GAME 4 2026 NBA FINALS

How the Spurs conceded the largest comeback in NBA Finals history

Quantum Sport Analytics

Marin Gautier, Hugo Stubler, Idriss Ben Mrad, Victor Helpiquet

August 2026

Contents

Chapter 1

Introduction

San Antonio led by 29 points. They lost.

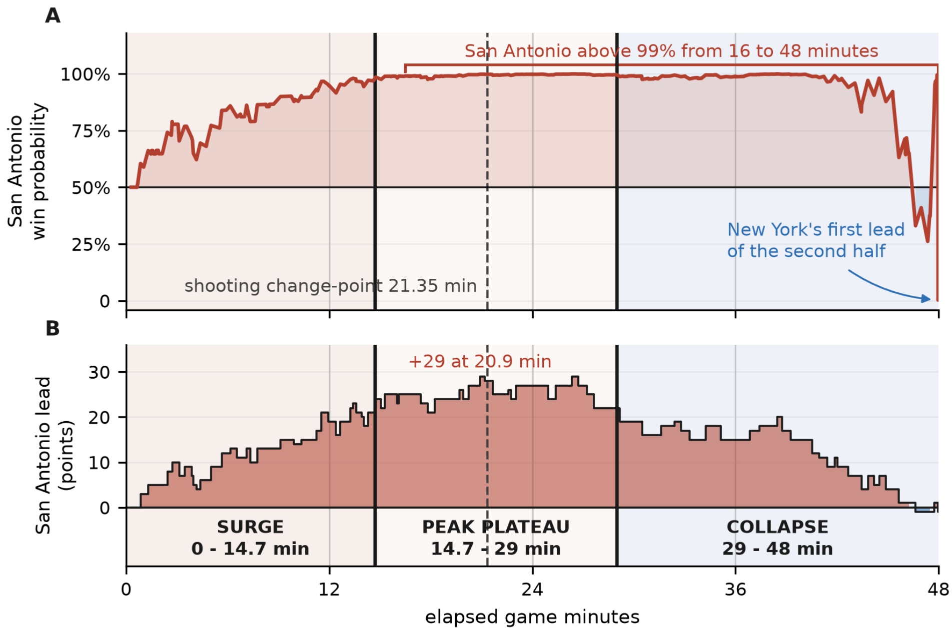

Figure 1.1: Official in-game win probability, and the San Antonio lead, minute by minute. The win probability is the classic model, driven by the score margin and the time remaining.

Figure 1.1 is the whole game on one line. San Antonio built a lead of 29 points by half-time (76–49), and from the sixteenth minute their win probability sat above 99 percent. Nobody comes back from that: no team in the history of the Finals had ever surrendered a lead that large. The probability did not fall below 50 percent until the final minute, and even then the finish contained a double reversal. San Antonio retook the lead 106–105 inside the last fifteen seconds; New York scored the winning basket with two seconds left. Final score 106–107. A game that was decided for three quarters was decided again twice in fourteen seconds.

A 29-point lead does not vanish at this level. So the question is not *what* happened, the box score says that much, but *why*: what broke in San Antonio's game, and what New York did to break it. This report is that analysis, built on player tracking.

The data. We built a computer-vision tracking pipeline that follows every player through the broadcast video (Fig. 1.2) and merged its output with the official NBA play-by-play. Every position on the floor is therefore tied to the official record of what happened.

Every metric is defined in one line where it first appears, and in full, with its formula, in the Glossary at the end.

Figure 1.2: The tracking pipeline, drawn back onto the broadcast.

1.1 Momentum

We cut the game into three phases, read straight from the scoreboard: the surge, the plateau, the collapse. The cut points are fixed once and used for every number that follows:

SURGE	0 – 14.7 minutes elapsed
PEAK PLATEAU	14.7 – 29.0
COLLAPSE	29.0 – 48.0

The data. Every unanswered run of six points or more, and the scoring inside each phase (Fig. 1.3). A run is a string of points by one team without an answer from the other, free throws included.

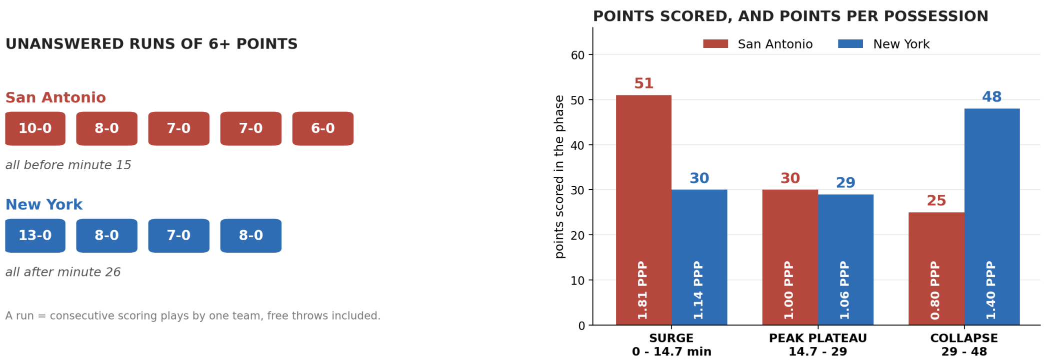

Figure 1.3: The unanswered runs, and the scoring inside each phase.

On the floor. Every San Antonio run comes before the break, every New York run after it. In the first phase San Antonio scored 1.81 points per possession and New York 1.14. In the last phase the numbers flip: 0.80 for San Antonio, 1.40 for New York. In between, the plateau is dead even, 30 to 29. The two teams traded places, and it happened at a precise moment.

The question. That date is the problem this report sets out to solve. A 29-point lead does not change hands by accident: something broke in San Antonio's game, and something in New York's changed to break it. The answer comes in two steps. Chapter 2 measures *what* broke at that date, in the performance metrics any analyst can compute. Chapter 3 then follows every player on the floor to show *how* it broke, and *why*.

Chapter 2

Performance Metrics: the Reflection of the Spurs Collapse

First, what broke. At the date of the swap, three performance metrics break in the same window and in the same direction: shot-making against shot quality, possession value, and turnovers. Between them they account for the result. This chapter establishes what collapsed; the next asks how and why.

2.1 The metrics used here

Four numbers carry this chapter (Fig. 2.1). Full derivations are in the Glossary.

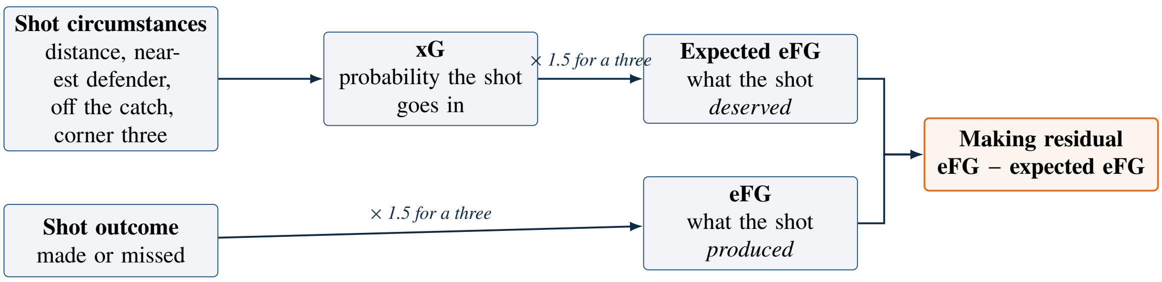

Figure 2.1: How the shooting metrics relate. Shot quality is what the shot deserved; shot-making is the gap between what it produced and what it deserved.

eFG (effective field goal percentage). Shooting percentage, with a made three counted as 1.5 makes because it is worth 1.5 times a two.

$$\mathrm{eFG} = \frac{\mathrm{FGM} + 0.5 \times 3\mathrm{PM}}{\mathrm{FGA}}$$

with FGM the shots made, 3PM the threes made, FGA the shots attempted. Example: 10 of 20, all twos, eFG = .500; 10 of 20, all threes, eFG = .750.

xG (expected goals), and expected eFG. The chance a shot goes in, judged on the situation only (distance, nearest defender, off the catch, corner three), never on the result. It is a statistical model of the kind used by Chang et al. (2014) [2]; the model and its coefficients are in the Glossary. Expected eFG weights it like eFG.

$$\mathrm{xG}_i = p\left(\mathrm{made} \mid \mathrm{situation}_i\right)$$

$$\mathrm{xeFG} = \frac{1}{\mathrm{FGA}} \sum_{i=1}^{\mathrm{FGA}} w_i\, \mathrm{xG}_i, \qquad w_i = \begin{cases} 1 & \text{two-point attempt} \\ 1.5 & \text{three-point attempt} \end{cases}$$

eFG is what the shots produced; expected eFG is what they deserved.

Shot-making. The number the whole report turns on: the gap between what the shots produced and what they deserved (in the Glossary: the making residual).

making residual = eFG − expected eFG

If positive: the team made more than its shots deserved. **If negative: fewer.**

Over the 34 second-half shots, that gap is worth about **ten shots that should have gone in and did not.**

EPV (expected possession value), and EPVA (EPV-Added). EPV is what the value of the possession at any instant, in points: about 1.0 with the ball at half-court, nearly 2.0 with an open man at the rim, 0 once the ball is lost. EPV-Added is how much a player raised or lowered that value while he had the ball:

$$\text{EPVA} = \sum_{\text{his touches}} \big[\text{EPV}(\text{release}) - \text{EPV}(\text{receive})\big] = \textsc{build} + \textsc{finish}$$

BUILD: value added while carrying or passing, whatever the shot does.

FINISH: the last act. A make cashes the value in; a miss or a turnover takes it to zero.

Example. Drive from the arc (1.0) to the rim (1.7), miss (0): BUILD +0.7, FINISH −1.7, EPVA −1.0.

2.2 Shot quality against shot-making

The phenomenon. A team's shooting efficiency can fall in only two ways:

1. **The shots get harder.** Shot quality falls: further out, more contested, off the dribble. This is about what the defence allows.

2. **The same shots stop going in.** Shot-making falls. This is about the shooter.

San Antonio's collapse is entirely the second: **same shots, stopped making them.**

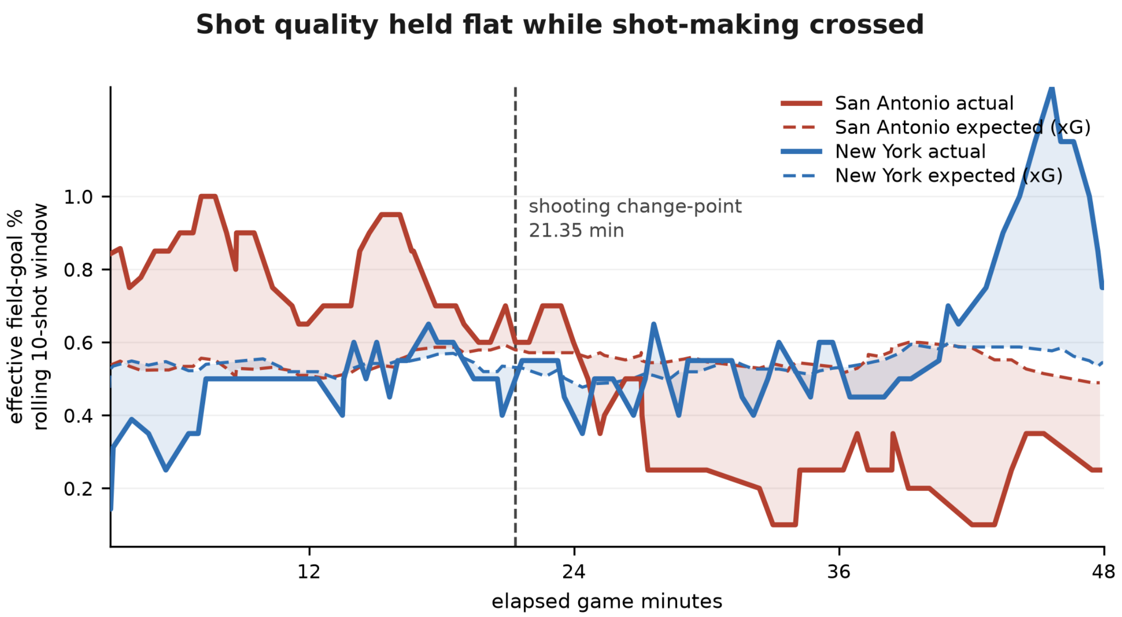

Figure 2.2: Shot quality held flat while shot-making crossed.

Figure 2.2 is the picture of the game. The dashed lines are shot quality (expected eFG): flat for both teams, all night. The solid lines are shot-making: San Antonio's starts far above what its shots deserved and ends far below it. In the first half they made seven shots more than their looks deserved; in the second, ten fewer (eFG .745, then .244). New York's does the opposite (.459 to .646), and the two lines cross just after half-time.

On the floor.

- **San Antonio's make rate: .650 before minute 21, .217 after.**
- **Largest decliner: Vassell,** not Wembanyama (Fig. 2.3). It is not that the wrong men were shooting.

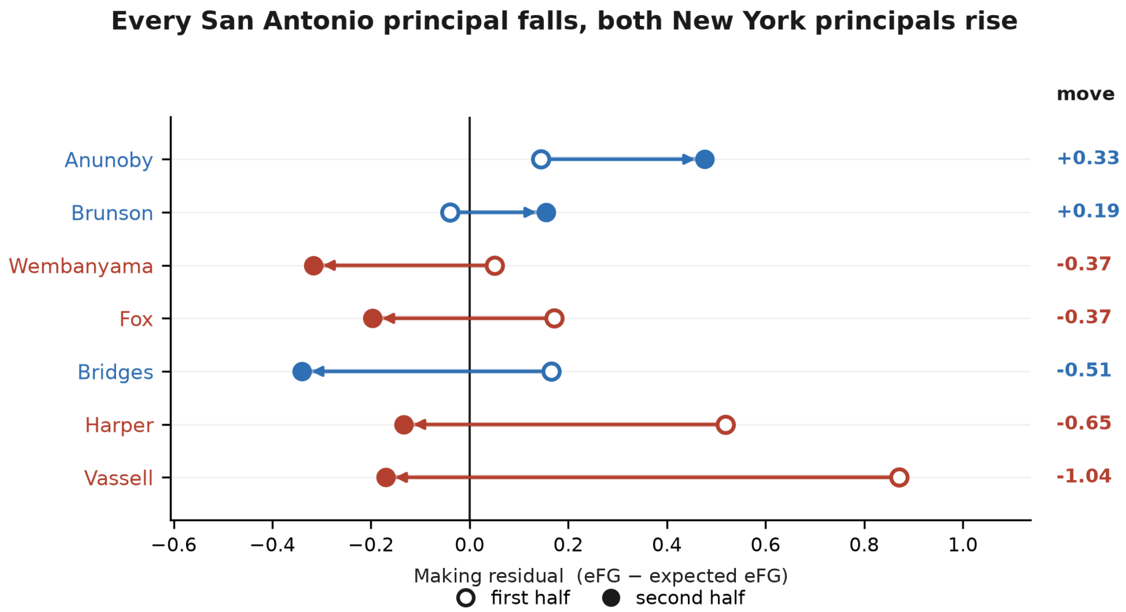

Figure 2.3: Shot-making (eFG minus expected eFG), player by player.

2.3 Possession value (EPV-Added)

The phenomenon. Measured in points added to the possessions they were involved in, the main men on each side trade places between the first phase and the last.

The data. EPV-Added by phase (Fig. 2.4), and the split into the two channels (Table 2.1).

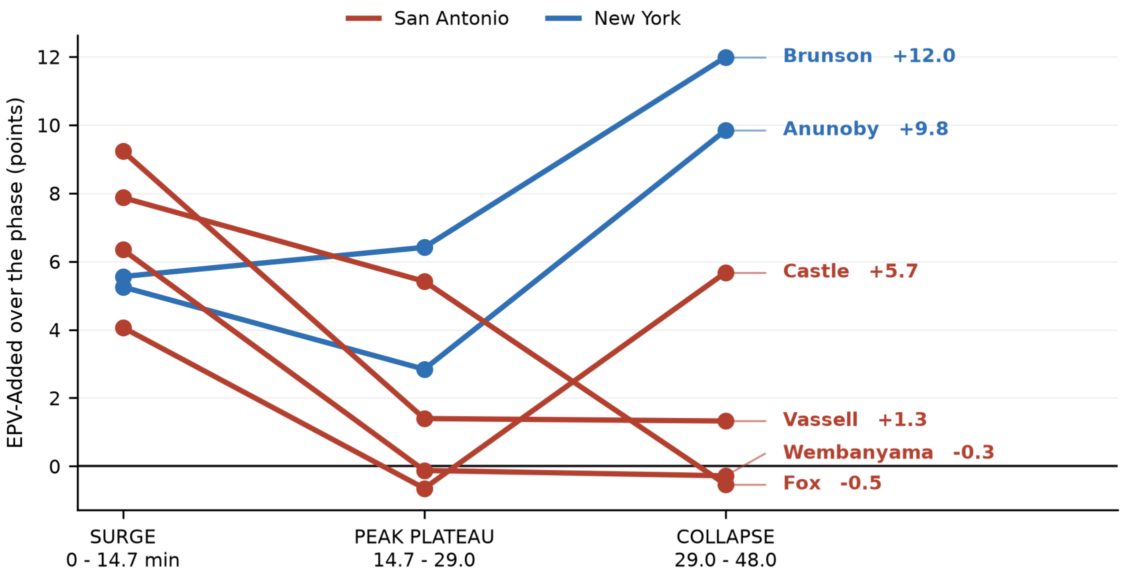

Figure 2.4: EPV-Added by phase.

Table 2.1: EPV-Added split into its two channels, first half → second half.

Player	BUILD	FINISH
Fox	5.62 → **6.09**	5.24 → **−4.23**
Wembanyama	0.22 → **1.40**	6.43 → **−0.38**
Vassell	0.55 → **1.03**	8.95 → **−0.95**
Harper	4.21 → **2.00**	8.08 → **−1.98**
Castle	4.09 → **3.86**	0.15 → **+1.70**

Green: up from the first half. Red: down.

On the floor.

- **Read the lines:** the Spurs start high and slide toward zero, Castle the lone exception; the two Knicks (Brunson, Anunoby) start low and climb. By the collapse, the players adding the most value are Knicks.
- **Wembanyama turns negative** in the collapse: the possessions that ended in his hands were worth less than when he got the ball, because the shots missed.
- **BUILD up, FINISH down** for Fox, Wembanyama and Vassell (Table 2.1). The machinery that created value kept working; the last act stopped paying.
- **A second way of measuring the same thing:** the value was still being created, and it stopped being cashed in.

2.4 Turnovers

The phenomenon. San Antonio did not turn the ball over at all for seventeen minutes, then gave it away twelve times.

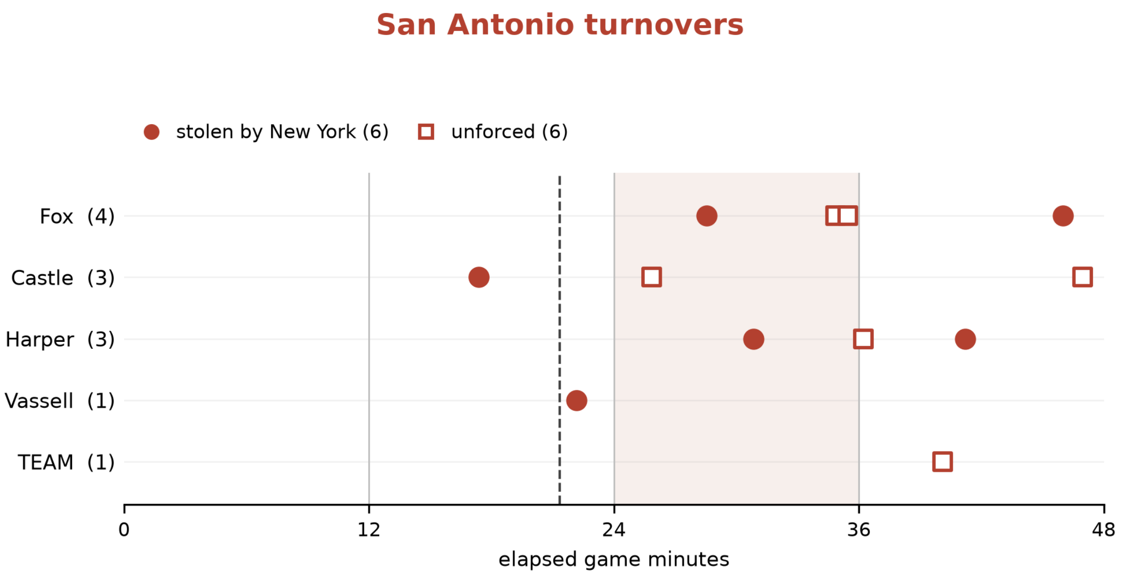

Figure 2.5: The twelve turnovers on the clock.

On the floor. Zero turnovers in the first quarter, two in the second, five in each of the third and fourth (Fox 4, Castle 3, Harper 3, Vassell 1, one shot-clock violation). Half are live-ball steals by New York, half are unforced. The rate peaks in the third quarter, the same window as the shooting break. Each turnover throws away a possession worth about 0.63 points on average, so twelve of them cost roughly seven and a half points, and each handed New York the ball moving forward. We do not say the turnovers caused the collapse; they are its most visible symptom.

What Explains the Remontada?

3.1 What the tracking data can decode

The shot-quality model of Chapter 2 says the shots did not change. So the cause of the collapse lies in everything that model does not read, and everything the tracking does.

What the shot-quality model sees: distance, nearest defender, off the catch, corner three. **Unchanged between halves.**

What only the tracking sees

one section each

- What New York changed, offensively and defensively §3.2
- Who carried the physical load, and who wore down §3.3
- Where the ball went before the shot §3.4
- Where the players stood at the shot §3.5
- How the contest at the rim changed §3.6

Verdict: this is where the collapse lives.

The next sections take these five one at a time. The answers come straight from measurement: where people stood, how far they ran, where the ball went, how close the defender was. Together, the five sections tell one story, and the Synthesis (Chapter 4) tells it without the numbers.

3.2 Towns, the key to New York's new defence and offence

3.2.1 The Towns effect

The phenomenon. One substitution changed the game. When Karl-Anthony Towns came on against Wembanyama, Wembanyama was pulled away from the rim, and New York's shooting efficiency nearly doubled.

The data. New York's shots with Wembanyama on the floor, with and without Towns (Fig. 3.1), and where Wembanyama stood in each case (Fig. 3.2).

	TOWNS OFF	TOWNS ON	
New York actual eFG	.400	.774	*nearly doubles*
Shot quality (expected eFG)	.453	.453	*identical*
Wembanyama's distance from rim	4.39 m	4.60 m	*+21 cm*
Shooter's nearest defender	1.81 m	2.03 m	*+22 cm*
New York offensive spread	4.65 m	5.05 m	*+40 cm*

New York shots with Wembanyama on the floor, split on whether Towns was also on.

Figure 3.1: New York with Wembanyama on the floor, Towns off against Towns on.

On the floor.

- **Same shots.** Shot quality is identical (.453 in both cases): New York did not get easier shots with Towns on.
- **Going in.** Actual eFG goes from .400 to .774, and this holds in both halves separately.
- **Wembanyama moved.** With Towns on the floor he defends 21 cm further from the rim (Fig. 3.2); the shooter has 22 cm more room, the offence is spread 40 cm wider. Small moves, and none of them nearly enough to explain a doubling of conversion.

> **The Towns effect is a conversion effect, not a shot-creation effect: the same shots, going in.**

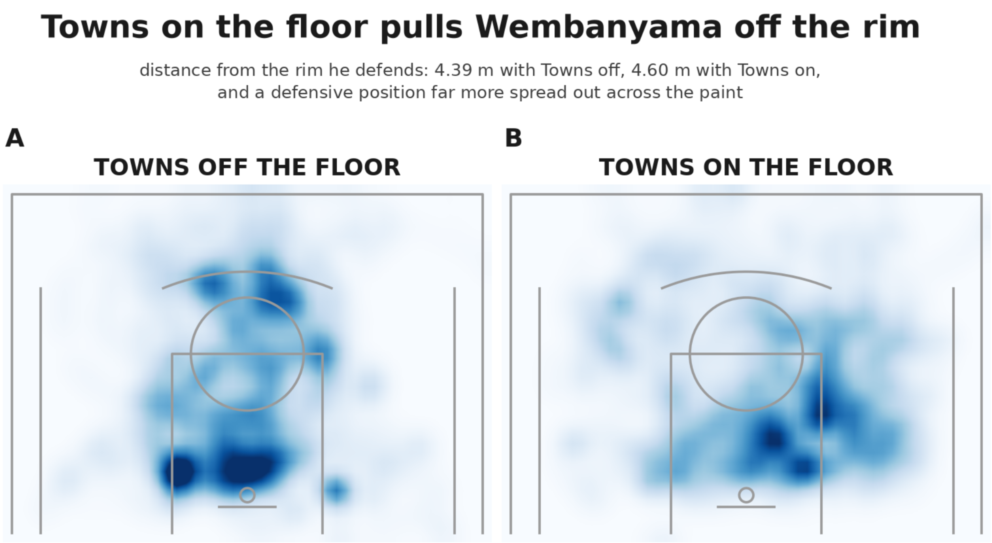

Figure 3.2: Where Wembanyama defended, with and without Towns: further from the rim, and far more spread out.

From an empty paint to a new way of attacking. The chain is short:

1. **Towns pulls Wembanyama out.** A centre who shoots threes has to be followed to the arc; the rim protector leaves the paint.

2. **The whole defence follows.** The other four Spurs step out with him to stay connected, and the paint is left empty.

3. **New York changes its attack.** With nobody at the basket, the Knicks stop settling for the perimeter and go through the middle with the pick-and-roll: Towns blocks the path of the man guarding the ball-handler (the *pick*), then cuts to the basket (the *roll*); the ball-handler drives into the space, or feeds Towns rolling to the rim (Fig. 3.3).

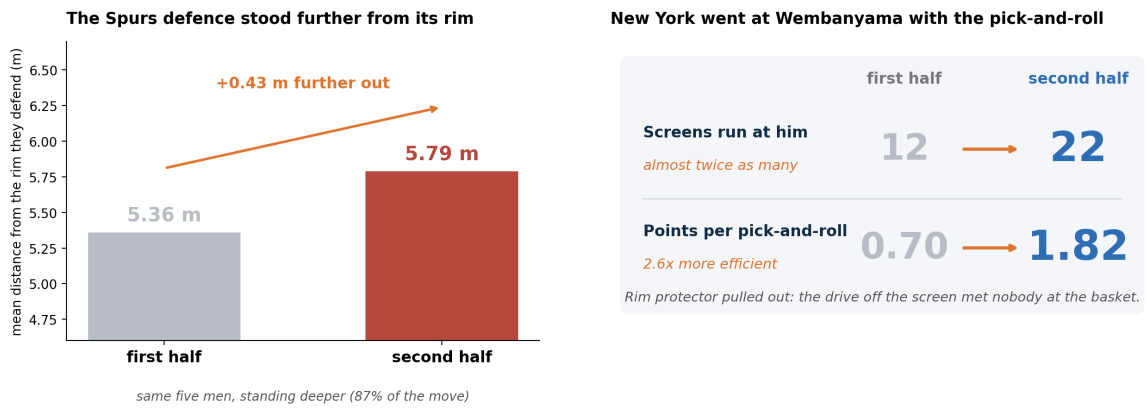

Figure 3.3: The San Antonio defence stood further out, and New York punished it with the pick-and-roll at Wembanyama.

- **Further out.** The Spurs defended 0.43 m further from their rim in the second half, and that is the same five men standing deeper, not a change of personnel.

- **More picks at Wembanyama.** A ball screen forces the rim protector to step out to the ball-handler or concede a jump shot. New York ran it at him 12 times in the first half, 22 in the second.

- **Far more efficient penetration.** Each pick-and-roll returned 0.70 points in the first half and 1.82 in the second: with Wembanyama pulled out, the drive off the screen met nobody at the basket.

Towns pulled Wembanyama out of the paint, the Spurs' defence followed him, and the way to the rim was open: New York switched to penetration through the pick-and-roll, ran it twice as often and scored 2.6 times more on every trip (0.70 to 1.82 points).

3.2.2 The defence on Wembanyama

The phenomenon. New York handed the Wembanyama assignment to Towns, and concentrated it on far fewer men.

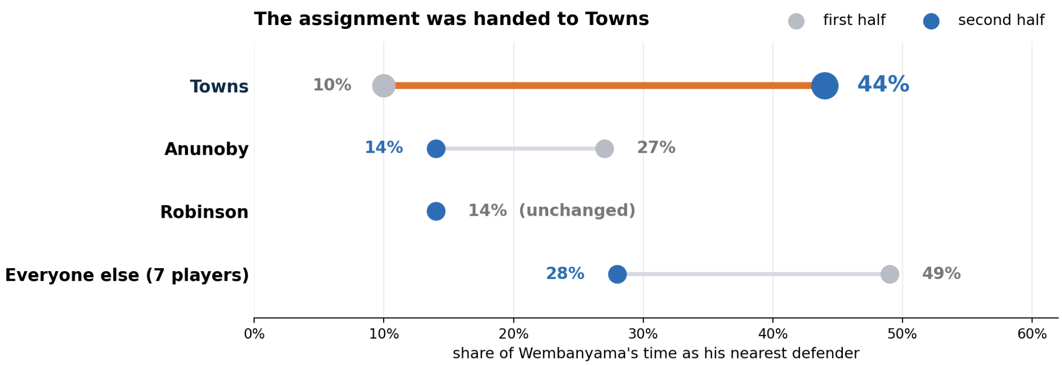

Figure 3.4: Who guarded Wembanyama, first half to second.

On the floor.

- **Towns takes over:** he was the closest man to Wembanyama 10 percent of the time in the first half, 44 percent in the second. The top three defenders go from 53 to 71 percent of his time.
- **Same pressure.** His nearest defender stood no closer, and the help arrived no more often: New York changed *who*, not *how hard*.
- **Anunoby freed up** for the other job, on Fox (Section 3.4.1).

3.3 The Spurs fatigue

Not an event in the game, but the condition underneath it: how the physical work was shared out. The rotation was set before the collapse and did not change during it.

Relative load. How much physical work a player has done compared with the men on the floor with him at the same moment. 1.0 means he has worked as hard as the average man out there; 1.6 means sixty percent harder. The load counts contact and holding position as well as running, which is why a rim protector can be the most loaded man on the floor while covering less ground than a guard.

How it is computed. From the tracking, we add up seven kinds of physical work for every player, twenty-five times a second: running, braking, contact, changing direction, holding position against a body, and carrying the ball.

Weights and formulas in the Glossary.

3.3.1 Who carried the load

The phenomenon. San Antonio played a short rotation, concentrated the work on few men, and their best player carried more than anyone on the floor.

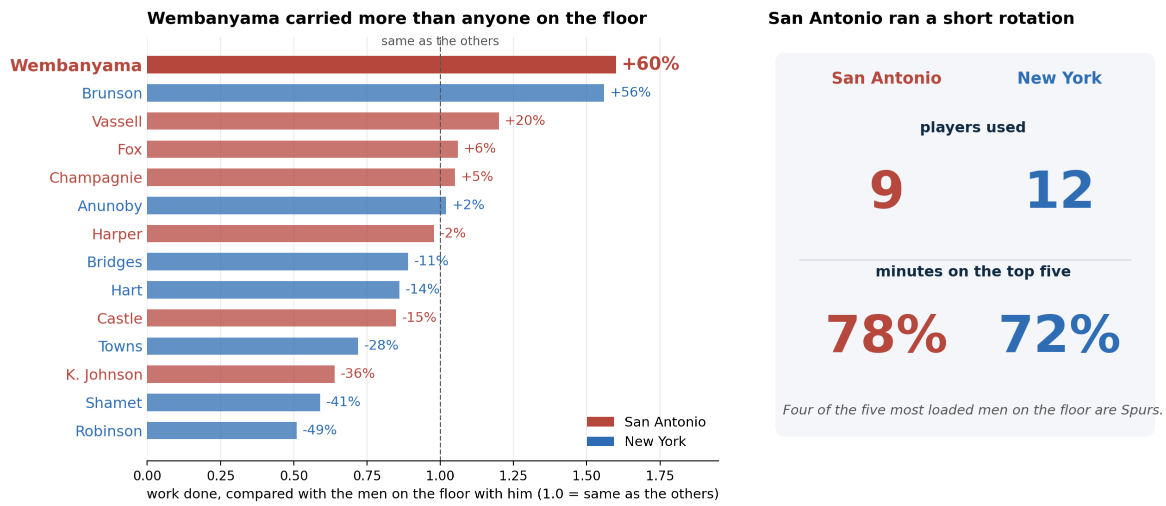

Figure 3.5: Second-half load, player by player, and the size of each rotation.

On the floor.

- **Wembanyama worked 60 percent harder than the men on the floor with him,** and four of the five hardest-working players are Spurs. Not only running: contact and holding position put him above every guard while covering less ground.
- **Nine men, three quarters of the minutes on five of them.** New York spread the same game across twelve: a deeper bench, and it paid off. The Knicks' starters came into the fourth quarter fresher than the men they were playing against.
- **Wembanyama's 43.9 minutes, against a season average of 29.2,** more total work than anyone in the game, a final unbroken shift of 16.6 minutes. He did not slow down; the fatigue lands on the last act: touch, balance, finishing.
- **The second-half load came from holding the ball.** In the first half the work was movement without it, cuts and screens that created the opening before the pass arrived. In the second, the tired men kept the ball in their hands and decided instead of reacting: more of the work was done with the ball in hand, less of it moving without it. The next section puts that switch on the clock.

3.3.2 Catch-and-shoot

The phenomenon. From the legs to the hands. The load above is the condition; here is where it shows. In the first half the ball barely stopped: it arrived and went up. In the second, on the same possessions, the same players held it and thought before releasing. The physical work did not change the shots San Antonio got; it changed how long they took to take them.

The data. Shot rhythm, measured from the ball's position (Fig. 3.6).

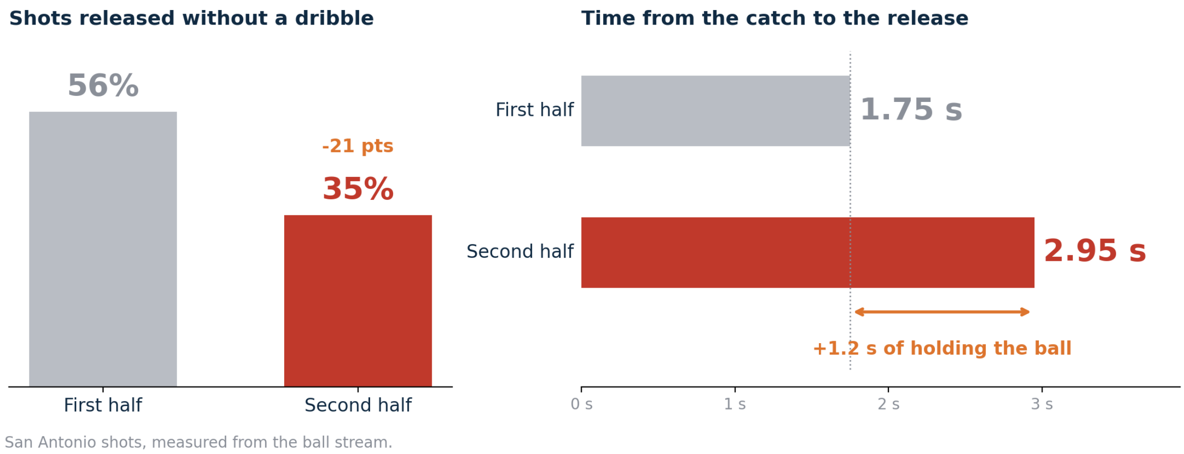

Figure 3.6: San Antonio's shot rhythm, first half to second: fewer shots off the catch, and 1.2 seconds more with the ball in hand before each release.

On the floor. The same shots from the same places, taken 1.2 seconds later by a player who has to decide rather than react. In the first half the ball arrived and went up; in the second it arrived and was considered. A shot-quality model grades where the shot was taken and who was near; neither changed, so it sees nothing. Tracking data goes a level deeper, to the second before the release, and finds one of the clearest reasons the same shots stopped going in.

3.4 Fox's physical collapse

Section 3.3 ended on a tired team holding the ball longer. The man holding it most was De'Aaron Fox, San Antonio's point guard. Two things hit him at the same time. He was already tired, and he now had the ball in his hands longer on every possession. And New York put its best defender, OG Anunoby, on him: taller, stronger and closer, so every dribble took more effort and every drive met a body. This section looks at who guarded him, how much less he moved, and why the ball still kept coming to him.

3.4.1 The Anunoby assignment

The phenomenon. New York's first tactical move was to put OG Anunoby on De'Aaron Fox: tried in the first quarter, made permanent from the end of the second.

The data. Who was Fox's nearest defender, on the clock (Fig. 3.7) and over the halves (Fig. 3.8).

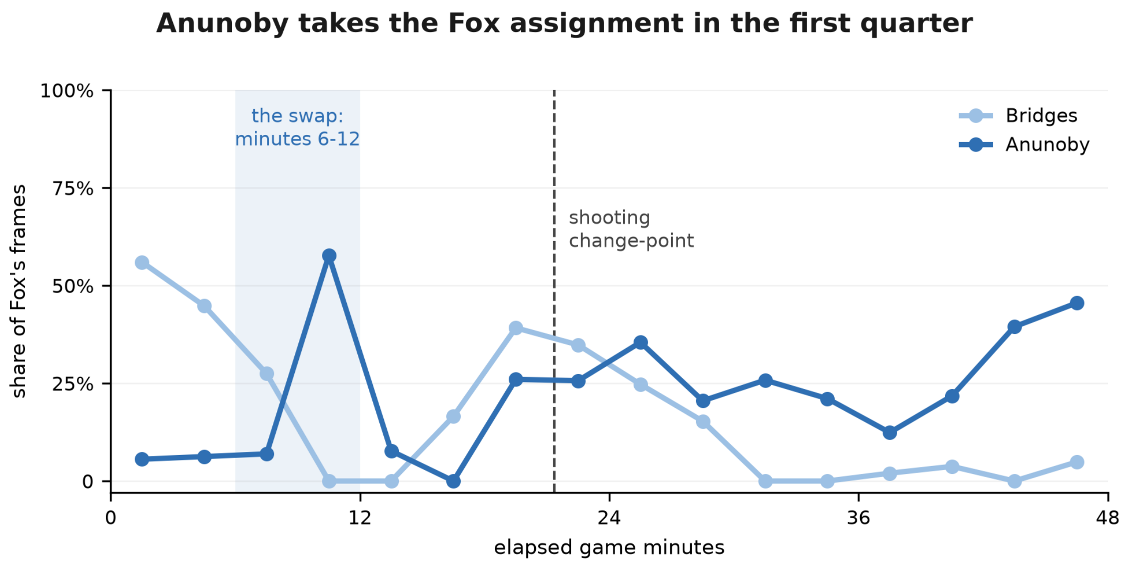

Figure 3.7: The Anunoby switch onto Fox: a first trial in the first quarter, then the permanent handover from the end of the second.

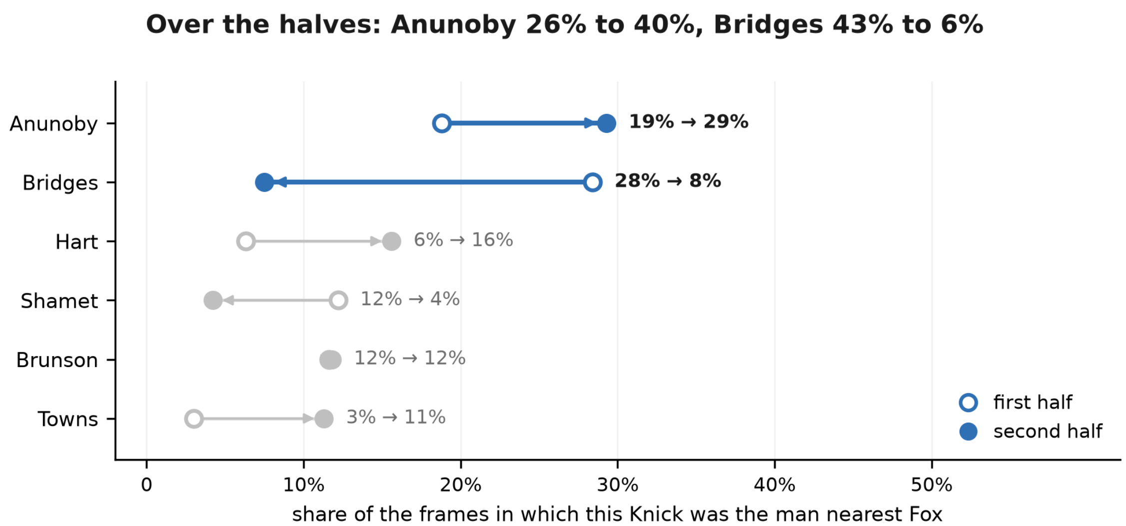

Figure 3.8: Who guarded Fox, half by half.

On the floor.

- **A first trial, minutes 6 to 12:** Anunoby's share of the time as Fox's nearest defender jumps to 58 percent, then falls back to zero. Bridges is still the man on the next stints.
- **The real switch, from minute 20:** Anunoby overtakes Bridges at the end of the second quarter, just before the shooting broke down, and never gives it back. Bridges is at zero from minute 30 on. Over the halves: Anunoby 26 to 40 percent, Bridges 43 to 6.
- **A standing condition** for everything after it: from the end of the second quarter, San Antonio's primary ball-handler is worked by New York's best perimeter defender rather than their second-best, on every possession.

The cost of that assignment shows up in Fox's legs first.

3.4.2 Fox's decline

The phenomenon. Fox is the player who lost the most between the two halves: less ground covered, lower speed, more than anyone else on the floor.

The data. The change from the first half to the second, for Fox and for the rest of the rotation, in distance covered per minute and in mean speed (Fig. 3.9).

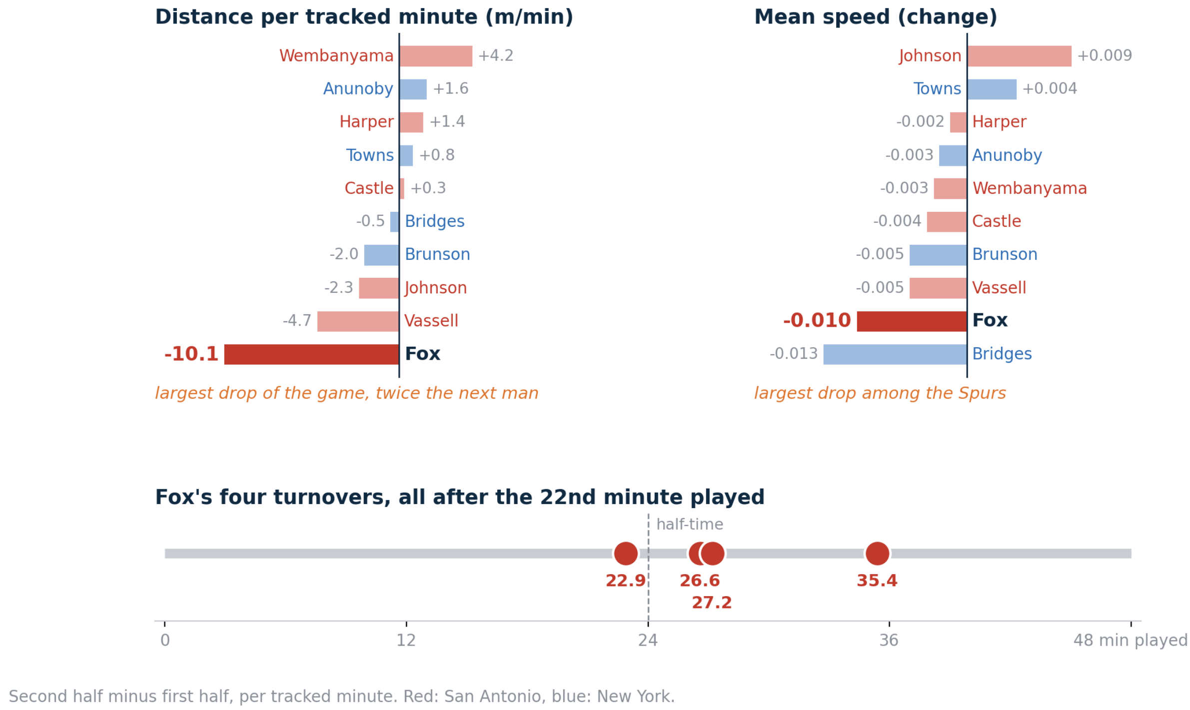

Figure 3.9: Fox's physical decline, first half to second: the largest drop in distance of the whole rotation, the largest drop in speed among the Spurs, and four turnovers once the decline had set in.

On the floor. Ten fewer metres in every minute he played, twice the drop of the next man, at lower speed. He stopped covering ground. For the player who initiates the offence, that is the difference between a defence that has to move and one that can stand still.

Locked in, so the ball stayed in his hands. While Fox was declining, the ball did not leave him: his share of the time San Antonio spent carrying it *rose*, from 24 to 37 percent. This is what a good defence does to a tired point guard. Anunoby took away the drive, the rest of New York stayed home on the shooters, and with less movement around him there was less to pass to. So he kept the ball, and kept it longer, on possessions that were no longer being built by five men. The next section measures how much of the collective was left.

3.4.3 Collective construction

The phenomenon. San Antonio kept generating shots from the same places, but through a far less collective process. Two ways of judging the same shot:

- **Shot quality: still good.** From where, and how open? Much the same positions, and much the same space, as in the first half.

- **How the shot was built: no longer good enough.** How many hands touched the ball, how far it travelled, how many passes, was the shooter set? On every one of these, the second half fell away.

The data. Every construction measure (Fig. 3.10), and the passing map, built on the passes that lead to a shot, made or missed [9] (Fig. 3.11). Passes before a shot, not assists: an assist is only recorded on a made shot, so an assist network thins automatically when a team stops scoring.

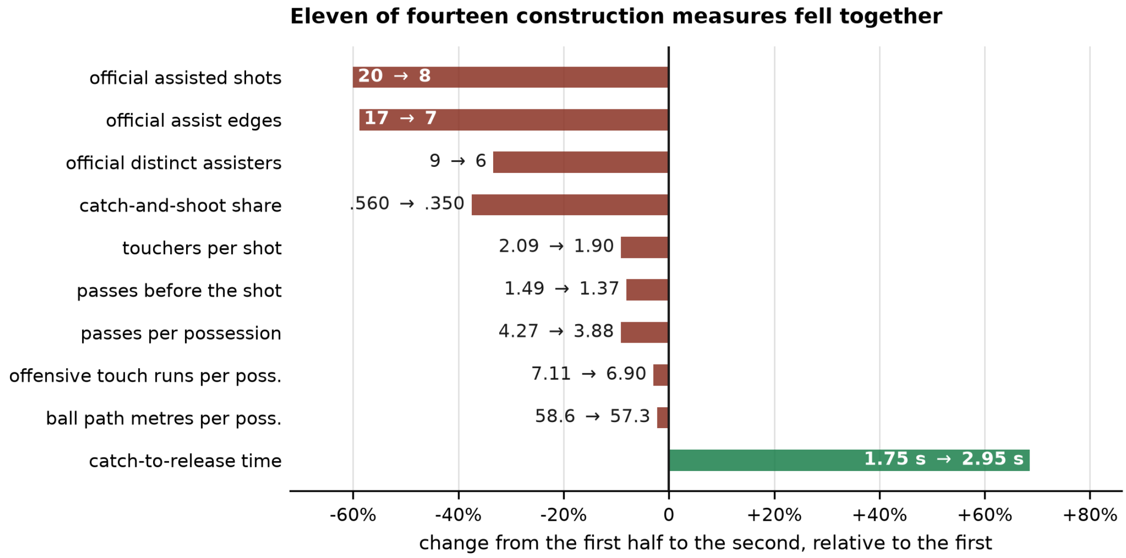

Figure 3.10: The construction measures, first half to second.

On the floor. Almost every construction measure went down together. San Antonio went from a well-distributed offence, where the ball passed through many hands along many routes, to something more static: fewer players involved, fewer options on each possession, the shooter less often set when the ball arrived. The passing map below shows the same thing in one picture: the same number of passes, but among fewer people along fewer routes. The ball kept moving; it stopped circulating.

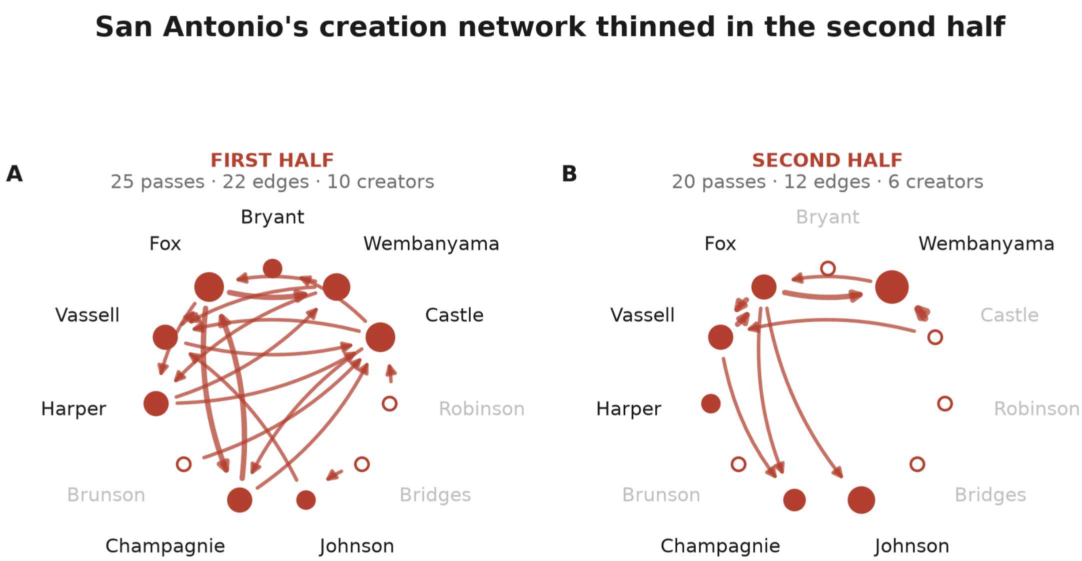

Figure 3.11: San Antonio's passing map: the passes that led to a shot, first half and second.

3.5 Wemby's shot concentration

3.5.1 Shot share against carry share

The phenomenon. The ball circulated less, the creators thinned out, and the shots had to go somewhere. They went to Wembanyama: his share of his team's shooting rose by half again, while the ball itself spent *less* time in his hands. With Fox keeping the ball longer, Wembanyama was no longer part of the construction, only of the finish: the ball reached him late, to shoot, not to create.

The data. His shots against his passes (Fig. 3.12).

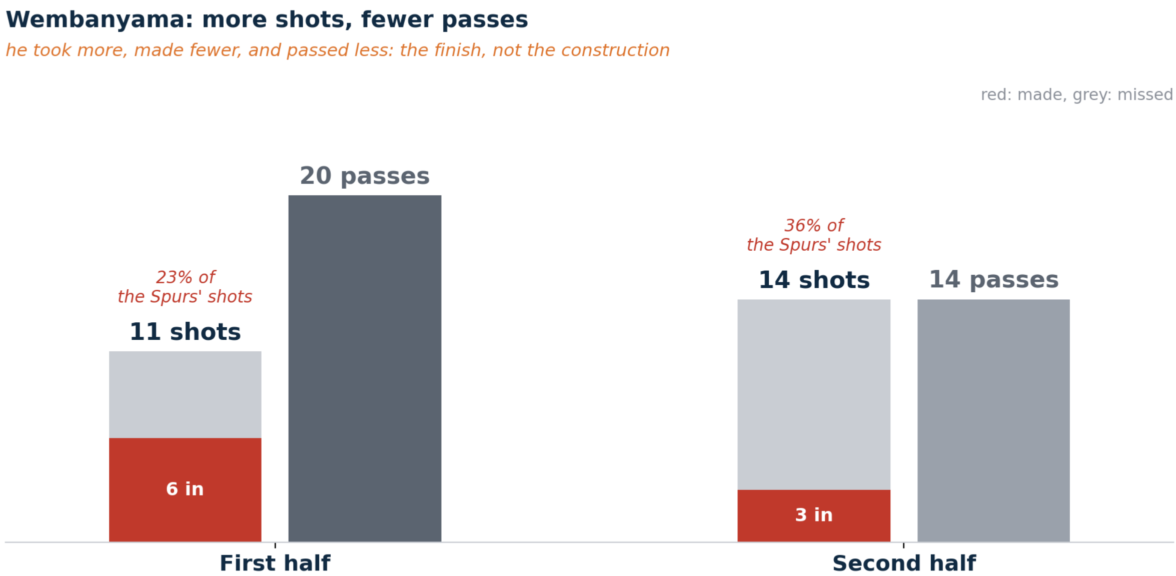

Figure 3.12: Wembanyama, first half to second: more shots, a larger share of his team's shooting, and fewer passes.

On the floor. More shots, fewer passes: the ball reached him to be shot, not to build with. Where the extra shots came from is the next question.

3.5.2 The withdrawal around him

The phenomenon. Wembanyama shot more because the rest of the team shot much less. This is not only an offence that chose to run through its best player; it is the men around him stepping back from their share of the shooting and leaving it to him.

The data. Every Spur's shot attempts, first half to second (Fig. 3.13).

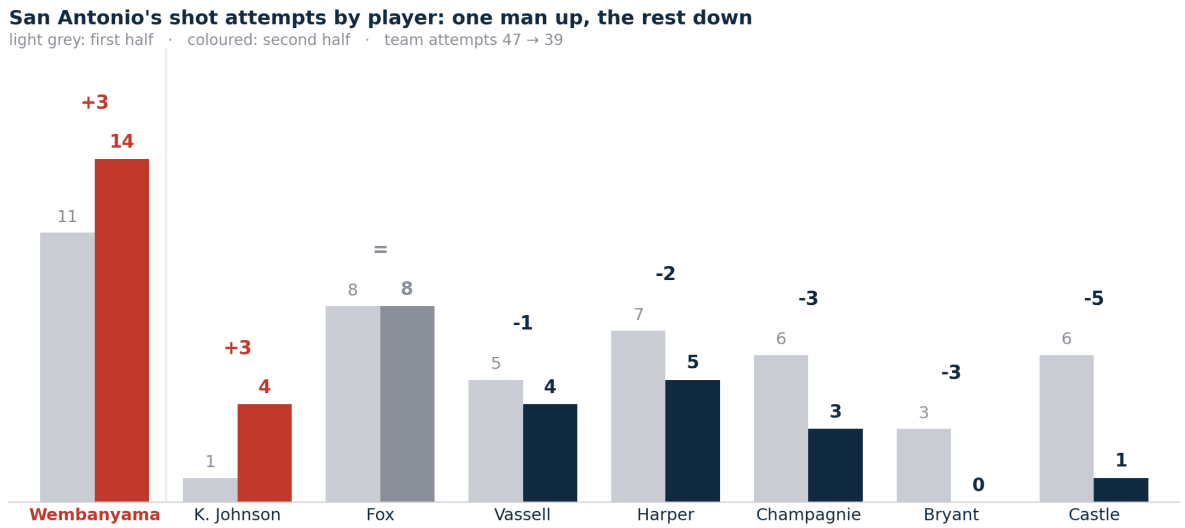

Figure 3.13: San Antonio's shot attempts by player, first half to second: Wembanyama and K. Johnson up, everyone else down.

On the floor. San Antonio reorganised around Wembanyama in the paint, and not by clearing space for him: the team closed in on him. Six of the eight Spurs who shot in the first half shot less in the second, and the tracking shows the whole team drawing in:

- **The team moved in:** on average, every Spur stood 26 cm closer to the rim.
- **Fewer men on the arc:** the share of Spurs beyond the three-point line while the possession was being set up fell from 76 to 68 percent.
- **Bodies around him:** when he shot, his nearest team-mate stood 1.28 m closer than in the first half.

The plan was clear: Wembanyama inside, the others close around him, the ball sent to him. But a crowd of team-mates brings a crowd of defenders. He took more shots from a shape built to feed him, with less room than before, and he missed them.

From the map to the rim. His shot map (Fig. 3.14) shows where the extra shots went: closer and closer to the basket. In the first half he was 3 of 3 at the rim, three of them dunks. In the second he was 1 of 7 there, and did not dunk once. The same man, the same spot, the opposite result. Why the rim stopped paying is the question of the next section, and the last link in the chain.

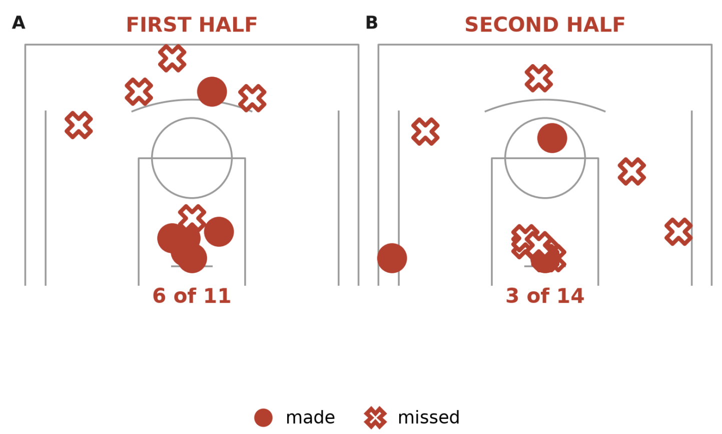

Figure 3.14: Wembanyama's shots, first half to second: more of them in the paint, and most of those missed.

3.6 The rim: a vicious circle

3.6.1 Finishing collapsed

The phenomenon. The Spurs kept getting to the basket and stopped converting there.

The data. Conversion within 4 feet of the rim, from the official shot chart (Fig. 3.15).

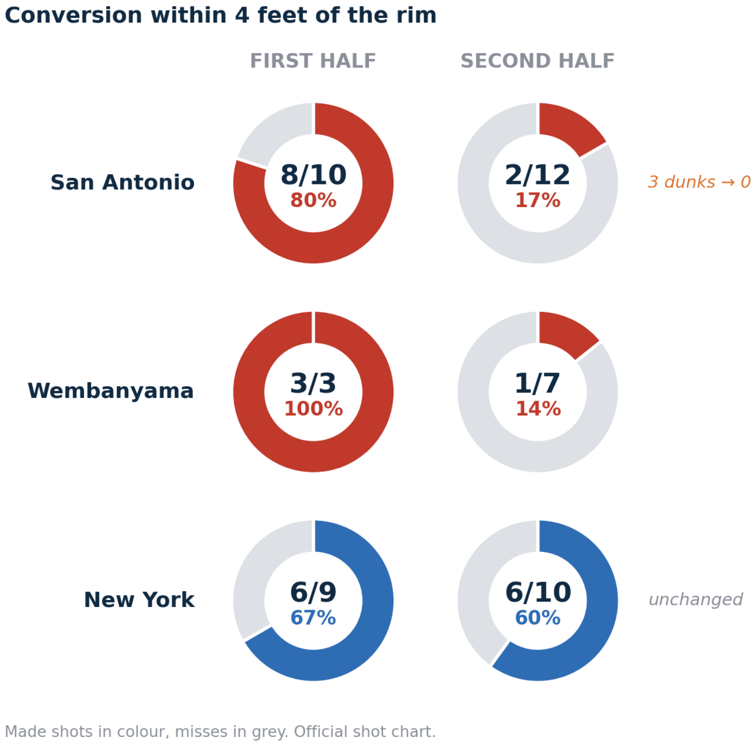

Figure 3.15: Conversion within 4 feet of the rim: San Antonio from 80 to 17 percent, Wembanyama from 3 of 3 to 1 of 7, New York unchanged.

On the floor. This is the key change between the two halves, and it needs no tracking to see: every shot is in the official play-by-play.

- **More attempts at the rim, 10 then 12.** San Antonio kept getting there.
- **Conversion divided by almost five, 80 to 17 percent.** Eight makes became two. Three dunks became none.
- **Wembanyama, 3 of 3 then 1 of 7.** The man the offence was now built to feed at the rim missed six of his seven finishes there.
- **New York, unchanged: 6 of 9 then 6 of 10.** Same rim, same half, same referees; only one team stopped finishing.
- **About a quarter of the whole drop in efficiency** between the halves sits inside these 4 feet.

Same players, same spot, same number of attempts, opposite outcome. Shot quality cannot explain it: a shot from under the basket is a good shot for anyone. What changed is the defence at that exact spot, and only there. San Antonio had crowded the paint around Wembanyama, and every team-mate who came in brought his defender with him: more bodies inside, more defenders at the rim. The tracking data can measure it: how close the defender was when the shot went up.

3.6.2 The defender arrived closer

The phenomenon. The closer the Spurs shot from the rim, the tighter the defence got. Across the whole floor New York pressed San Antonio no harder than in the first half. Inside 8 feet the defender was already closer; inside 4 feet he was almost on the shooter, half a metre nearer than before.

The data. The nearest defender's distance at the moment of the shot, zone by zone (Fig. 3.16), and on the floor (Fig. 3.17).

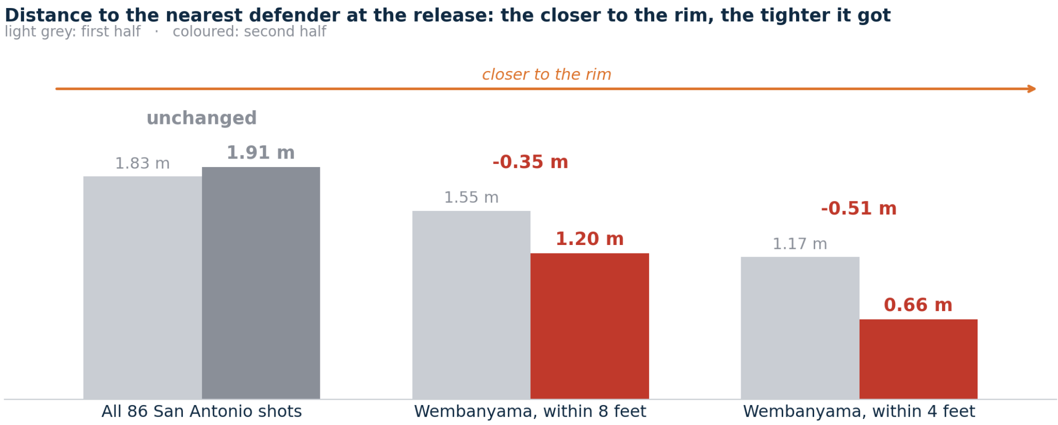

Figure 3.16: Distance to the nearest defender at the moment of the shot, first half to second: unchanged across the floor, half a metre tighter at the rim.

On the floor. Fig. 3.16 is the whole story of the rim. New York did not press harder as a team: on all 86 Spurs shots the nearest defender stood where he had stood in the first half. But the closer the shot came to the basket, the more the second half differs from the first: 35 cm tighter on Wembanyama inside 8 feet, 51 cm tighter inside 4 feet, where he went 3 of 3 and then 1 of 7.

The circle. What sets the contest at the rim is team-mate proximity, in both halves: bodies near the shooter mean defenders near the shooter. That closes the loop with Section 3.5. The offence reorganised to feed Wembanyama inside; feeding him inside pulled his team-mates in around him; team-mates in around him brought their defenders with them; and the extra body at the rim is exactly what killed the finishing. Fig. 3.17 shows where: the red zone, where the defender arrived closer, is the paint, and it is deepest on the right-hand side of it, exactly where Wembanyama took most of his shots.

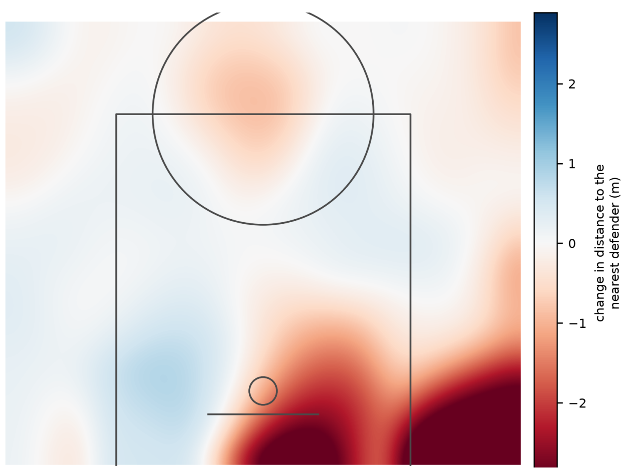

Figure 3.17: Where the defender arrived closer in the second half: the paint, and above all its right-hand side, where Wembanyama took most of his shots.

Chapter 4

Synthesis

The story of the game, without the numbers, runs like this.

What collapsed. A little past the midpoint of the game, San Antonio stopped scoring, and the turnovers arrived in the same window. Classical shot statistics show no particular change in the quality of the shots they were taking. It takes the full tracking of every player on the floor to reveal the hidden patterns underneath, and the tactical changes that explain them. That is what follows.

New York's first key: Towns. The substitution that changed the game. A centre who shoots threes pulled Wembanyama out of the paint, the Spurs' defence stepped out with him, and New York went straight through the space it opened with the pick-and-roll. At the other end Towns took the Wembanyama assignment himself. San Antonio spent the second half solving, on both sides of the floor, a problem they had not had in the first.

A tired team, and a change of rhythm. San Antonio played nine men and put three quarters of the minutes on five of them; New York spread the same game across twelve, and that deeper bench sent fresher men into the fourth quarter. Wembanyama carried more physical load than anyone on the floor and never sat down again. The fatigue did not slow the legs; it changed the rhythm. In the first half the ball arrived and went up. In the second the same players held it a second longer before releasing, deciding instead of reacting.

New York's second key: Anunoby on Fox. Tried in the first quarter, made permanent from the end of the second: New York's best perimeter defender on San Antonio's point guard, on every possession. By the second half Fox was the most physically diminished man on the floor, and the ball did not leave him: locked in, he kept it longer, on possessions that were no longer being built by five men.

Why did a team this good stop playing as a team? Not because it forgot how. Three measured things met at once. Its creator was taken away by one matchup, so the ball stayed in one pair of hands and the passing network thinned: fewer creators, fewer routes, the same passes among fewer people. Its rhythm slowed with its legs, so the shooter was less often set when the ball arrived. And its bench chose a shape, Wembanyama inside and the others drawn in around him, that concentrated the shooting on one man.

The trap at the rim. The shots had to go somewhere and they went to Wembanyama, more of them in the paint. The team-mates who came in to feed him brought their defenders with them: unchanged pressure across the floor, half a metre tighter at the rim. The man now taking most of the shots was finishing in the most crowded part of the floor, and the easiest shots in the game stopped falling: 3 of 3 at the rim, then 1 of 7.

> **Two New York decisions, Towns and Anunoby, met a short, tired San Antonio rotation. The creator was locked, the rhythm slowed, the shape closed in around one man, and the defender arrived at the rim. The box score only records that the shots stopped falling. Following every player, every second, is what lets us read why.**

Chapter 5

Further Analysis

Two directions in which this analysis can grow, both already within reach of the pipeline built here.

New York's side. This report is an analysis of a collapse, so New York appear mostly as the agent acting on San Antonio. The same machinery, pointed at their creators and their adjustments, would say whether they anticipated the collapse or followed it.

The full series. Everything here rests on a single game, so the smallest samples can only be described. Running the same pipeline over all five games would let each player serve as his own control across opponents and rest patterns, the one design that separates fatigue from matchup. It is the highest-value extension available, and it is mostly data collection rather than new method.

Appendix A

Glossary

Every metric used in the report, in full. Nothing here is needed to follow the story; it is here so that any number in it can be reproduced.

Pick-and-roll. A two-man action. The screener stands still in the path of the defender guarding the ball-handler (the *pick*), then cuts to the basket (the *roll*). The defence must choose: the screener's man steps out to the ball, and the roller is free at the rim; or he stays back, and the ball-handler drives or shoots. Against a rim protector it is the standard way to drag him out of the paint.

Effective field goal percentage (eFG).

$$\text{eFG} = \frac{\text{FGM} + 0.5 \times \text{3PM}}{\text{FGA}}$$

FGM is field goals made, 3PM is three-pointers made, FGA is attempts [1]. The 0.5 exists because a three-pointer is worth 1.5 times a two-pointer: crediting a made three with 1.5 makes puts both shot types on one scale, so eFG is proportional to points per field-goal attempt (points per FGA = 2× eFG). Free throws are excluded. Per shot we write $w_i = 1$ for a two-point attempt and 1.5 for a three, and $\text{eFG}_i = w_i$ if the shot went in and 0 if it did not; team eFG is the mean of eFG_i.

Expected goals (xG), expected eFG, and the making residual. A defender-aware logistic model in the Chang (2014) family [2], fitted in metres:

$$\begin{aligned}
\text{logit}(p) &= b_0 + b_1\,\text{dist}_\text{m} + b_2 \min(\text{def_dist}_\text{m},\, 3) + b_3\,\text{catch_and_shoot} + b_4\,\text{corner_three} \\
p &= \frac{1}{1 + e^{-\text{logit}(p)}}
\end{aligned}$$

with dist_m the shot distance in metres, def_dist_m the nearest defender's distance at the validated release (capped at 3 m, beyond which extra space stops helping), and the last two terms binary indicators. The coefficients are $b_0 = +0.140$, $b_1 = -0.157$ per metre, $b_2 = +0.169$ per metre, $b_3 = +0.150$, $b_4 = +0.080$. Expected eFG applies the same weighting to the model, $\text{xeFG}_i = w_i\, p_i$; the making residual is $\text{eFG}_i - \text{xeFG}_i$, and summed over shots, makes minus expected makes is in units of shots.

Expected possession value (EPV) and EPV-Added (EPVA). EPV is the number of points the possession is worth at an instant, given where the ball is and how closely it is defended [3]. Our implementation (*epv_lite*) is a deterministic function of ball location and nearest-defender distance, evaluated on a 0.2 s tick. EPV-Added is what a player adds to that value while he is involved:

$$\text{EPVA(player)} = \sum_{\text{his involvements}} \big[\, \text{EPV(end)} - \text{EPV(start)} \,\big]$$

denominated in points, so the totals are directly comparable with the scoreboard. It splits exactly into EPVA = BUILD + FINISH: BUILD is the value added while carrying and moving the ball, measured from the tracking feed tick by tick, independent of whether the shot went in; FINISH is the value realised or lost when the possession ends in the player's hands, dominated by the outcome of the shot.

The z-score. A z-score subtracts the mean of a reference group and divides by that group's standard deviation, restating a quantity as how far above or below the group a player sits: zero is exactly average,

+1 is one standard deviation above. Our reference group is always the players on the floor at that instant, which is what makes quantities in different units addable and removes the clock from the comparison.

Relative load. A player's raw cumulative load rises with the clock, so comparing raw totals between halves compares stopwatches. Relative load removes the clock by measuring each player against the men he is actually on the floor with:

$$\mathrm{rel\_z}(i,t) = \frac{F_{\mathrm{cum}}(i,t) - \mathrm{mean}_j\, F_{\mathrm{cum}}(j,t)}{\mathrm{sd}_j\, F_{\mathrm{cum}}(j,t)}$$

with j running over the players on court at instant t. Raw cumulative load correlates with elapsed time at +0.90 to +0.97; normalised this way the correlation falls below 0.09 in absolute value. It measures how a coach distributed the work, not how long the game had been going.

Metabolic power. Running converted into watts per kilogram without needing the player's mass, by the equivalent-slope method of di Prampero et al. [4] and Osgnach et al. [6]: accelerating on the flat is energetically equivalent to running at constant speed up a slope.

$$\begin{aligned}
ES &= a_{\mathrm{fwd}}/g && \text{(equivalent slope)}\\
EM &= \sqrt{ES^2+1} && \text{(equivalent mass)}\\
C(ES) &= (155.4\,ES^5 - 30.4\,ES^4 - 43.3\,ES^3 \\
&\qquad + 46.3\,ES^2 + 19.5\,ES + 3.6) \times EM \times K_T && \text{(energy cost, J/kg/m)}\\
P_{\mathrm{met}} &= C(ES) \times v && \text{(W/kg)}
\end{aligned}$$

with a_{fwd} the forward acceleration, v the speed and $K_T = 1$ on hardwood; the energy-cost polynomial $C(ES)$ follows Minetti et al. [5].

The load index L^*. Seven load streams, each z-scored across the on-court cohort, then weighted:

$$\begin{aligned}
L^*(t) = {} & 1.00\, z(P_{\mathrm{met}}) && \text{metabolic power (the anchor)}\\
& + 0.80\, z(L_{\mathrm{dec}}) && \text{eccentric braking, } \max(0, -a) \times v\\
& + 0.70\, z(L_{\mathrm{con}}) && \text{contact / defender proximity and collisions}\\
& + 0.50\, z(L_{\mathrm{cod}}) && \text{change of direction, } v \times |\mathrm{d}(\mathrm{heading})/\mathrm{d}t|\\
& + 0.40\, z(L_{\mathrm{pl2d}}) && \text{2-D PlayerLoad, horizontal jerk}\\
& + 0.35\, z(L_{\mathrm{iso}}) && \text{isometric holding, } \mathbf{1}[v < 1.0 \text{ and nearest_opp} < 1.5]\\
& + 0.30\, z(L_{\mathrm{ball}}) && \text{on-ball load, has_ball}
\end{aligned}$$

Cumulative load comes in two forms, a leaky integrator with a recovery constant and a non-recovering sum:

$$\begin{aligned}
F_{\mathrm{dose}}(t) &= F_{\mathrm{dose}}(t-\mathrm{d}t)\, e^{-\mathrm{d}t/\tau} + \max(L^*(t), 0)\,\mathrm{d}t, \qquad \tau = 25 \text{ min}\\
F_{\mathrm{cum}}(t) &= \textstyle\sum_{\mathrm{ticks}} \max(L^*, 0)\,\mathrm{d}t
\end{aligned}$$

Relative load is built on F_{cum}. The weights are documented anchors from the sports-science literature, and the index deliberately counts more than running: contact and isometric holding are work a distance-only measure never sees [7].

Convex hull, and offensive spread. The convex hull is the area of the smallest polygon containing all five players: how much floor the offence occupies [8]. Offensive spread is the mean distance of the five from their own centroid.

References

[1] J. Kubatko, D. Oliver, K. Pelton, and D. Rosenbaum, "A starting point for analyzing basketball statistics," *Journal of Quantitative Analysis in Sports*, vol. 3, no. 3, Art. 1, 2007. [Primary.] Effective field goal percentage and possession estimation (Poss = FGA − OREB + TOV + 0.44 × FTA); the eFG definition and every possession denominator in this report.

[2] Y.-H. Chang, R. Maheswaran, J. Su, S. Kwok, T. Levy, A. Wexler, and K. Squire, "Quantifying shot quality in the NBA," in *MIT Sloan Sports Analytics Conference*, 2014. [Secondary.] The defender-aware logistic shot-quality form, and the 3 m cap beyond which additional separation stops helping.

[3] D. Cervone, A. D'Amour, L. Bornn, and K. Goldsberry, "A multiresolution stochastic process model for predicting basketball possession outcomes," *Journal of the American Statistical Association*, vol. 111, no. 514, pp. 585–599, 2016; arXiv:1408.0777. [Primary.] Expected possession value and EPV-Added. Our *epv_lite* is a deliberately reduced implementation: canonical EPV conditions on score and clock, which our inputs do not carry.

[4] P. E. di Prampero, S. Fusi, L. Sepulcri, J. B. Morin, A. Belli, and G. Antonutto, "Sprint running: a new energetic approach," *Journal of Experimental Biology*, vol. 208, pp. 2809–2816, 2005. [Primary.] The equivalent-slope method: accelerating on the flat is energetically equivalent to running up a slope.

[5] A. E. Minetti, C. Moia, G. S. Roi, D. Susta, and G. Ferretti, "Energy cost of walking and running at extreme uphill and downhill slopes," *Journal of Applied Physiology*, vol. 93, no. 3, pp. 1039–1046, 2002. [Primary.] The energy-cost polynomial $C(ES)$ evaluated verbatim in the Glossary.

[6] C. Osgnach, S. Poser, R. Bernardini, R. Rinaldo, and P. E. di Prampero, "Energy cost and metabolic power in elite soccer: a new match analysis approach," *Medicine & Science in Sports & Exercise*, vol. 42, no. 1, pp. 170–178, 2010. [Primary.] The team-sport application of metabolic power to match tracking data. Terrain constant $K_T = 1.0$ on hardwood.

[7] A. Martinez-Serrano et al., "Dismissing the idea that basketball is a 'contactless' sport," 2024. [Secondary.] 18.6 percent of contacts produce no kinematic displacement at all, which is why L^* carries contact and isometric terms: a running-only index scores a defensive stance or a held seal as rest.

[8] R. Metulini, M. Manisera, and P. Zuccolotto, "Modelling the dynamic pattern of surface area in basketball," *Journal of Quantitative Analysis in Sports*, vol. 14, no. 3, 2018. [Secondary.] Five-player convex-hull area, the floor-size measure in Section 3.5.2.

[9] J. H. Fewell, D. Armbruster, J. Ingraham, A. Petersen, and J. S. Waters, "Basketball teams as strategic networks," *PLoS ONE*, vol. 7, no. 11, e47445, 2012. [Secondary.] Ball-movement networks, the lineage behind the creation network in Section 3.4.3.

[10] J. B. Gelbach, "When do covariates matter? And which ones, and how much?" *Journal of Labor Economics*, vol. 34, no. 2, pp. 509–543, 2016. [Primary.] The conditional decomposition used as the exact engine behind every attribution in Chapter 3: order-invariant, with exact grouping of correlated features.

[11] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. [Secondary.] SHAP, and the additivity property that lets a change in mean SHAP decompose a between-half shift in the target's own units.

[12] D. N. Politis and H. White, "Automatic block-length selection for the dependent bootstrap," *Econometric Reviews*, vol. 23, no. 1, pp. 53–70, 2004. [Primary.] Block-length selection for the dependent bootstrap: with lag-1 autocorrelation at 0.917, resampling whole possessions rather than frames is the only valid choice.