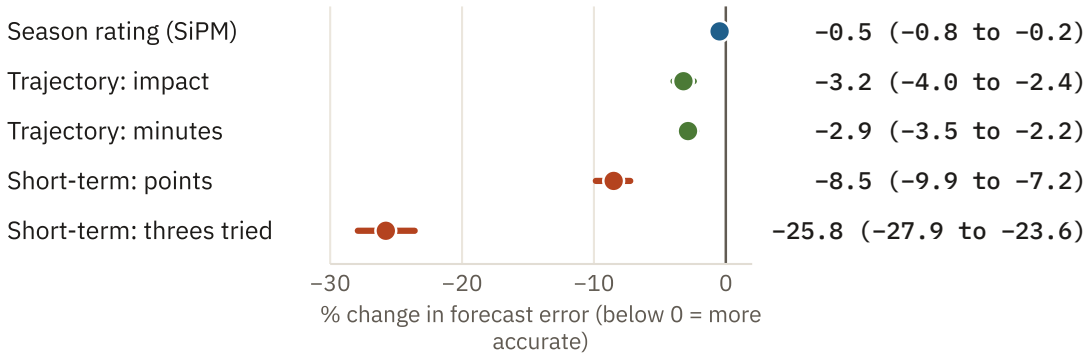
FIGURE 4

Preseason games carry real signal, most for rookies, and counting them makes every player forecast more accurate

A. What preseason adds to next season's forecast



B. Forecast error, a preseason minute counted as 0.7 of a regular-season minute



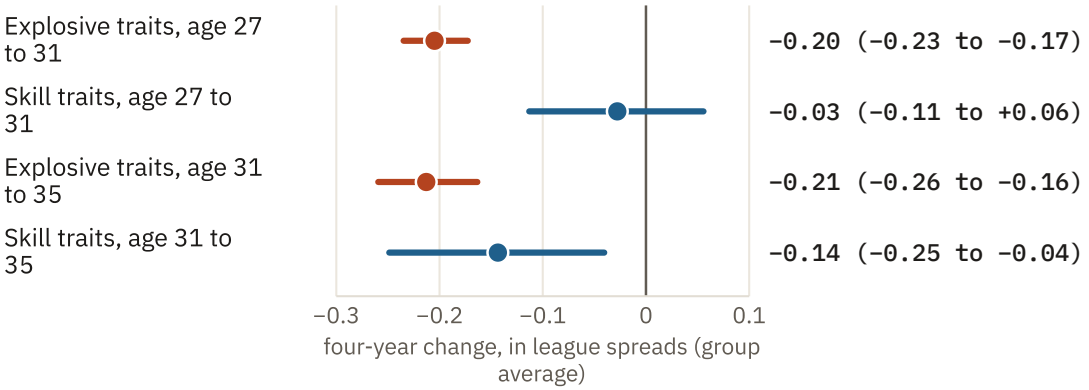
Preseason box scores say the most about rookies. For a rookie's box-score impact they add 0.14 of R^2 to a forecast built on draft slot and age; for returning players the gain is small but real. Counted at 0.7 of a regular-season minute, preseason makes all three player forecasts more accurate in walk-forward tests, and the short-term box-score forecast improves on all 14 stats. This preseason, the largest upward flag among returning players belongs to the Bulls' Leonard Miller: +1.54 on his box-score impact forecast.

R^2: the share of next season's variation a forecast explains. Rookies are compared with draft slot and age; returning players with last season, age and years in the league. Walk-forward: each season is forecast from earlier seasons only. Bar: 90% interval, players resampled.

Trajectories, in more detail

FIGURE 5

From 27 to 31, explosive traits fall while skills barely do; after 31 both fall

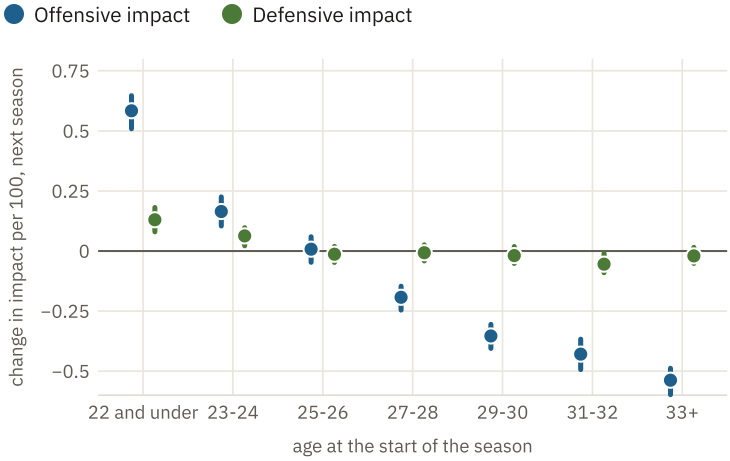


A skill-based game loses less through the late 20s, but not past the early 30s. From 27 to 31 the explosive group falls 0.20 league spreads and the skill group 0.03, a gap of −0.18 (−0.27 to −0.08); from 31 to 35 the explosive group falls 0.21 and the skill group 0.14, a gap not distinguishable from zero.

Dot: estimate. Bar: 90% interval. Explosive group: rim share, dunks, transition share, free-throw rate, blocks, offensive rebounds. Skill group: three-point %, free-throw %, mid-range %, assists, ball security.

FIGURE 6

Offense rises fast when young and falls from the late 20s; measured defense barely moves

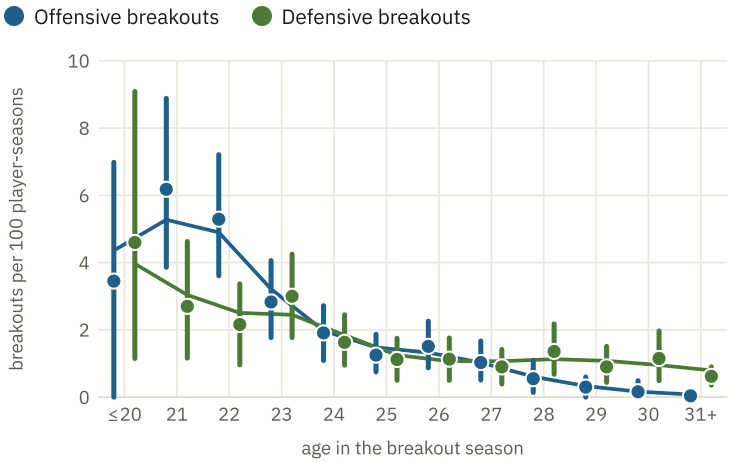


Defense is not the first thing to go, though the measure may be too blunt to see it go. Offensive impact rises 0.58 points per 100 a season at 22 and under, stops rising by 25-26 and falls 0.54 a season at 33 and over. Measured defense moves far less: +0.13 when young and −0.02 when old.

Dot: average change from one season to the next. Bar: 90% interval. Every player with 1,000+ minutes in the earlier season who plays 250+ in the next counts once; players likely to drop out are weighted up, using only what was known that season.

FIGURE 7

Breakouts come young; defensive ones come about two years later than offensive ones

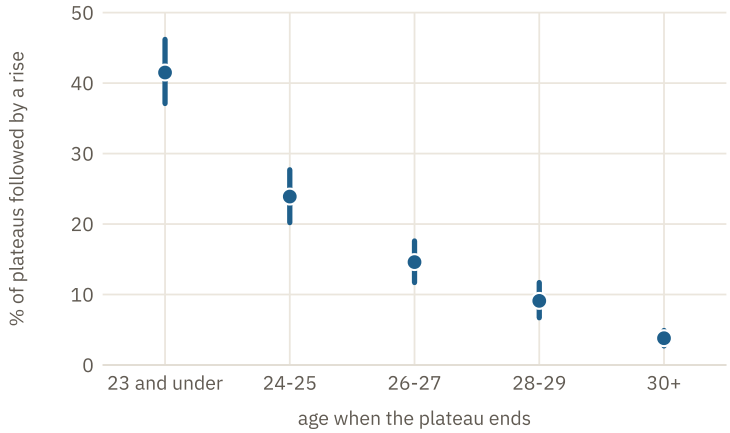


Do not write off a young player's defense at 23. 52% of offensive breakouts come at 23 or under, against 33% of defensive ones (median age 23.9 against 25.7). Most of a breakout is not forecast: for an overall-impact breakout the forecast expected a median 18% of the rise, though the definition builds in some surprise. Young high picks are about five times over-represented among those breakouts, and a change of team is under-represented (0.53 times).

A breakout is a one-season rise in the top 2.5% against the player's own past. Dot: rate. Bar: 90% interval.

FIGURE 8

A plateau at 23 is often a pause; at 30 it is usually the level



The chance of a rise after a plateau falls from 42% when it ends at 23 or under to 4% at 30 and over. Of 2,184 plateaus, 335 were followed by a rise larger than the forecast expected, often on another part of the game such as usage or creation.

Plateau: two or more flat seasons. Dot: share. Bar: 90% interval.

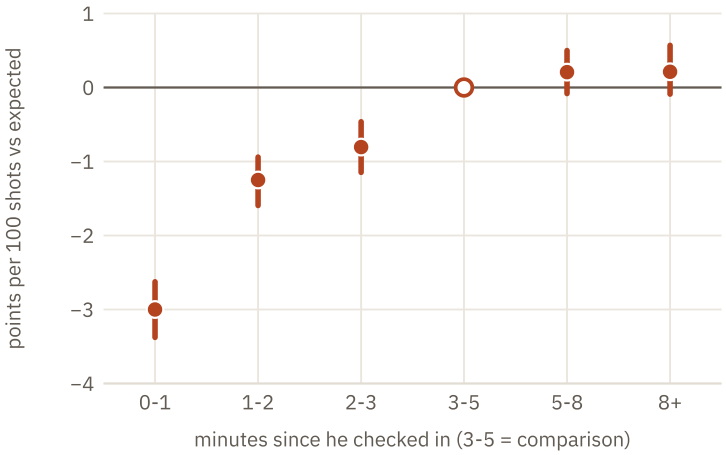
A related caution, for players 22 and under: a young player on a strong team can look worse than he is. Each extra point per 100 in his four teammates' combined prior-season rating lowers his measured impact by 0.125 per 100 (0.078 to 0.168).

From the front office to the bench

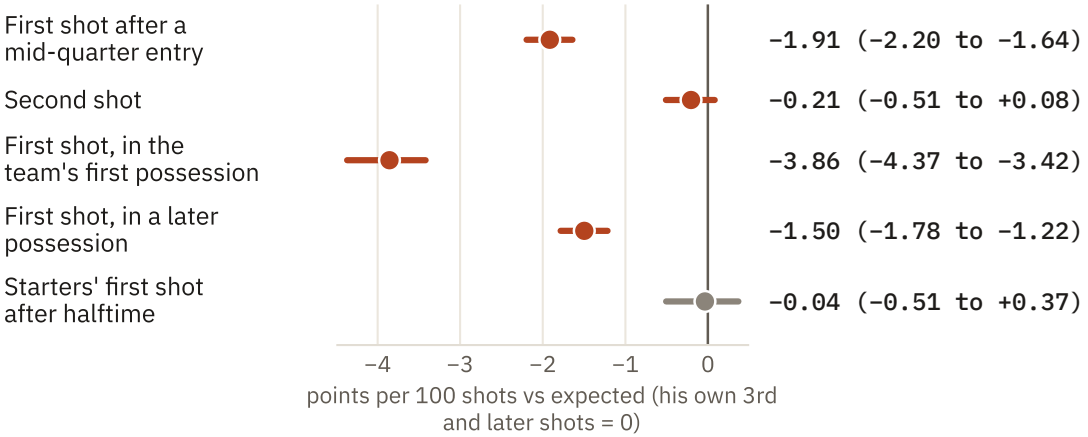
FIGURE 9

A player's first shot after checking in is his worst, especially in the team's first possession

A. By minutes since entry



B. By shot order and possession



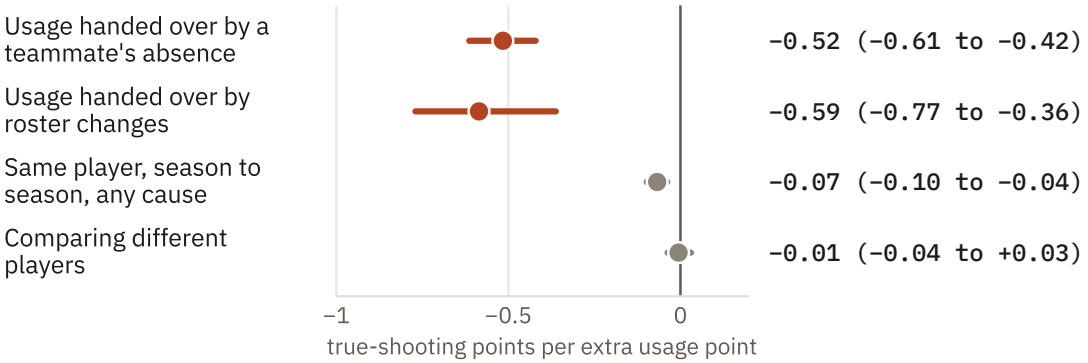
A player's first shot after a mid-quarter entry is worth 1.91 fewer points per 100 shots than his own later shots, after accounting for shot type, location and his skill. The first minute is worst (about 3 points), and the dip is gone after about three minutes. It is worst in the team's first possession after he enters (−3.86), against −1.50 in a later possession. The effect is small: about 0.3 points a game for a team. Moving that first-possession shot to someone else could recover 0.06 to 0.10 a game. Whether to act on it is the coaching staff's call.

Every regular-season shot 1996-97 to 2025-26. Dot: estimate. Bar: 90% interval. The recovery is 0.06 a game if his first shot moves to a later possession, and 0.10 if an average shot replaces it. The possession split is not tested on fresh data, and coaches may call plays for players they expect to be guarded poorly. Starters coming out of halftime, after a warm-up, show no dip. A player's own dip does not repeat from season to season, so it is no use for rating players.

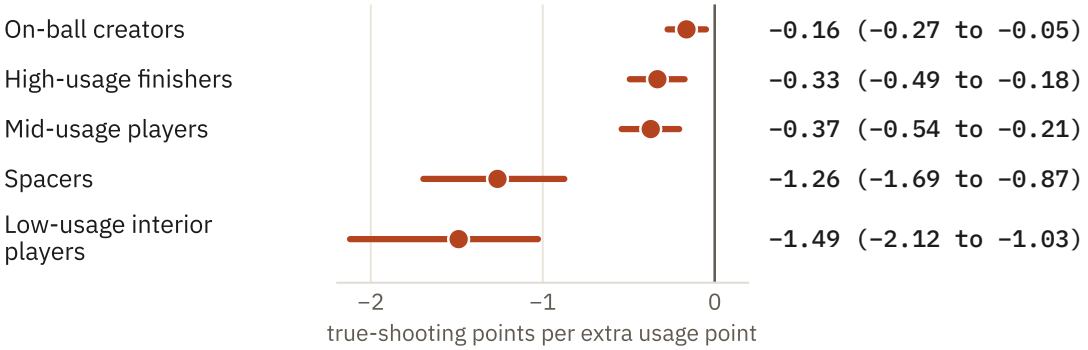
FIGURE 10

When a teammate's absence hands a player a bigger role, his true shooting falls, mostly from losing the teammate who set him up

A. The cost, by how it is measured



B. Who pays it (teammate-absence test)

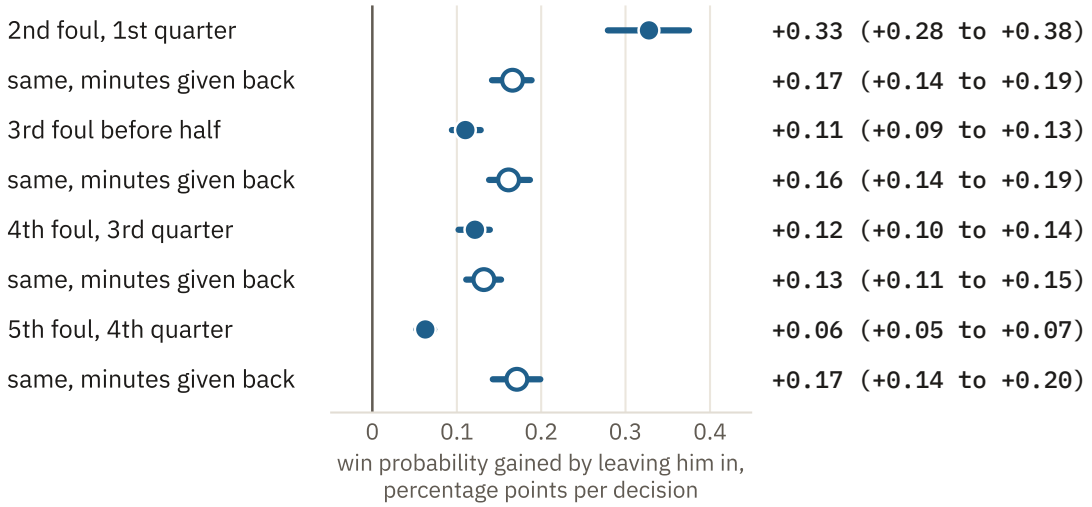


When a teammate's absence or a roster change hands a player a bigger role, his true shooting falls 0.52 to 0.59 points per usage point. Comparing players shows almost no cost, because coaches give the ball to the players who can carry it. In the teammate-absence test, most of the cost is losing the teammate who set him up; with that passing held fixed, the extra shots cost little: −0.06 (−0.21 to +0.08). The roster-change test cannot split the cost this way. On-ball creators pay little; spacers and interior players pay a lot. Usage is a resource to fit together as each player's efficiency develops.

Dot: estimate. Bar: 90% interval. Teammate-absence test: 1996-97 to 2025-26. Roster-change test: 2022-23 to 2025-26.

FIGURE 11

Leaving a starter in after an early second foul is worth about a third of a point of win probability; coaches lose about 0.3 wins a season by sitting

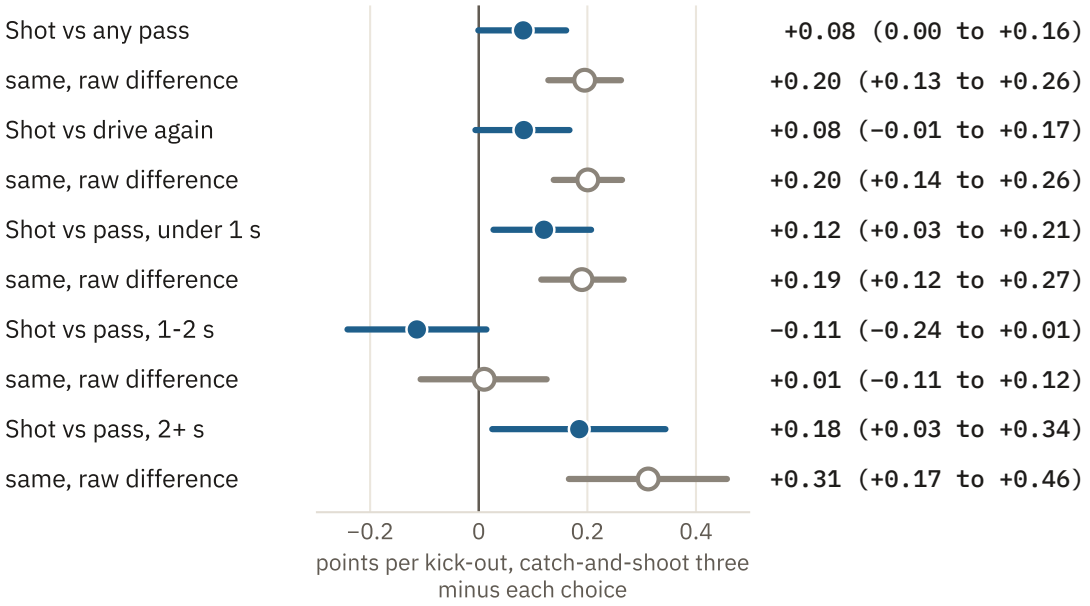


Leaving a starter in usually pays. After a second foul in the first quarter it adds +0.33 points of win probability (+0.28 to +0.38), or +0.17 once the minutes coaches give back later are counted; sitting is right only when the substitute is rated better than the starter. Coaches' actual choices give up 0.30 wins per team per season (0.27 to 0.32), almost all by sitting.

Value: the win value of a point of margin, times the starter's edge over the teammates who replace him, times the extra minutes he plays. 103,235 trouble events, 1996-97 to 2025-26. Billy Donovan's leave-in rate is close to the league average. Bar: 90% interval, players resampled; it covers sampling, not the model's assumptions.

FIGURE 12

After a kick-out, the catch-and-shoot three beats passing, but by less than half the raw gap once the situation is held fixed



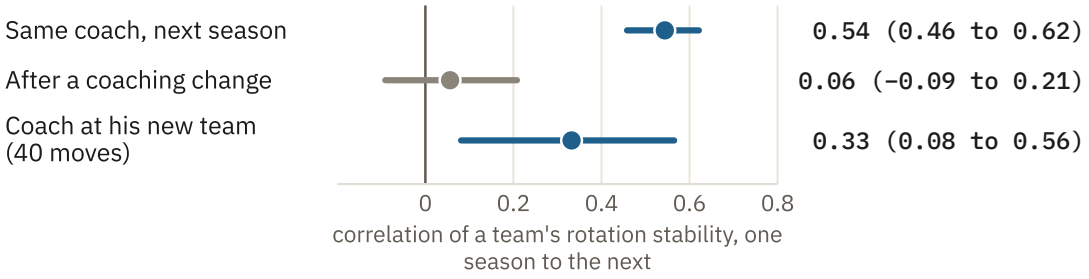
The catch-and-shoot three is still the best choice after a kick-out, but by less than the raw numbers suggest. Its raw edge over any pass, +0.20 points per kick-out, falls to +0.08 (0.00 to +0.16) once the catcher's space, the defenders, the clocks and the score are held fixed. A pass held one to two seconds scores about as well as the shot, and always following the model's best choice would add only +0.03 points (−0.03 to +0.10).

From player tracking of 631 games, October 2015 to January 2016: 136,735 perimeter catches, 7,243 of them after a drive. The estimate combines a model of who chooses what with a model of the points each choice produces. Filled: situation held fixed; hollow: raw difference. Bar: 90% interval, games resampled.

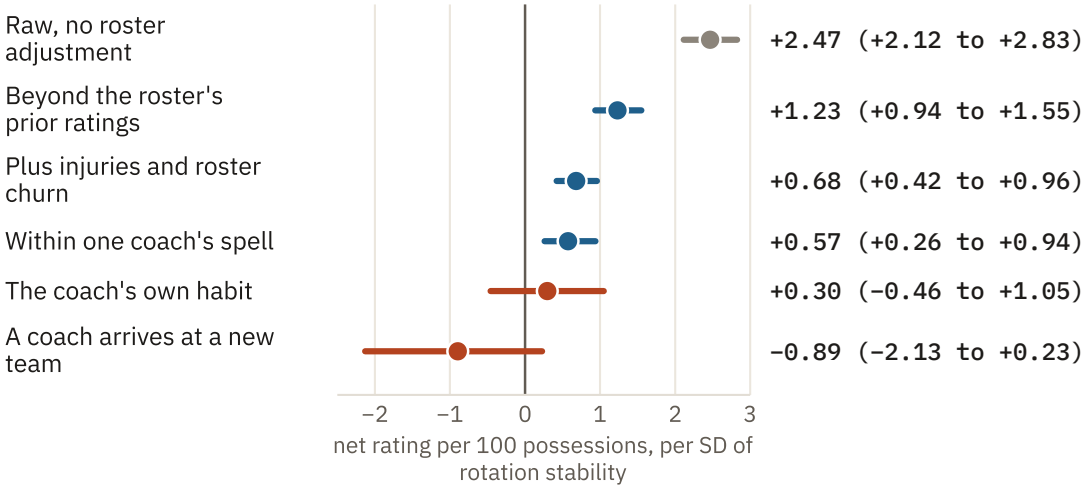
FIGURE 13

Rotation stability is a coach's habit that follows him to new teams; stable seasons go with better results, but nothing shows the habit causes them

A. Stability follows the coach



B. Stability and results



Rotation stability belongs to the coach, and nothing shows it costs young players. With the same coach it carries over at r = 0.54; after a coaching change, 0.06. Steadier teams beat their roster's expectation by +0.68 per 100 per SD (+0.42 to +0.96), but the designs that isolate the coach find nothing, so healthy, winning teams likely settle their rotations rather than the reverse. Billy Donovan's Bulls sat about at the league average (0.02 below it on a 0-to-1 scale).

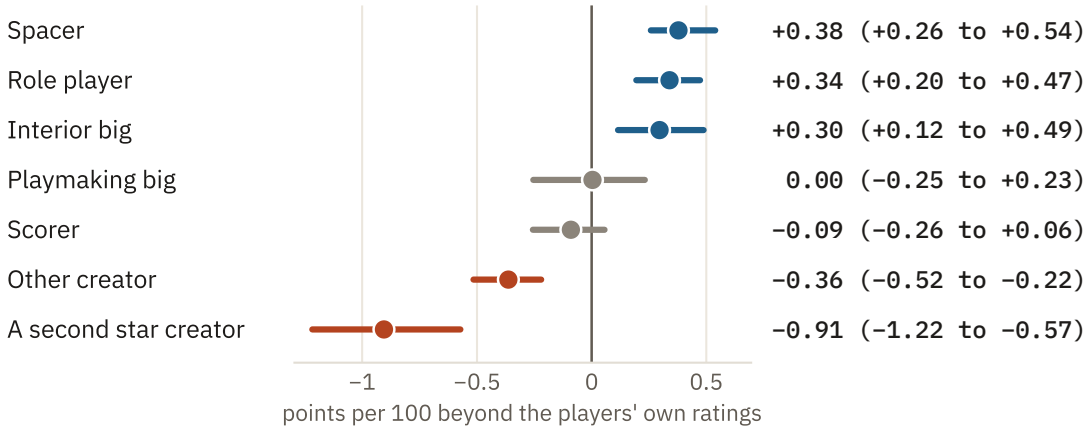
Stability compares each rotation player's minute-by-minute pattern across every pair of games: 1 = the same minutes every night, 0 = random (league average 0.25). Players 23 and under change by −0.15 BPM per SD of team stability (−0.26 to −0.06), the opposite sign to a development cost. 390 team-seasons, 2013-14 to 2025-26. Bar: 90% interval, teams resampled.

Fit: building around a star creator

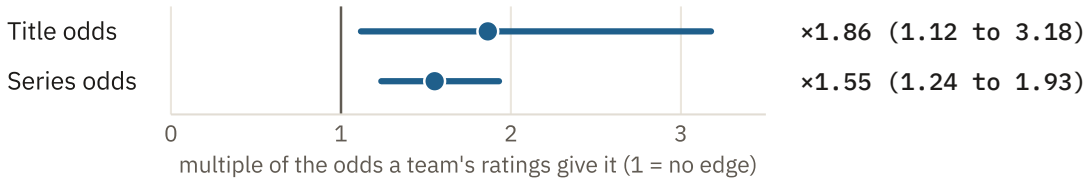
FIGURE 14

Spacers and interior bigs fit next to a star creator; two star creators are worth less than the sum of their ratings

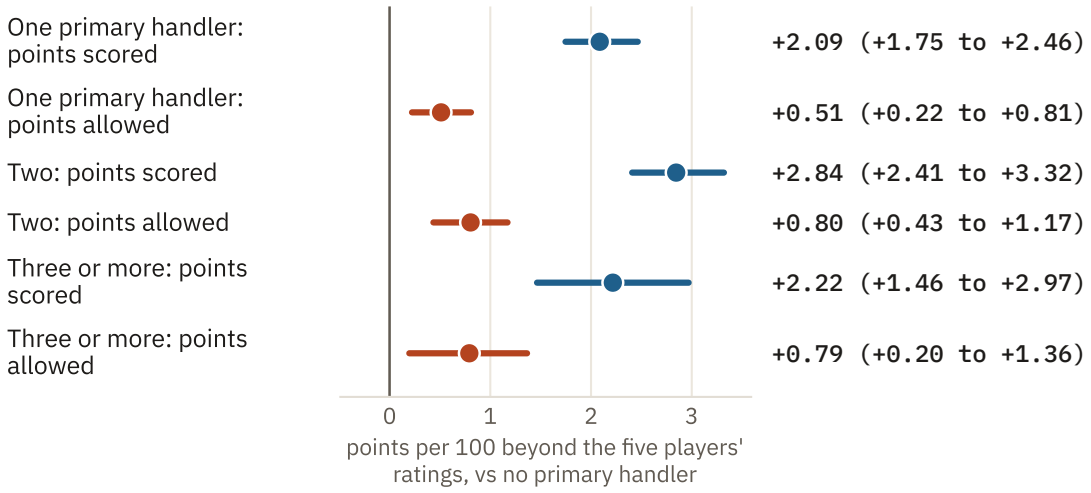
A. Next to a star creator, each type adds



B. Title and series odds with a lead option who shoots at least league average



C. Stacking primary ball handlers: the second adds little on offense and costs in full on defense



Next to a star creator, spacers add 0.38 points per 100 beyond their own ratings and interior bigs 0.30; two star creators together are worth 0.9 per 100 less than their ratings add up to. Pair effects held up in later seasons at 1.21× their size (0.69 to 1.71). Teams with a lead option who shoots at least league average won 32 titles; their ratings and playoff brackets predicted 26.7. As a group they were already favorites, so multiplying their odds by 1.86 adds about five titles over 46 seasons. Without the shooting condition the multiplier is ×1.46 (0.85 to 2.67), which cannot be told apart from no edge. A forward test on the 2008-26 playoffs was too small to confirm the gain. Panel C: lineups with one primary ball handler score 2.1 more per 100 than their ratings predict; a second handler adds only 0.8 more, while the points allowed keep rising (+0.51, then +0.80). Sharing the ball does not buy the defensive cost back.

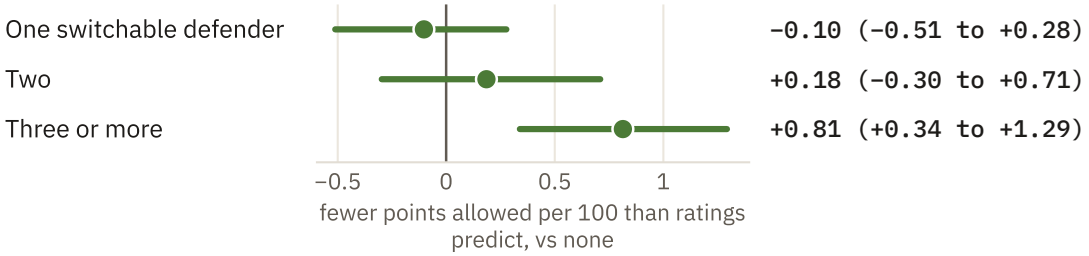
Star creator (panel A): usage at least one standard deviation above average and offense in the season's top 10%. Lead option (panel B): a regular-minutes player whose league usage rank is at most half the number of teams (15th with 30 teams) and

whose true shooting is at or above league average. Dot: estimate. Bar: 90% interval. Pairs: lineups 1996-97 to 2025-26. Titles: playoff teams 1980-81 to 2025-26.

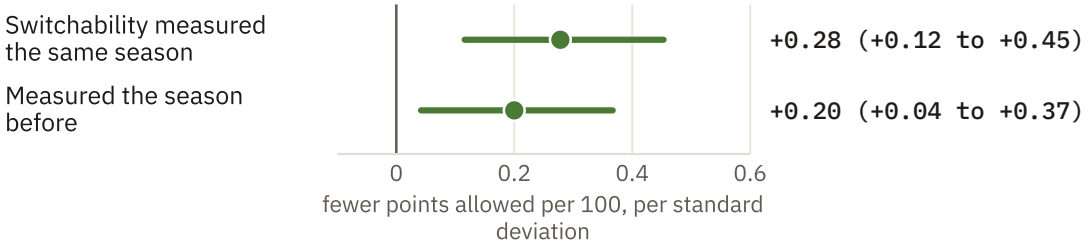
FIGURE 15

Lineups with more switchable defenders allow fewer points than their ratings predict

A. By number of switchable defenders



B. Per standard deviation of lineup switchability



Most of the gain comes with the third switchable defender. Per standard deviation of switchability, lineups allow 0.28 fewer points per 100 than the five players' defensive ratings predict (0.12 to 0.45), and last season's switchability still predicts it. The intervals are wide.

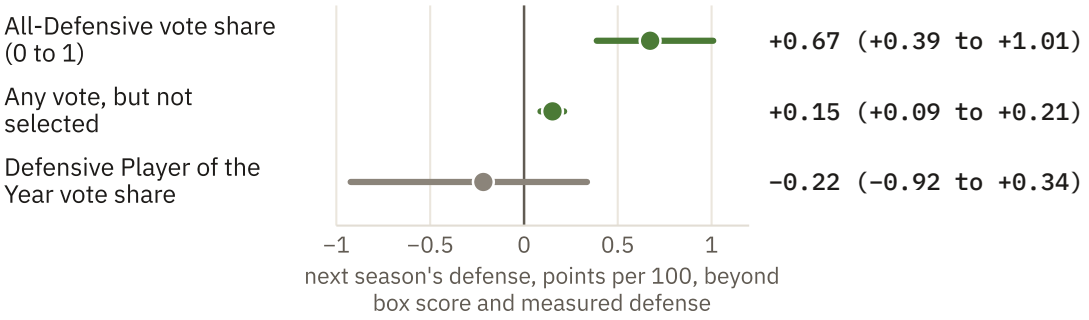
Switchability: how widely the heights of the players a defender guards vary, beyond what his usual assignment predicts; a stable trait (reliability 0.72). NBA.com matchup data 2017-18 to 2024-25. Dot: estimate. Bar: 90% interval.

Defense, and the limits of public data

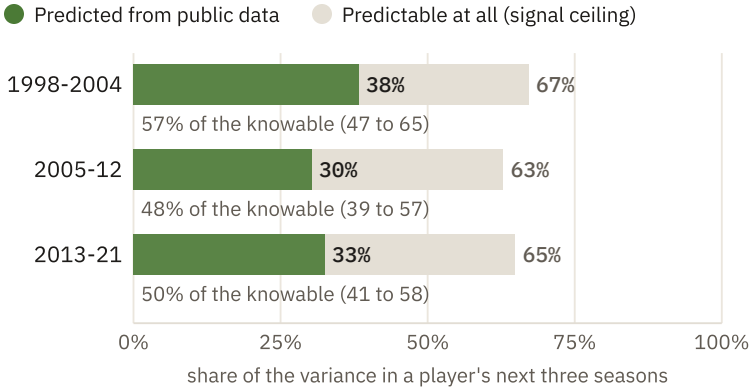
FIGURE 16

Award votes add defensive signal the numbers miss; public data captures about half of the signal in a player's future

A. What award votes add beyond box score and measured defense



B. How much public data can know



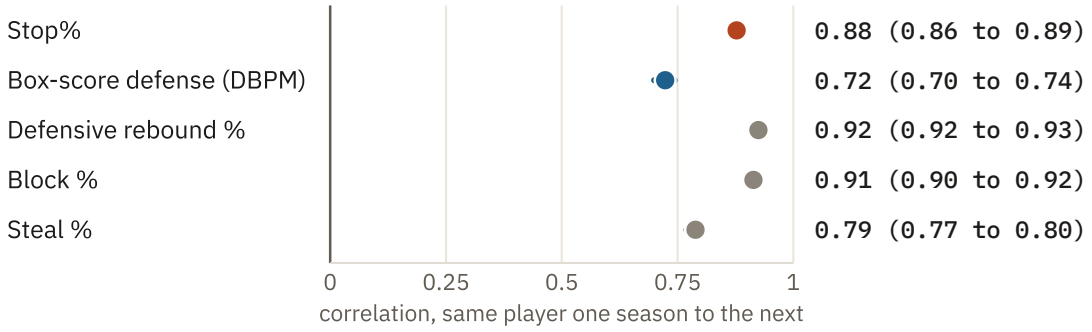
All-Defensive votes predict a player's defense next season beyond his box score and measured defense: +0.67 per 100 for a full share, and +0.15 for any vote short of selection. Public data explains 30 to 38% of a player's next three seasons of impact, and 63 to 67% of that impact is signal rather than noise, so public data captures about half of the signal. Better models of public data can close some of the gap; the rest needs data teams hold: tracking, medical, practice and scouting. In the draft, when draft order and a public stats model disagree, draft order wins 56% of the time on four-year value (51 to 62%).

A: 1996-97 to 2024-25; dot: estimate, bar: 95% interval. B: RAPM averaged over the next three seasons; the ceiling is the share of that average that is signal rather than noise, and it has no interval.

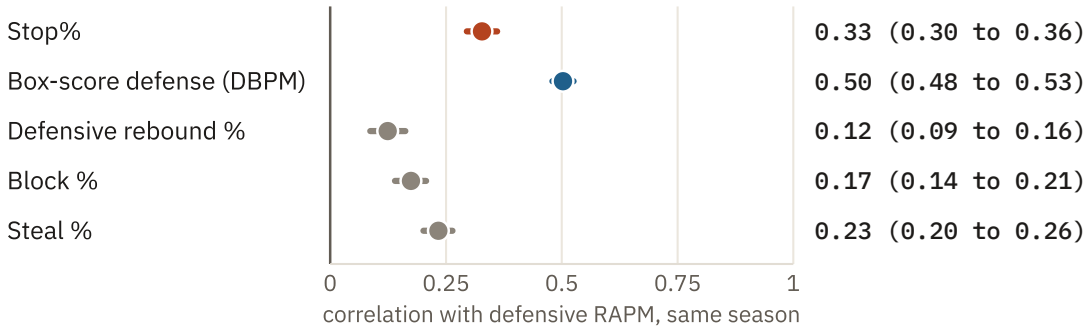
FIGURE 17

Stop% is stable because it is mostly rebounds and blocks; it tracks defensive impact less well than box-score defense and adds nothing to it

A. Stable from season to season



B. Tracks measured defensive impact



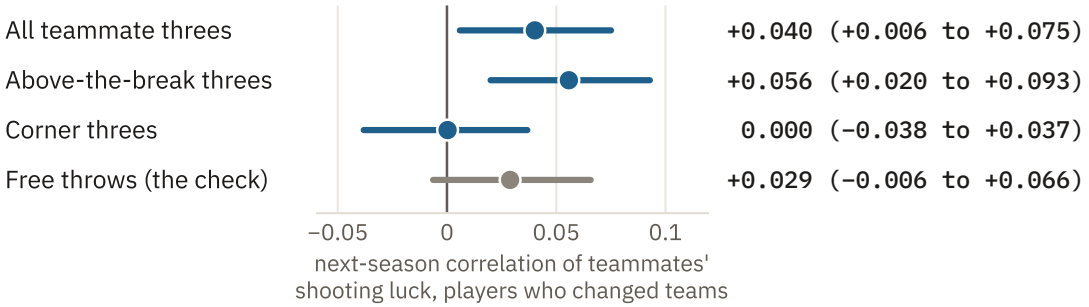
Stop% tracks measured defense less well than box-score defense (0.33 against 0.50) and adds nothing to it. Added to box-score defense, it changes next-season prediction error by −0.01% (−0.06% to +0.05%). It is as stable as the rebound and block rates it is built from: 85% of the 20 career leaders are big men, against 23% of all 819 players with 5,000+ minutes since 2003-04.

Stop%: Dean Oliver's individual stop rate, rebuilt from box scores; about 60% of it is the team's forced misses and turnovers, shared out by minutes. Player-seasons with 1,000+ minutes, 2003-04 to 2025-26; out-of-sample test fit on 2003-13 and scored on 2014-24. Bar: 90% interval, players resampled.

Central questions, tested

FIGURE 18

Gravity: a small effect, clearest on above-the-break threes

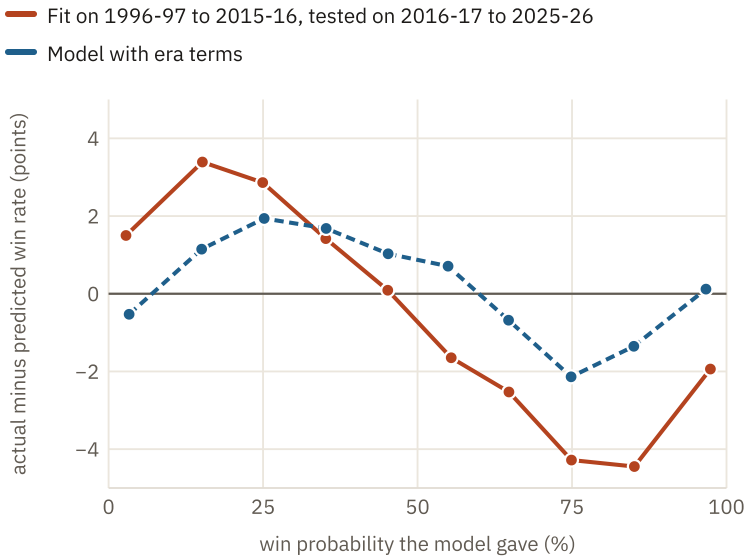


The evidence for gravity is suggestive, not settled. When a player changes teams, his new teammates' threes beat their norms in the same direction they did on his old team: weakly for all threes, more clearly above the break, not in the corners. The all-threes correlation (0.040) barely clears what random teammates produce by chance (0.038), and the free-throw check, where gravity should not show, is smaller (+0.029) but not zero, which weakens the result. A typical player's true effect is about 0.47 percentage points on each teammate three, but one season of the measure is about 98% noise. Several seasons, or tracking data, would settle it.

Dot: correlation. Bar: 90% interval. Free throws are the check because nobody guards them.

FIGURE 19

Leads are less safe than they used to be: an older win-probability model is overconfident on recent games



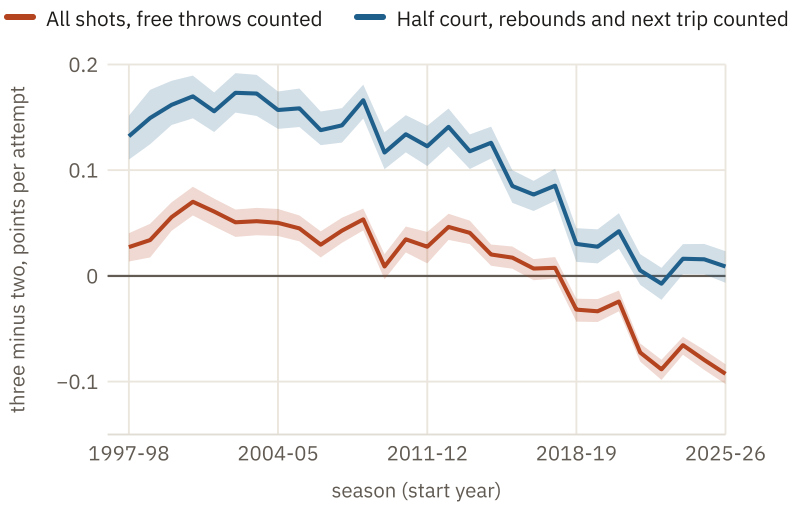
Leads are less safe than they used to be, so late-game decisions should use a model fit on the current era. A model built from score, clock and possession on 1996-97 to 2015-16 games is too confident on newer games: when it gave a team 85%, that team won 81%. Its odds need pulling back toward even by about a sixth (a slope of 0.83, where 1.00 means no correction). The model with era terms is closer but still off by about 2 points: when it gave a team 25%, that team won 27%, and when it gave 75%, it won 73%. Each season on its line is scored by a model fit on the other seasons, so it has already seen the current era.

Each point is one of ten bins of predicted win probability. 0 = perfectly calibrated; below 0 on the right means favorites won less often than the model said.

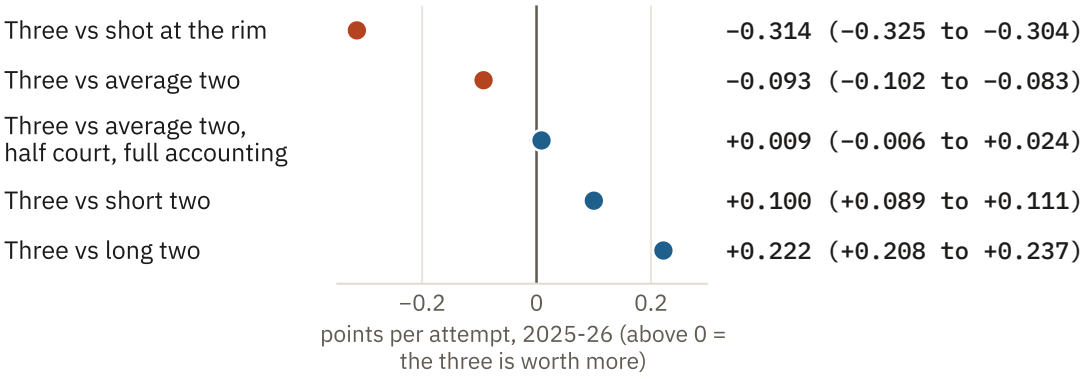
FIGURE 20

Threes versus twos: in the half court, fully counted, they are even, and the balance can be tracked every season

A. Three minus two, every season



B. By kind of two, 2025-26

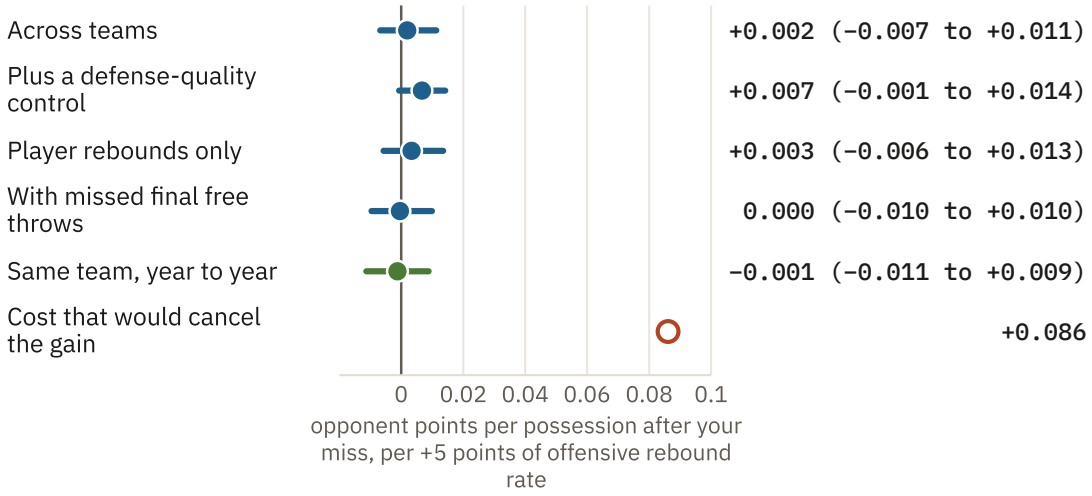


In the half court, once offensive rebounds and the opponent's next trip are counted, threes and twos are even (+0.009, −0.006 to +0.024): NBA teams have found the balance. Across all shots, counting free throws but not rebounds, the average two returns 0.09 points per attempt more than the average three, so it can look as if teams take too many threes; leaving out putbacks and fast breaks removes about 60% of that lead. Shots at the rim are the only twos worth more than a three: get to the rim, and do not trade threes for midrange twos.

Band and bar: 90% interval. Half court excludes transition, putbacks, heaves and garbage time.

FIGURE 21

Crashing the offensive glass costs almost nothing in transition: each 5 points of offensive rebound rate nets about 6.6 points per 100 misses

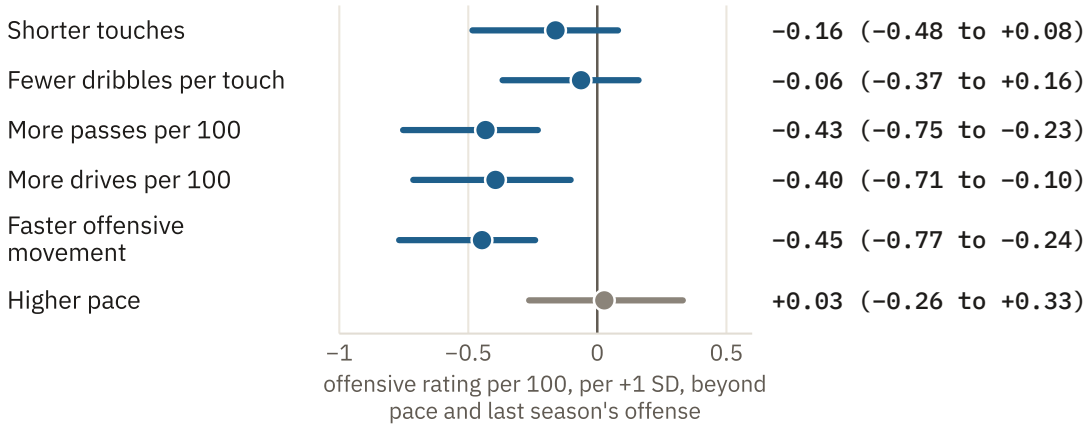


Crashing the offensive glass barely costs anything in transition. For every 5 points of offensive rebound rate a team adds, opponents score +0.002 points per possession after its misses (−0.007 to +0.011), far below the +0.086 that would cancel the extra rebounds; counting the opponent's next possession, the net gain is +6.6 points per 100 misses (5.6 to 7.5). Offensive rebound rate on misses has risen every season, from 26.0% in 2020-21 to 29.9% in 2025-26.

664,302 missed field goals in 180 team-seasons of play-by-play, 2020-21 to 2025-26. Bar: 90% interval, franchises resampled. Blind to how many players crash; that needs tracking.

FIGURE 22

Team pace habits are stable, but playing faster with the ball does not buy offense



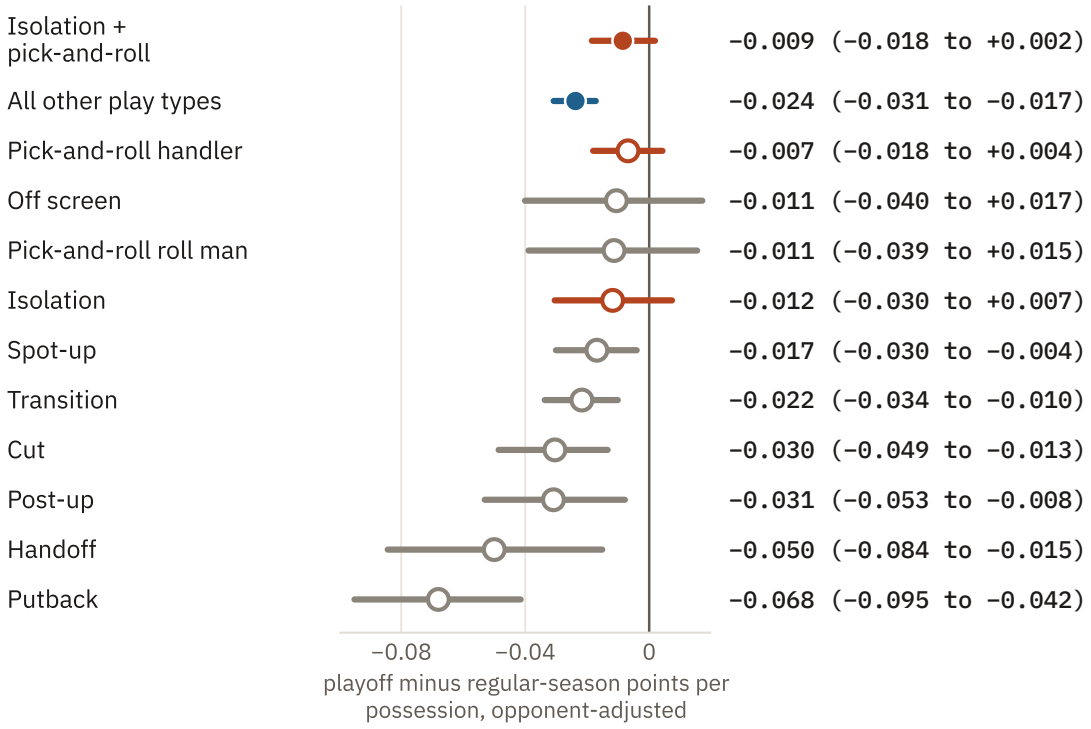
Playing faster with the ball does not buy offense. Touch length and passes per 100 are stable team habits (r = 0.63 each, year to year), more than pace (0.55), but more passes, more drives and faster movement each go with slightly worse offense, about −0.4 per 100 per SD, and pace shows nothing (+0.03). The likely reason runs the other way: teams with a star creator hold the ball and score well.

390 team-seasons, 2013-14 to 2025-26, from NBA.com team tracking. Six habits were tested, so one chance result is possible. Bar: 90% interval, franchises resampled.

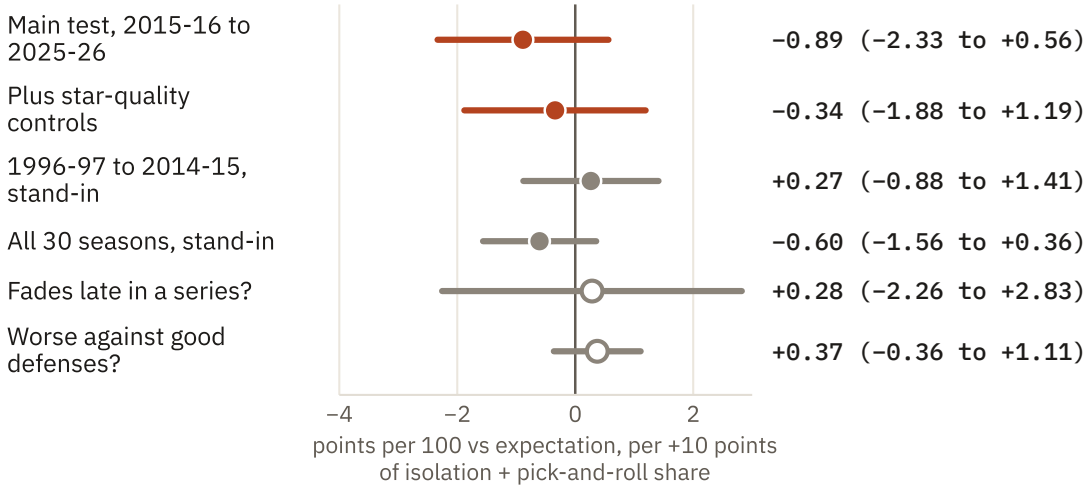
FIGURE 23

No measurable playoff penalty for an isolation and pick-and-roll style: those plays hold up better than others, and series results show no clear cost

A. The same players' plays, regular season to playoffs



B. Playoff games and series



Isolation and pick-and-roll plays hold up in the playoffs better than other plays. From the regular season to the playoffs, a player's isolation and pick-and-roll plays change by −0.009 points per possession (−0.018 to +0.002), against −0.024 for his other plays. A team whose isolation and pick-and-roll share is 10 points above its opponent's does −0.89 per 100 against expectation (−2.33 to +0.56), which rules out a series-win penalty larger than about 13 points; two secondary versions lean negative with intervals that just exclude zero, so a small penalty remains possible.

Synergy play types, 2015-16 to 2025-26. Stand-in: each team's share of unassisted made shots (r = 0.76 with the Synergy share). A: 90% interval, players resampled; hollow dots are single play types. B: 90% interval, allowing games from one team-season to be correlated.

How ideas get tested

820	304	82	434
questions tested	answered no	mixed	answered yes

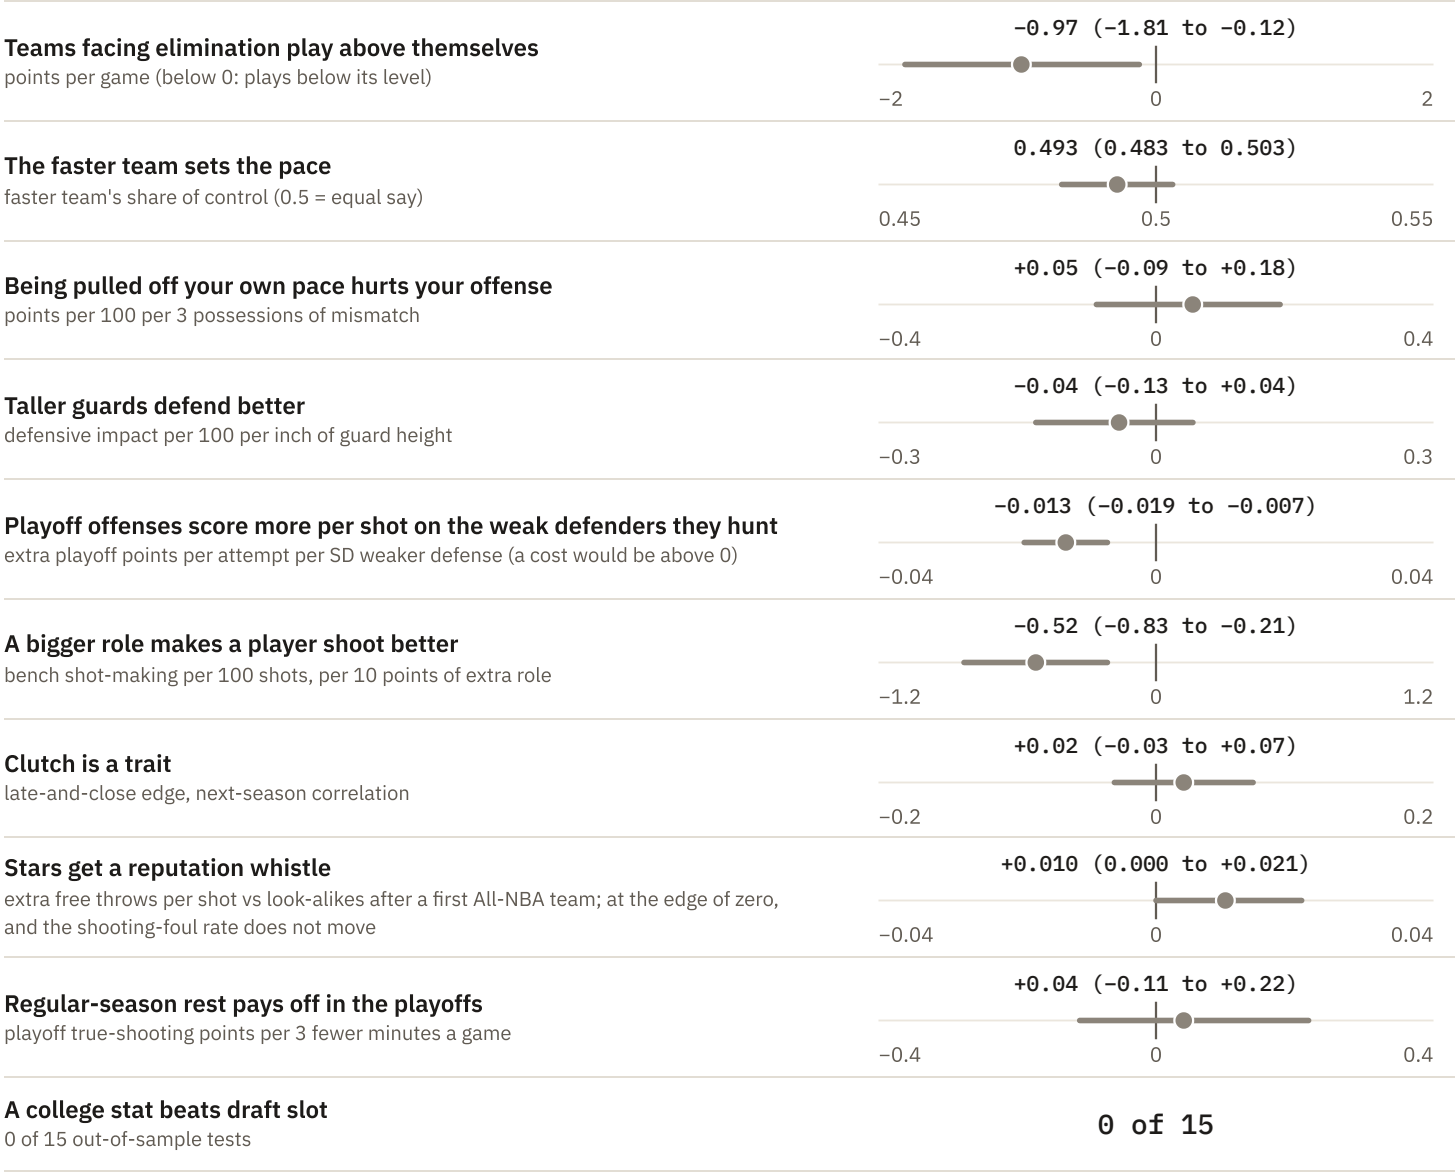
- Forecasts are scored on seasons the model never saw.
- Most numbers carry a 90% interval.
- With this many tests, some "yes" answers will be luck. Weigh single results with that in mind.

FIGURE 24

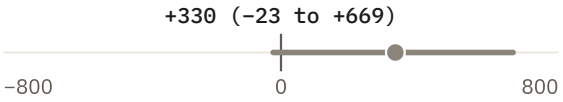
Tested and dropped: popular ideas the data does not support

Each idea sounds right and failed a fair test: the dot sits at or near no effect, or on the wrong side. Three point clearly the wrong way: teams facing elimination play slightly below their level, a bigger role slightly lowers bench players' shot-making, and hunted defenders are not scored on more efficiently. Length still matters on defense: longer players defend better, even though taller guards do not.

Dot: estimate. Bar: 90% interval. Each row has its own scale; the vertical line marks no effect (equal say for pace).



College shooters struggle to make it in the NBA
extra NBA minutes in the first three seasons vs draft slot, shooters minus other picks, 2018-23 classes



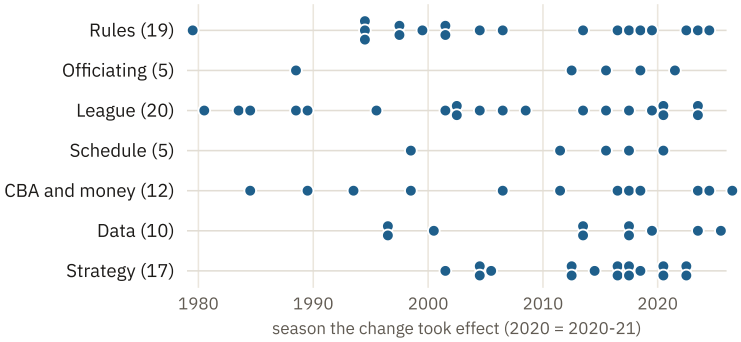
Style uniqueness predicts stardom
change in All-NBA forecast precision (no interval saved)



FIGURE 25

Qualitative history, quantified: every change to the game, dated

88 changes to the game (87 dated), 1979-80 to 2026-27



Any trend over time is read against what changed in its window: rule changes, officiating memos, CBA terms, data-feed changes and strategy shifts, each dated. On the web version, hover over a dot to see the change.

When a provider quietly changes its data, the inputs get re-checked

After the project's September pull, NBA.com re-issued its 2024-25 player-tracking data on its newer Hawk-Eye definitions, with no public note. A routine analysis caught it as a sudden jump in one season; the 24 affected files were re-pulled with backups, and the results built on them were re-run. No headline finding flipped: the shot-making rating still tracks the NBA's own Hawk-Eye shot-making (correlation 0.71 before, 0.70 after). One stability measure moved (drive fouls drawn, year to year: 0.74 to 0.57) because the new files leave the free-throw fields nearly empty; those fields are now blanked or flagged.

The idea the project is named after did not hold

Singular started from a hunch: the NBA, more than other sports, runs on singular talents, so a player whose style is unlike anyone else's should be likelier to become a star. Tested on future All-NBA players, style uniqueness slightly hurt the forecast, and a machine-learned style profile did worse.

That version of the idea could not be proven with public data alone. But the best players in the league plainly each have a unique game, so the nature of top talent remains a guiding principle. There is something there, and it is what I want to keep pursuing.

TERMS USED

Per 100 possessions	points scored or allowed per 100 trips down the floor, so fast and slow teams compare fairly.	**Net rating**	points scored minus points allowed per 100 possessions.
True shooting	shooting efficiency that counts threes and free throws.	**Usage**	the share of a team's plays a player finishes (shot, free throws or turnover) while he is on the floor.
RAPM	a player's impact measured from the lineups he played in, with teammates and opponents accounted for.	**90% interval**	the range the true value plausibly falls in; a bar that crosses zero means "cannot tell".
League spread	one standard deviation across players, so traits on different scales can share a chart.		

Prepared October 2026. Intervals are 90% unless a figure says otherwise.