



Dr. Austin Brown & Prof. Michael Frankel

ABSTRACT

Background

- Re-injury rates among professional athletes remain concerning, with NBA players facing elevated risk upon return to play after initial injury
- Traditional return-to-play protocols rely on subjective assessments and physical benchmarks, but may miss subtle performance decline indicating incomplete recovery
- Statistical process control methods, specifically Exponentially Weighted Moving Average (EWMA) control charts, have successfully predicted pitcher re-injury in Major League Baseball by monitoring performance deterioration
- Net rating, a comprehensive metric reflecting point differential while a player is on court, provides objective measure of player impact and may reveal performance decline preceding re-injury

Objective

- Adapt EWMA control chart methodology from baseball to NBA basketball for post-injury performance monitoring
- Assess model's ability to provide advance warning before re-injury occurrence

METHODS

Data Source

- hoopR R package: Official NBA statistics for 2024-25 season
- Injury tracking: players missing games due to injury (excluded rest, suspension, personal reasons, G-League)

Injury Definitions & Player Selection

- Initial injury:** First occurrence of 2+ consecutive missed games due to injury
- Re-injury:** Any subsequent injury causing 2+ consecutive missed games after return
- Inclusion criteria:** (1) initial injury of 2+ games, (2) return to play, (3) 10+ consecutive games post-return

Net Rating Calculation

- Net rating = (Team Points - Opponent Points) $\times$ 100 / Team Possessions while player on court
- Higher values = stronger impact; lower values = declining effectiveness

EWMA Control Chart Implementation







- EWMA statistic:** weighted average emphasizing recent performance
- Lambda (λ) smoothing:**
 - higher = faster response, more volatility
- L (control limit):**
 - higher = fewer false alarms, more missed detections
- Out-of-control signal:** EWMA falls below lower control limit (performance significantly below league average)

Parameter Optimization

- Maximized F1-score while maintaining specificity $\geq$ 30%
- Optimal parameters: $\lambda = 0.25$, $L = 0.6$

Games Saved & Team Analysis

- Games saved: re-injury game - first signal game
- Team analysis: Pearson correlation and regression examined injuries vs. season wins



Metric	Value
Sensitivity	71 %
Specificity	30.6 %
Accuracy	46.2 %
Precision	39.3 %
F1-Score	0.506
Mean Games Saved	19.8 games
Median Games Saved	17.5 games
Re-Injuries Detected Early	22 of 58 re-injuries

Injuries vs Wins	
Correlation (r)	-0.178
P-value	0.428
Basic Model R^2	0.032
Injury Type Model R^2	0.079

RESULTS

Model Performance (Figure 4)

- Tuned parameters through grid search: $\lambda = 0.25$ (strong weight on recent games) and $L = 0.6$ (tight control limits).
- These settings produced high sensitivity (71.0%) and low specificity (30.6%), meaning the model successfully identified most re-injury patterns but also triggered frequent false alerts.
- Overall F1 = 0.506, indicating balanced though moderate effectiveness; recall was prioritized over precision.
- Trade-off: highly responsive to subtle declines in performance, which improves early detection but increases the number of false positives.

Early Warning Capability

- 22 of 58 re-injuries were detected before they occurred ($\approx 38\%$ of all injured players, $\approx 71\%$ recall among re-injured group).
- Average lead time: ≈ 20 games before the next absence (min = 2, max = 63).
- This window gives medical staff time to adjust playing time, adjust workloads, or add additional rehab.

Examples

- Jerami Grant (Figure 5A):** Model detected a sustained drop in net rating beginning two games after return, nine games before re-injury, representing a clear early warning.
- Caleb Martin (Figure 5B):** Model failed to trigger before re-injury; his performance remained stable and within limits throughout, illustrating an undetected case.

Team-Level Analysis (Figure 3)

- The analysis showed a weak negative correlation between total injuries and team wins ($r = -0.18$, $p = 0.43$).
- Teams with higher injury counts tended to follow the downward trend more consistently, while those with fewer stints added greater variability.
- This pattern suggests a potential relationship between injury burden and performance, but the effect is muted by limited data and uneven representation across teams.
- With additional seasons or a larger dataset, this relationship would likely become clearer and statistically stronger.

DISCUSSION

- Key Findings:** The control chart method detected re-injury risk in 71% of cases with an average 20-game advance warning. Although the false alarm rate was high (69%), the model showed that performance-based monitoring can reveal players at elevated risk before another absence. Team results appeared mostly unaffected by injury counts, likely due to limited sample size.
- Why This Matters:** The findings show that statistical tracking can highlight potential re-injury risk before medical symptoms appear, giving staff time to adjust workloads, modify rehab, or intervene to prevent setbacks.
- Limitations:** Net rating fluctuates by team context, making it difficult to isolate individual decline. The model only detects performance-related deterioration and misses sudden or contact injuries. A single season and 58 re-injured players limit generalization.
- Next Steps:** Future work should test more stable metrics such as usage rate or efficiency rating. Expanding to multiple seasons and testing in real time with team staff would help confirm whether early warnings reduce re-injuries.