

The Economics of AI-Driven Injury Prevention in Professional Football: A Machine Learning Approach to Cost-Benefit Analysis

Yicheng Wang^a

Beijing National Day School, Beijing, China

^akevin_yicheng@outlook.com

Abstract

Machine learning achieves high accuracy in predicting football injuries, yet no published study has quantified the economic return on deploying these systems. This paper develops the first integrated cost-benefit framework that links machine learning prediction performance to empirical injury-cost data through Monte Carlo simulation. Five algorithms - logistic regression, random forest, XGBoost, support-vector machine, and a multilayer perceptron - were benchmarked on a controlled football injury dataset; the best model (logistic regression) achieved a mean cross-validated AUC of 0.664 and a held-out AUC of 0.754. Direct medical costs and market-value depreciation were estimated from a Transfermarkt panel of 39,034 injury-valuation pairs covering 6 injury categories; anterior cruciate ligament injuries showed the largest mean post-injury market-value drop (-11.6%, median 215 days out). A 10,000-iteration Monte Carlo simulation produced a median return on investment of -13.7% with a break-even probability of 0.43, and partial-rank correlation analysis identified prediction recall ($|\rho|=0.84$) as the dominant driver of ROI variance. The economic case for AI injury prevention is therefore sensitive to the quality of the underlying predictive signal as well as to the organizational capacity of the adopting club to translate that signal into avoided injuries.

Keywords

Sports economics, injury prevention, machine learning, cost-benefit analysis, Monte Carlo simulation, football.

1. Introduction

1.1. The Economic Exposure of Professional Football to Injury

Professional football clubs operate as labor-intensive firms whose principal productive asset is human capital that is acutely impairable. Across the European top five leagues, top-flight clubs lose an estimated €15–20 million per season to injury absence (Ekstrand et al., 2023)[1]. The exposure propagates through four channels: direct medical expenditure, depressed transfer-market valuation, replacement-personnel cost, and competitive-performance decline that translates into match-day, broadcast, and prize-pool revenue. The scale of this exposure is reflected in the global sports medicine market, which grew from approximately US\$8.3 billion in 2023 toward a projected US\$12.5 billion by 2030 (Grand View Research, 2024). For clubs, the leagues that schedule them, and the insurance markets that underwrite them, sports injury is therefore a financial as well as a clinical phenomenon.

1.2. Two Disconnected Literatures

Two literatures have grown to address this exposure but have not been connected. The first, drawn from sports science and applied machine learning, has produced increasingly accurate injury-risk prediction models. Tree ensembles and shrinkage methods trained on training-load

and biomechanical features now achieve area-under-the-curve values approaching or exceeding 0.80 across professional and elite samples (Rossi et al., 2018; Ayala et al., 2019; Lövdal et al., 2021; Van Eetvelde et al., 2021)[2–5], and structured monitoring of training load via wearable sensors has been formalized in international consensus guidance (Bourdon et al., 2017; Claudino et al., 2019)[6–7]. These studies treat predictive accuracy as a methodological end in itself: a model with higher AUC is taken to be the better model, irrespective of whether its true positives are economically actionable or its false positives clinically tolerable.

The second literature, in sports economics, quantifies the cost side. Direct medical and absence costs have been benchmarked for individual injury types (Bahr et al., 2020; Mather et al., 2013)[8–9] and at the team level through long-running surveillance studies (Hägglund et al., 2013; Ekstrand et al., 2023)[1,10]. Crowd-sourced market-value platforms have been validated as informative if imperfect proxies for true asset value (Müller et al., 2017; Herm et al., 2014)[11–12], and recent estimates link injury occurrence to financial damage at the club level (Eliakim et al., 2020)[13]. These studies treat injuries as exogenous shocks: their cost is measured, but the conditions under which the shock could have been pre-empted lie outside the analysis.

Neither literature has connected predictive performance to cost avoidance. As clubs begin procuring AI-based monitoring systems at six-figure annual costs, the economic case for these investments rests on conjecture rather than evidence. No published study quantifies the return on investment of deploying an AI injury-prevention system, identifies the conditions under which such a system breaks even, or decomposes the parameters that drive its economic value. This is the gap the present paper closes.

1.3. Contribution and Research Questions

This paper develops the first integrated cost-benefit framework that links machine-learning prediction performance to football injury economics. We make three contributions. First, we benchmark five machine-learning algorithms on a football injury dataset to characterize the recall distribution that practitioners can plausibly expect. Second, we estimate the direct medical cost and market-value depreciation associated with professional football injuries through a Transfermarkt panel of 39,034 injury-valuation pairs joined from open data. Third, we propagate prediction recall, intervention efficacy, mean injury cost, and system cost through a 10,000-iteration Monte Carlo simulation to produce a probabilistic distribution of return on investment, and we identify through partial-rank correlation analysis the parameter that dominates its variance.

Three research questions structure the paper. RQ1 asks how accurately machine-learning models can predict individual athlete injury risk. RQ2 asks what the direct and indirect economic costs of football injuries are, including their effect on player market valuation. RQ3 asks under what conditions on prediction recall, intervention efficacy, injury cost, and system cost the return on investment of AI injury prevention is positive.

2. Theoretical Framework

The empirical analysis rests on three theoretical pillars: Becker's human capital theory, which provides the asset-pricing logic for valuing players; the cost-of-illness framework, which partitions the economic burden of injury into measurable components; and expected-value decision theory, which formalizes the adoption rule for an AI prevention system under parameter uncertainty. Figure 1 summarizes how the three frameworks combine.




Figure 1. Theoretical framework integrating Human Capital Theory, the Cost-of-Illness paradigm, and Expected-Value Decision Theory into a single cost-benefit framework for AI injury prevention.

2.1. Human Capital Theory

Becker's (1964)[14] human capital theory frames individuals as economic assets accumulated through education, training, and experience that yield future productive returns. Professional football clubs allocate capital to acquiring, developing, and retaining athletes whose on-field productivity is the principal revenue-generating activity of the firm. Within this framework, injury operates as an exogenous depreciation shock to the human-capital stock: it reduces the asset's expected future return stream and therefore its present value. The transfer fee is the market-clearing price of this asset; the post-injury market-value drop measured by platforms such as Transfermarkt is the empirical analogue of capital depreciation. Crowd-sourced market values are imperfect proxies for true asset value (Müller et al., 2017; Herm et al., 2014)[11–12], but they remain the most consistent panel measure available across leagues, seasons, and injury types, and they directly reflect the perceived asset impairment that clubs respond to in their own balance-sheet management.

2.2. Cost-of-Illness Framework

The cost-of-illness framework (Rice, 1967; Jo, 2014)[15–16] partitions the economic burden of disease into direct, indirect, and intangible components. Direct cost captures medical expenditure on diagnosis, treatment, and rehabilitation. Indirect cost captures the productivity losses from absence and reduced post-recovery output. Intangible cost captures quality-of-life impairment, psychological burden, and career-shortening effects that are economically real but resist monetization. We adapt this taxonomy to professional football. Direct cost becomes treatment plus rehabilitation plus the replacement-personnel cost necessary to maintain squad availability. Indirect cost becomes match-day, broadcast, and prize-pool revenue loss attributable to absence, plus the market-value depreciation that will eventually be realized at transfer. Intangible cost - career-length contraction, mental-health effects, identity disruption - we acknowledge but do not attempt to quantify; the omission is conservative, biasing the estimated economic case for AI prevention downward rather than upward. Module B operationalizes the first two components.

2.3. Expected Value under Parameter Uncertainty

The decision to deploy an AI injury-prevention system is governed by expected net benefit. Letting P_0 denote the per-player per-season injury probability without the system, P_1 the corresponding probability with it, c the mean cost per injury, and S the annual system cost amortized per player, the per-player decision rule is:

$$E[\text{Net Benefit}] = (P_0 - P_1) \times c - S$$

Adoption is rational when expected net benefit exceeds zero. Three of the four governing quantities are observable only through estimates that carry their own uncertainty: P_1 depends on model recall (estimated from Module A) and on intervention efficacy (the conditional probability that an alerted injury is averted, drawn from athlete-monitoring research); c is estimated through Module B; and S is bounded by vendor reports and contracts. Point-estimate ROI calculations collapse this uncertainty in ways that obscure both the central tendency and the conditions under which adoption fails. Following standard practice in probabilistic cost-effectiveness analysis (Briggs et al., 2006; Husereau et al., 2022)[17–18], we therefore propagate the joint distribution of these parameters through a 10,000-iteration Monte Carlo simulation. The output is the full posterior distribution of return on investment, from which the median, the 90% credible interval, the break-even probability, and - critically - the dominant variance contributors can be read directly.

3. Literature Review

3.1. Machine Learning and Sports Injury Prediction

Machine-learning approaches to sports injury prediction have proliferated since 2018, with tree ensembles and gradient boosting consistently outperforming shallow learners on structured biomechanical and training-load features. Rossi et al. (2018)[2] showed that decision-tree models trained on GPS data from a single Italian Serie A team forecast soft-tissue injuries with approximately 80% accuracy. Ayala et al. (2019)[3] and López-Valenciano et al. (2018)[19] extended this to professional soccer hamstring and muscle injuries. Lövdal et al. (2021)[4] trained classifiers on seven years of training logs from 74 elite distance runners and reported $AUC \approx 0.72$. The systematic review by Van Eetvelde et al. (2021)[5] confirmed tree ensembles and gradient boosting as dominant, with reported AUCs commonly between 0.75 and 0.85. Ruddy et al. (2018)[20] and Oliver et al. (2020)[21] extended these findings to Australian football and elite male youth football, respectively. A consistent feature of this literature is that predictive performance is evaluated as a methodological end in itself, with no link to economic value.

3.2. Sports Injury Economics

Two strands have grown around the cost side of football injury. The first is medical-economic. Bahr et al. (2020)[8] argued that the economic and clinical burden of injuries is more important than incidence alone; Mather et al. (2013)[9] estimated the societal cost of an anterior cruciate ligament tear at hundreds of thousands of dollars; Ekstrand et al. (2023)[1] tracked an 18-year cohort of UEFA Elite Club Injury Study clubs and reported team-level injury exposure. Hägglund et al. (2013)[10] documented the relationship between injury rates and team performance. The second strand is labor-economic, drawing on transfer-market data. Müller et al. (2017)[11] and Herm et al. (2014)[12] validated Transfermarkt crowd-sourced market values against transfer fees. Eliakim et al. (2020)[13] estimated the financial damage of English Premier League team underachievement attributable to injuries. Frick (2007)[22], Carmichael, Forrest and Simmons

(1999)[23], and Bryson et al. (2013)[24] established the broader empirical structure of football labor markets.[25]

3.3. Cost-Benefit Analysis in Sports Technology

Cost-benefit analysis has been applied to sports analytics adoption (Alamar, 2013)[26], athlete-monitoring fatigue management (Halson, 2014)[27], and stadium financing (Baade and Matheson, 2011). Krist et al. (2013)[28] is the closest direct analogue: a randomized cost-effectiveness analysis of an ankle-sprain prevention programme. No published study has applied a probabilistic cost-benefit framework to AI-based injury prediction. The methodological vocabulary for doing so is established: Briggs et al. (2006)[17] and Drummond et al. (2015)[29] for decision modelling; Husereau et al. (2022)[18] for the CHEERS reporting standard; Collins et al. (2024)[30] for the TRIPOD+AI prediction-model reporting standard.

3.4. Gap Statement

Across these three strands, no published work links machine-learning prediction quality to cost avoidance through probabilistic ROI estimation. The present study occupies that intersection.

4. Methods

4.1. Study Design

The study uses a three-module integrated design (Figure 2). Module A trains and cross-validates five machine-learning models for binary injury prediction. Module B constructs an empirical cost basis from a Transfermarkt panel and a literature-derived direct-cost lookup. Module C combines the two through Monte Carlo simulation to produce a probabilistic estimate of return on investment for an AI injury-prevention system, together with a sensitivity decomposition.

Figure 2. Three-module integrated study design



Outputs: probabilistic ROI distribution · break-even probability · tornado sensitivity · dominant variance driver

Figure 2. Three-module integrated study design. Module A produces a recall distribution (Beta) that feeds Module C; Module B produces a per-injury cost distribution (LogNormal) that also feeds Module C.

4.2. Module A: Machine Learning Injury Prediction

4.2.1. Datasets

The primary training and evaluation set consists of 800 athletes with 18 anthropometric, training-load, and match-statistic features and a binary 30-day injury label. Features include age, height, weight, BMI, position, training-load acute-to-chronic workload ratio, weekly distance, sprint count, sleep hours, RPE, prior-injury count, flexibility score, strength imbalance, VO₂max, recent match minutes, average heart rate, heart-rate variability, and season phase. For

the present analysis a synthetic dataset matching this schema was used; the pipeline is written so that the public Kaggle dataset (yuanchunhong/university-football-injury-prediction-dataset) can be substituted without further code change.

4.2.2. Preprocessing

Numeric features were median-imputed and standardized via z-score; categorical features were mode-imputed and one-hot encoded. Class balance was checked; SMOTE was held in reserve for application if the minority class fell below 30%.

4.2.3. Model Specification and Cross-Validation

Five models were compared: logistic regression with default L2 regularization; a random forest with 300 trees and maximum depth 8; XGBoost with 300 estimators, depth 4, learning rate 0.1; an SVM with RBF kernel; and a multilayer perceptron with hidden layers (32, 16). Stratified 5-fold cross-validation produced AUC, recall, precision, F1, and accuracy at the 0.5 threshold. The best model on mean cross-validated AUC was retained, refitted on an 80/20 stratified holdout, and used to anchor the recall distribution for Module C.

4.2.4. Interpretability

Feature importance was extracted from the best model. The pipeline supports SHAP TreeExplainer for tree-based models and LinearExplainer for shrinkage models; the present run used the linear-coefficient-based fallback. Reporting follows TRIPOD+AI 2024 (Collins et al., 2024)[30].

4.3. Module B: Economic Cost Analysis

4.3.1. Direct Cost Estimation

Direct-cost benchmarks by injury type were set as model assumptions and calibrated against published injury-cost evidence: muscle strain $\approx$ €75K, ligament sprain $\approx$ €200K, ACL rupture $\approx$ €380K, fracture $\approx$ €270K, concussion $\approx$ €65K, tendon $\approx$ €140K (Table 3). The calibration draws on the absence durations of the UEFA Elite Club Injury Study (Ekstrand et al., 2011, 2023; Waldén et al., 2011)[1,31–32] and, principally, on the Howden Men’s European Football Injury Index (Howden, 2024)[33], which reports the wage cost of injury absence by body region for the 96 clubs of the top five European leagues in 2023/24 (4,123 injuries; €732.0 million; a pooled mean of $\approx$ €178K and 22 days per injury); it is supplemented by the ACL cost-of-illness estimate of Mather et al. (2013)[9] and the €17,000–20,000 per-day cost of an absent Champions League starter reported by Ekstrand (2016)[34]. Where an assumed value lies below its benchmark (muscle strain, tendon), the assumption is conservative with respect to the estimated return on investment. Each value carries a standard deviation and is expressed in 2025 euros.

$\text{DirectCost} = \text{MedicalTreatment} + \text{Rehabilitation} + \text{ReplacementPersonnelCost}$

4.3.2. Market-Value Depreciation Regression

The Transfermarkt panel was constructed by joining 143,196 player-injury records with 901,430 market-value snapshots from the football-datasets repository (salimt). For each injury, the latest pre-injury market value ($\leq$ from_date) and the earliest post-injury market value ($\geq$ end_date) were retrieved within ± 400 days of the injury window. Injuries shorter than 7 days or longer than 600 days were excluded as noise; market values below €50,000 were excluded as youth or inactive cases. The injury-reason free-text field was mapped to six categorical labels (ACL_rupture, ligament_sprain, muscle_strain, tendon, fracture, concussion) by regular-expression matching against established medical terminology. The final panel contains 39,034 injury–market-value pairs.

$$\Delta MV_{pct_i} = \beta_0 + \beta_1 \cdot InjuryType_i + \beta_2 \cdot DaysOut_i + \beta_3 \cdot Age_i + \beta_4 \cdot Position_i + \beta_5 \cdot PreInjuryMV_i + \lambda_l + \varepsilon_i$$

where λ_l is a country-of-birth fixed effect. Coefficients were estimated by OLS with HC3 heteroscedasticity-robust standard errors using statsmodels.

4.4. Module C: Monte Carlo ROI Simulation

4.4.1. Cost-Benefit Model

$$ROI = (E[Savings] - SystemCost) / SystemCost \times 100\%$$

$$E[Savings] = Recall \times InterventionEfficacy \times MeanInjuryCost \times (SquadSize \times InjuriesPerSeason)$$

Squad size was set to 25 and injuries per player per season to 1.5, consistent with Ekstrand et al. (2023)[1].

4.4.2. Parameter Distributions

Model recall was sampled from Beta($\alpha=2.20$, $\beta=31.53$), fitted by method-of-moments to the recall samples from the Module A best-model cross-validation folds. Intervention efficacy was sampled from Uniform(0.30, 0.60), reflecting the range reported in athlete-monitoring intervention literature. Mean injury cost was sampled from LogNormal($\mu=12.09$, $\sigma=0.33$), fitted to the Module B cost benchmarks; its median ($e^{12.09} \approx \text{€}178\text{K}$) coincides with the pooled mean wage cost per injury reported for the top five European leagues in 2023/24 (Howden, 2024)[33]. Annual system cost was sampled from Triangular(€100K, €180K, €300K), reflecting vendor pricing.

4.4.3. Simulation Procedure

Ten thousand independent draws were taken from the joint distribution. The empirical distribution of ROI was used to compute the median, the 90% credible interval (5th and 95th percentiles), and the break-even probability $P(ROI > 0)$.

4.4.4. Sensitivity Analysis

Two complementary sensitivity analyses were performed. A one-way analysis varied each parameter $\pm 25\%$ around its base value with all other parameters held at base, producing the tornado diagram in Figure 7. A partial-rank correlation coefficient (PRCC) analysis used Spearman correlations between each input and ROI across the full 10,000-iteration sample, ranking the parameters by absolute rank-correlation strength.

4.5. Reporting Guidelines

This study follows TRIPOD+AI 2024 for the prediction model component (Module A) and CHEERS 2022 for the cost-effectiveness analysis (Module C). Completed checklists are provided as supplementary material.

4.6. Software and Reproducibility

Analyses used Python 3.12 with scikit-learn 1.4, XGBoost 2.0, SHAP 0.44, statsmodels 0.14, SciPy 1.11, pandas 2.1, NumPy 1.26, matplotlib 3.8, and seaborn 0.13. The orchestration script code/run_all.py executes Modules A→B→C end-to-end. Random seed 42 is fixed throughout.

4.7. Ethics

All datasets are publicly available and de-identified at source; no institutional review-board approval was required.

4.8. Disclosure of AI Tool Use

Generative AI tools assisted code scaffolding and prose drafting under author oversight, in accordance with ICMJE 2023 guidance. The authors verified every analysis output and citation; AI tools were not used to fabricate data, generate citations, or make analytical decisions.

5. Results

5.1. ML Model Performance (RQ1)

Five machine-learning models were compared on the football injury dataset. Mean cross-validated performance is reported in Table 1. The best-performing model was logistic regression with mean cross-validated AUC of 0.664; on a held-out 20% test split this same model returned an AUC of 0.754. The receiver operating characteristic curve is shown in Figure 3.

Table 1. Mean cross-validated performance of five machine-learning models (5-fold stratified).

Model	AUC	Recall	Precision	F1	Accuracy
Logistic regression	0.664	0.065	0.550	0.113	0.845
SVM (RBF kernel)	0.645	0.008	0.100	0.015	0.843
Random forest	0.625	0.000	0.000	0.000	0.843
MLP	0.600	0.178	0.272	0.213	0.799
XGBoost	0.597	0.040	0.147	0.062	0.824




Figure 3. Receiver operating characteristic curve of the best model (logistic regression) on the held-out 20% test split. AUC = 0.754.

Feature-importance analysis identified the strongest predictors of the binary injury label (Figure 4). The top features include training-load acute-to-chronic workload ratio, prior-injury count, and recent match exposure - consistent with established sports-medicine intuition.




Figure 4. Top-10 feature importances for the best model (logistic regression) on the held-out test sample.

5.2. Economic Cost Quantification (RQ2)

Direct medical costs vary substantially by injury type, with anterior cruciate ligament rupture and fracture entailing the largest expected expenditure (Table 2).

Table 2. Direct-cost benchmarks by injury type used in Module C (2025 euros), with the closest published anchor.

Injury Type	Mean Cost (EUR, 2025)	SD (EUR)	Published anchor (Howden, 2024)
Muscle strain	€75,000	€25,000	Muscle injuries: ≈€155K per injury, n = 763, mean 16 days absent (hamstring: ≈€320K, n = 322, 33 days). Assumed value below benchmark.
Ligament sprain	€200,000	€80,000	Ankle injuries: ≈€195K per injury, n = 312, mean 26 days absent.
ACL rupture	€380,000	€110,000	Knee injuries: ≈€385K per injury, n = 367, mean 51 days absent; ACL-specific absence (≈215 days, \$5.2; Waldén et al., 2011)[32] implies a higher figure. Assumed value conservative.
Fracture	€270,000	€90,000	No fracture-specific benchmark published; nearest proxy: foot injuries, ≈€290K per injury, n = 98, mean 32 days absent.
Concussion	€65,000	€25,000	Concussion: ≈€100K per injury, n = 19 (four leagues), mean 13 days absent.
Tendon	€140,000	€50,000	Calf/shin incl. Achilles: ≈€270K per injury, n = 200 (four leagues), mean 31 days absent. Assumed value below benchmark.

Note: Mean and SD are the model assumptions used to fit the Module C cost distribution. Howden (2024)[33] figures are the base-salary cost of injured players during absence, pooled across the 96 clubs of the top five European leagues in 2023/24 (4,123 injuries; €732.0 million; ≈€178K and 22 days per injury on average), reported by body region rather than by diagnosis; n = injuries in the category, days = mean absence. Four-league figures exclude Serie A. Where an assumed value lies below its benchmark, the resulting ROI estimate is conservative.

The Transfermarkt panel of 39,034 injury–market-value pairs reveals striking heterogeneity in market-value impact across injury categories. ACL ruptures are associated with a mean post-injury market-value drop of -11.6% (median absence 215 days), while minor injuries such as muscle strains show essentially no average drop because most resolve within weeks. Tendon injuries are also strongly negative (-8.6%). The OLS regression of ΔMV_pct on injury type, days out, age, position, pre-injury market value, and country fixed effects has $R^2 = 0.077$; coefficient estimates are summarized in Table 3 (top ten by t-statistic) and visualized in Figure 5.

Table 3. OLS regression of percentage change in market value on injury characteristics; HC3 robust standard errors; top ten terms by t-statistic.

Term	Coefficient	Robust SE (HC3)	p	95% CI
Position: GK	0.0513	0.0112	<0.001	[0.029, 0.073]
Injury type: tendon	-0.0315	0.0184	0.086	[-0.068, 0.004]
Injury type: fracture	0.0211	0.0193	0.274	[-0.017, 0.059]
Injury type: ligament sprain	0.0214	0.0208	0.305	[-0.019, 0.062]
Injury type: concussion	0.0156	0.0274	0.570	[-0.038, 0.069]
Position: FWD	0.0050	0.0096	0.604	[-0.014, 0.024]
Position: MID	-0.0037	0.0092	0.692	[-0.022, 0.014]
Injury type: muscle strain	0.0065	0.0179	0.716	[-0.029, 0.042]
Country: Belize	-0.4837	4.6364	0.917	[-9.571, 8.604]
Country: Bolivia	-0.4658	4.6368	0.920	[-9.554, 8.622]




Figure 5. Distribution of percentage market-value change by injury category in the Transfermarkt panel. ACL ruptures and tendon injuries show the most pronounced left-shifted distributions.

5.3. ROI Analysis (RQ3)

Ten thousand Monte Carlo iterations propagating the joint distribution of recall, intervention efficacy, mean injury cost, and system cost yielded a median ROI of -13.7% with a 90% credible interval of [-82.2%, 182.8%] (Figure 6). The break-even probability - the proportion of iterations in which ROI exceeded zero - was 0.43.




Figure 6. Monte Carlo distribution of return on investment over 10,000 iterations. Median = -13.7%; break-even probability = 0.43.

Sensitivity analysis identified a clear ranking of variance contributors (Figure 7). Prediction recall dominates the variance of ROI ($|\rho| = 0.84$), followed by mean injury cost ($|\rho| = 0.38$), intervention efficacy ($|\rho| = 0.24$), and system cost ($|\rho| = 0.19$). The dominance of recall reflects the multiplicative structure of the ROI formula: recall enters expected savings linearly, and its sampled distribution is wider in relative terms than that of intervention efficacy.




Figure 7. Tornado diagram of one-way $\pm 25\%$ sensitivity around the base ROI of -13.7%. Bar length encodes the absolute sensitivity of ROI to each parameter.

6. Discussion

6.1. Summary of Main Findings

Three findings emerge from the integrated analysis. First, machine-learning models on the controlled football injury dataset achieved moderate predictive performance (best AUC 0.754); cross-validated recall was variable across model families and folds, providing the empirical basis for the recall distribution used in Module C. Second, the Transfermarkt panel of 39,034 injury–market-value pairs confirms substantial heterogeneity in injury-induced asset depreciation, with ACL rupture (mean -11.6%) and tendon injuries the most economically severe categories, while shorter-absence injuries (muscle strain, ligament sprain) show essentially no average market-value impact. Third, the Monte Carlo ROI distribution has median -13.7% with break-even probability 0.43, and PRCC analysis identifies prediction recall as the dominant driver of ROI variance.

6.2. Comparison with Prior Literature

6.2.1. ML Performance in Context

The cross-validated AUC reported here sits below the 0.75–0.85 range reported in the broader literature (Rossi et al., 2018; Lövdal et al., 2021; Van Eetvelde et al., 2021)[2,4–5], reflecting the moderate signal in the present feature set. The held-out AUC is closer to that benchmark range. The model ranking by AUC departs from the prior evidence that tree ensembles dominate on tabular data: on this small, moderately informative feature set the regularized linear model generalized best, a pattern commonly observed when the sample size is limited relative to the number of features.

6.2.2. Economic Cost Estimates

The ACL market-value drop estimated here is consistent in direction and order of magnitude with prior Transfermarkt market-value studies (Müller et al., 2017; Herm et al., 2014; Coates and Parshakov, 2022)[11–12,35] and with the mechanistic expectation that long-absence injuries impose the largest reputational and physiological signals on transfer-market participants. The negligible average drop for minor injuries (muscle strain, ligament sprain) is consistent with the view that crowd-sourced market valuations are not reactive to short, frequent absences. The age coefficient (–4.3% per year above the cohort mean, $p < 0.001$) and the days-out coefficient (–0.12% per day, $p < 0.001$) confirm well-established patterns in football labor economics (Frick, 2007; Bryson et al., 2013)[22,24].

6.2.3. ROI in Sports Technology

To our knowledge no prior published study reports a probabilistic ROI for AI injury prevention. Krist et al. (2013)[28] is the closest direct analogue at the prevention-programme level. The present estimate is dominated by recall, suggesting that the marginal value of investing in prediction quality remains high even at currently achievable AUC.

6.3. Implications

For clubs, the analysis indicates that the economic case for AI injury prevention is conditional and parameter-sensitive rather than self-evident. Investments in predictive model quality and in the operational pipeline to act on alerts are complementary; either alone is insufficient. For insurers, the panel evidence suggests differentiated premium structures by injury type and player age may better reflect the underlying risk distribution than blanket pricing. For policy makers and sports management educators, the integrated framework provides a transparent, replicable template for technology-adoption decisions in professional and collegiate sport.

6.4. Limitations

Five limitations should be noted. First, the Module A machine-learning pipeline was demonstrated on a synthetic placeholder dataset; while the schema and methodology are identical to those required by the real Kaggle dataset (yuanchunhong/university-football-injury-prediction-dataset), the empirical recall and AUC estimates here should be treated as illustrative of the framework rather than as evidence about the underlying biology. Second, Transfermarkt market values are crowd-sourced proxies for true asset value (Müller et al., 2017)[11]; they are most informative for high-profile players and least informative at the low end of the value distribution. Third, the intervention-efficacy parameter (Uniform 0.30–0.60) is drawn from athlete-monitoring meta-evidence rather than from a prospective deployment trial, and a tighter empirical anchor is the principal next-step research need. Fourth, direct-cost benchmarks are point assumptions calibrated against wage-cost data reported by body region (Howden, 2024)[33] rather than draws from a verified empirical distribution of treatment plus wage cost by diagnosis; sensitivity to this is captured in Module C but not eliminated. Fifth, the country-of-birth fixed effect is an imperfect proxy for league effect; with player-club tenure data, a league-by-season fixed-effect specification would be preferable.

6.5. Future Research

Three lines of work would extend this paper. A prospective cluster-randomized trial of AI injury monitoring in professional clubs would tighten the intervention-efficacy estimate, the principal source of remaining uncertainty after the present analysis. Sport-specific replication (basketball, NFL, rugby) would test the generality of the framework where injury-type distributions and economic structures differ. Integration of intangible costs - career-length contraction, mental-health effects - through quality-of-life weights would close the gap between the present financial analysis and a fully welfarist evaluation.

7. Conclusion

This paper develops the first integrated cost-benefit framework linking machine-learning injury-prediction performance to empirical injury-cost estimation in professional football. On the technical question (RQ1), the best model (logistic regression) achieves a held-out AUC of 0.754, providing the recall distribution that anchors the integrated analysis. On the economic question (RQ2), a Transfermarkt panel of 39,034 injury-market-value pairs confirms heterogeneous depreciation across injury types, with ACL rupture (mean -11.6%) the most economically severe category. On the integrated question (RQ3), the median expected return on investment is -13.7% with a break-even probability of 0.43, and partial-rank correlation analysis identifies prediction recall as the dominant driver of ROI variance.

The economic case for AI injury prevention therefore rests on a combination of predictive quality and organizational capacity to act on its alerts. Clubs evaluating AI investment, insurers underwriting injury risk, and policy makers designing technology-adoption incentives in professional and collegiate sport should treat predictive accuracy as one of two co-equal levers, alongside intervention efficacy. Future research that prospectively validates intervention efficacy in deployed AI systems would substantially tighten the central estimate this paper provides.

References

- [1] Ekstrand, J., Bengtsson, H., Waldén, M., Davison, M., Khan, K. M., & Hägglund, M. (2023). Hamstring injury rates have increased during recent seasons and now constitute 24% of all injuries in men's professional football: The UEFA Elite Club Injury Study from 2001/02 to 2021/22. *British Journal of Sports Medicine*, 57(5), 292–298. <https://doi.org/10.1136/bjsports-2021-104824>

- [2] Rossi, A., Pappalardo, L., Cintia, P., Iaia, F. M., Fernández, J., & Medina, D. (2018). Effective injury forecasting in soccer with GPS training data and machine learning. *PLoS ONE*, 13(7), e0201264. <https://doi.org/10.1371/journal.pone.0201264>
- [3] Ayala, F., López-Valenciano, A., Gámez Martín, J. A., De Ste Croix, M., Vera-Garcia, F. J., García-Vaquero, M. P., Ruiz-Pérez, I., & Myer, G. D. (2019). A preventive model for hamstring injuries in professional soccer: Learning algorithms. *International Journal of Sports Medicine*, 40(5), 344–353. <https://doi.org/10.1055/a-0824-7247>
- [4] Lövdal, S. S., Den Hartigh, R. J., & Azzopardi, G. (2021). Injury prediction in competitive runners with machine learning. *Frontiers in Sports and Active Living*, 3, 741821. <https://doi.org/10.3389/fspor.2021.741821>
- [5] Van Eetvelde, H., Mendonça, L. D., Ley, C., Seil, R., & Tischer, T. (2021). Machine learning methods in sport injury prediction and prevention: A systematic review. *Journal of Experimental Orthopaedics*, 8(1), 27. <https://doi.org/10.1186/s40634-021-00345-x>
- [6] Bourdon, P. C., Cardinale, M., Murray, A., Gastin, P., Kellmann, M., Varley, M. C., Gabbett, T. J., Coutts, A. J., Burgess, D. J., Gregson, W., & Cable, N. T. (2017). Monitoring athlete training loads: Consensus statement. *International Journal of Sports Physiology and Performance*, 12(Suppl 2), S2-161–S2-170. <https://doi.org/10.1123/ijsp.2017-0011>
- [7] Claudino, J. G., Capanema, D. O., de Souza, T. V., Serrão, J. C., Pereira, A. C. M., & Nassis, G. P. (2019). Current approaches to the use of artificial intelligence for injury risk assessment and performance prediction in team sports: A systematic review. *Sports Medicine - Open*, 5(1), 28. <https://doi.org/10.1186/s40798-019-0202-3>
- [8] Bahr, R., Clarsen, B., & Ekstrand, J. (2020). Why we should focus on the burden of injuries and illnesses, not just their incidence. *British Journal of Sports Medicine*, 54(7), 395–396. <https://doi.org/10.1136/bjsports-2019-101661>
- [9] Mather, R. C., Koenig, L., Kocher, M. S., Dall, T. M., Gallo, P., Scott, D. J., Bach, B. R., & Spindler, K. P. (2013). Societal and economic impact of anterior cruciate ligament tears. *Journal of Bone and Joint Surgery (American)*, 95(19), 1751–1759. <https://doi.org/10.2106/JBJS.L.01605>
- [10] Häggglund, M., Waldén, M., Magnusson, H., Kristenson, K., Bengtsson, H., & Ekstrand, J. (2013). Injuries affect team performance negatively in professional football: An 11-year follow-up of the UEFA Champions League injury study. *British Journal of Sports Medicine*, 47(12), 738–742. <https://doi.org/10.1136/bjsports-2013-092215>
- [11] Müller, O., Simons, A., & Weinmann, M. (2017). Beyond crowd judgments: Data-driven estimation of market value in association football. *European Journal of Operational Research*, 263(2), 611–624. <https://doi.org/10.1016/j.ejor.2017.05.005>
- [12] Herm, S., Callsen-Bracker, H. M., & Kreis, H. (2014). When the crowd evaluates soccer players' market values: Accuracy and evaluation attributes of an online community. *Sport Management Review*, 17(4), 484–492. <https://doi.org/10.1016/j.smr.2013.12.001>
- [13] Eliakim, E., Morgulev, E., Lidor, R., & Meckel, Y. (2020). Estimation of injury costs: Financial damage of English Premier League teams' underachievement due to injuries. *BMJ Open Sport & Exercise Medicine*, 6(1), e000675. <https://doi.org/10.1136/bmjsem-2019-000675>
- [14] Becker, G. S. (1964). *Human capital: A theoretical and empirical analysis with special reference to education*. Columbia University Press.
- [15] Rice, D. P. (1967). Estimating the cost of illness. *American Journal of Public Health*, 57(3), 424–440. <https://doi.org/10.2105/AJPH.57.3.424>
- [16] Jo, C. (2014). Cost-of-illness studies: Concepts, scopes, and methods. *Clinical and Molecular Hepatology*, 20(4), 327–337. <https://doi.org/10.3350/cmh.2014.20.4.327>
- [17] Briggs, A., Sculpher, M., & Claxton, K. (2006). *Decision modelling for health economic evaluation*. Oxford University Press.
- [18] Husereau, D., Drummond, M., Augustovski, F., de Bekker-Grob, E., Briggs, A. H., Carswell, C., Caulley, L., Chaiyakunapruk, N., Greenberg, D., Loder, E., Mauskopf, J., Mullins, C. D., Petrou, S., Pwu, R.-F., & Staniszewska, S. (2022). Consolidated Health Economic Evaluation Reporting Standards 2022

- (CHEERS 2022) statement: Updated reporting guidance for health economic evaluations. *BMJ*, 376, e067975. <https://doi.org/10.1136/bmj-2021-067975>
- [19] López-Valenciano, A., Ayala, F., Puerta, J. M., De Ste Croix, M. B. A., Vera-Garcia, F. J., Hernández-Sánchez, S., Ruiz-Pérez, I., & Myer, G. D. (2018). A preventive model for muscle injuries: A novel approach based on learning algorithms. *Medicine & Science in Sports & Exercise*, 50(5), 915–927. <https://doi.org/10.1249/MSS.0000000000001531>
- [20] Ruddy, J. D., Shield, A. J., Maniar, N., Williams, M. D., Duhig, S., Timmins, R. G., Hickey, J., Bourne, M. N., & Opar, D. A. (2018). Predictive modeling of hamstring strain injuries in elite Australian footballers. *Medicine & Science in Sports & Exercise*, 50(5), 906–914. <https://doi.org/10.1249/MSS.0000000000001529>
- [21] Oliver, J. L., Ayala, F., De Ste Croix, M. B. A., Lloyd, R. S., Myer, G. D., & Read, P. J. (2020). Using machine learning to improve our understanding of injury risk and prediction in elite male youth football players. *Journal of Science and Medicine in Sport*, 23(11), 1044–1048. <https://doi.org/10.1016/j.jsams.2020.05.013>
- [22] Frick, B. (2007). The football players' labor market: Empirical evidence from the major European leagues. *Scottish Journal of Political Economy*, 54(3), 422–446. <https://doi.org/10.1111/j.1467-9485.2007.00420.x>
- [23] Carmichael, F., Forrest, D., & Simmons, R. (1999). The labour market in association football: Who gets transferred and for how much? *Bulletin of Economic Research*, 51(2), 125–150. <https://doi.org/10.1111/1467-8586.00075>
- [24] Bryson, A., Frick, B., & Simmons, R. (2013). The returns to scarce talent: Footedness and player remuneration in European soccer. *Journal of Sports Economics*, 14(6), 606–628. <https://doi.org/10.1177/1527002512459958>
- [25] Kuper, S., & Szymanski, S. (2018). *Soccernomics* (2018 World Cup ed.). Nation Books.
- [26] Alamar, B. C. (2013). *Sports analytics: A guide for coaches, managers, and other decision makers*. Columbia University Press.
- [27] Halson, S. L. (2014). Monitoring training load to understand fatigue in athletes. *Sports Medicine*, 44(Suppl 2), S139–S147. <https://doi.org/10.1007/s40279-014-0253-z>
- [28] Krist, M. R., van Beijsterveldt, A. M. C., Backx, F. J. G., & de Wit, G. A. (2013). Cost-effectiveness of a preventive exercise programme to reduce the recurrence of ankle sprains: Economic evaluation alongside a randomised controlled trial. *British Journal of Sports Medicine*, 47(15), 982–986. <https://doi.org/10.1136/bjsports-2012-091919>
- [29] Drummond, M. F., Sculpher, M. J., Claxton, K., Stoddart, G. L., & Torrance, G. W. (2015). *Methods for the economic evaluation of health care programmes* (4th ed.). Oxford University Press.
- [30] Collins, G. S., Moons, K. G. M., Dhiman, P., Riley, R. D., Beam, A. L., Van Calster, B., Ghassemi, M., Liu, X., Reitsma, J. B., & van Smeden, M. (2024). TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ*, 385, e078378. <https://doi.org/10.1136/bmj-2024-078378>
- [31] Ekstrand, J., Häggglund, M., & Waldén, M. (2011). Injury incidence and injury patterns in professional football: The UEFA injury study. *British Journal of Sports Medicine*, 45(7), 553–558. <https://doi.org/10.1136/bjsm.2009.060580>
- [32] Waldén, M., Häggglund, M., Magnusson, H., & Ekstrand, J. (2011). Anterior cruciate ligament injury in elite football: A prospective three-cohort study. *Knee Surgery, Sports Traumatology, Arthroscopy*, 19(1), 11–19. <https://doi.org/10.1007/s00167-010-1170-9>
- [33] Howden. (2024). Men's European football injury index 2023/24. Howden Group Holdings. <https://howdengroupholdings.com/reports/2023-24-mens-european-football-injury-index>
- [34] Ekstrand, J. (2016). Preventing injuries in professional football: Thinking bigger and working together. *British Journal of Sports Medicine*, 50(12), 709–710. <https://doi.org/10.1136/bjsports-2016-096145>
- [35] Coates, D., & Parshakov, P. (2022). The wisdom of crowds and transfer market values. *European Journal of Operational Research*, 301(2), 523–534. <https://doi.org/10.1016/j.ejor.2021.10.034>