A Bayesian Forecast of Mexico's Group A Outlook in the 2026 World Cup

Douglas Fabian Paris Merin
dmerin3@gatech.edu
April, 2026

Abstract—In football, results often are a combination of skill and luck (randomness), making point predictions difficult and misleading. The project approaches this by building a Bayesian model, capturing the uncertainty of the most popular sport. Combining recent match data, FIFA ranking points, and a Poisson-based goals model, the project estimates and simulates outcomes. Three model variants are explored, a standard Poisson model, one with match-level random effects, and a Dixon–Coles adjustment for low-scoring matches. These models are compared using diagnostics, posterior predictive checks, and PSIS-LOO. The results are expressed as probabilities rather than point estimates, providing a more realistic view of how football matches unfold.

1 INTRODUCTION

The 2026 FIFA World Cup will be the largest ever, with 48 teams divided into 12 groups, making it interesting from a modeling perspective. Mexico will host and compete in group A, along with South Africa, Korea Republic, and Czechia. The objective of the project is to, with Bayesian statistics, estimate group A's outlook, with special focus in Mexico's national team. Instead of producing a single, fixed prediction, we are able to output probabilities with uncertainty. This is useful because in football, even if a team is stronger on average than another the matches often result in not so expected scores.

This project, despite offering a wide range of options, centers around three primary questions.

1. How likely is Mexico to win Group A? Estimate with posterior uncertainty shown as a highest-density interval, Mexico's performance in group A.

2. How likely is Mexico to play the Round of 16 in Mexico City? In the official bracket, Match 92 (Winner Match 79 vs Winner Match 80) is played at Mexico City Stadium. Therefore, Mexico reaches that Mexico City Round of 16 if it wins

Match 79, implying Mexico wins Group A.

3. How likely is Mexico to reach the quarterfinals? This is simulated from every possible Mexico path: first, second, or qualified third in Group A. The "fifth game" is the psychological and sporting barrier that the Mexican National Football Team has not managed to overcome in World Cups played away from home (since 1986), since it represents reaching the Quarter-Finals (5th game)

Figure 1—2026 World Cup Groups. Source: ESPN Deportes.

2 DATA

The project uses two data sources obtained from two separate Kaggle discussions. One source is an international football match dataset containing historical match results, including match date, home team, away team, goals scored by each team, and whether the match was played at a neutral venue (mikesicarioleighton, 2026). The remaining source is FIFA ranking data over time, specifically FIFA ranking points for national teams (mikesicarioleighton, 2026).

Additionally, the participating teams for the world cup 2026 were obtained directly from the FIFA website. The 48 teams were sorted alphabetically and given a unique id.

The match data was restricted to teams participating in the World Cup and matches between January 1st, 2025, and March 31st, 2026, transformed into team-1 versus team-2 format. The data was transformed into team-1 versus team-2 format, for each match the dataset contains:

- match date
- team names
- goals scored by team 1 and team 2,
- numeric team identifiers,
- indicators for whether team 1 or team 2 had home or host advantage (1.0 for home advantage, 0.0 otherwise).

	date	home_team	away_team	t1_score	t2_score	team_1	team_2	H_1	H_2
0	2025-01-22	Uzbekistan	Jordan	0.0	0.0	47	25	0.0	0.0
1	2025-03-20	Canada	Mexico	0.0	2.0	8	27	0.0	0.0
2	2025-03-20	USA	Panama	0.0	1.0	45	32	1.0	0.0
3	2025-03-20	Croatia	France	2.0	0.0	11	18	1.0	0.0
4	2025-03-20	Netherlands	Spain	2.0	2.0	29	40	1.0	0.0
5	2025-03-20	Brazil	Colombia	2.0	1.0	6	9	1.0	0.0

***Figure* 2**—Transformed match data.

FIFA rankings is a standard ordinal ranking of FIFA points by team. Points are used instead of rankings, as they provide a continuous measure of team "strength" that is more suitable for modeling. Points were standardized before entering the model, by having mean zero and a variance 1, keeps the model stable and makes the prior for the ranking coefficient easier to interpret.

	country_full	total_points	rank_date	team_id	z_points
0	Spain	2165.0	2026-04-01	40	2.172882
1	Argentina	2113.0	2026-04-01	1	1.877757
2	France	2082.0	2026-04-01	18	1.701817
3	England	2020.0	2026-04-01	17	1.349937
4	Brazil	1984.0	2026-04-01	6	1.145620

***Figure* 3**—Transformed rankings data.

3 MODELING

Football scores are naturally modeled as int count data, there are no half or negative goals. For this reason, the baseline model uses a Poisson likelihood for goals scored by each team. This approach is commonly used, as goals are nonnegative integer numbers and most matches involve relatively low scores. Early work by Maher (1982) shows that a Poisson model provides a reasonably accurate description of football scores, especially when team-specific parameters are included (Maher, 1982).

Having match $m = 1, \ldots, n$ is played between team $i[m]$ and team $j[m]$. Let y_{m1} and y_{m2} be the goals scored by each team. The general structure is

$$y_{m1} \sim \text{Poisson}(\lambda_{m1}) y_{m2} \sim \text{Poisson}(\lambda_{m2})$$

The scoring rates are modeled on the log scale. In the baseline specification,

$$\log \lambda_{m1} = \mu + \theta_{i[m]} - \theta_{j[m]} + \gamma H_{m1}$$
$$\log \lambda_{m2} = \mu + \theta_{j[m]} - \theta_{i[m]} + \gamma H_{m2}$$

Common formulations like Maher (1982) and Dixon-Coles (1997) include attacking and defense rates of team i, j however by not having this detail in the data, the project follows this standard structure through generalized "strength" differences, plus a home advantage effect (Maher, Dixon and Coles, 1982, 1997).

Here, μ is the baseline scoring level, θ_k is the strength of team k, γ captures home or host advantage, and H_{m1}, H_{m2} indicate whether either team receives that advantage. The inclusion of a home advantage term is also established in the literature and has been shown to be statistically significant (Maher, 1982).

Team strength is modeled using FIFA ranking points:

$$\theta_k = \beta z_k$$

β controls how strongly FIFA ranking points translate into latent team strength. This specification makes the model simple and stable. Teams with higher FIFA points tend to have higher estimated strength, but the actual match results still update the parameters through the likelihood.

The priors chosen for the baseline model are weakly informative normal distributions, to stabilize estimation while remaining relatively ranged.

$$\mu \sim N(0, 2^2), \gamma \sim N(0, 1^2), \beta \sim N(0, 1^2)$$

3.1 Model 1: Baseline Poisson Model

Model 1 is the model described above. It assumes that, conditional on the scoring rates, FIFA points and home advantage, two goal counts are independent Poisson outcomes. This model is simple and easy to interpret, so it serves as the baseline.

The model was sampled using NUTS with 4 chains, 6,000 tuning draws, and 4,000 posterior draws per chain with target accept=0.95 . Diagnostics indicated a stable fit: there were no divergences, all $\hat{R} = 1.0$, and acceptable effective sample sizes for all parameters. Here is the model summary:

	mean	sd	hdi_3%	hdi_97%
mu	0.150	0.068	0.024	0.279
gamma	0.070	0.114	-0.149	0.281
beta	0.298	0.054	0.194	0.397

3.2 Model 2: Poisson model with match-level random effect

Model 2 extends Model 1 by adding a shared match-level random effect u_m to both teams' scoring rates:

$$\log \lambda_{m1} = \mu + \theta_{i[m]} - \theta_{j[m]} + \gamma H_{m1} + u_m,$$
$$\log \lambda_{m2} = \mu + \theta_{j[m]} - \theta_{i[m]} + \gamma H_{m2} + u_m.$$

The random shared match effect with a scale parameterization is used:

$$u_m = \sigma_u u_m^*, \qquad u_m^* \sim N(0, 1)$$

Adding the prior:

$$\sigma_u \sim \text{HalfNormal}(1)$$

This model allows some matches to be more open or closed. For example, a match-level effect can increase both teams' expected goals in a dynamic game or decrease both expected goals in a defensive game.

The model was sampled using NUTS with 4 chains, 6,000 tuning draws, and 4,000 posterior draws per chain with target accept=0.99 . Diagnostics indicated a stable fit: there were no divergences, all $\hat{R} = 1.0$, and acceptable effective sample sizes for all parameters.

Model 2 suggests a meaningful match-level random effect. The posterior mean of σ_u was ~ 0.36, with posterior standard deviation ~ 0.10, indicating non-ignorable match-to-match variation besides team strength and host effect. Here is the model summary:

	mean	sd	hdi_3%	hdi_97%
mu	0.082	0.083	-0.074	0.234
gamma	0.064	0.117	-0.157	0.280
beta	0.307	0.055	0.201	0.407
sigma_u	0.363	0.102	0.171	0.557

Having a look at the trace plot, the chains mix well, overlapping across posterior densities. This is consistent with the numerical stats and diagnostics.

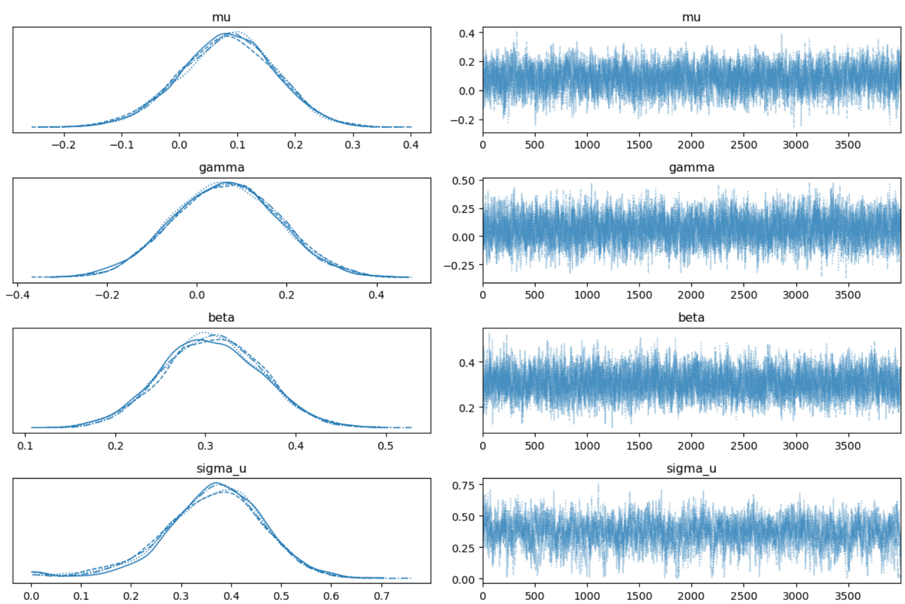

Figure 4—Model 2 Trace Plot.

3.3 Model 3: Poisson model with Dixon-Coles correction

Model 3 keeps the same scoring-rate specification as Model 1, but modifies the likelihood using a Dixon-Coles style correction for low-score outcomes (Dixon and Coles, 1997). The corrected likelihood is

$$p(x, y) = \text{Poisson}(x; \lambda_{m1})\text{Poisson}(y; \lambda_{m2})\tau(x, y; \lambda_{m1}, \lambda_{m2}, \rho),$$

Where τ adjusts the probabilities of the low-score outcomes $(0,0)$, $(1,0)$, $(0,1)$, and $(1,1)$. This correction is motivated by the fact that independent Poisson models often struggle to match the observed frequency of draws and other low-scoring results in football data, a feature emphasized by Dixon and Coles (1997) (Dixon and Coles, 1997).

The prior used for the correction parameter is

$$\rho \sim N(0, 0.1^2). \tag{1}$$

This model is explored because football, particularly in low-scoring matches are known to exhibit small deviations from the independent Poisson assumption, (Dixon and Coles, 1997).

The model was sampled using NUTS with 4 chains, 6,000 tuning draws, and 4,000 posterior draws per chain with target accept=0.99 . Diagnostics indicated a stable fit: there were no divergences, all $\hat{R} = 1.0$, and acceptable effective sample sizes for all parameters. Here is the model sumary:

	mean	sd	hdi_3%	hdi_97%
mu	0.151	0.067	0.020	0.273
gamma	0.069	0.111	-0.144	0.275
beta	0.299	0.054	0.197	0.400
rho	-0.058	0.072	-0.190	0.078

4 MODEL COMPARISON

The models were evaluated using sampling diagnostics, posterior predictive checks, and Pareto-smoothed importance sampling leave-one-out cross-validation (PSIS-LOO). As previously presented, models sampled without divergences, acceptable effective sample sizes, and unit $\hat{R}$ values.

PSIS-LOO provides an estimate of out-of-sample predictive performance by approximating leave-one-out cross-validation using importance sampling (ArviZ Developers, 2024). This method evaluates how well each model predicts each observation when that observation is left out of the fitting process.

For PSIS-LOO, the comparison was done at the match level. For Models 1 and 2, the match-level log-likelihood was computed as the sum of both teams' goal log-likelihoods. For Model 3, the custom Dixon–Coles likelihood was used directly.

Model	rank	elpd_loo	p_loo	elpd_diff	weight	se	dse	warning
Model 2	0	-410.07	35.85	0.00	0.97	12.80	0.00	False
Model 3	1	-412.78	4.00	2.71	0.03	13.81	2.56	False
Model 1	2	-412.83	3.58	2.76	0.00	13.84	2.58	False

The model comparison favored Model 2. Its expected log point-wise predictive density (ELPD) was slightly better than the other two models. The difference was not large, but maintaining a relatively simple structure while allowing additional match-level variation, model 2 was selected for the group-stage simulation. Model 1 and model 3 performed similarly, with no sufficient evidence that Dixon-Coles correction add value in this project.

Figure 5 complements the PSIS-LOO results by showing how model 1 and 2 behave in non-neutral low score outcomes . While the ELPD differences across models are modest, the plot helps illustrate where each stands in specific scorelines.

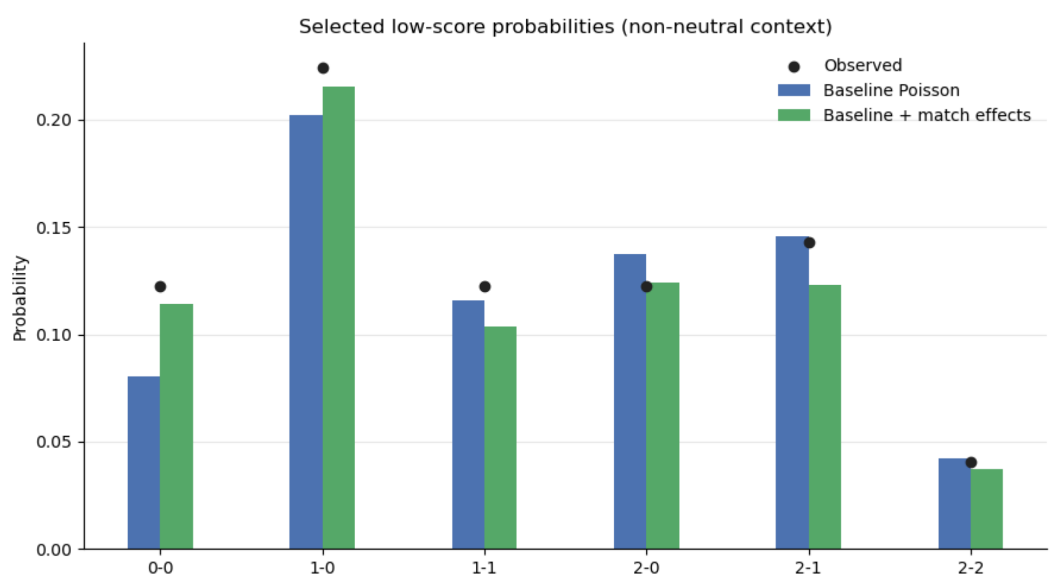

Figure 5—Observed frequencies and posterior predictive probabilities for non neutral low score outcome

5 SIMULATION

After selecting Model 2, the model was used to simulate through the Round of 16, allowing us to identify quarterfinal participants. More precisely, the methodology employed is as follows:

1. Simulate all group-stage matches for all 12 groups.
2. Rank every group.
3. Select best 3rd place teams.
4. Build Round of 32 using FIFA's Regulations Annex C. Fédération Internationale de Football Association, 2026
5. Simulate all Round of 32 matches.
6. Simulates all Round of 16 matches.
7. Record Round of 16 winners (Quarter-Finals participants).

Each simulated group was ranked using points, goal difference, goals scored, head-to-head performance among tied teams, and FIFA ranking points as a final deterministic tie-breaker. This is a simplified version of the official FIFA group ranking rules, but it keeps the simulation focused and functional.

For the knockout stage, if a simulated match ended in a draw, the winner was selected at random. This approximates the uncertainty of extra time and penalty shootouts without adding a separate penalty model.

For each posterior draw, the model parameters are fixed and then new match outcomes were generated by drawing a fresh match-level random effect and then drawing goals from the Poisson distributions. This process captures three types of uncertainty:

- uncertainty in the model parameters,
- match-to-match randomness,
- randomness in the actual goal outcomes.

The simulation used 1,000 posterior draws and 500 full tournament simulations per draw, producing 500,000 simulated tournament paths. For a tournament event A, the posterior predictive probability was estimated as

$$\hat{p}(A) = \frac{1}{SR} \sum_{s=1}^{S} \sum_{r=1}^{R} \mathbf{1}\{A_{sr}\},$$

where S is the number of posterior draws, R is the number of simulations per posterior draw, and $\mathbf{1}\{A_{sr}\}$ is an indicator for whether event A occurred in tournament simulation r under posterior draw s.

5.1 Mexico's Group A performance

Figure 6 shows the posterior distributions of Mexico's finishing probabilities in Group A, with a 94% HDI. The distributions reflect both parameter uncertainty and match-level randomness, Mexico finishing first is important because it places Mexico into Match 79 in the Round of 32 (played in Mexico city). Finishing second sends Mexico to Match 73, while some third-place outcomes can send Mexico to Match 82 or Match 74 depending on the full set of third-place qualifiers.

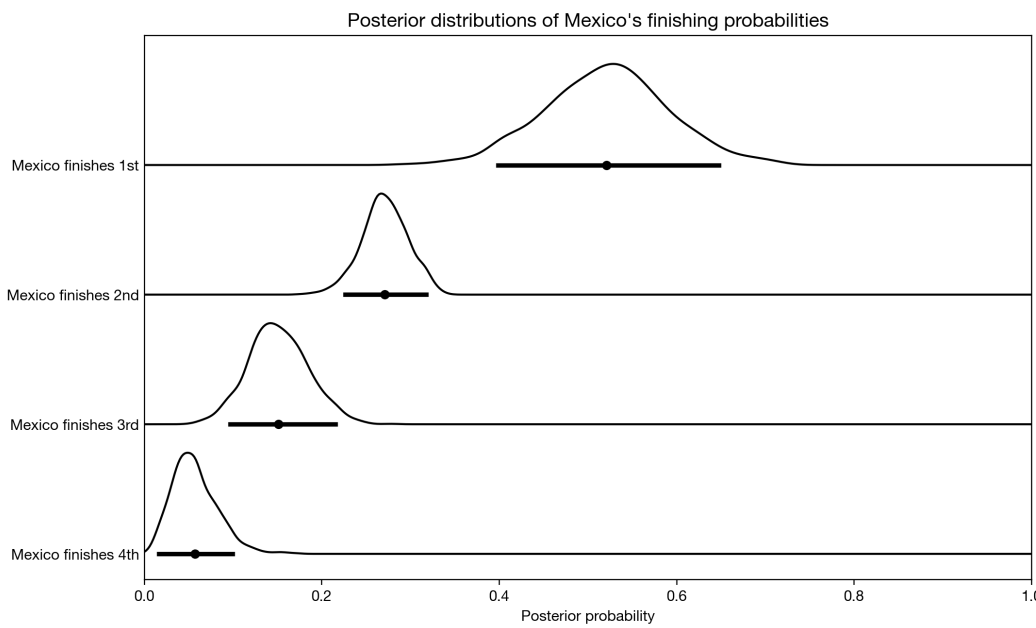

Figure 6—Posterior distributions of Mexico's performance probabilities in Group A.

Mexico has the highest probability of finishing first, with posterior mean 52% and a 94% credible interval of [40%, 65%], followed by a 27% probability of finishing second. The probability of finishing third or fourth is substantially lower. The outlook for South Africa is less favorable, with a low probability of advancing. While Korea Republic and Czechia appear to be closely competing for the second spot.

Rank	Mean Prob.	HDI (3%)	HDI (97%)
1	52.1%	39.6%	65.0%
2	27.1%	22.4%	32.0%
3	15.1%	9.4%	21.8%
4	5.7%	1.4%	10.2%

Table 1—Posterior summaries of Mexico's performance probabilities in Group A.

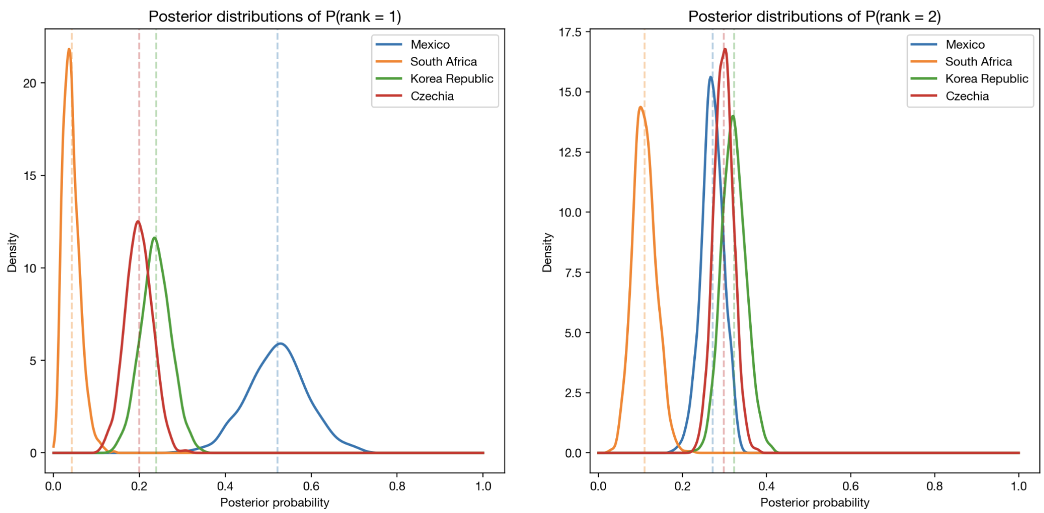

Figure 7—Posterior distributions of Group A's performance probabilities.

5.2 Round of 16 in Mexico City - Match 92

The 2026 bracket creates a specific path leading to Mexico City, Match 92 -Round of 16-. In the bracket, Match 92 is played between the winner of Match 79 and the winner of Match 80. Therefore, Mexico reaches Match 92 if it wins Group A, enters Match 79, and then wins Match 79.

The simulation estimates that Mexico plays Match 79 with mean probability 52.1%. Mexico reaches Match 92 with mean probability 31.8%, with a 94% HDI of $[22.8\%, 41.2\%]$. This means that the Mexico City Round of 16 path is not the most likely outcome, but it is also not a rare tail event under this model.

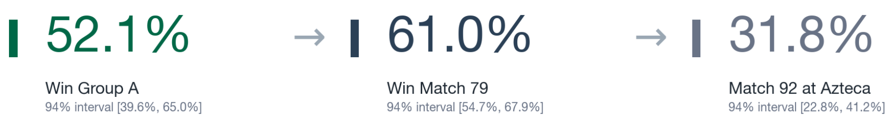

Figure 8—Road to Match 92.

Conditional on Mexico reaching Match 92, the opponent is the winner of Match 80. Because Match 80 depends on the winner of Group L and one of the qualified third-place teams, the possible opponent distribution is wide. In the simulations where Mexico reached Match 92, England was the most common opponent, followed by Croatia.

These probabilities should be interpreted conditionally. For example, England appears as Mexico's Match 92 opponent in 38.6% of the simulations in which Mexico reaches Match 92, not in 38.6% of all tournament simulations. Same for Match 79 Opponents.

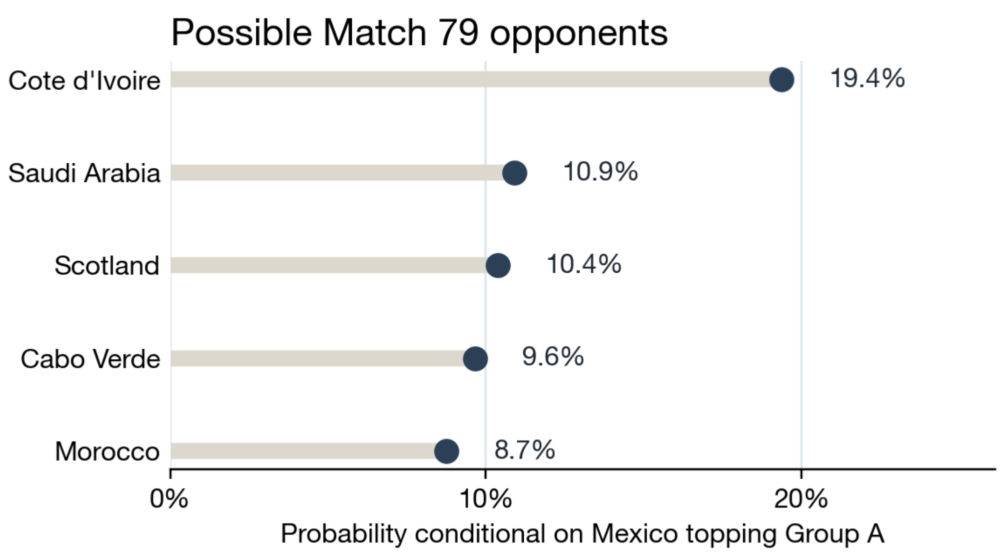

Figure 9—Possible Match 79 opponents

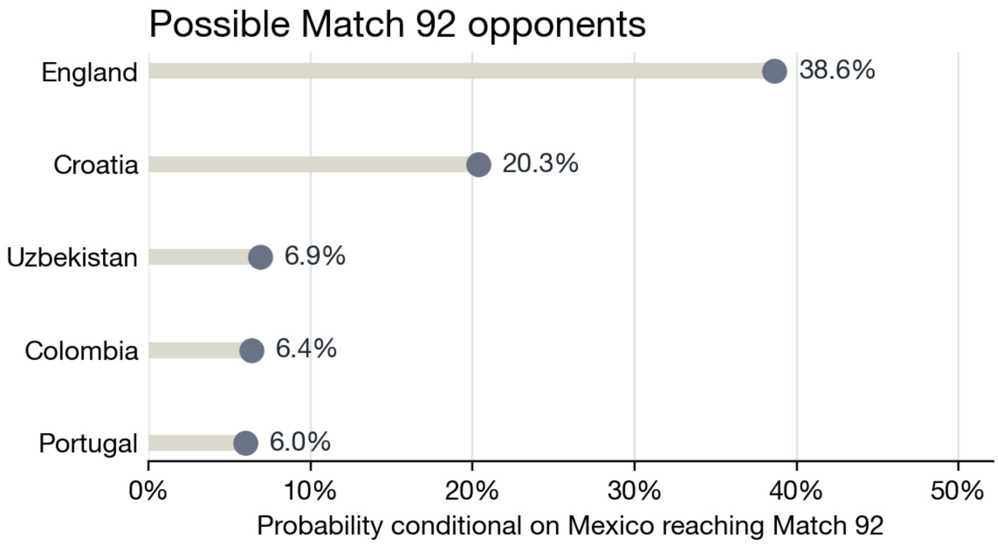

Figure 10—Possible Match 92 opponents

5.3 Quarter-Finals

In recent Mexican football folklore, a discussion often focuses on reaching the "fifth game" (quarter-finals). The "fifth game" is the psychological and sporting barrier that the Mexican National Football Team has not managed to reach quarter finals in World Cups since 1986. In this new 48-team format, the same milestone corresponds to reaching the sixth game. In this simulation, Mexico reaches the quarterfinals with posterior mean probability 25% and a 94% HDI of [21.2%, 29.0%].

This probability combines all possible Mexico paths, not only the Group A winner path. Mexico can reach the quarterfinals by winning its Round of 32 match and then winning its Round of 16 match from any bracket position. For this reason, the quarterfinal probability is lower than the Round of 16 probability, but it is not limited only to the Mexico City route.

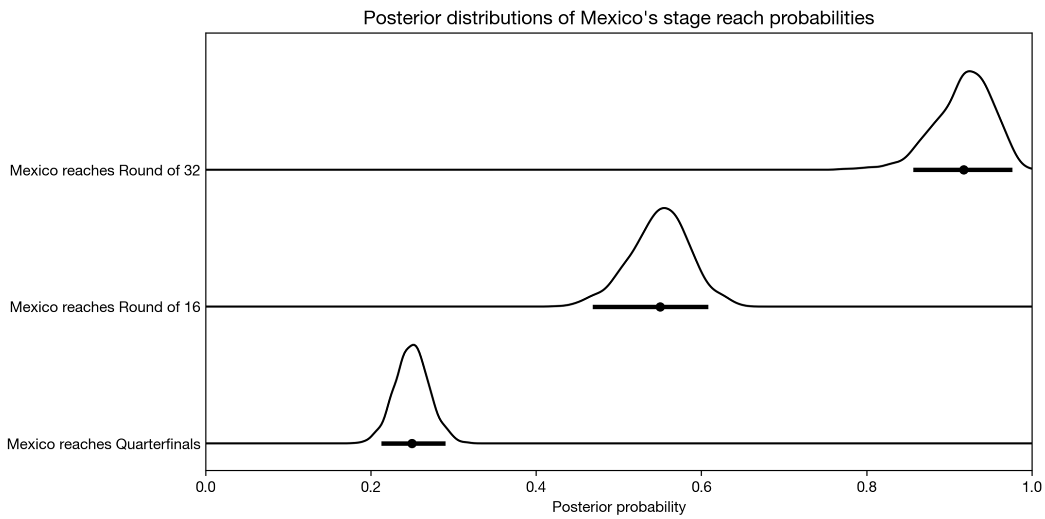

Figure 11—Mexico's stage probabilities.

Mexico is heavily favored to advance out of Group A and has a plausible path to a Mexico City Round of 16, but each additional round adds meaningful uncertainty. The Bayesian simulation makes this explicit: the result is not a single bracket prediction, but a distribution of possible tournament paths under the selected model assumptions.

6 CONCLUSION

Tackling the project with a Bayesian approach provides results that are not just a single prediction. Instead, it produces a distribution over Mexico's (and others) possible tournament paths and stage-reaching probabilities. This makes the results closer to real football environment, Mexico may have a central estimated probability of winning the group or advancing, but there is also uncertainty around that estimate.

Event	Mean Prob.	HDI (3%)	HDI (97%)
Qualifies to Round of 32	91.7%	85.6%	97.6%
Plays Match 79	52.1%	39.6%	65.0%
Reaches Round of 16	54.9%	46.8%	60.8%
Reaches Match 92	31.8%	22.8%	41.2%
Reaches quarterfinals	25.0%	21.2%	29.0%

Table 2—Posterior predictive probabilities for Mexico's main tournament milestones.

This project used Bayesian modeling to estimate the outlook of Mexico in the 2026 World Cup. While focused on Mexico, it can be used for other national team's analysis and predictions with uncertainty. The selected model was a points informed Poisson model with a match-level random effect. FIFA ranking points provided a simple and interpretable belief of team strength, while the match-level effect allowed individual matches to vary behavior.

7 REFERENCES

[1] ArviZ Developers (2024). *arviz.loo — Pareto-smoothed importance sampling leave-one-out cross-validation*. https://python.arviz.org/projects/stats/en/stable/api/generated/arviz_stats.loo.html. Accessed April 14, 2026.

[2] Dixon, Mark J. and Coles, Stuart G. (1997). "Modelling Association Football Scores and Inefficiencies in the Football Betting Market". In: *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 46.2, pp. 265–280.

[3] Fédération Internationale de Football Association (Mar. 2026). *Regulations for the FIFA World Cup 26™: 11 June–19 July 2026*. PDF document.

[4] Maher, M. J. (1982). "Modelling association football scores". In: *Statistica Neerlandica* 36.3, pp. 109–118.

[5] mikesicarioleighton (2026a). *FIFA World Ranking*. Kaggle dataset discussion. Accessed April 07, 2026. URL: https://www.kaggle.com/datasets/cashncarry/fifaworldranking/discussion/687166.

[6] mikesicarioleighton (2026b). *International Football Results from 1872 to 2017*. Kaggle dataset discussion. Accessed April 07, 2026. URL: https://www.kaggle.com/datasets/martj42/international-football-results-from-1872-to-2017/discussion/687168.