Who Will (Most Likely) Win the 2018 FIFA World Cup?

Achim Zeileis

https://eeecon.uibk.ac.at/~zeileis/

2018 FIFA World Cup prediction

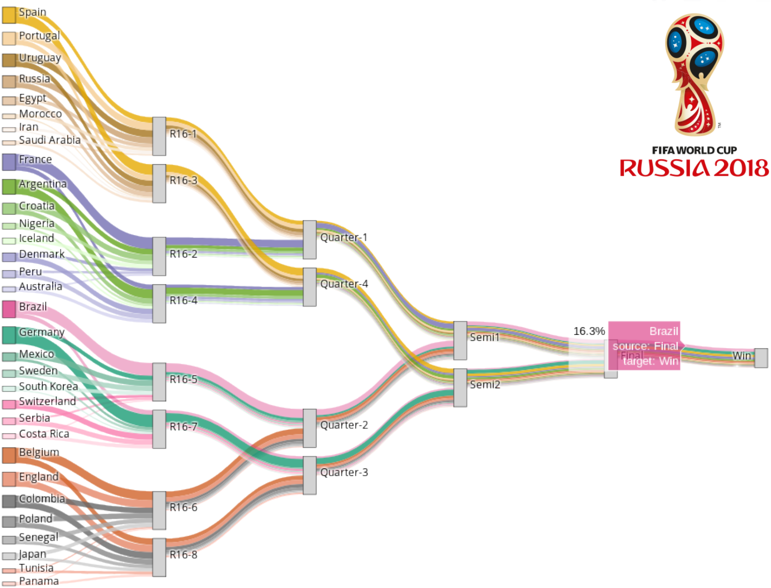

Source: Zeileis, Wikipedia

2018 FIFA World Cup prediction

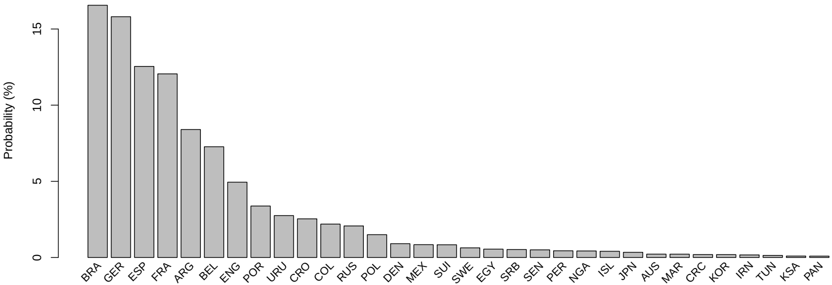

- Tournament forecast based on bookmakers odds.
- Main results: Brazil and Germany are the top favorites with winning probabilities of 16.6% and 15.8%, respectively.
- Top favorites are most likely to meet in the final (5.5%), then with odds very slightly in favor of Brazil (50.6% winning probability).

Bookmakers odds

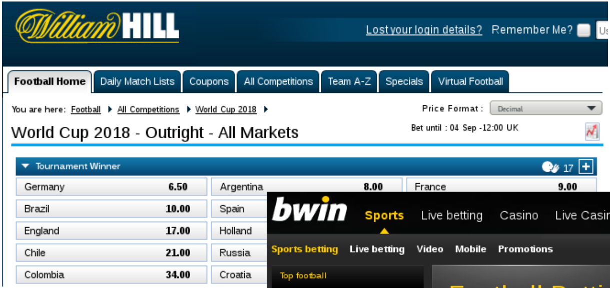

Source: williamhill.com, bwin.com

Bookmakers odds: Motivation

Forecasts of sports events:

- Increasing interest in forecasting of competitive sports events due to growing popularity of online sports betting.
- Forecasts often based on ratings or rankings of competitors' ability/strength.

In football:

- Elo rating.
  - Aims to capture relative strength of competitors yielding probabilities for pairwise comparisons.
  - Originally developed for chess.
- FIFA rating.
  - Official ranking, used for seeding tournaments.
  - Often criticized for not capturing *current* strengths well.
  - June 2018: Decision to change calculation to be more similar to Elo.

Bookmakers odds: Motivation

Alternatively: Employ bookmakers odds for winning a competition.

- Bookmakers are “experts” with monetary incentives to rate competitors correctly. Setting odds too high or too low yields less profits.
- Prospective in nature: Bookmakers factor not only the competitors abilities into their odds but also tournament draws/seedings, home advantages, recent events such as injuries, etc.
- Statistical “post-processing” needed to derive winning probabilities and underlying abilities.

Bookmakers odds: Statistics

Odds: In statistics, the ratio of the probabilities for/against a certain event,

$$odds = \frac{p}{1-p}.$$

Bookmakers odds: Statistics

Odds: In statistics, the ratio of the probabilities for/against a certain event,

$$odds = \frac{p}{1-p}.$$

Illustrations:

- Even odds are “50:50” ($= 1$).
- Odds of 4 correspond to probabilities $4/5 = 80\%$ vs. $1/5 = 20\%$.

Bookmakers odds: Statistics

Odds: In statistics, the ratio of the probabilities for/against a certain event,

$$odds = \frac{p}{1-p}.$$

Illustrations:

- Even odds are “50:50” ($= 1$).
- Odds of 4 correspond to probabilities $4/5 = 80\%$ vs. $1/5 = 20\%$.

Thus: Odds can be converted to probabilities and vice versa.

$$\begin{aligned} p &= \frac{odds}{odds+1} \\ 1-p &= \frac{1}{odds+1} \end{aligned}$$

Bookmakers odds: Quoted odds

Quoted odds: In sports betting, the payout for a stake of 1.

Bookmakers odds: Quoted odds

Quoted odds: In sports betting, the payout for a stake of 1.

Fair bookmaker: Given the probability p for the event the bookmaker could set

$$\textit{quoted odds} = \frac{1-p}{p} + 1.$$

Bookmakers odds: Quoted odds

Quoted odds: In sports betting, the payout for a stake of 1.

Fair bookmaker: Given the probability p for the event the bookmaker could set

$$\textit{quoted odds} = \frac{1-p}{p} + 1.$$

Expected payout: Wins and losses cancel out each other.

$$p \cdot \frac{1-p}{p} - (1-p) \cdot 1 = 0.$$

Bookmakers odds: Quoted odds

Quoted odds: In sports betting, the payout for a stake of 1.

Fair bookmaker: Given the probability p for the event the bookmaker could set

$$\textit{quoted odds} \; = \; \frac{1-p}{p} + 1.$$

Expected payout: Wins and losses cancel out each other.

$$p \cdot \frac{1-p}{p} - (1-p) \cdot 1 \; = \; 0.$$

Thus: "Naive" computation of probability

$$p \; = \; \frac{1}{\textit{quoted odds}}.$$

Bookmakers odds: Quoted odds

Illustration: Quoted odds for bwin obtained on 2018-05-20.

Team	Quoted odds	"Naive" probability
Brazil	5.0	0.200
Germany	5.5	0.182
Spain	7.0	0.143
France	7.5	0.133
	⋮	
Saudi Arabia	501.0	0.002
Panama	1001.0	0.001

Bookmakers odds: Quoted odds

Illustration: Quoted odds for bwin obtained on 2018-05-20.

Team	Quoted odds	"Naive" probability
Brazil	5.0	0.200
Germany	5.5	0.182
Spain	7.0	0.143
France	7.5	0.133
	⋮	
Saudi Arabia	501.0	0.002
Panama	1001.0	0.001

Problem: Probabilities of all 32 teams sum to $1.143 > 1$.

Bookmakers odds: Adjustment

Reason: Bookmakers do not give honest judgment of winning chances but include a profit margin known as “overround”.

Simple solution: Adjust quoted odds by factor 1.143 so that probabilities sum to 1.

Bookmakers odds: Adjustment

Reason: Bookmakers do not give honest judgment of winning chances but include a profit margin known as “overround”.

Simple solution: Adjust quoted odds by factor 1.143 so that probabilities sum to 1.

Team	Adjusted odds	Probability
Brazil	5.71	0.175
Germany	6.28	0.159
Spain	8.00	0.125
France	8.57	0.117
	⋮	

Bookmakers odds: Overround

Refinement: Apply adjustment only to the odds, not the the stake.

$$\textit{quoted odds}_i \;=\; \textit{odds}_i \cdot \delta + 1,$$

- where $\textit{odds}_i$ is the bookmaker's "true" judgment of the odds for competitor i,
- δ is the bookmaker's payout proportion (overround: $1 - \delta$),
- and $+1$ is the stake.

Bookmakers odds: Overround

Winning probabilities: The adjusted $odds_i$ then corresponding to the odds of competitor i for losing the tournament. They can be easily transformed to the corresponding winning probability

$$p_i = \frac{1}{odds_i + 1}.$$

Determining the overround: Assuming that a bookmaker's overround is constant across competitors, it can be determined by requiring that the winning probabilities of all competitors (here: all 32 teams) sum to 1: $\sum_i p_i = 1$.

Bookmakers odds: 2018 FIFA World Cup

Data processing:

- Quoted odds from 26 online bookmakers.
- Obtained on 2018-05-20 from `http://www.bwin.com/` and `http://www.oddschecker.com/`.
- Computed overrounds $1 - \delta_b$ individually for each bookmaker $b = 1, \dots, 26$ by unity sum restriction across teams $i = 1, \dots, 32$.
- Median overround is 15.2%.
- Yields overround-adjusted and transformed winning probabilities $p_{i,b}$ for each team i and bookmaker b.

Modeling consensus and agreement

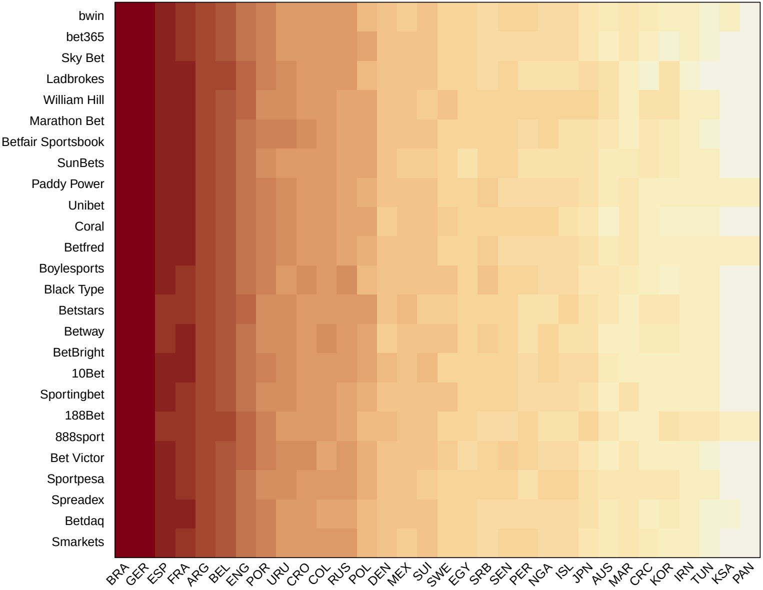

Modeling consensus and agreement

Goal: Get consensus probabilities by aggregation across bookmakers.

Straightforward: Compute average for team *i* across bookmakers.

$$\bar{p}_i = \frac{1}{26} \sum_{b=1}^{26} p_{i,b}.$$

Modeling consensus and agreement

Goal: Get consensus probabilities by aggregation across bookmakers.

Straightforward: Compute average for team i across bookmakers.

$$\bar{p}_i \;=\; \frac{1}{26} \sum_{b=1}^{26} p_{i,b}.$$

Refinements:

- Statistical model assuming for latent consensus probability p_i for team i along with deviations $\varepsilon_{i,b}$.
- Additive model is plausible on suitable scale, e.g.,

$$\mathrm{logit}(p) \;=\; \log\left(\frac{p}{1-p}\right).$$

Modeling consensus and agreement

Model: Bookmaker consensus model

$$\text{logit}(p_{i,b}) \; = \; \text{logit}(p_i) \; + \; \varepsilon_{i,b},$$

where further effects could be included, e.g., group effects in consensus logits or bookmaker-specific bias and variance in $\varepsilon_{i,b}$.

Modeling consensus and agreement

Model: Bookmaker consensus model

$$\text{logit}(p_{i,b}) \;=\; \text{logit}(p_i) \;+\; \varepsilon_{i,b},$$

where further effects could be included, e.g., group effects in consensus logits or bookmaker-specific bias and variance in $\varepsilon_{i,b}$.

Analogously: Methodology can also be used for consensus ratings of default probability in credit risk rating of bank b for firm i.

Modeling consensus and agreement

Here:

- Simple fixed-effects model with zero-mean deviations.
- Consensus logits are simply team-specific means across bookmakers:

$$\widehat{\text{logit}(p_i)} \;=\; \frac{1}{26}\sum_{b=1}^{26} \text{logit}(p_{i,b}).$$

- Consensus winning probabilities are obtained by transforming back to the probability scale:

$$\hat{p}_i \;=\; \text{logit}^{-1}\left(\widehat{\text{logit}(p_i)}\right).$$

- Model captures 98.7% of the variance in $\text{logit}(p_{i,b})$ and the associated estimated standard error is 0.184.

Modeling consensus and agreement

Team	FIFA code	Probability	Log-odds	Log-ability	Group
Brazil	BRA	16.6	−1.617	−1.778	E
Germany	GER	15.8	−1.673	−1.801	F
Spain	ESP	12.5	−1.942	−1.925	B
France	FRA	12.1	−1.987	−1.917	C
Argentina	ARG	8.4	−2.389	−2.088	D
Belgium	BEL	7.3	−2.546	−2.203	G
England	ENG	4.9	−2.957	−2.381	G
Portugal	POR	3.4	−3.353	−2.486	B
Uruguay	URU	2.7	−3.566	−2.566	A
Croatia	CRO	2.5	−3.648	−2.546	D
		⋮			

Abilities and tournament simulations

$$\Pr(i \text{ beats } j) \;=\; \pi_{i,j}$$

$$= \frac{\mathit{ability}_i}{\mathit{ability}_i + \mathit{ability}_j}$$

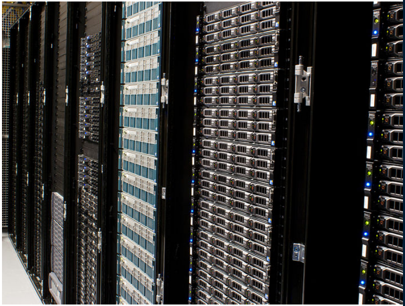

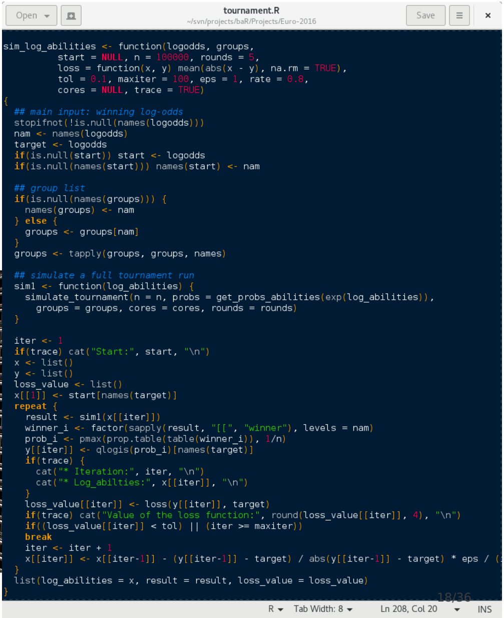


```r
sim_log_abilities <- function(logodds, groups,
          start = NULL, n = 100000, rounds = 5,
          loss = function(x, y) mean(abs(x - y), na.rm = TRUE),
          tol = 0.1, maxiter = 100, eps = 1, rate = 0.8,
          cores = NULL, trace = TRUE)
{
  ## main input: winning log-odds
  stopifnot(!is.null(names(logodds)))
  nam <- names(logodds)
  target <- logodds
  if(is.null(start)) start <- logodds
  if(is.null(names(start))) names(start) <- nam

  ## group list
  if(is.null(names(groups))) {
    names(groups) <- nam
  } else {
    groups <- groups[nam]
  }
  groups <- tapply(groups, groups, names)

  ## simulate a full tournament run
  sim1 <- function(log_abilities) {
    simulate_tournament(n = n, probs = get_probs_abilities(exp(log_abilities)),
      groups = groups, cores = cores, rounds = rounds)
  }

  iter <- 1
  if(trace) cat("Start:", start, "\n")
  x <- list()
  y <- list()
  loss_value <- list()
  x[[1]] <- start[names(target)]
  repeat {
    result <- sim1(x[[iter]])
    winner_i <- factor(sapply(result, "[[", "winner"), levels = nam)
    prob_i <- pmax(prop.table(table(winner_i)), 1/n)
    y[[iter]] <- qlogis(prob_i)[names(target)]
    if(trace) {
      cat("* Iteration:", iter, "\n")
      cat("* Log_abilties:", x[[iter]], "\n")
    }
    loss_value[[iter]] <- loss(y[[iter]], target)
    if(trace) cat("Value of the loss function:", round(loss_value[[iter]], 4), "\n")
    if((loss_value[[iter]] < tol) || (iter >= maxiter))
    break
    iter <- iter + 1
    x[[iter]] <- x[[iter-1]] - (y[[iter-1]] - target) / abs(y[[iter-1]] - target) * eps / (
  }
  list(log_abilities = x, result = result, loss_value = loss_value)
}
```


Source: Wikipedia, Zeileis

Abilities and tournament simulations

Further questions:

- What are the likely courses of the tournament that lead to these bookmaker consensus winning probabilities?
- Is the team with the highest probability also the strongest team?
- What are the winning probabilities for all possible matches?

Motivation:

- Tournament draw might favor some teams.
- Tournament schedule was known to bookmakers and hence factored into their quoted odds.
- Can abilities (or strengths) of the teams be obtained, adjusting for such tournament effects?

Abilities and tournament simulations

Answer: Yes, an approximate solution can be found by simulation when

- adopting a standard model for paired comparisons (i.e., matches),
- assuming that the abilities do not change over the tournament.

Model: Bradley-Terry model for winning/losing in a paired comparison of team i and team j.

$$\Pr(i \text{ beats } j) = \pi_{i,j} = \frac{ability_i}{ability_i + ability_j}.$$

Abilities and tournament simulations

"Reverse" simulation:

- If the team-specific $\mathit{ability}_i$ were known, pairwise probabilities $\pi_{i,j}$ could be computed.
- Given $\pi_{i,j}$ the whole tournament can be simulated (assuming abilities do not change and ignoring possible draws during the group stage).
- Using "many" simulations (here: 1,000,000) of the tournament, the empirical relative frequencies $\tilde{p}_i$ of each team i winning the tournament can be determined.
- Choose $\mathit{ability}_i$ for $i = 1, \dots, 32$ such that the simulated winning probabilities $\tilde{p}_i$ approximately match the consensus winning probabilities $\hat{p}_i$.
- Found by simple iterative local search starting from log-odds.

Abilities and paired comparisons

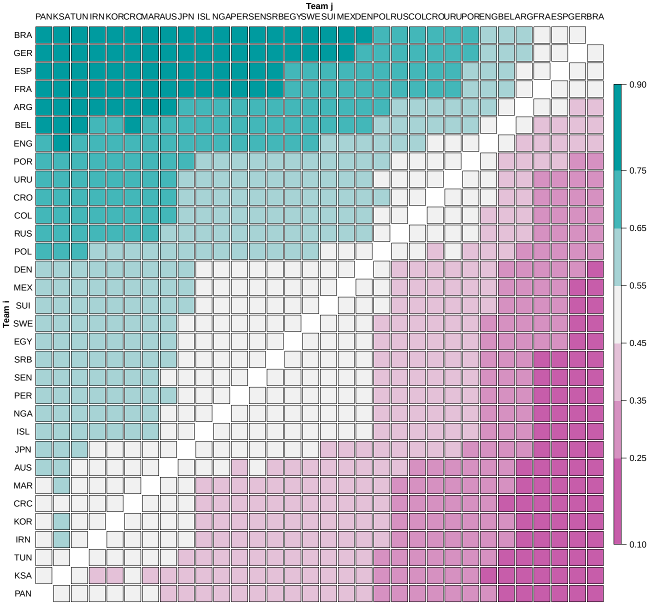

Tournament simulations: Survival curves

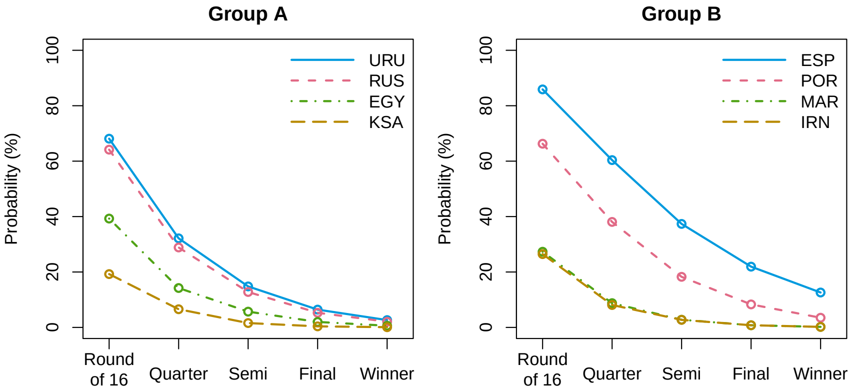

Tournament simulations: Survival curves

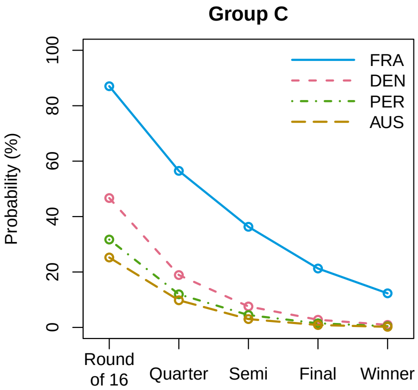

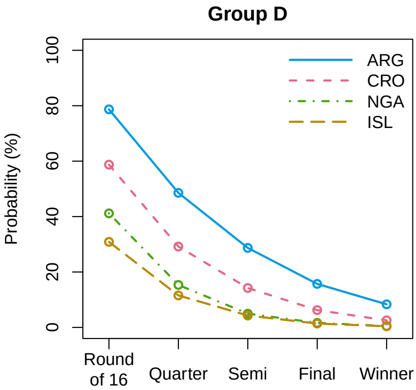

Tournament simulations: Survival curves

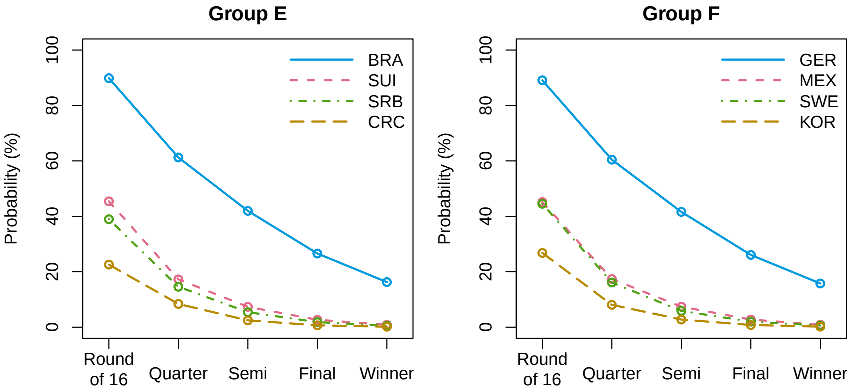

Tournament simulations: Survival curves

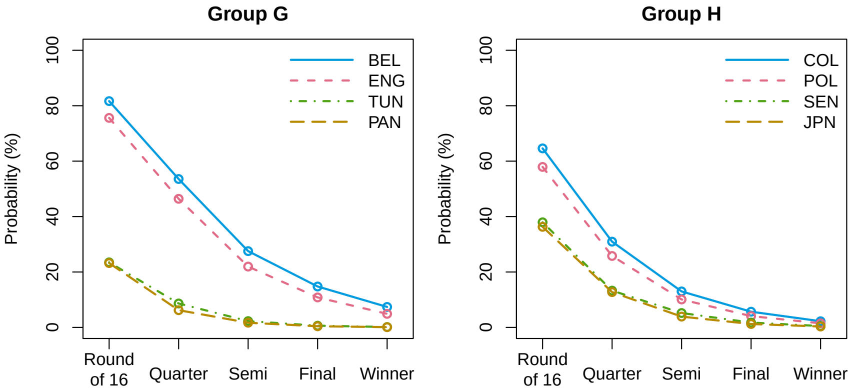

Outcome verification

Source: Spiegel.de

Outcome verification

Illustration: Check results for UEFA Euro 2016.

Question: Was the bookmaker consensus model any good?

- Ex post the low predicted winning probability for Portugal (4.1%) seems wrong.
- However, they profited from Spain's and England's poor performances in the last group stage games.
- And they only won 1 out of 7 games in normal time.
- Even in the final Gignac might as well have scored a goal instead of hitting the post in minute 92...

Problems:

- Just a single observation of the tournament and at most one observation of each paired comparison.
- Hard to distinguish between an unlikely outcome and systematic errors in the predicted (prob)abilities.

Outcome verification

Possible approaches:

- Compare forecasts with the observed tournament ranking (1 POR, 2 FRA, 3.5 WAL, 3.5 GER, ...).
- Benchmark against Elo and FIFA ratings.
- Note that the Elo rating also implies ability scores based on which pairwise probabilities and “forward” simulation of tournament can be computed:

$$ability_{Elo,i} \;=\; 10^{Elo_i/400}.$$

- Check whether pairwise probabilities roughly match empirical proportions from clusters of matches.

Outcome verification: Ranking

Spearman rank correlation of observed tournament ranking with bookmaker consensus model (BCM) as well as FIFA and Elo ranking:

BCM (Probabilities)	0.523
BCM (Abilities)	0.436
Elo (Probabilities)	0.344
Elo	0.339
FIFA	0.310

Outcome verification: BCM pairwise prob.

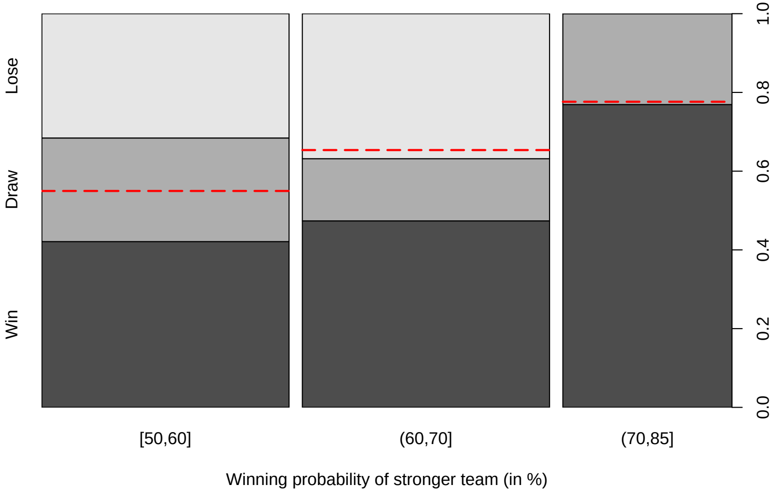

Outcome verification: Elo pairwise prob.

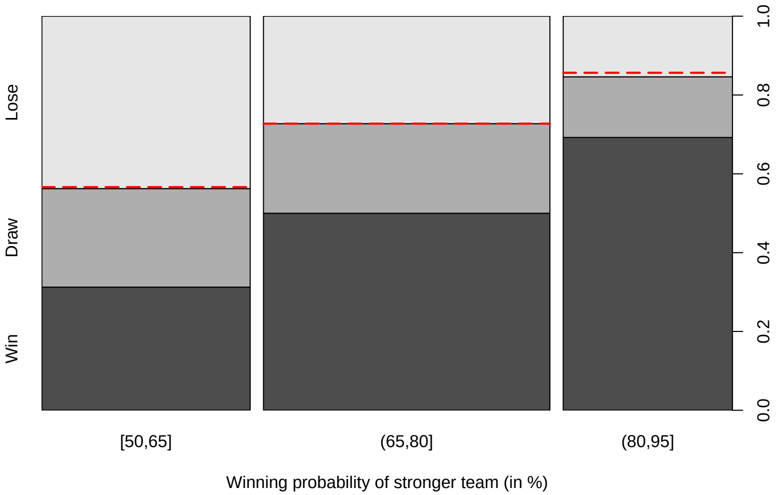

Outcome verification: BCM abilities

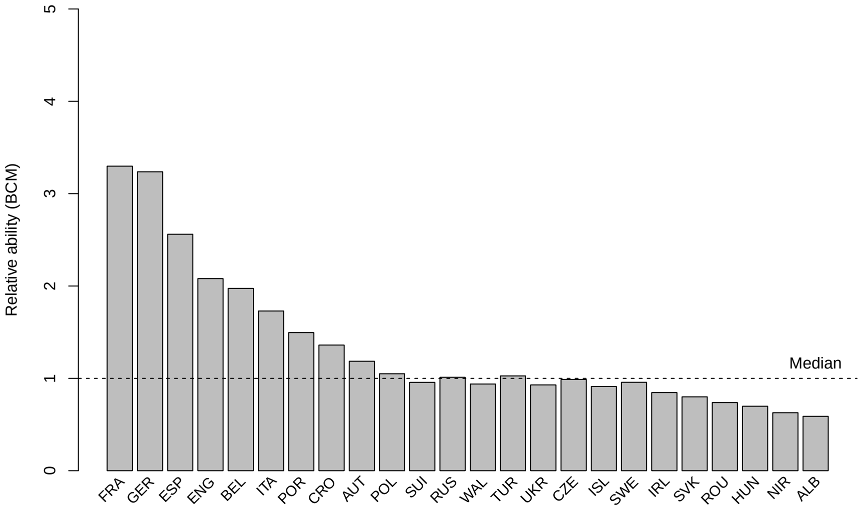

Outcome verification: Elo abilities

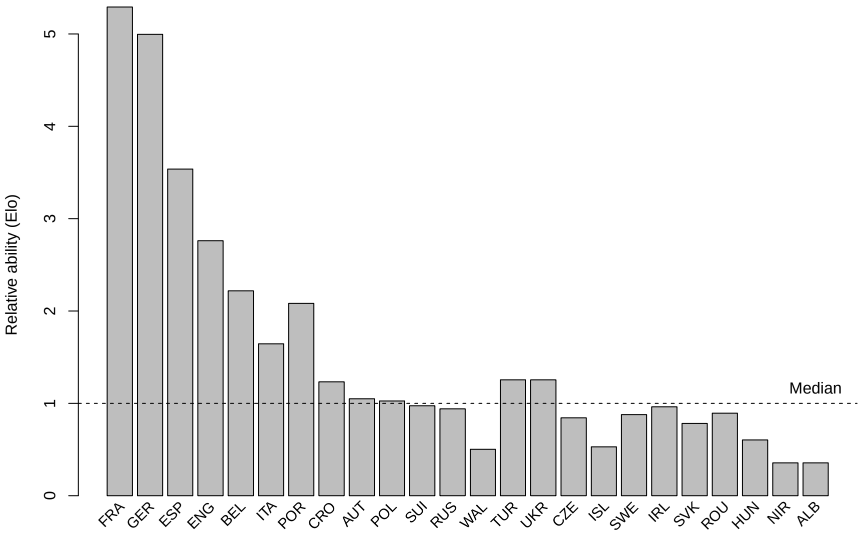

Discussion

Summary:

- Expert judgments of bookmakers are a useful information source for probabilistic forecasts of sports tournaments.
- Winning probabilities are obtained by adjustment for overround and averaging on log-odds scale.
- Competitor abilities can be inferred by post-processing based on pairwise-comparison model with “reverse” tournament simulations.
- Approach outperformed Elo and FIFA ratings for the last UEFA Euros and correctly predicted the final 2008 and winner 2012.

Limitations:

- Matches are only assessed in terms of winning/losing, i.e., no goals, draws, or even more details.
- Inherent chance is substantial and hard to verify.

References

Zeileis A, Leitner C, Hornik K (2018). “Probabilistic Forecasts for the 2018 FIFA World Cup Based on the Bookmaker Consensus Model.” *Working Paper 2018-09*, Working Papers in Economics and Statistics, Research Platform Empirical and Experimental Economics, Universität Innsbruck.
URL `http://EconPapers.RePEc.org/RePEc:inn:wpaper:2018-09`.
Blog: `https://bit.ly/fifa-forecast`.

Zeileis A, Leitner C, Hornik K (2016). “Predictive Bookmaker Consensus Model for the UEFA Euro 2016.” *Working Paper 2016-15*.
URL `http://EconPapers.RePEc.org/RePEc:inn:wpaper:2016-15`.

Leitner C, Zeileis A, Hornik K (2011). “Bookmaker Consensus and Agreement for the UEFA Champions League 2008/09.” *IMA Journal of Management Mathematics*, **22**(2), 183–194. `doi:10.1093/imaman/dpq016`.

Leitner C, Zeileis A, Hornik K (2010). “Forecasting Sports Tournaments by Ratings of (Prob)abilities: A Comparison for the EURO 2008.” *International Journal of Forecasting*, **26**(3), 471–481.
`doi:10.1016/j.ijforecast.2009.10.001`.