

THE SHANAHAN ERA / 2017–2025 + 2026 UPDATE

The 49ers injury problem

A real pattern. An uncertain explanation.

1st

Era injury burden


+40%

Versus the NFL average

8 of 9

Seasons in the worst 10

San Francisco has carried an unusually persistent injury burden. It leads our full-era, season-normalized comparison, narrowly ahead of Arizona. Removing 2020 and 2021 still leaves the 49ers fourth-highest.



Each grey dot is one NFL team-season. Red tracks SF. A value of 1.0x equals that season's league average. Sources [1–7]; calculations described on pages 5 and 13.

The strongest predictors of new public-report appearances: earlier playing usage and previous injury-report history. Their importance does not establish that the 49ers overwork players, recruit unusually fragile players, or have inferior medical care.

The shareable conclusion: the burden is real; the data do not isolate a single cause. Arizona is close, 2023 was healthy, and a reporting model cannot explain every serious injury.

Research snapshot: September 25, 2026, Pacific time. Full seasons: 2017–2025. Early 2026 is presented separately. Independent analysis using public sources; no team medical records.

01 / THE MEASUREMENT

Three questions, three measures

Counting injuries and measuring their consequences are different tasks.

Question	Measure used	What it can tell us
How much did injuries disrupt the team?	Adjusted games lost (AGL)	Lost availability among expected starters and important contributors.
How often were players listed unavailable?	Physical-condition Out/Doubtful report entries	A supplementary reporting count; repeated weeks can be the same injury.
Which features predict the next report?	A new report appearance after a game	A conditional reporting forecast among players who played and were previously unlisted.

Why AGL is the main benchmark

Football Outsiders/FTN counts missed games and fractional losses for players who play while listed as injured. It focuses on expected starters, their injury replacements and key situational players. A reserve’s absence may not count; an All-Pro is not given extra talent weight simply for being an All-Pro. AGL is a team-availability measure, not a count of new diagnoses. [1–8]

A serious blind spot in weekly reports

Players placed directly on injured reserve can disappear from the weekly dataset. Our audit found no season-report rows for Nick Bosa or Solomon Thomas in 2020, or Brandon Aiyuk in 2024. A missing next-week row can therefore coexist with a catastrophic injury. The model and report-entry counts cannot replace AGL.


Scope boundary: all historical comparisons and modeling use regular seasons. Playoff injuries are excluded. George Kittle’s Achilles injury in January 2026 belongs to the *2025 postseason*, and is absent from 2025 regular-season AGL. [22]

We recovered all 32 teams for nine completed seasons: **288 AGL observations**. Weekly reports, snaps, player IDs and schedules support a separate player-level analysis. The 2026 feed has partial current-week information, so it is not mixed into historical rankings. [9–12]

02 / THE TIME SERIES

2023 was the exception

High injury burden recurred across most of the era.



AGL divided by the same-season mean across all 32 teams. COVID absences excluded in 2020 and 2021. Source [1–7].

Season	SF AGL	Most-injured rank	Relative to NFL
2017	91.6	10 of 32	1.31×
2018	97.2	7 of 32	1.31×
2019	95.8	6 of 32	1.26×
2020	161.6	1 of 32	2.27×
2021	119.5	4 of 32	1.48×
2022	93.9	9 of 32	1.18×
2023	34.5	29 of 32	0.50×
2024	141.2	1 of 32	1.83×
2025	115.0	5 of 32	1.40×


The two most extreme SF seasons were **2020 (2.27× the league mean)** and **2024 (1.83×)**. In 2023 the team was fourth-healthiest by AGL, at roughly half the league average. In 2025 it was fifth-most injured, not the league’s worst.

Ranks above run from 1 = most injury burden to 32 = least. Publisher tables generally use the reverse convention. 2017 uses its original published methodology; see the sensitivity analysis and source notes.

03 / ALL 32 TEAMS

First, but barely clear of Arizona

A long-running problem shared with other franchises.




Mean of annual AGL / league-mean AGL, 2017–2025. Equal weight per season. 1.0x is average. Sources [1–7].

SF: 1.396x. Arizona: 1.383x. The small gap is not evidence that San Francisco belongs in a unique category. Arizona actually has a slightly worse average annual rank, while SF has the higher burden on our magnitude-based measure.

04 / ROBUSTNESS

The finding survives tougher checks

The rank changes; the excess burden persists.



Labels show excess above the same-year NFL average and rank for highest burden. All use equal-weight annual ratios. Sources [1–7].

Removing the extraordinary 2020 season lowers SF to third. Removing both pandemic seasons puts it fourth, still **26% above the league average**. The most recent three seasons place it seventh, reflecting the unusually healthy 2023 season.

Check	Result	Interpretation
Season resampling	1.10×–1.69×	95% bootstrap interval across the nine observed seasons; the resampled-season score varies.
SF minus Arizona	–0.39× to +0.43×	Paired season-resampling interval includes zero by a wide margin.
Conventional boxplot rule	Not flagged	SF 1.396× is below the 1.432× upper Tukey fence across 32 era averages.


Answer to “truly an outlier?” SF is at the high end of a persistent burden distribution. It is not uniquely separated under a standard boxplot rule. “Most burdened in this compiled era” is better supported than “uniquely cursed.”

Alternate measure: mean AGL per 16 scheduled games also ranks SF first overall, and second when 2020 is excluded. Different weighting of high-injury league seasons changes a close rank. Nine seasons offer limited inferential precision; resampling does not measure source error.

05 / CHANCE AND SELECTION

An extreme team will always exist

A reference simulation shows why selection matters.



100,000 simulations independently reshuffle the 32 observed team burdens within each season, then average across nine seasons. Red is SF's observed 1.396x.

0.22%

One preselected team

6.47%

Any of 32 teams

Under this simplified reshuffling model, a team chosen in advance reaches SF's burden about 0.22% of the time. But if we look across the league and report whichever team is highest, the chance is about 6.47%. That second comparison is less dramatic.

This is not a test proving “bad luck” or disproving it. The simulation deliberately removes persistent team differences. Real rosters, prior injuries and reporting practices carry over between seasons. It also does not correct for trying many historical windows.

Why persistence still matters


Eight of nine SF seasons land among the ten most burdened teams. That is a descriptive fact worth investigating. However, team selection, shared personnel, diagnosis mix and measurement revisions prevent a simple probability calculation from identifying the cause.

Monte Carlo p-values use $(\text{exceedances} + 1) / (\text{simulations} + 1)$. The maximum-over-32 calculation handles team selection under this reference model, not all choices made after observing the history.

06 / WHERE THE BURDEN LANDED

The injured position changes

2025 was dominated by edge rushers and receivers.



49ers 2025 positional AGL. The nine categories sum to 115.0. This measures lost contributor availability, not new injury counts. Source [8].

28.5

Edge-rusher AGL

27.1

Receiver AGL

0.4

Running-back AGL

Edge rushers contributed the most positional AGL in the league; receivers were second-highest. Together they represented **48.3% of SF's 2025 AGL**. The running-back group fell from **32.2 AGL in 2024 to 0.4 in 2025**. [6,8]

This changing concentration fits a mixture of persistent vulnerability and volatile injury events. It does not fit a simple claim that one position or a run-heavy offense alone explains the entire era.

The burden crosses both sides of the ball: **2024 offense 65.0 / defense 76.2; 2025 offense 60.2 / defense 54.8**. The burden is not confined to offense; this split does not identify playcalling's contribution. [6,7]

Quality check: a source inconsistency affects the Chargers' positional totals. The SF row reconciles exactly; no cross-team positional averages using that inconsistent row appear here.

07 / THE PREDICTIVE MODEL

A useful signal, with modest accuracy

We tested the model on seasons it had not seen.

41,858

Held-out player transitions



9.07%

New-report base rate

0.663

Held-out ROC AUC

Task: among players who took a snap and had no injury-related report before that game, predict a physical injury/condition report before their team’s next game. A new listing can be a recurrence or disclosure of an existing problem; it is not necessarily new tissue damage.



Training 2017–2022; model choice on 2023; refit through 2023; final evaluation on 2024–2025. Each calibration point is a tenth of held-out predictions. Sources [9–12].

Gradient boosting beat a fixed logistic-regression candidate on validation average precision. On unseen seasons, average precision was **0.151 versus 0.091** for a constant baseline. AUC was **0.663** (game-resampling interval 0.654–0.671). AUC 0.5 is chance and 1.0 is perfect. This provides limited ranking ability, not reliable individual injury forecasting.


Essential limit: this outcome misses direct-to-IR injuries and excludes already-injured nonparticipants. It models public reporting in a selected active-player sample, not total team injury burden.

Brier score: 0.0802 versus 0.0824 baseline; lower is better. Severe Out/Doubtful sensitivity: AUC 0.636, average precision 0.040 versus 0.023 baseline. Bootstraps condition on the fitted model and do not capture all repeated-player dependence.

08 / FEATURE IMPORTANCE

Usage and history carry most signal

The leading two feature groups are effectively tied.



Grouped permutation importance on 2024–2025: decline in average precision when a group is shuffled. Larger drops mean more predictive information. Whiskers are 95% game-resampling intervals, conditional on a fitted model and one permutation. Sources [9–12].

Earlier usage: 2.79 AP points; prior-report history: 2.61. Their uncertainty intervals overlap. Features are correlated and scores are not additive. These are not percentages of injuries caused or explained.

The usage result mostly reflects snap share in an earlier game, rather than accumulated 28-day snap volume. That can encode player role, exposure and previous availability. The index game’s snaps were excluded because an injury in that game could shorten them and distort the relationship.

Schedule/recovery: 0.61 AP points. Most of that signal comes from time until the next game. Report timing, recovery and practice opportunities can all contribute. It is not proof that short rest causes injuries.


Age, home/away and surface add little here. This does not establish that age or surfaces are harmless: the outcome is coarse, important cases are missing, and features overlap. In the severe-report sensitivity, history leads, followed by usage and position.

Age score rounds to 0.01 AP points; exact 0.0148. Surface –0.0238: negligible, consistent with noise. Global permutation importance measures predictive dependence, not effects of changing a feature or causes of SF’s excess.

09 / WHAT THIS MEANS FOR SF

History is plausible. Overuse is unproven.

League-wide feature importance alone cannot explain a team’s excess.



Unadjusted held-out rates among eligible player transitions, 2024–2025. Prior appearances can repeat the same injury. Bin sizes: 16,420; 11,794; 9,228; 4,416. Sources [9–12].

Players with six or more injury-related report appearances in the prior year had a **16.1%** next-report rate, versus **5.6%** with none. This association is compatible with recurrence and persistent vulnerability; it also includes reporting habits and underlying conditions.

Held-out eligible sample	SF	Other 31 teams
Mean age	27.79	26.98
Mean report appearances, prior year	2.40	2.04
Mean snaps in prior 28 days, excluding index game	96.25	99.16
No prior-year report appearances	32.3%	39.4%

SF’s selected sample had more report history and was older, but had **slightly fewer measured earlier snaps**. This supports examining player history and roster continuity; it does not demonstrate unusually heavy workload or prove a risky recruitment strategy.


A model excluding team identity predicts a **9.48%** new-report rate for SF; the observed rate is **10.39%** (133 / 1,280 transitions). That modest reporting residual cannot explain the much larger AGL burden, because the outcomes and populations differ.

These profiles are conditional on participation and a clean report, not whole-roster averages. SF observed-rate game-resampling interval: 8.87%–11.97%; no claim of a statistically distinct causal team effect is made.

10 / MEDICAL EVIDENCE

Plausible mechanisms have evidence

The evidence is league-wide; SF-specific attribution remains open.



Separate observational comparisons and injury types; do not compare them as shares of one causal explanation. IRR = injury rate ratio. Primary studies [13,14].

Prior injury and recurrence

An NFL hamstring study recorded 2,075 injuries, of which 38.4% were recurrent. That supports the biological plausibility of history as a risk marker. It does not mean previously injured players have a 38.4% higher risk than uninjured players, or that SF returned anyone too soon. [15]

Preparation and workload progression

A 2026 study using NFL medical and tracking data associated gradual or varying early-camp practice schedules with fewer muscle strains: IRR 0.81 (0.68–0.95). After league-wide acclimation changes and education, camp strains fell from 288 in 2021 to 186 in 2023. This supports auditing preparation; the bundled before/after design cannot isolate a single cause or implicate SF. [13]

Surface and movement demands

A 2012–2016 all-team study found a 16% higher lower-extremity injury rate per play on synthetic surfaces. SF's home schedule is coded as natural grass; away exposure and specific field properties still matter. Our weak surface importance score cannot settle that clinical question. [12,14]

Snapshots omit speed, cuts, contact and practice. NFL medical data show exposure differences by play type; a single-team wearable study supports examining abrupt workload changes but cannot establish a universal safe threshold. [16,24]

11 / COMPETING EXPLANATIONS

Which stories survive the evidence?

A ranked research agenda is more defensible than assigning blame.

Explanation	Assessment	What would establish more
Prior report history / recurrence	Strongest SF-specific measured clue. History predicts reporting and SF’s eligible players had more of it.	Medical history before acquisition, severity, recurrence and exposure; compare recruitment choices.
Role, exposure and preparation	Plausible. Earlier usage predicts reports; medical studies support gradual preparation.	Practice tracking, high-speed load, contact, rehab progression and adherence.
Older roster	SF sample is older, but age adds little in this model.	Whole-roster exposure-adjusted, diagnosis-specific analysis.
Field surface	Relevant to some lower-extremity injuries; no demonstrated SF attribution.	Actual injury location, mechanism, footwear and field measurements.
Travel / short rest	Limited or inconsistent supporting NFL evidence.	Prospective exposure, sleep/recovery measures and specific schedule comparisons.
Coach / medical staff	Public data cannot isolate responsibility.	Consistent clinical records, exposure and credible comparison around interventions.

A gamebook study found fewer recorded in-game injuries on four-day rest, while another multi-season study found no travel-distance association. Neither proves short rest is protective or travel harmless. They weaken a confident claim that these obvious-sounding factors explain SF’s history. [17,18]

2026, as of September 25: the final Week 3 report lists Bosa, Height, Robinson and Thompson Jr. out, with Evans questionable. Bosa’s calf injury occurred during a practice warmup; Robinson and Thompson had game-related ankle injuries. Different exposure settings already matter. [21]

The 2026 download includes weeks 1–3 but its SF Week 3 statuses were still blank. The official final team report is more current. There is no completed 2026 AGL season; no annualized claim or full-season league rank is warranted.

12 / METHODS AND AUDIT

How to reproduce the findings

The accompanying data and code preserve the analytical choices.

Benchmark construction

For team t and season s , $R(t,s) = \text{AGL}(t,s) / \text{mean AGL}(\text{all 32 teams},s)$. The era score is the unweighted mean of R over the selected seasons. This neutralizes season length and common league-wide changes; it does not equal a count of injuries. COVID absences are excluded in both 2020 and 2021. Source revisions are recorded per row.

The 2017 original release uses an older AGL methodology. This is a compilation of public vintages, not a uniformly reprocessed historical dataset. Omitting 2017 leaves SF first. A separate per-game check uses 16 games through 2020 and 17 thereafter; AGL counts the canceled 2022 BUF–CIN game, as its publisher instructs.

Player-level data and cleaning

Join injury reports to schedules and snap participation using player IDs, team, season and week. Harmonize relocated team names. Match 99.91% of snap rows to player IDs. Keep a physical injury/condition component across four injury-text fields; remove explicit illness, COVID, rest and noninjury reasons plus unknown placeholders. Canceled-game report rows are excluded from this player-level analysis.

The final cohort has 178,854 eligible transitions. Features use prior report history, age, position, earlier snaps, index-game surface/home status, rest intervals and week. Earlier workload excludes the index game. Players must nevertheless participate in that game, so this is a conditional analysis, not a deployable pregame medical risk score.

Validation and uncertainty

Choose between two fixed models using 2023 average precision, refit through 2023 and evaluate on 2024–2025. No test-season tuning. Grouped importance uses 20 shuffles; performance uses 500 game-cluster bootstrap draws. Importance intervals use 300 game-cluster draws conditional on the first shuffle. Repeated players and teams remain a source of dependence; model-selection uncertainty is omitted.

Checks that matter

Coverage: full 2025 has 544/544 team-games with report rows. Earlier seasons have seven unmatched empty team reports, excluded from modeling; empty does not automatically mean healthy. Severe-status, no-team-feature and no-2017 sensitivities preserve modest performance. Raw report counts also put SF first: 441 Out/Doubtful entries in 149 games (2.96/game), versus 2.12 across the league. This is supplementary reporting evidence, not 441 injuries.

No medical records, training loads, diagnoses, acquired-player risk counterfactuals or injury-site records are available. Public concussion studies have missed substantial recorded cases; that validation finding must not be applied as a numeric completeness estimate to our all-condition dataset. [20]

13 / SOURCE REGISTER

Data sources

Publisher tables, archived originals and reproducible public datasets.

- [1] [Football Outsiders: 2017 AGL \(archived original\)](#)
- [2] [Football Outsiders: 2019 AGL, with revised 2018 column](#)
- [3] [Football Outsiders: 2020 AGL, COVID-excluded table](#)
- [4] [Football Outsiders: 2022 AGL, COVID-excluded 2021 column](#)
- [5] [FTN: 2023 AGL, with revised 2022 comparison](#)
- [6] [FTN: 2024 AGL; 49ers injury burden and revised 2023](#)
- [7] [FTN: 2025 AGL, with revised 2024 comparison](#)
- [8] [FTN: 2025 AGL totals by position](#)
- [9] [nflverse: injury/practice report releases](#)
- [10] [nflverse: snap-count releases](#)
- [11] [nflverse: player IDs and birth dates](#)
- [12] [nflverse: schedules and game context](#)

Every AGL row retains its source URL, archive URL, COVID convention and revision note. The reproducibility bundle includes downloaded CSV snapshots, hashes, scripts, model diagnostics and source notes.

AGL values: Football Outsiders / FTN. Injury reports and snaps: nflverse, from underlying public reporting. Analyses, normalizations, simulations and charts were calculated for this report. Accessed September 25, 2026 (Pacific).

14 / SOURCE REGISTER

Research and current context

Primary studies and official team reporting.

- [13] [Herzog et al. \(2026\): gradual preseason acclimation](#)
- [14] [Mack et al. \(2019\): surface and lower-extremity injury](#)
- [15] [Bodendorfer et al. \(2021\): NFL hamstring recurrence](#)
- [16] [Herzog et al. \(2023\): lower-extremity strains and exposure](#)
- [17] [Perez et al. \(2020\): Thursday games and injury rates](#)
- [18] [Angileri et al. \(2023\): travel, timing and NFL injuries](#)
- [19] [Mack et al. \(2020\): lower-extremity injury incidence](#)
- [20] [Matava et al. \(2023\): validity of public injury research](#)
- [21] [49ers: final Week 3 injury report, September 25, 2026](#)
- [22] [49ers: Kittle Achilles injury in the 2025 postseason](#)
- [23] [NFL: 2016 injury-report definition changes](#)
- [24] [Li et al. \(2020\): wearable workload and soft-tissue injury](#)

Responsible social framing: “The 49ers’ injury burden has been persistently high. Prior injury history and playing usage help predict new report appearances, but public data cannot tell us who or what caused the team’s excess.”