Psychology of Sport and Exercise 15 (2014) 538–547

Contents lists available at ScienceDirect

Psychology of Sport and Exercise

journal homepage: www.elsevier.com/locate/psychsport

Bayesian networks for unbiased assessment of referee bias in Association Football

Anthony Costa Constantinou [a, b, *], Norman Elliott Fenton [a], Liam Jackson Hunter Pollock [c]

[a] Risk and Information Management Research Group, School of Electronic Engineering and Computer Science, Queen Mary, University of London, London E1 4NS, UK

[b] Forensic Mental Health Research Unit, Centre for Psychiatry, Wolfson Institute of Preventive Medicine, Barts and The London School of Medicine and Dentistry, Queen Mary, University of London, London EC1M 6BQ, UK

[c] Biological and Experimental Psychology Research Group, The School of Biological and Chemical Sciences, Queen Mary, University of London, London E1 4NS, UK

ARTICLE INFO

Article history:
Received 28 January 2014
Received in revised form
27 May 2014
Accepted 27 May 2014
Available online 10 June 2014

Keywords:
Causal modelling
Crowd effect
Home advantage
Officiating bias
Soccer

ABSTRACT

Objectives: To assess referee bias with respect to fouls and penalty kicks awarded by taking explanatory factors into consideration.
Design: We present a novel Bayesian network model for assessing referee bias with respect to fouls and penalty kicks awarded. The model is applied to the 2011-12 English Premier League season.
Method: Unlike previous studies, the model takes into consideration explanatory factors which, if ignored, can lead to biased assessments of referee bias. For example, a team may be awarded more penalties simply because it attacks more, not because referees are biased in its favour. Hence, we incorporate causal factors such as possession, time spent in the opposition penalty box, etc. prior to estimating the degree of penalty kicks bias.
Results: We found fairly strong referee bias, based on penalty kicks awarded, in favour of certain teams when playing at home. Specifically, the two teams (Manchester City and Manchester United) who finished first and second appear to have benefited from bias that cannot be fully justified by the explanatory factors. Conversely Arsenal, a team of similar popularity and wealth and who finished third, benefited least of all 20 teams from referee bias at home with respect to penalty kicks awarded.
Conclusions: Among our conclusions are that, in contrast to many previous studies, being the home team does not in itself result in positive referee bias. More importantly, the model is able to explain significant discrepancies of penalty kicks bias into non-significant after accounting for the explanatory factors.

Introduction

The notion that referees in Association Football (hereafter referred to simply as *football*) are biased towards certain teams or in certain contexts is widely accepted by football pundits and supporters. In fact, whether or not such bias exists is an area of increasing interest that attracts the attention of researchers from the domains of sport science, psychology, statistics and computer science.

Irrespective of the true underlying causes, there is no doubt that ‘playing at home’ has a significant impact on a team's success. This home advantage effect has been extensively studied (Anders & Rotthoff, 2012; Constantinou & Fenton, 2013; Courneya & Carron, 1992; Hirotsu & Wright, 2003; Nevill & Holder, 1999; Pollard, 1986; 2006; Pollard & Pollard, 2005; Poulter, 2009). Numerous explanatory factors have been proposed for home advantage. The crowd effect is normally suggested as one of the most important factors (Agnew & Carron, 1994; Dohmen, 2008; Downward & Jones, 2007; Goumas, 2012; Nevill, Newell, & Gale, 1996, 1999; 2002) and is said to occur to a greater extent in leagues in which home crowds are more hostile and vociferous (Anders & Rotthoff, 2012). Other proposed factors include the travelling effect (Clarke & Norman, 1995), the familiarity with the playing grounds (Neave & Wolfson, 2003; Pollard, 2006), as well as referees themselves who are said to

* Corresponding author. Risk and Information Management Research Group, School of Electronic Engineering and Computer Science, Queen Mary, University of London, London E1 4NS, UK. Tel.: +44 790 3292 836.
E-mail address: anthony@constantinou.info (A.C. Constantinou).

http://dx.doi.org/10.1016/j.psychsport.2014.05.009
1469-0292/

favour home teams on the basis of penalty kicks, free kicks, yellow/red cards and/or extra time data (Boyko, Boyko, & Boyko, 2007; Buraimo, Forrest, & Simmons, 2010; Dawson, Dobson, Goddard, & Wilson, 2007; Dohmen, 2008; Downward & Jones, 2007; Goumas, 2012; Nevill et al., 1996, 1999; 2002; Sutter & Kocher, 2004). However, the degree of influence of referee decisions relative to the overall home advantage effect has not been extensively studied.

It is apparent that the literature tends to indicate with strong belief that referee decisions favour the home team. However, some researchers (Page & Page, 2010) have questioned this outcome and expressed their uncertainty as "*it could be the case that these biases do not manifest themselves into significant differences in terms of the overall performance of a team*" (Page & Page, 2010); the increased number of fouls, yellow cards, red cards, penalties and so on in favour of the home team might simply be the result of the home team performing better than the away team. For example, if the home team is in control of the ball (possession) more often than not, then we would expect it to be awarded more fouls and penalties, and less yellow and red cards relative to the opponent, on the basis that its control of possession will lead to it being on the receiving end of more tackles. We should also expect a higher proportion of these to be committed nearer to the opponent's goal, as greater possession also tends to correspond to a marked territorial advantage. We agree that the kind of explanatory causal factors proposed in (Page & Page, 2010) must be incorporated into any study of referee bias.

Hence, in this paper we present a novel Bayesian network (BN) model developed for referee bias analysis in football. It is the most comprehensive attempt to date to include within-game explanatory variables in order to justify the observed discrepancies between fouls and penalty kicks awarded between adversaries prior to formulating beliefs about referee bias. Although previous attempts have been made to control within-game events such as shots, fouls and corners (Dohmen, 2008; Goumas, 2012), this paper integrates a number of important additional variables which are required for formulating a causal network model, specifically for penalty kicks awarded.

The paper is organised as follows: Section The model describes the BN model, Section Results and discussion discusses the results and Section Concluding remarks and future research provides our concluding remarks.

The model

In this section we describe the BN model which was developed using the AgenaRisk BN tool (Agena Ltd., 2013). The tool was chosen because of its ability to properly incorporate continuous variables, without any constraint (like Normality), and without the need for static discretisation. This is achieved through its dynamic discretisation algorithm (Neil, Marquez, & Fenton, 2010). Details about the role of qualitative judgements and how inference is done are provided in (Fenton & Neil, 2012; Fenton, Neil, & Caballero, 2007; Neil et al., 2010).

The data used to inform priors and provide observations for each of the teams is available online at (WhoScored?.com, 2012), although the data for number of penalties awarded was manually recorded by a member of the research team from bbc.co.uk/football. However, the data is limited in the sense that, instead of having the value for each explanatory factor for each team in each match, we only have the averaged values for a set of match instances (namely match instances at home, away, and overall). With this limitation in place we have to make distributional assumptions based on expert judgement.

The data limitation also affects our ability in performing accurate simulation for estimating penalty kicks awarded. Specifically, for a proper simulation we want to know, for example, the percentage of time spent in the opposition penalty box (while in possession of the ball) relative to the overall percentage of possession for each individual match, rather than the average values over a number of match instances. Since we have a known average rate, distributional assumptions such as the *~Poisson distribution*, which expresses the probability of a given number of events occurring in a fixed interval of time, help us in addressing these issues by also keeping the model simple (more details in the subsections that follow). The drawback is that uncertainty is increased, since we are estimating those values for each match.

The model is constructed on the basis of two components as illustrated by the model topology in Fig. 1. Component 1 (described in Section Component 1) measures the referee bias over all fouls awarded, while Component 2 (described in Section Component 2) measures the referee bias over fouls awarded within the opposition penalty box (effectively penalty kicks). All the technical information required for developing the model (by following the model topology presented in Fig. 1) are provided in Table B.1.

The model is used to assess the referee bias for each case at home, away, and overall. While it is possible that there is some dependency between the two biases, our analysis assumes that they are independent; implying that the bias for penalty kicks awarded is only measured based on penalty kicks predicted and observed, and the same applies for the free kicks bias.

Component 1

This component simply assumes that the fouls awarded in a game are a consequence of a team's ability with respect to the following attributes (each corresponding to a node in the model):

1. *Possession*: percentage of time the team is in control of the ball (we assume *Truncated ~Normal* distribution);
2. *Pass accuracy*: the percentage of successful passes (i.e. those that reach a team mate, and we assume *Truncated ~Normal* distribution);
3. *Aerial duels*: the percentage of aerial duels won (we assume *Truncated ~Normal* distribution);
4. *Dribbles*: the average number of times, per match instance, a player manoeuvres the ball around a player of the opposing team (we assume *~Poisson* distribution);
5. *Interceptions*: the average number of times, per match instance, a player intercepts a pass made by a player of the opposing team (we assume *~Poisson* distribution).

Accordingly, we use the above five observable variables as predictors, in a naive Bayesian classification framework, for the latent variable *True fouls awarded* (predicted average per match instance) for a team at the specified ground (we assume *~Poisson* distribution). Subsequently, the referee bias is simply inferred by measuring the discrepancy[1] in distributions between *predicted* (*True fouls awarded* node) and observed (*Fouls awarded* node) fouls

[1] While the common practise is to let the *true* value be the parent of the *observed* value, we chose to model this relationship in an inverse manner. This is due to the naive Bayesian assessment performed; i.e. if we had followed the common practise then the *true* value would had been predicted (up to a degree, depending on how the *bias* node is defined) given the *observed* value (as it happens with all of the other factors in the naive Bayes framework). The way we chose to model this (i.e. not following the common causal practise) certainly keeps the *true* value constant, and the discrepancy between *true* and *observed* values is fully explained in the *bias* node.