The Queta Problem

Why Neemias Queta Is a Top-50 Player, and Why 2026–27 Proves It

GRAVITY Research Note No. 2 — a prediction paper

July 10, 2026

Abstract

Our July 2026 ranking of all 644 NBA players (GRAVITY) placed Neemias Queta — undrafted-adjacent, twice-waived, the 39th pick of 2021 — at **#41**, ahead of household names and a full tier above his public reputation. This paper treats that placement as a testable scientific claim rather than a quirk. We audit it three ways: we shrink the model's own shiniest input (his +8.6 on/off) using an empirical noise model built from 184 player-pairs, and the case survives; we run a comparables engine over every statistically similar big-man season since 2020 (33 of them — the Gobert / Robert Williams / Jarrett Allen family) and find a 39% base rate of a top-50-caliber follow-up, before adjusting for Queta's unusually secure situation; and we stress the one real threat, Boston's signing of Mitchell Robinson — noting that the Celtics answered it themselves by extending Queta for $56 million *afterward*. The paper ends with a falsifiable forecast: under transparent scenario weights, GRAVITY puts the probability that Queta finishes 2026–27 as a top-50 player at roughly **75%**, and top-25 at roughly **40%**. We will score this prediction in July 2027.

1 The claim, and why it sounds wrong

Every ranking system produces a few outputs its authors flinch at. In our July 2026 edition of "The 644," the flinch was #41: Neemias Queta, ranked above Ivica Zubac, Rudy Gobert, Jaren Jackson Jr., and roughly four hundred players with larger reputations. The public model of Queta is a hustle big — a two-way-contract success story, fifth big on a title roster, fine. The statistical model of Queta is a starting center on a 56-win team who posted one of the most efficient high-volume finishing seasons in basketball, protected the rim, and led all Boston regulars in filtered on/off impact.

When a model and the crowd disagree this sharply, one of them is wrong, and papers that only ever defend their model are advertising. So this note does three things: it explains where #41 comes from, it attacks the number with the two strongest counterarguments we know how to formalize (noise and regression), and it converts what survives into a prediction with a scoring date. One relevant institution has already taken a side. On July 6, Boston signed Mitchell Robinson — a direct competitor for Queta's minutes. Days later, they signed **Queta to a $56 million extension anyway**.

2 How the rating works (the short version)

GRAVITY asks two questions: *how good is a player per minute on the floor* (box impact, efficiency relative to offensive burden, creation, defense, and garbage-time-filtered on/off from Cleaning the Glass), and *how much can you count on him* (games, role size, documented injuries). The two

multiply. Everything is expressed in z-scores against the league-average minute: 0 average, +1 clearly good, +2 elite; the final scale sets 50 = average, 60+ = quality starter, 80+ = franchise engine (Jokić tops 2025–26 at 97.8). The full specification is in Appendix A; the seven-season historical variant (GRAVITY-RS) drops the on/off term, which is not uniformly available for early seasons.

3 Who is Neemias Queta, statistically

The first Portuguese-born NBA player; the 39th pick in 2021 out of Utah State; two two-way contracts; waived by Sacramento; a garbage-time cameo on Boston's 2024 title team. Then Porziņģis was traded, Horford and Kornet left, and in 2025–26 coach Joe Mazzulla handed him the starting job: 76 games, 75 starts, 25.3 minutes per game, 10.2 points, 8.4 rebounds, 1.3 blocks — career highs across the board.

Table 1: Queta's NBA career. IMPACT is per-minute value in z units (GRAVITY-RS, league-standardized within each season); Pct is his percentile among 1,000+ minute players; Ctr-Pct among centers only.

Season	Team	G	MP	TS%	WS/48	BPM	IMPACT	Pct	Ctr-Pct
2021–22	SAC	15	120	.495	.051	−5.1	−0.46	24	9
2022–23	SAC	5	29	.607	.149	−3.4	−0.10	46	13
2023–24	BOS	28	333	.662	.259	+1.5	+0.50	77	61
2024–25	BOS	62	863	.674	.179	−0.7	+0.03	54	28
2025–26	BOS	76	1926	.674	.224	+2.8	**+0.85**	**88**	**77**
2026 playoffs	BOS	7	152	.767	.230	+1.2	—	—	—

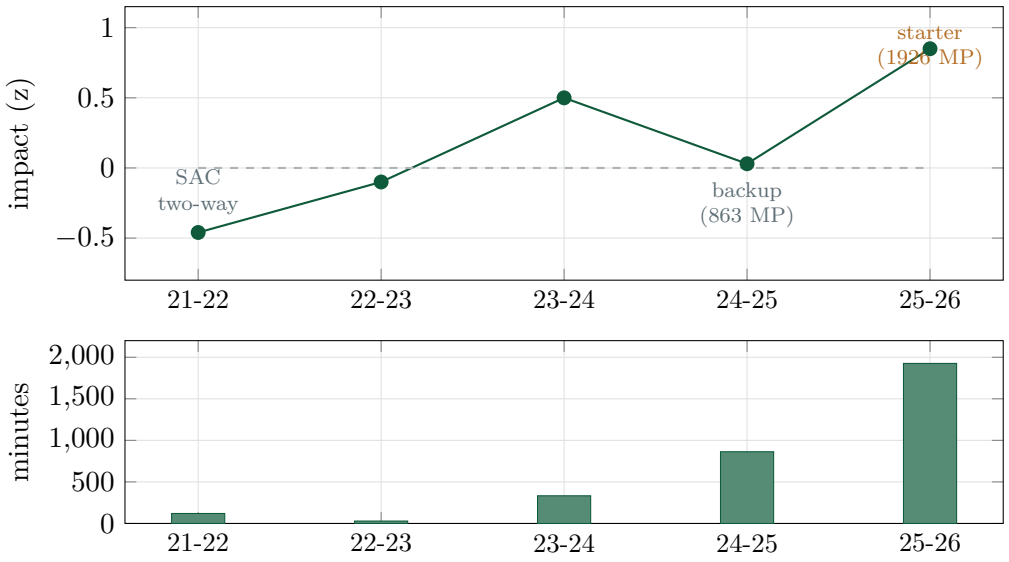

Figure 1: Five seasons: per-minute impact (top) and the role that finally arrived (bottom). The 2024–25 dip is real — his box value that year was ordinary — and the paper does not hide it.

Two mechanical details explain why the role took five years and why it should persist. First, the classic young-big constraint — fouling — has receded on schedule: his foul rate went from the 6th percentile in 2023–24 (a foul machine, unplayable in long stretches) to the 24th, to the 31st.

Coaches could not keep him on the floor before; in 2025–26 they could. Second, his efficiency did not bend as the load grew.

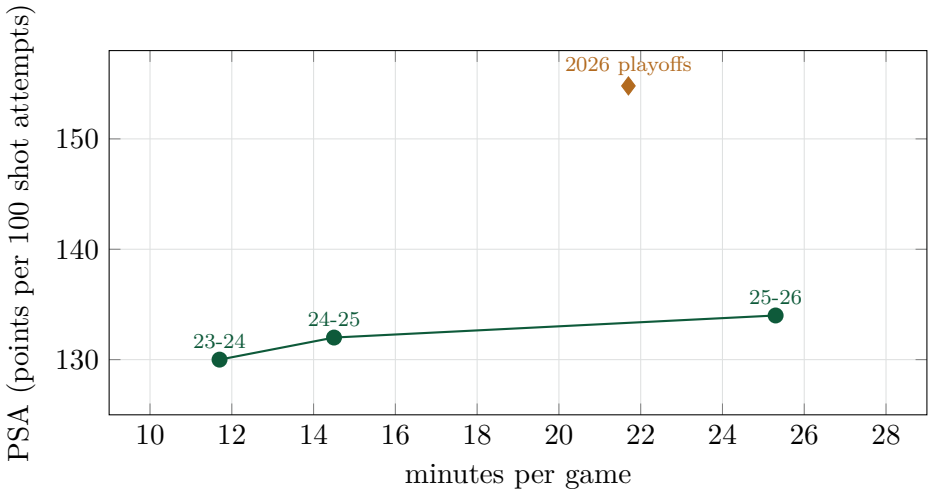

Figure 2: The anti-fluke chart. CTG scoring efficiency (PSA) against role size, Boston seasons. The normal pattern is a downward slope — bench bigs feast on easy minutes and regress when the role doubles. Queta's role grew sevenfold and the efficiency line went *up*: 130 → 132 → 134, all 79th–87th percentile, with a 154.8 (97th percentile) playoff series on top.

Table 2: The 2025–26 profile in CTG percentiles (among bigs). This is what #41 is made of: elite finishing at real volume, elite offensive rebounding, real rim protection, and the league's 17th-best filtered on/off among 1,500+ minute players.

Component	Value	Percentile
Scoring efficiency (PSA)	134.0	84th
eFG%	65.4%	86th
Rim accuracy	75%	85th
Offensive rebound rate	12.3%	79th
Block rate	2.7%	79th
Defensive rebound rate	19.9%	74th
Filtered on/off (raw)	+8.6	#17 of 158 league-wide
Turnover rate	12.6%	55th
Foul rate	4.4%	31st (up from 6th in '24)
Free-throw %	70.1%	41st

In plain terms: Queta finally got starter minutes at 26 and his efficiency — which usually falls when a bench big's role doubles — went up instead. He finishes like a top-10 center, crashes the offensive glass, blocks shots, and stopped fouling himself off the floor. That's the whole trick: there is no empty-calorie stat propping him up.

4 Stress test I: shrink the shiny number

The skeptic's first target should be that +8.6 on/off — Boston was 8.6 points per 100 possessions better with Queta on the floor (garbage time excluded). On/off numbers are notoriously noisy, so

we measured the noise rather than argue about it. For all 184 players with 1,000+ CTG minutes in both 2024–25 and 2025–26, we correlated their on/off across the two seasons (Figure 3):

$$r = 0.317. \tag{1}$$

That is the honest, ugly truth about on/off: even filtered, only about a third of it is signal. The empirical-Bayes response is to shrink toward the mean by exactly that factor:

$$\widehat{\text{on/off}}_{\text{true}} \;=\; r \times 8.6 \;=\; +\mathbf{2.7}. \tag{2}$$

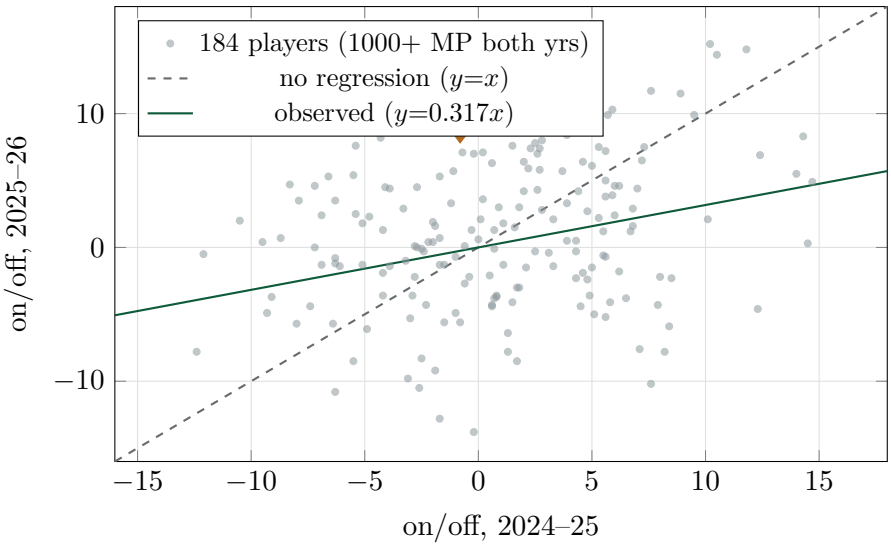

Figure 3: How much of an on/off number survives a year? The cloud says: about a third. The amber diamond is Queta — his 2024–25 on/off (a −0.8 in a 768-minute backup role, too few minutes even to qualify for the cloud) carried almost no information; his 2025–26 value is large but must be shrunk by (1) before you believe it.

This is the paper attacking its own best exhibit, so state the result plainly: properly shrunk, Queta's lineup impact is +2.7 per 100 — solidly positive, no longer spectacular. Does #41 survive? Yes, because on/off is only 14% of the GRAVITY impact blend (Appendix A); the placement is carried by the box profile of Table 2, which involves hundreds of shots and rebounds rather than lineup arithmetic, and by Table 1's 88th-percentile composite. Rerunning his 2025–26 impact with the on/off input forced to +2.7 moves his score from 58.8 to ≈ 57.6 — from #41 to roughly #48. The claim "top-50 player" does not depend on believing the shiny number.

In plain terms: On/off stats are two-thirds noise — we measured it. Discount Queta's +8.6 all the way down to +2.7 and he's still a top-50 player, because the ranking was never resting on that number in the first place.

5 Stress test II: what happened to everyone like him

The second attack: *every* backup-turned-starter big has one shiny efficient season; most regress. Rather than argue by anecdote, we built the cohort. From six seasons of league data (2019–20 through 2024–25), we extracted every big-man season matching Queta's 2025–26 statistical

silhouette — usage ≤ 19, WS/48 $\geq .190$, TS $\geq .630$, block rate $\geq 1.8\%$, 900–2,500 minutes, age ≤ 28 — and followed each into the next season. There are 33 of them, and the family is exactly who you'd hope: Gobert, Robert Williams, Jarrett Allen (four times), Zubac, Hartenstein, Gafford, Claxton, Kessler, Duren — and, three times, Mitchell Robinson.

Table 3: Selected comparables (of 33): elite-efficiency, low-usage big seasons and what happened next. Full criteria in text; IMPACT in z units, Pct = next-season percentile among 1,000+ minute players.

Year	Player	Age	MP	IMPACT	Next MP	Next IMPACT	Next Pct
2020–21	Robert Williams	23	985	+1.66	1804	+1.48	95
2024–25	Jalen Duren	21	2034	+1.02	1976	+1.60	97
2019–20	Rudy Gobert	27	2333	+1.12	2187	+1.38	94
2023–24	Daniel Gafford	25	1815	+0.89	1226	+1.32	95
2024–25	Jarrett Allen	26	2296	+1.26	1519	+0.88	88
2023–24	Isaiah Hartenstein	25	1896	+0.87	1590	+0.86	84
2021–22	Isaiah Hartenstein	23	1216	+1.34	1626	−0.16	42
2022–23	Walker Kessler	21	1703	+0.89	1493	+0.24	67
2022–23	Nic Claxton	23	2271	+0.96	2116	+0.38	73
2021–22	Mitchell Robinson	23	1848	+0.75	1591	+0.63	82
2019–20	Mitchell Robinson	21	1412	+1.02	853	+0.34	69
2020–21	Enes Freedom	28	1758	+0.39	411	−0.23	38
All 33: median next-season percentile **75th**; ≥80th: **45%**; ≥85th: **39%**							
median impact regresses +0.89 → +0.47; median minutes change **−146** (roles more often shrink than grow)							

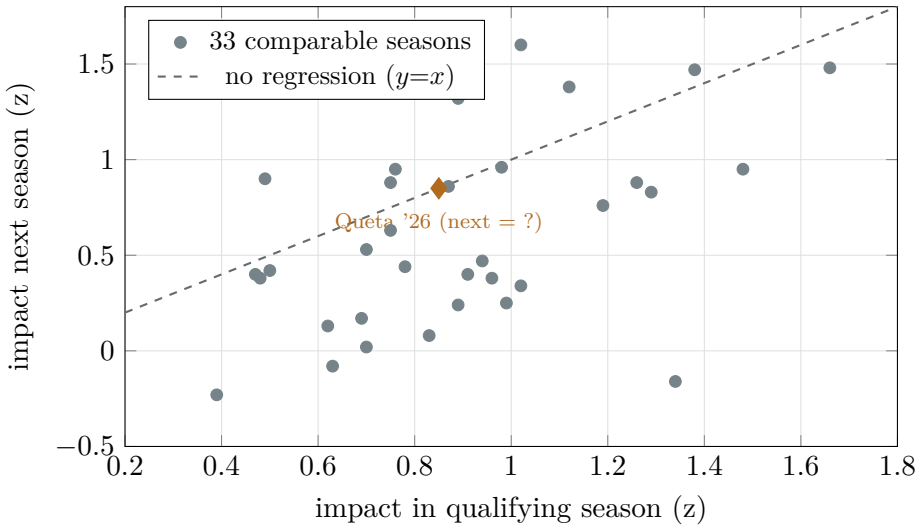

Figure 4: Regression to the mean, drawn honestly. Most dots sit below the dashed line: seasons like this usually give some value back. But the floor is high — three-quarters of the cohort remained above the league median, and two in five delivered a top-50-caliber (85th+ percentile) follow-up.

The base rates give the skeptic real ammunition — the median comparable gave back about half his impact z, and minutes *shrank* more often than grew — and they give the bull real support: 39% of the cohort posted a top-50-caliber next season, and the family's floor is remarkably high. Two situational facts then push Queta above the cohort median. Most comparables were backups fighting for a role (Kornet, Jackson-Davis, Clarke); Queta is the *incumbent starter* on a 56-win

team. And no comparable had his team commit $56M in July. Base rate 39%, situation-adjusted: higher.

In plain terms: We found the 33 players statistically most like 2025–26 Queta over the last six years. Most gave a little back the next season — but 2 in 5 delivered a top-50-level year, and unlike most of them, Queta enters next season with the starting job, a new contract, and a 56-win team around him.

6 The Robinson question

The strongest objection isn't statistical, it's transactional: on July 6 Boston signed Mitchell Robinson — the champion Knicks' rim protector — for three years and $47.4M. If Robinson takes 24 minutes a night, the bear case writes itself.

Three replies, in declining order of hardness. *First*, the revealed preference: Boston extended Queta for $56M days *after* the Robinson deal, and beat reporting projects Queta to open as the starter, with something like an even split behind him. Front offices do not pay both centers starter money to bench one. *Second*, Robinson's own availability record is the worst in the comparables family — he appears in Table 3 three times precisely because he keeps posting elite fragments (853, 768, 1,591 minutes follow-ups); planning Boston's center minutes around Robinson's health has broken repeatedly for a decade. *Third*, the tide: Jayson Tatum returns for a full season (his 16-game 2025–26 sample and the seven-game playoff series he played at an 8.9 BPM say the Achilles took less than feared), which raises every Celtic's context — and Queta's playoff series alongside him (.767 TS) was his best basketball of the year. The honest version of the Robinson signing is a minutes tax on the bear case, not a refutation: it lowers Queta's floor scenario from "starter by default" to "starter by performance," which — given Section 5 — is a bet the data supports.

7 The prediction

Now the arithmetic. The two-year GRAVITY machinery (Appendix A) converts any assumed 2026–27 performance level and workload into a score; the top-50 cutoff in the July 2026 distribution is **57.4**. We fix four scenarios spanning the realistic range, weight them transparently, and let the pipeline speak.

Table 4: Four futures, weighted. Impact is 2026–27 per-minute value in z units (his 2025–26 was +0.97 with the on/off input, +0.85 without); the blend automatically carries 2025–26 forward as the trailing season.

Scenario	Weight	Impact	MP	GP	Score	Top-50?	Reading
Bust: loses the job	15%	+0.30	1000	50	53.8	no	Robinson takes over
Bear: even timeshare	25%	+0.75	1450	62	58.1	*barely*	regression + split
Base: holds the job	40%	+1.00	2000	74	62.3	yes (#23–27)	2025–26 continues
Bull: consolidation	20%	+1.25	2300	76	66.9	yes (#10–13)	comps' top tercile
$P(\text{top-50}) \approx 0.25{\times}0.6 + 0.40 + 0.20 = \mathbf{0.75}$ (bear counted 60%, being borderline)							
$P(\text{top-25}) \approx 0.40{\times}0.5 + 0.20 = \mathbf{0.40}$							

Note what makes the claim robust: it does not require a leap. The base case is simply *2025–26 happening again* — same job, same efficiency, ordinary regression — and it lands at #23–27, comfortably inside the claim. The prediction fails only if Robinson takes the job outright or the

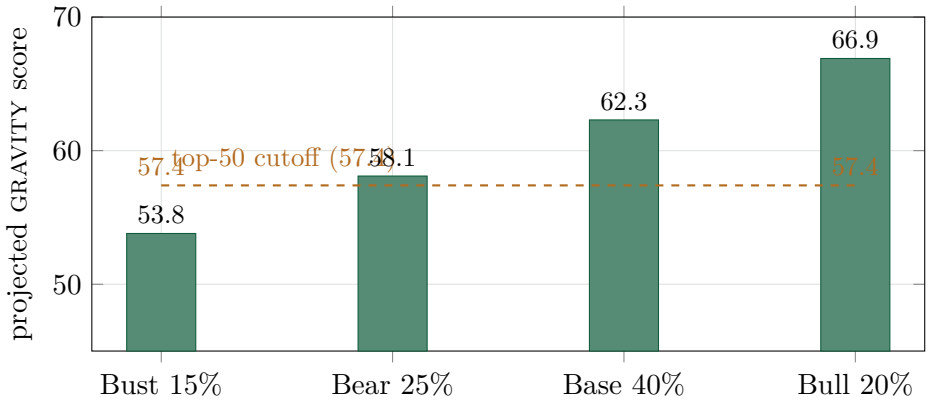

Figure 5: The bet, drawn. Three of four scenarios clear the top-50 line; the weighted probability is about 75%.

season-scale regression lands in the worst tercile of the comparables. **Falsification criteria, on the record:** in the July 2027 edition of "The 644," computed by this same specification, Queta at #50 or better scores as a win; #51–70 as a push with explanation owed; below #70 as a clean miss.

In plain terms: The bet: about 75% that Queta is a top-50 player in next July's ranking, and 2-in-5 that he's top-25. It's not a moonshot claim — it just requires this season to mostly repeat. We wrote down in advance exactly what counts as being wrong.

8 Conclusion

The Queta problem dissolves under measurement. A #41 ranking that looks like a model glitch turns out to rest on the most trustworthy statistics in the sport — hundreds of finished possessions at 85th-percentile efficiency, elite offensive rebounding, real rim protection — and to survive the two best attacks we could mount on it: shrinking his on/off by an empirically measured noise factor ($r = 0.317$) still leaves a top-50 profile, and the 33-player comparables cohort puts a 39% base rate on a top-50-caliber follow-up before counting his incumbency, his new contract, or Tatum's return. Boston's own front office, holding better information than any public model, signed the same conclusion for $56 million. The market will grade us in twelve months.

A The GRAVITY specification (compact)

Identical to Research Note No. 1 and the July 10, 2026 "The 644" ranking (v3). For season y, minutes-weighted moments define clipped z-scores $z_k(i) = \mathrm{clip}((x_k(i) - \mu_k)/\sigma_k, \pm 4)$. Components: efficiency-under-load $(\mathrm{TS} - \mu_{\mathrm{TS}}) \cdot 100 \cdot (\mathrm{clip}(\mathrm{USG}, 8, 36)/20)^{0.7}$; creation $\mathrm{AST}\% - 0.55\,\mathrm{TOV}\%$; load $\mathrm{USG} + 0.4\,\mathrm{AST}\%$; stocks $2\,\mathrm{STL}\% + 1.2\,\mathrm{BLK}\%$; rim pressure; volume (pts/100). Skill blend $S = 0.27 z_{\mathrm{vol}} + 0.19 z_{\mathrm{eff}} + 0.22 z_{\mathrm{crea}} + 0.09 z_{\mathrm{load}} + 0.10 z_{\mathrm{stk}} + 0.05 z_{\mathrm{reb}} + 0.08 z_{\mathrm{rim}}$. Season impact

$$I_y = \left[0.44 z_{\mathrm{BPM}} + 0.12 z_{\mathrm{WS/48}} + 0.14 z(\Delta_{\mathrm{on/off}} \min(1, m/2400)) + 0.30 S\right] \sqrt{\min(1, m/800)}, \quad (3)$$

blended across two seasons with minute-proportional weights (0.75 on the recent year), shrunk toward replacement $\rho = -1.9$ with credibility $c = \sqrt{\min(M, 2200)/2200}$, adjusted for playoff

evidence and age, then discounted by availability $A = 0.78 + 0.22\min(1, (0.7g_y + 0.3g_{y-1})/55)$, role size $\Phi = 0.62 + 0.38\min(1, \text{mpg}^*/32)$, and documented injuries $J \in [0.78, 1]$. Final scale $G = 50 + 15[\rho + (\tilde{R} + \pi + \alpha - \rho)A\Phi J]$. The historical variant GRAVITY-RS (used for Table 1 and the comparables engine) deletes the on/off term and renormalizes to $[0.533z_{\text{BPM}} + 0.133z_{\text{WS/48}} + 0.333S']\sqrt{\min(1, m/800)}$.

Data and forecast registration

Basketball-Reference league tables 2019–20 through 2025–26 (≈4,400 player-seasons); Cleaning the Glass subscriber database (filtered on/off, shot zones, foul and rebounding rates; player page 4902), accessed July 10, 2026. The on/off noise estimate uses all 184 players with 1,000+ CTG minutes in both 2024–25 and 2025–26. No external rankings or projections were consulted. Prediction registered July 10, 2026; to be scored against the July 2027 "The 644."

References

[1] Basketball-Reference.com, league tables, accessed July 10, 2026.

[2] Cleaning the Glass, player and leaders tables (subscription). https://cleaningtheglass.com/stats/player/4902

[3] "Why Neemias Queta is the Celtics' starting center for 2025–26," Yahoo Sports / NBC Sports Boston, October 2025.

[4] "Celtics sign Neemias Queta to $56 million extension after Mitchell Robinson addition," Yahoo Sports, July 2026.

[5] "Mitchell Robinson joining Celtics on 3-year, $47.4M deal," NBA.com / Hoops Rumors, July 6, 2026.

[6] GRAVITY Research Note No. 1, "The Rise and Fall of Ja Morant," July 10, 2026.