

The constructor problem

-- why basketball value doesn't travel --

Javier Romanelli -- last updated in August 2026

Introduction

Every trade deadline my league melts down over the same argument. A deal happens, someone posts the two stat lines side by side, and everyone fights about who won. I've been doing fantasy analytics for about fifteen years and it took me embarrassingly long to notice the argument itself is broken. It assumes a player's numbers belong to him. That 25 a game is a property of the guy, like his height, and it ships wherever he goes. Almost everything I've seen (playing, projecting, reading the actual research) says otherwise. This writeup is the case.

The cleanest frame comes from a different sport. In Formula 1 nobody pretends the driver does it alone: a model of 65 years of race results found the constructor -- the car and the team -- explains about 86% of the variance in results, and the driver about 14%, replicated at roughly 88% by an independent Bayesian re-analysis. Basketball is nowhere near that lopsided (Part 4 gets into how far), but we talk about players as if the ratio ran the other way. A box score is a joint measurement. One player, one role, one scheme, four specific teammates, an opponent, some luck. We collapse all of that into a line and hand it to the player.

And this is not a message-board problem. Teams commit hundreds of millions on these numbers, and every fantasy season is won or lost on them. So: Part 1 is what the research shows. The rest follows one player who carries the whole argument. Basketball-Reference lists "The System" among James Harden's real nicknames, which tells you where this is going. The punchline at the end: the same fact tells a fantasy manager and a GM to do opposite things.

Part 1 -- the car and the driver

Three findings anchor everything. The first is that **spillovers are large and unpriced**. Using possession-level NBA data, Arcidiacono, Kinsler & Price (2017) separated each player's ability to produce from his ability to raise the production of the four teammates around him. A one-standard-deviation increase in spillover ability improves team success about 63% as much as the same increase in direct productivity -- helping others produce is worth nearly two-thirds of producing. Their model also shows a player's value formally varies with his teammates -- value is a function, not a scalar -- and that salaries price own production only, leaving the spillover term unpriced: a documented, standing market inefficiency.

The second is that **skill curves are real and measured** (Figure 1). Dean Oliver proposed that every player's efficiency declines as his usage rises; Goldman & Rao estimated the curves empirically and found that they slope downward for nearly everyone, that guards run flatter than big men, and -- the fact that quietly settles most star debates -- that higher-scoring players have flatter curves in general. Stars aren't the players who score a lot; they're the players whose efficiency barely bends under load. The same literature shows an optimal offense equalizes *marginal*, not average, efficiency, so the best player should be used less than greedy logic suggests; Skinner's "Price of Anarchy in Basketball" derives the Braess corollary that removing a key player can improve an offense -- the Ewing Theory, formalized.

The third is that **context as thin as a contract date moves the numbers**. Across 230+ NBA and MLB players, scoring statistics rise in the contract year and then fall below the pre-contract baseline after signing (in the NBA, PER shows the same bump-then-dip). If a signature date bends the box score, imagine what a scheme change does.




Figure 1 | The load is assigned; the slope is owned. (a) Harden's scoring load, every season from 2009-10 to 2025-26, colored by role: bench (red), No. 1 option (blue), co-star (light blue). (b) True shooting over the same span. The load tripled and came back down; the efficiency sat in a 58-66% band for sixteen straight seasons. Data: [Basketball-Reference](#).

What the database says

Those findings are other people's work, so I built a database to test them directly -- every NBA player-season since 1947 (Basketball-Reference, via [Datta's open dataset](#)): 8,125 consecutive-season pairs with real minutes since usage tracking began in 1978 (the full database ships alongside). First, selection hides the curve: within careers, usage and efficiency correlate *positively*, because coaches hand possessions to players when they're playing well -- but across the 101 natural experiments where a team's 27%+-usage top option departed, the teammates forced to absorb 3+ extra usage points lost about 0.17 points of true shooting per point, double the cost of expansions players chose (Figure 2). Second, the curve is asymmetric: movers who shed 3+ usage points recovered about 0.14 TS per point -- the Durant effect is a league-wide regularity, and shedding load pays more than adding it costs. Third, no relocation tax: changing teams at a constant role cost nothing (+0.1 $\pm$ 0.2 TS, $n = 695$). Whatever a trade destroys, it is not the player's shooting




Figure 2 | What one usage point costs. (a) Year-over-year change in true shooting against change in usage for all 8,125 consecutive-season pairs (1,000+ minutes both years, age-adjusted), 1978-2025, binned. The red point is teammates forced to expand after a 27% +-usage top option left. (b) The per-point price, by how the change happened; whiskers are ± 1 s.e. Computed in the companion database.




Figure 3 | How much of a team carries over has never changed. Season-to-season correlation of franchise win percentage, by decade, 1950-2025. Free agency and the salary cap are marked. The 2020s point is partial and includes two COVID seasons.

Part 2 -- the system called Harden

In 2011-12, Harden was Oklahoma City's sixth man: 16.8 points in 31.4 minutes on 10.1 attempts, off the bench in 60 of 62 games, at a 21.6% usage rate for a 47-19 Finals team. The market priced the role as the player -- RotoWire's preview that summer found it "hard to imagine him ever becoming a 20-point per game scorer with the Thunder." A reasonable statement about a role, mistaken for one about a player.

Houston handed him the keys in October 2012 and he was immediately -- not after a development arc -- a 25.9-point All-NBA player on 17.1 attempts at 29.0% usage, eventually scaling to 36.1 points per game at a 40.5% usage rate in 2018-19. Same driver, new car -- the possessions simply hadn't existed in his old role. There's a reason "The System" stuck.

And Harden isn't an isolated case -- the league runs this experiment every season, and it's repeatable enough that I think of it as the **vacated-usage effect** (Table 1): Gilgeous-Alexander (10.8 ppg Clippers rookie → MVP, scoring champion, Finals MVP once handed a franchise), Brunson (Dallas sixth man, let walk → New York No. 1, All-NBA, ultimately a ring), Siakam (16.9 → 22.9 the season Kawhi's usage vacated), Ingram (18.3 → 23.8 and Most Improved upon leaving LeBron's orbit), Oladipo (15.9 beside record-usage Westbrook → 23.1, All-NBA). Jerami Grant said the quiet part out loud, signing with Detroit because the GM promised him a "franchise player" role -- and doubled his scoring overnight at 26. So, stated carefully: the set of players who could average roughly 20-25 points *at some efficiency* if handed primacy on a bad team is far larger than the set we call stars. What the role cannot supply is the flat curve -- as Grant's usage kept climbing, his efficiency plummeted. The role sets the volume; the player supplies the slope.

Player	Before	After
James Harden	OKC sixth man, 2011-12: 16.8 ppg on 10.1 FGA, 21.6% usage (Sixth Man of the Year)	HOU No. 1 option, 2012-13: 25.9 ppg on 17.1 FGA, 29.0% usage (All-NBA); later 36.1 ppg at 40.5% usage
Shai Gilgeous-Alexander	LAC rookie third guard, 2018-19: 10.8 ppg	OKC franchise player, 2024-25: 32.7 ppg, scoring title, MVP, champion
Jalen Brunson	DAL sixth man, 2020-21: 12.6 ppg; let walk in 2022	NYK No. 1 option, 2022-23: 24.0 ppg; All-NBA the next year
Pascal Siakam	TOR fourth option, 2018-19: 16.9 ppg (champion)	TOR No. 1 after Kawhi left, 2019-20: 22.9 ppg, All-NBA
Brandon Ingram	LAL in LeBron's orbit, 2018-19: 18.3 ppg	NOP primary option, 2019-20: 23.8 ppg, Most Improved, All-Star
Victor Oladipo	OKC beside record-usage Westbrook, 2016-17: 15.9 ppg	IND No. 1 option, 2017-18: 23.1 ppg, All-NBA, steals leader
DeMar DeRozan	SAS, 2020-21: 21.6 ppg, priced as a declining veteran at 31	CHI touch-heavy No. 1, 2021-22: 27.9 ppg, MVP-ballot season at 32
Jerami Grant	DEN role forward, 2019-20: 12.0 ppg in 26.6 min	DET "franchise player" role, 2020-21: 22.3 ppg overnight, at an efficiency cost
Jordan Poole	GSW third guard, 2022-23: 20.4 ppg	WAS No. 1 option, 2023-24: 17.4 ppg at career-low efficiency on a 15-67 team — the role without the slope
Kevin Durant	OKC No. 1 option, 2015-16: 28.2 ppg, 30.6% usage, .634 TS	GSW co-star, 2016-17: 25.1 ppg, 27.8% usage, .651 TS (then a career high)

Player	Before	After
Russell Westbrook	OKC co-star, 2015-16: 23.5 ppg, 31.6% usage	OKC solo engine, 2016-17: 31.6 ppg, record 41.7% usage, MVP; TS unchanged at .554
Chris Bosh	TOR franchise player, 2009-10: 24.0 ppg, 10.8 rpg	MIA third option, 2010-11: 18.7 ppg, 8.3 rpg
DeMarcus Cousins	NOP franchise center, 2017-18: 25.2 ppg (pre-Achilles)	GSW fifth option, 2018-19: 16.3 ppg
Andrew Wiggins	MIN No. 1 option, 2016-17: 23.6 ppg on thin efficiency	GSW fourth option, 2021-22: 17.2 ppg, All-Star starter, champion
Wilt Chamberlain	PHI scoring engine, 1965-66: 33.5 ppg, a seventh straight scoring title, MVP	PHI hub of Hannum's offense, 1966-67: 24.1 ppg on a record 68% FG, 68-13, champion — told to shoot less
Oscar Robertson	CIN's entire offense for a decade: nine straight All-NBA first teams, six assists titles	MIL co-star beside Kareem, 1970-71: 19.4 ppg, 8.2 apg, first title at 32
Moses Malone	HOU one-man franchise, 1981-82: 31.1 ppg, 14.7 rpg, MVP	PHI beside Erving, 1982-83: 24.5 ppg, MVP again, Finals MVP, champion
Steve Nash	DAL co-star, 2003-04: 14.5 ppg, 8.8 apg	PHX engine of D'Antoni's system, 2004-05: 15.5 ppg, 11.5 apg, MVP; team +33 wins

Table 1 | The player held constant, the role changed, 1966 to 2025. Data: [Basketball-Reference](#) and [FiveThirtyEight](#); player pages hyperlinked in the text.

Part 3 -- one trade, two experiments

The cleanest paired observation in the dataset is the 2016 Durant decision, which ran both arms of the trial at once. Durant's usage fell from 30.6% to 27.8% in Golden State's ecosystem, his scoring fell from 28.2 to 25.1 -- and his true shooting *rose* to a then-career-high .651. Westbrook absorbed the vacuum to a record 41.7% usage rate, a 31.6-point triple-double MVP season with his true shooting unchanged at .554 -- his curve was flat too, just at league-average altitude. [FiveThirtyEight](#) documented the divergence side by side (Figure 4): one trade, two directions, and both stat lines moved without either player changing.

	KEVIN DURANT		RUSSELL WESTBROOK	
	2015-16 SEASON	2016-17 SEASON	2015-16 SEASON	2016-17 SEASON
True shooting	63.4%	65.0%	55.4%	54.2%
Usage rate	30.6%	27.6%	31.6%	41.8%
Isolation rate	14.9%	11.4%	12.5%	19.1%
Turnovers/game	3.5	2.3	4.3	5.4

2016-17 stats are through Feb. 8

Figure 4 | [FiveThirtyEight](#)'s side-by-side of the 2016 split -- the isolation-rate and turnover rows are what the standard box score hides. The 2016-17 columns run through Feb. 8; by season's end Westbrook's true shooting had recovered to .554, even with 2015-16, and his usage settled at 41.7%. Source: [Herring, FiveThirtyEight \(2017; archived\)](#).

The compression arm repeats wherever stars join stacked teams: Bosh went from 24.0 points and 10.8 rebounds as Toronto's franchise player to 18.7 and 8.3 as Miami's third option -- and even his rebounding drop was partly the "[Bargnani effect](#)": his Toronto boards were inflated by playing beside one of the worst-rebounding seven-footers ever -- even rebounds depend on who's standing next to you. The same logic decides titles: Dallas's 2011 zone held the reigning MVP to 17.8 points per game, down from 26.7 -- and the same player averaged 28.6 in the next June's Finals and won. In 2017, the "inevitable" Warriors trailed 78-55 with Kawhi Leonard at 26 points in 23 minutes before one closeout ended his series; a sweep followed. And in 2018, a 65-win machine led the conference finals 3-2 when Chris Paul's hamstring went, then missed 27 consecutive threes in Game 7. A lot of that is just variance. Which is the uncomfortable part: the wins and losses we grade players on are laundering hamstrings and coin flips.

It cuts the other way too. Even the most stacked constructor of the era kept nearly losing. The season before Durant arrived, the 73-9 Warriors -- the best regular-season machine ever assembled -- blew a 3-1 Finals lead, with Green suspended for Game 5. With Durant, they were one hamstring and twenty-seven missed threes from losing that 2018 series. A year later the dynasty ended in the Finals when Durant's Achilles and Thompson's ACL tore days apart. In 2021, Durant's Nets superteam went out in a second-round Game 7 by

the width of his sneaker on the three-point line. None of this is new: the 1969 Lakers assembled Chamberlain, West, and Baylor and lost Game 7 to Boston anyway -- West is still the only Finals MVP from the losing team. Seven games between near-equal machines comes down to tendons, suspensions, and shooting luck. Then we turn those margins into permanent rankings.

Part 4 -- where the thesis breaks

The strongest counterargument first. Lopez, Matthews & Baumer (2018) modeled a decade of games across the four major leagues and found the NBA has the **largest talent dispersion and the least game-level randomness**: the best NBA team at home is essentially always a 60%+ favorite -- a determinism baseball and hockey never approach. And basketball's individual effects are the largest in team sports: when LeBron left an otherwise largely intact Cleveland roster, the team fell 61-21 to 19-63, the worst single-season collapse in NBA history; Boston's Garnett-and-Allen summer produced the +42 mirror image. And this is not a modern artifact -- Figure 5 plots every franchise's season-over-season swing from 1950 through 2025, and swings of this size trace to single arrivals and departures in every era.




Figure 5 | Forty-win swings are as old as the league. Every franchise's season-over-season win change per 82 games, 1950-2025 (1,620 team seasons; under-50-game seasons excluded). The extremes trace to single arrivals and departures in every era. Computed from Basketball-Reference team records; shortened seasons scaled to 82 games.

I think this sharpens the thesis instead of killing it. Team strength being persistent is exactly what a constructor-dominant sport looks like; the stable, legible unit is the machine. And look at who actually produces the forty-win swings: players whose signature skills are the *spillover* kind. They tend to be systems themselves. This is what Ben Taylor formalized as “portability” in Thinking Basketball: shooting, rim gravity, screening, defense, and low-maintenance playmaking compound as teammates improve, while high-volume on-ball creation travels worst, because better teams have less of that one resource to spare. The other side of the boundary: Jordan Poole scored 20.4 as Golden State's third guard, was handed a franchise's worth of possessions in Washington, and his scoring went *down*, at career-low efficiency, on a 15-67 team. The role handed him all the volume he wanted. The efficiency underneath wasn't there.

Part 5 -- two chairs, opposite trades

In the fantasy chair, you are long the box score. The contaminated measurement is the commodity, so context is the whole edge. The industry consensus is that projected minutes and usage beat talent priors, and professional DFS tooling treats minutes as the single most important column, with usage as the multiplier. Possessions are conserved, so vacated usage is the closest thing fantasy has to a law of physics -- every

breakout in Part 2 was visible months early on the depth chart. The caveat (the industry says it themselves): opportunity is necessary, not sufficient. top-50 minutes with no curve underneath converts to nothing.

In the GM's chair, you are long wins, so the same effects flip into contamination. Ask what a player averages at the usage *your* system can actually allocate, not what he averaged where possessions were free. Buy the spillovers the salary market provably doesn't price. Expect trades to underdeliver -- of 47 significant midseason trades across five seasons, only 4 measurably improved the team -- and budget for integration: accumulated shared experience is a productive asset that a trade destroys on both ends. And price the scheme, not coaching reputation: only 14 of 62 coaches studied measurably moved player performance; Nash-into-D'Antoni turned 29-53 into 62-20 while most coaching changes move nothing.

One exception, at full strength: for the one or two constructor-grade players per era, "acquire at nearly any price" is correct -- the ± 42 -win swings belong to the players who carry the context with them. The thesis governs roster spots two through ten; spot one belongs to the exceptions. And the inversion is the fun part: fantasy managers and GMs are trading the *same mispricing from opposite sides* (fantasy is the only place the "empty stats" guy is priced right). In my opinion a mispricing that funds profitable strategies on both sides is not a hot take. It is structure.

Closing

Formula 1 gives out two trophies, driver and constructor, and keeps the ledgers separate. Basketball gives out both too (the MVP is the drivers' title; the championship is basically the constructors' title) and then merges the ledgers anyway. We count rings on fingers, turn the constructor award into a driver stat, and argue about drivers for decades. None of this means stars don't matter. Part 4 says the opposite. It means most of what we argue about (stat lines, trade grades, legacies) gets measured inside a context the player didn't pick and mostly can't take with him.

Once you see it, a lot of loud debates get quiet. Stop asking what a guy averaged. Ask what the role allowed, what the curve underneath looks like, and which of his skills survive a move. Same questions whether you are setting a lineup or running a front office. The box score was hiding all three.