Behavioural Public Policy (2022), 6: 2, 237–260
© *Cambridge University Press* doi:10.1017/bpp.2019.9 First published online 29 April 2019

A machine learning perspective on responsible gambling

ARMAN HASSANNIAKALAGER
University of Bath, Bath, UK
PHILIP W.S. NEWALL*
University of Warwick, Coventry, UK

Abstract: Gamblers are frequently reminded to 'gamble responsibly'. But these qualitative reminders come with no quantitative information for gamblers to judge relative product risk in skill-based gambling forms. By comparison, consumers purchasing alcohol are informed of product strength by alcohol by volume percentage (ABV%) or similar labels. This paper uses mixed logistic regression machine learning to uncover the potential variation in soccer betting outcomes. This paper uses data from four bet types and eight seasons of English Premier League soccer, ending in 2018. Outcomes across each bet type were compared using three betting strategies: the most-skilled prediction, a random strategy and the least-skilled prediction. There was a large spread in betting outcomes, with, for example, the per-bet average loss varying by a factor of 54 (from 1.1% to 58.9%). Gamblers' losses were positively correlated with the observable betting odds across all bets, indicating that betting odds are one salient feature that could be used to inform gamblers about product risk. Such large differences in product risk are relevant to the promotion of responsible gambling.

Submitted 19 September 2018; revised 9 March 2019; accepted 19 March 2019

Introduction

Consumers purchasing gambling and alcohol products receive similar advice. 'Please gamble responsibly' is seen perhaps as frequently as advice to 'please drink responsibly'. But how should this advice be followed? With alcohol, the familiar alcohol by volume percentage (ABV%) tells consumers about product strength, and this information is further broken down into the number of units of alcohol per serving on the majority of alcohol sold in the UK (Portman Group, 2017). In gambling, a similar percentage determines

* Correspondence to: Philip W.S. Newall, Applied Psychology, WMG, University of Warwick, Coventry, CV4 7AL, UK. Email: philip.newall@warwick.ac.uk

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

product strength in terms of statistical risk, but this advice is given to consumers more inconsistently. Electronic gambling machines represent one UK gambling product that comes with product risk information, where by law statistical risks are presented as (Parke *et al.*, 2016):

$$\text{'return-to-player'}\% = 100 - \text{house edge}$$

And the return-to-player percentage is described in practice as, for example, 'This game has an average percentage payout of 97.3%'. The return-to-player corresponds to the long-run proportion of every £100 wagered that gets returned to gamblers as prize money. The gambling operator keeps the remainder as the 'house edge'. House edge will be used here in contrast to return-to-player, as a higher house edge represents a higher average risk (Woolley *et al.*, 2013), mirroring the risks of high-ABV% alcohol. In European roulette, for example, the house edge on all bets is 2.7% (return-to-player of 97.3%), and this forms the basis for many roulette-based electronic gambling machine games. Qualitative research suggests that many regular gamblers fail to see UK electronic gambling machine return-to-player warning labels (Collins *et al.*, 2014), and that some gamblers do not understand the return-to-player percentage (Collins *et al.*, 2014; Harrigan *et al.*, 2017; Rowe *et al.*, 2017). Despite these current failings, evidence from the behavioural risk communication literature suggests that this information could be communicated to more people through the use of graphical risk representations (Garcia-Retamero & Cokely, 2017), while the tobacco literature suggests that making these warning labels more prominent should also boost effectiveness (Hammond, 2011) . But despite these significant failings, many UK gambling products are marketed with even less product risk information than is currently given on return-to-player warning labels.

Soccer is the UK's national sport, and soccer betting is now the most popular form of gambling with British problem gamblers (Gambleaware, 2017). The most common advice featuring on soccer bets, either at point of purchase or on marketing communications, is currently the generic, 'When the fun stops, stop'. Absolutely no advice is given on relative product risk, even as these risks may vary substantially, in contrast to return-to-player or ABV% labels. When consumers see a soccer bet, either in a physical bookmaker, on a bookmaker's website or in a marketing communication, they typically receive only one piece of risk information: the betting odds. For example, a bet on the correct score in soccer such as 'Arsenal to win by three goals to one, 25-to-1' means that gamblers will win £25 of profit for every £1 bet if the team Arsenal wins by three goals to one (Cortis, 2015). But each potential payout must be multiplied by its probability of happening in order to calculate a

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

gambler's long-term returns. In European roulette, potential payout and event probability happen to be perfectly negatively correlated, so that the house edge is a constant 2.7%. Some psychologists have recently proposed that payout and probability are in fact negatively correlated in many real-world environments, resulting in roughly constant average expected returns (Pleskac & Hertwig, 2014). However, as we show next, the house edge can fluctuate wildly across different bet types in soccer betting.

Bookmakers rarely share past betting data with researchers (Chagas & Gomes, 2017). But the *average* house edge can be estimated indirectly in sports betting, based on bookmakers' odds (Cortis, 2015), under the assumption that bookmakers offset the risks of different bets (Stark & Cortis, 2017). Each soccer match ends in one of three salient outcomes: home win, draw or away win. Bets on these outcomes are the oldest bets that can be made in soccer (Forrest, 2008) and will be called 'home–draw–away' bets here. These events are high-profile and are relatively few, so it should be easy for gamblers to shop around and find the bookmaker with the best home–draw–away odds, especially with odds comparison sites such as oddschecker.com. In addition, the house edge on home–draw–away bets has decreased in recent years (Constantinou & Fenton, 2013; Buhagiar *et al.*, 2018). For example, a study found home–draw–away edges of 4.5% on average over the 2014 World Cup (Newall, 2015), compared to an average house edge of 10.3% in data from the 1990s (Kuypers, 2000). But bookmakers are increasingly innovating new bet types (Newall, 2015, 2017; Newall *et al.*, 2018). Correct score bets, for example, pay off if a gambler can predict the more specific final scoreline. The same 2014 World Cup study found a far higher house edge of 21.9% on correct score bets (Newall, 2015), which is proportionally not much lower than the 26% house edge found on correct score bets from a study in the 1990s (Dixon & Pope, 2004). The house edge in soccer depends a lot on the type of bet chosen and can be more than eight times the 2.7% European roulette house edge on a per-bet basis. This is important for gamblers, as a lower house edge allows gamblers to benefit from a lower 'price' of gambling, and thereby to be able to afford a higher number of gambling experiences from some fixed budget allocated for gambling (Woolley *et al.*, 2013).

Perhaps one of the unique attractions of sports betting is that 'skilful' bettors can beat these average house edges (Kaunitz *et al.*, 2017). Sports betting is a skill-based form of gambling, whereas roulette is luck-based as there is nothing that can be legally done to beat the 2.7% house edge (Thorp, 1998). The reality, however, is that many sports fans are not especially skilled at predicting sports outcomes. Sports fans are overconfident, even when the financial stakes are high (Simmons & Massey, 2012), and they are subject to wishful

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

thinking (Babad, 1987). In addition, soccer experts are not necessarily better than non-experts at predicting outcomes (Andersson *et al.*, 2005, 2009). Prediction skill means that soccer bettors can plausibly receive returns that are both better and worse than the averages summarised in the previous paragraph. This paper will therefore use machine learning to investigate the total potential impact of prediction skill across four soccer betting types. Although it could be difficult to directly make sports bettors more skilful, it might be easier to alert sports bettors to observable product features, such as the level of betting odds, which are correlated with product risk. For example, with alcohol, it is reasonable to guess that the ABV% of a bottle of whiskey is higher than that of a bottle of wine, which is higher than that of a bottle of beer. Any salient attribute of sports bets that is correlated with the returns that gamblers receive could similarly be used to communicate product risk non-numerically.

Risks are relatively easy to estimate in electronic machine gambling, where the chance of each outcome is predetermined. This makes the house edge fixed and the winning probabilities and payoffs for gamblers easy to calculate via probability theory and Monte Carlo simulation. It is harder to calculate event probabilities in sports betting, where each match is unique and the outcomes are governed partly by chance and partly by skill (Hill, 1974). This is one potential reason for the inconsistent labelling of product risk on UK gambling products. But consider what could happen if a gambler sees a warning label on an electronic gambling machine, but not on a soccer bet. It is possible the gambler might assume that the unlabelled soccer bet poses less risk, while actually, soccer bets can be over eight times riskier on a per-bet basis. But while it is harder to calculate product risk for sports bets, the variable nature of sports betting risks also provides a unique opportunity to reduce the harms associated with sports betting.

Gambling-related harm increases as gamblers' total losses increase (Markham *et al.*, 2014, 2016). But gamblers' losses are not fixed in sports betting, and they depend on the relative skill levels across gamblers and bookmakers. At present, British bookmakers are creating soccer bets with ever-increasing potential payoffs, which is the only piece of risk information currently communicated to gamblers (Newall, 2015, 2017; Newall *et al.*, 2018). This matters because these high-potential-payoff bets have higher house edges than traditional home–draw–away bets. But changing the behaviour of even a small number of gamblers, such as by making them more aware of how potential payoff combines with event probability to determine the house edge, could incentivise bookmakers to compete on 'price' – the house edge (Woolley *et al.*, 2013). House edges have decreased for home–draw–away bets since the late 1990s (Kuypers, 2000; Constantinou & Fenton, 2013;

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

Buhagiar *et al.*, 2018), demonstrating that bookmakers can compete on price. Any decreases in the house edge mean then even gamblers who do not change their behaviour in any way can also benefit from resulting decreases in their gambling losses.

Gambling-related harm is seen as a public health problem by many researchers (Korn & Shaffer, 1999; Livingstone & Woolley, 2007; Orford, 2010; Markham & Young, 2015; Lopez-Gonzalez *et al.*, 2017). Any intervention to reduce gambling's public health costs must consider the trade-off between individual freedom and protecting vulnerable consumers. At one end of the trade-off spectrum is the freedom-preserving choice to 'do nothing', while stronger interventions such as 'change the default' and 'use tax incentives' lie increasingly towards the other end (Nuffield Council on Bioethics, 2007). Providing information on product risk, such as via warning labels, is the first and most freedom-preserving public health intervention on this spectrum. We have used the examples of alcohol and electronic gambling machine warning labels to highlight an omission of product risk information for sports bets. However, any new warning label should be trialled experimentally and in field trials before being rolled out at a population level to ensure consumer understanding and to mitigate the risk of any behavioural backfiring (Stibe & Cugelman, 2016). Any gambling warning label should also ideally leverage insights from the behavioural risk communication literature (Garcia-Retamero & Cokely, 2017), which current UK electronic gambling machine labels do not (Collins *et al.*, 2014; Rowe *et al.*, 2017). In addition, if warning labels are not effective at protecting consumers, then stronger public health interventions should also be considered (Nuffield Council on Bioethics, 2007). These are all important goals for future research. This paper tackles a more fundamental issue in soccer betting: risks must be understood before they can be communicated.

This paper will use machine learning to investigate how product risk varies across the largest dataset of past betting odds and results available to us. The aim is to uncover information relevant to the concept of 'responsible gambling': "If gambling is to be conceptualized as a leisure activity to be engaged in as an individual choice … then it is fundamental that the individual is presented with all relevant information, in a timely fashion, in order to make an informed choice" (Parke *et al.*, 2015, p. 31). This will be operationalised through three research questions:

RQ1. What is the variation in prediction skill across four key soccer bet types?

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

Prediction skill matters to the extent that some gamblers can achieve better returns than others. Our machine learning model will make bets across different soccer matches and bet types. We have data for home–draw–away, correct score and two other bet types (explained in the 'Data' section below). The model will do this with three strategies in order to mimic the variability across human sports bettors. Model performance will then be compared across bet types and strategies using the strategies' arithmetic average returns for comparison. If a gambler approaches a given bet type with a given strategy, this percentage gives the amount of staked money lost on average *per bet*.

Some human sports bettors try to get the highest return possible by inspecting recent performance and developing knowledge about the sport. This strategy will be mimicked by a strategy we call 'Most-skilled'. The existence of skill can be demonstrated via showing an improvement in performance over a strategy that requires no skill (or knowledge). Therefore, to check the validity of our machine learning model, a second strategy called 'Random' will be generated. For RQ1, Random places a bet on *every* available outcome within a bet type. For RQ1, this generates more precise estimates of what a truly random betting strategy will return than just selecting one bet at random. It should be the case that the Most-skilled strategy yields higher returns than Random. However, if it is possible to beat a random strategy, then it is also possible to do worse than a random strategy. But there should be limits to how bad a strategy can do, as even the most unlikely events do eventually happen in sports betting (such as Leicester FC winning the 2015/2016 English Premier League). Our machine learning model will mimic the worst returns a human might get via the 'Least-skilled' strategy. For RQ1, we hypothesise that the returns will be distributed within each bet type as: Most-skilled > Random > Least-skilled. Prediction skill likely matters more for some bet types than others, and this will also be investigated in RQ1.

RQ2. A gambler starts with £100 and bets 10% of their 'bankroll' on each game. How long will their stake take on average to fall to £50 (the 'half-life')?

RQ1 provides the best estimates of the returns from placing a single bet under the three strategies previously discussed. But many gamblers place sequences of bets from a given starting amount of dedicated gambling money, called a 'bankroll'. Average arithmetic returns provide an incomplete picture for such sequences of gambles (Chen & Ankenman, 2006). RQ2 attempts to mirror the returns of the three strategies (Most-skilled, Random and Least-skilled), as measured by the number of bets before half an initial amount of money is lost, on average. This should provide a more accurate measure of what repeat gamblers could experience. This research question is

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

motivated by the statistical concept of the 'gambler's ruin', which shows how a gambler with negative expected returns will eventually lose all of their money (Mohan, 1955; Harik *et al.*, 1999). The 10% proportion was chosen heuristically by us, in the absence of us having any human betting data to compare against (Chagas & Gomes, 2017). Larger-percentage bet sizes can compress the differences between betting strategies; for example, with a bet of 50%, one losing bet hits the 'half-life' definition. Smaller bet sizes can do the opposite, as gamblers' bankrolls asymptote towards lasting forever.

As with RQ1, it should be the case that, within each bet type, the strategies have half-lives of: Most-skilled > Random > Least-skilled. For RQ2, Random involves randomly picking one bet per bet type. This is because betting on every outcome, as in RQ1, creates a portfolio of offsetting bets that loses a small amount of money with certainty, due to the house edge, leading to an artificial boost in the bankroll half-life (and no gambler would enjoy losing a small amount of money with certainty). RQ2 therefore investigates how gamblers' bankrolls evolve across different betting strategies and bet types.

RQ3. Which specific outcomes within each bet type do the Most-skilled and Least-skilled prediction strategies tend to choose?

Machine learning should ideally provide insights to help human decision-makers. Humans often overestimate the likelihood of low-probability events (Kahneman & Tversky, 1979). This tendency has been reflected in sports betting via the 'favourite–longshot' bias, where longshots with low winning chances yield higher than average losses for gamblers than bets on favourites (Vaughan Williams, 1999). The favourite–longshot bias has been observed in large-dataset investigations of home–draw–away soccer odds (Constantinou & Fenton, 2013; Buhagiar *et al.*, 2018). Ideally, the Least-skilled strategy will look like a biased human sports bettor, while the Most-skilled strategy will bet differently. Looking at recurring features of the specific bets chosen by the Most-skilled and Least-skilled strategies could also help inform strategies intended to educate gamblers about observable bet attributes that are correlated with product risk.

Model

Machine learning models come in varying degrees of complexity. Innovations in machine learning research often focus on new techniques for processing large amounts of data. Our intention with this paper was not to innovate new machine learning techniques, but to apply established prediction models from the operations research literature to the topic of responsible gambling.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

Better machine learning models will always be possible. However, any improvement in prediction should not necessarily bias our estimates of different betting strategies' *relative* risks. Our key innovations are to consider the total impact of skill via comparing the Most-skilled and Least-skilled strategies and to use the most extensive, freely available dataset of soccer betting odds (four bet types over eight seasons). We encourage other researchers to attempt to replicate these results with different datasets and models.

Data

We wanted to train the model on as many bet types as possible, given the ever-increasing number of ways to bet on soccer. However, long-running, freely available datasets were required, limiting the number of bet types that could be studied. Historical odds were available for the following four betting types: correct score, home–draw–away, over/under and Asian handicap. Data for eight seasons of English Premier League soccer (the highest English club competition) betting odds and results from 2010/2011–2017/2018 were taken from oddsportal.com.

Summary information is provided in Table 1. Of the four bet types, correct score bets stand out as an outlier in Table 1 with 25 potential outcomes. Although there is in theory no limit on the number of goals scored by both teams in a soccer match, reliable data were available only on the 24 most common scorelines. All other potential scorelines were included as a 25th catch-all outcome in the model. Home–draw–away bets have three potential outcomes (home win, draw and away win). The over/under is a bet on the number of goals scored: to be either between zero and two goals or to be three or more goals (two potential outcomes). The Asian handicap performs a skill adjustment on the two teams, effectively removing the potential of a draw (e.g., Arsenal to win by 1.5 goals or more, therefore giving two potential outcomes).[1] 'Complex' bet types with many potential outcomes, such as correct score bets, have previously been associated with high house edges (Ayton, 1997) and an overestimation by soccer fans of average event probability (Newall, 2017).

The model was also given in-game statistics of recent team performance (explained at greater length in the 'Inputs' section below). These data were downloaded from football-data.co.uk.

1 There are a small number of instances where the Asian handicap was set as a discrete number of goals; for example, Arsenal to win by one goal or more). Then, if Arsenal do in fact win by exactly one goal, a third outcome is created where the bet neither wins nor loses – a 'push'. In this instance, the original stake is returned to the gambler.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

Table 1. Example of betting odds data

Betting market	Example	Number of potential outcomes
Correct score	'Arsenal to win 1–0'	25
Home–draw–away	'Arsenal to win'	3
Over/under 2.5 goals	'Over 2.5 goals to be scored in the match'	2
Asian handicap	'Arsenal to win by 1.5 goals or more'	2

Note: Although there is potentially an infinite number of potential correct scores, reliable betting odds were only available for the 24 most-likely scorelines. Therefore, we include these 24 scorelines (either team scoring up to four goals, excluding 4–4) while adding a 25th catch-all outcome for any other scoreline.

Model specifications

Machine learning models approach new datasets like a cautious gambler, learning patterns between current data and future outcomes before making any predictions (the learning phase). The first three seasons were solely used to learn from the data.[2] Prediction quality on these learned associations was then assessed only in a new, unseen part of the data (the prediction phase). From the fourth season on, the model strategies placed (simulated) bets and received hypothetical returns. Because old data become gradually less informative, each season of prediction was based only on the previous three seasons of learning.

An emphasis on predicting new data is a hallmark of the machine learning approach to statistical testing. Less emphasis is placed on interpreting *p*-values, as in traditional null hypothesis statistical testing. The main check on model validity is that the Most-skilled and Least-skilled model strategies should respectively perform better and worse than Random in the prediction phase. Some researchers have argued that the social sciences could benefit from placing greater emphasis on predicting new data and less on explaining past data, as is currently done via null hypothesis statistical testing (Yarkoni & Westfall, 2017).

Prediction-phase performance was available for five seasons in total: 2013/2014–2017/2018. Prior to running any analyses, a decision was made to

2 Machine learning does not have strong methodological norms regarding what proportion of the data to use for learning and prediction. Even for similar research questions, splits as variable as 90–10 (Harvey & Liu, 2015) and 50–50 (Bajgrowicz & Scaillet, 2012) have been observed. Given the limits on our time series (eight seasons), we chose a conservative ratio of three learning-phase seasons for every prediction-phase season to reduce potential problems from more extreme ratios. Other learning/prediction-phase splits were not tested.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

discard the first five games of each season due to potentially large between-season variations in team performance. This leaves a total sample size of 1650 matches to compare model performance on (330 matches per season). Each prediction is based on 990 learning-phase observations (three seasons). The required number of learning-phase observations to obtain 95% statistical power is 380 for the low-complexity bet types (home–draw–away, over/under 2.5 goals and Asian handicap) and 950 for the high-complexity correct score bet type (Hsieh *et al.*, 1998).

Appropriate machine learning model choice depends on the underlying data's statistical properties. Some soccer forecasting models were developed for individual matches, which is typically accounted for by multinomial logistic regression models, given the three main outcomes of a soccer match (e.g., Koning, 2000; Willoughby, 2002; Khazaal *et al.*, 2012). But the present model also forecasts bet outcomes, which depend on the specific teams' characteristics, where a conditional logistic regression is more appropriate (Smith *et al.*, 2009). Each of these features needs to be accounted for separately. Therefore, we used a combination of multinomial and conditional (mixed) logistic regression to capture the full scope of the data (McFadden, 1974; McFadden & Train, 2000; Lessmann *et al.*, 2012).

Inputs

The machine learning model, like an expert gambler, learns associations between publicly available information (predictors) and soccer outcomes. Given the applied nature of our research questions, our aim was not to test the predictive power of previously unexplored inputs, but to only use inputs established in the previous literature to have positive predictive performance. While better models could always be designed, our relatively simple approach should not necessarily bias our estimates of different betting strategies' relative risks. It might seem that adding more inputs should boost predictive power. However, adding more inputs can increase the risk of 'overfitting', where chance random associations are discovered and erroneously predicted to recur in the prediction phase, hence leading to poor prediction-phase performance (Yarkoni & Westfall, 2017).

Our first model input was an average of relevant match betting odds, using an average of around 30 bookmakers per data point. Betting odds have been shown to successfully predict outcomes in the soccer forecasting literature (Goddard & Asimakopoulos, 2004; Štrumbelj & Šikonja, 2010; Štrumbelj, 2014). Betting odds tend to fare even better than other prediction methods, such as Elo rating systems, which were originally designed for chess ratings (Hvattum & Arntzen, 2010; Leitner *et al.*, 2010). As is the standard in the

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

literature, these odds were averaged to reduce random forecast error, and they were normalised to remove the effect of the house edge[3] (Hvattum & Arntzen, 2010; Leitner *et al.*, 2010).

Gamblers will often factor in recent performance. Therefore, the model was provided with in-game statistics summarising each team's performance over the previous five matches. Based on successful in-game predictors from the previous literature, we chose the cumulative number of points earned – three for a win, one for a draw and zero for a loss (Goddard & Asimakopoulos, 2004; Goddard, 2005) – and the number of goals scored and conceded (Dixon & Robinson, 1998; Oberstone, 2009; Baio & Blangiardo, 2010; Angelini & De Angelis, 2017). Predictors that have been used in the previous literature but were not included here include the number of corners won (Andersson *et al.*, 2009), shots on target (Oberstone, 2009), recent injuries (Constantinou & Fenton, 2017) and disciplinary bookings (Titman *et al.*, 2015).

We combined the estimate of the likelihood of each event happening from the machine learning model with the potential payoff size to create an estimate of the expected return from each bet with different strategies. The Most-skilled strategy chose the bet with the highest long-run expected return within each bet type for each match, while the Least-skilled strategy chose the bet with the lowest long-run expected return.

Results

In response to RQ1, the arithmetic average percentage returns varied over both the four bet types and the three betting strategies used, as summarised in Figure 1. Figure 1 presents the results of different hypothetical bettors, so it is more accurate to present the returns as experienced by each betting strategy, rather than as the bookmaker's overall house edge. The negative returns experienced by each strategy in Figure 1 can be interpreted as showing that predictions were never able to overcome the bookmaker's house edge. A return of –3.3% is equivalent to a house edge of 3.3%, or equivalently a return-to-player of 96.7%. In each bet type, it can be seen that the Most-skilled prediction performed best, the Least-skilled strategy performed worst and Random was in the middle. Therefore, in each bet type, (positive) betting skill improved performance, while (negative) betting skill worsened performance. Prediction skill mattered most in with correct score bets. With correct score bets, the Least-skilled strategy returned –58.9% on average, Random –34.3% and the

3 The odds used to calculate bet payoffs were only averaged, effectively mirroring the returns that a bettor who chooses bookmakers randomly would face.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

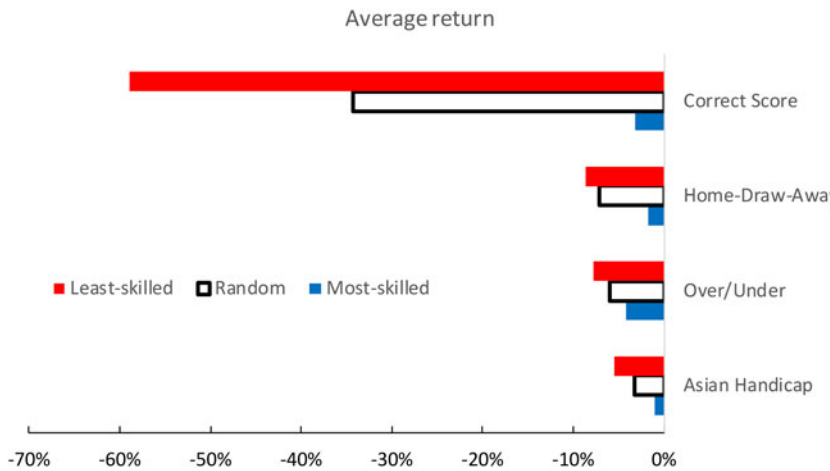


Figure 1. Prediction skill variability over each bet type and strategy

Most-skilled strategy –3.3% (returns as experienced by the model). No other betting strategy experienced returns worse than –10% per bet (even the Least-skilled strategies). Additionally, in the other three bet types, the per-bet difference in returns across the three betting strategies was 7% or under. Prediction skill always mattered, but especially with correct score bets (where there were many potential outcomes).

Next, we investigated the variation in relative betting performance. Comparing the relative losses can indicate the potential magnitude of savings a gambler could make by shifting towards lower-'price' bets. Table 2 shows each loss as a multiple of the lowest-loss performance, which was Most-skilled Asian handicap, at –1.1%. The variation in relative performance was large. The Least-skilled correct score strategy lost 54.0 times more per bet than the lowest-loss performance. High relative losses were not unique to correct score bets, with, for example, the Least-skilled home–draw–away strategy losing 8.0 times more per bet than the lowest-loss performance.

The half-life measured the number of bets before an average of 50% of starting wealth was lost (betting 10% of remaining wealth per match). Half-life variation is summarised in Figure 2, in response to RQ2. Similarly to Figure 1, the Most-skilled strategy had the longest half-life across each of the four bet types and the Least-skilled strategy had the shortest half-life. Correct score bets had half-lives of 10.8 to 7.0 bets, making it an outlier compared to the other three bet types. The other three bet types had Least-skilled half-lives of 22.2 (home–draw–away), 52.2 (over/under) and 64.4 bets (Asian handicap), all lasting longer than the Most-skilled correct score strategy.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

Table 2. Relative per-bet losses (in percentages) standardised so that the lowest loss strategy (Asian handicap, Most-skilled) = 1

Strategy	Correct score	Home–draw–away	Over/under	Asian handicap
Least-skilled	54.0	8.0	7.2	5.0
Random	31.5	6.5	5.5	2.9
Most-skilled	3.0	1.6	3.8	1.0 (–1.09%)

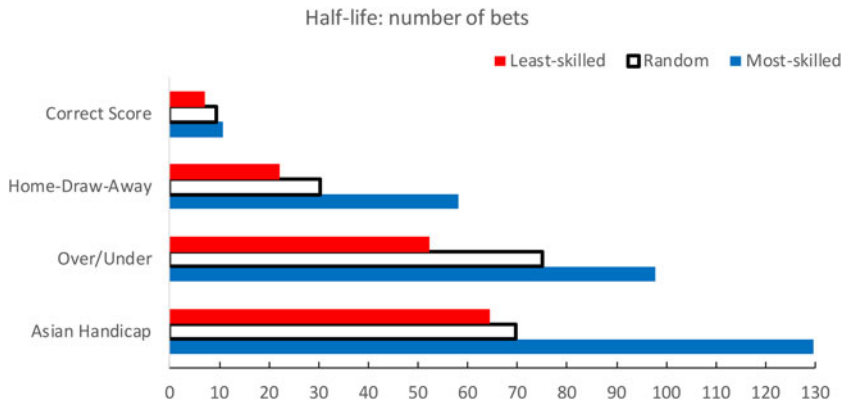


Figure 2. Half-life variability over each bet type and strategy

Table 3. Relative half-lives (in number of bets) standardised so that the lowest half-life strategy (correct score, Least-skilled) = 1

Strategy	Correct score	Home–draw–away	Over/under	Asian handicap
Least-skilled	1.0 (7.0 bets)	3.2	7.5	9.2
Random	1.3	4.3	10.7	9.9
Most-skilled	1.5	8.3	14.0	18.5

A comparison of relative half-lives in Table 3 highlights the large differences. The longest half-life strategy (Asian handicap, Most-skilled, 129.6 bets) lasted 18.5 times longer than the shortest half-life strategy (correct score, Least-skilled, 7.0 bets). There was again a large spread in betting outcomes, as measured by the half-life. All correct score betting strategies had short half-lives because this is the most volatile bet type.

Next, we investigated the outcomes chosen by the Most-skilled and Least-skilled prediction strategies across each bet type. Figure 3 provides an

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

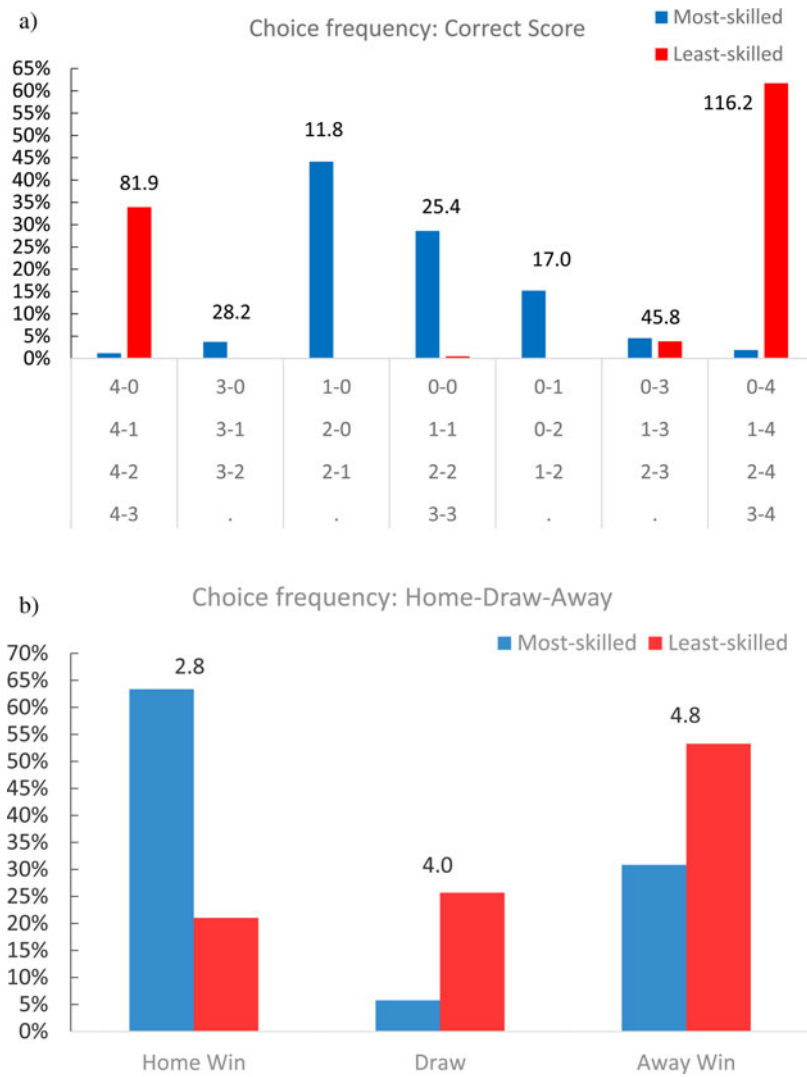


Figure 3. Most- and Least-skilled specific choice frequencies across each bet type: (a) correct score; (b) home–draw–away; (c) over/under; and (d) Asian handicap. The decimal odds are shown above each x-axis category, showing the total average payoff from a successful bet of $1. Similar scorelines were grouped together in (a), although the model estimates frequencies for each scoreline separately. By convention, scorelines are written as 'home team score–away team score'

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

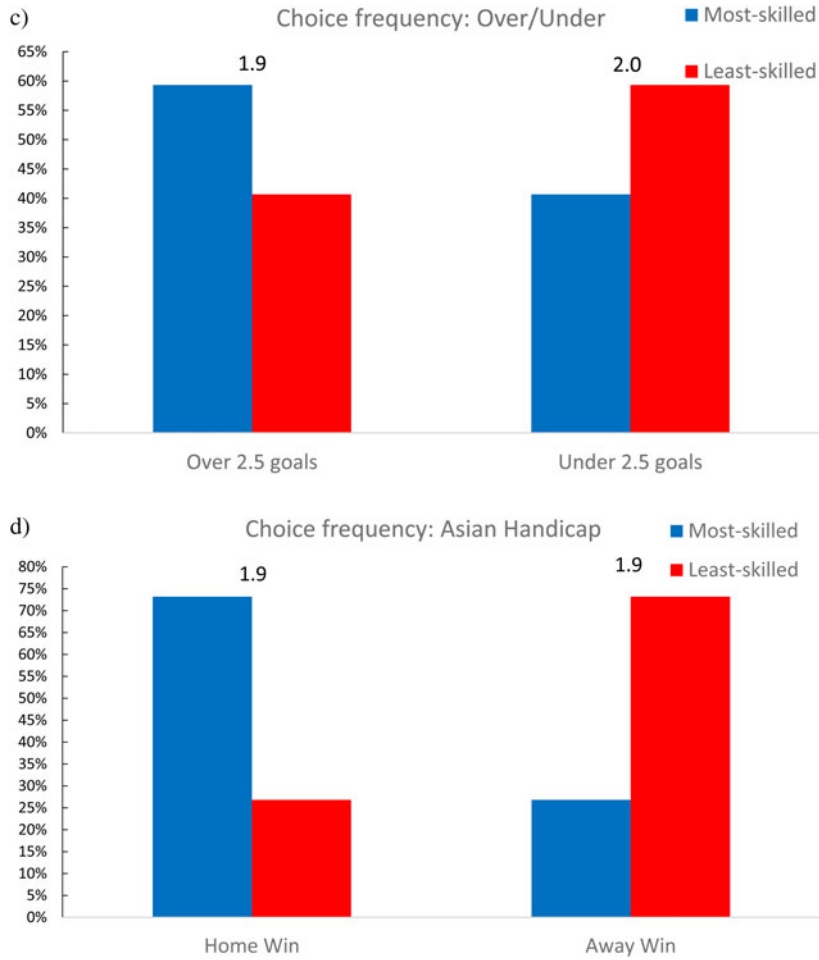


Figure 3. (Continued)

answer to RQ3. There are clear differences, as the Most-skilled strategy preferred low-scoring home wins with correct score bets (Figure 3(a)) and home wins with home–draw–away bets (Figure 3(b)). Figure 3 also provides the average decimal odds across each outcome shown on the x-axis. The decimal odds represent the total payoff from a successful bet of $1 (Cortis, 2015). The higher the decimal odds, the more extreme a longshot the bet is. In soccer, the team playing at home is usually the favourite to win (due to home advantage). By convention, the home team's score is written first, so a scoreline of 2–0 translates to the home team scoring twice and not conceding.

Across each bet type, the Most-skilled strategy usually chose likely outcomes with low decimal odds. For example, the home team winning 1–0, 2–0 or 2–1

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

with correct score bets was chosen 44.1% of the time (average decimal odds = 11.8), and the home team winning in home–draw–away was chosen 63.4% of the time (average decimal odds = 2.8). The Most-skilled strategy also favoured the home team winning in the Asian handicap and over 2.5 goals in the over/under. By contrast, the Least-skilled strategy chose the most extreme outcomes. This is demonstrated best with correct score bets, where the Least-skilled strategy places more than 95% of its bets across either the away team winning and scoring four goals (61.7%, average decimal odds = 116.2) or the home team winning and scoring four goals (33.9%, average decimal odds = 81.9). (Each scoreline is estimated separately by the models; we have grouped similar scorelines together for simplicity.) The Most-skilled strategy tended to bet on favourites, while the Least-skilled strategy tended to bet on lower-probability longshots, reflecting the established human bias of overweighting small probabilities (Kahneman & Tversky, 1979).

Discussion

The 'Introduction' section called soccer betting a "skill-based gambling form." In support of this claim, it was found that prediction skill mattered in each of the four soccer bet types. Most strikingly, it was found that gamblers' losses varied by a factor of 54 (from 1.1% to 58.9%). To return to the ABV% comparison from the 'Introduction' section, this is like the difference between a reduced-alcohol lager and a strong whiskey. Meanwhile, the 'half-life' number of bets before a starting bankroll was reduced by half varied by a factor of 18.5 (from 7.0 to 129.6 bets). The large number of outcomes makes every correct score bet a longshot compared to the three other bet types. Usually starving but occasionally feasting is a poor way to maintain wealth on average over time. This suggests that responsible gamblers should limit their exposure to volatile bet types such as correct score bets. The Least-skilled strategy tended to bet on lower-probability longshots, reflecting the established human bias of overweighting small probabilities (Kahneman & Tversky, 1979). But how else might these results be relevant to the topic of responsible gambling in humans?

The wide variation in gamblers' losses here contrasts with European roulette, where each possible bet has a constant house edge of 2.7%. This variability in soccer betting only increases the amount of information that should be given to gamblers, in keeping with responsible gambling principles (Parke *et al.*, 2015). But at present, UK soccer bets are marketed with only the betting odds and generic advice such as 'When the fun stops, stop'. By contrast, the more informative house edge is currently communicated to electronic machine gamblers as the 'return-to-player' percentage (Parke *et al.*, 2016).

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

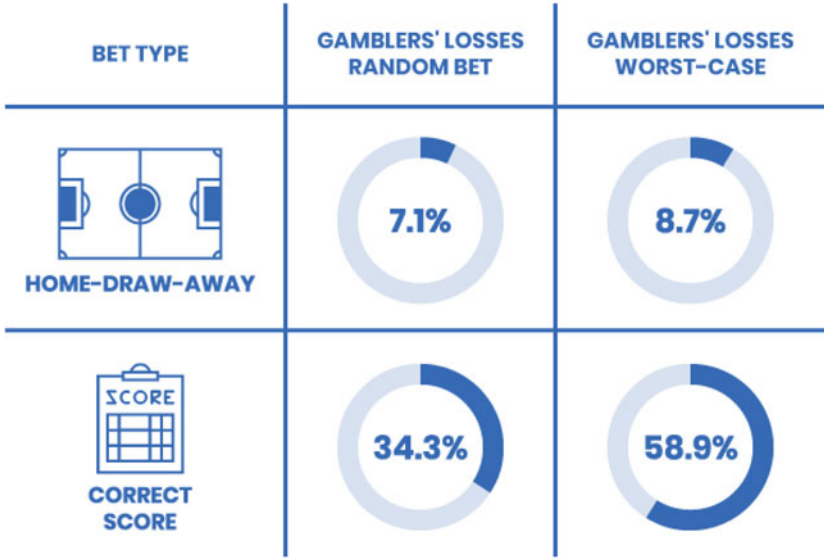


Figure 4. Example of warning labels for home–draw–away and correct score bet types. The 'worst-case' column corresponds to the returns of the Least-skilled prediction strategy

In addition, from April 2019, a £2 wagering limit will be placed on UK electronic gambling machines (Casey, 2018). By comparison, soccer betting, which is the most popular form of gambling with British problem gamblers (Gambleaware, 2017), has seen neither element of consumer protection.

Providing more information to soccer bettors could be seen as an informational 'nudge' (Thaler & Sunstein, 2008), as an example of 'boosting' (Hertwig, 2017) or as an example of consumer education (Peters, 2017). Figure 4 provides an example of how more information could be given for home–draw–away and correct score bets – the two bet types from this paper that have been seen in observational studies of UK gambling advertising (Newall, 2015, 2017; Newall *et al.*, 2018). Figure 4 provides a prototype example of how gamblers' losses when using the Random and Least-skilled strategies could be communicated to soccer bettors. In keeping with the spirit of ABV% labels, it might be wise not to present too much information to consumers, such as the Most-skilled strategy returns, which could plausibly cause confusion given that sports fans are prone to wishful thinking (Babad, 1987).

However, Figure 4 presents only estimates based on our historical simulation of eight seasons of data, the largest dataset available to us. Because the risks of sports betting are not fixed, it is important to present gamblers with the most

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

up-to-date information. We did not have access to actual past betting data for this paper, and so we were forced to use the Random strategy as a proxy for gamblers' average losses. Replacing this figure with gamblers' past average losses for a given bet type over some recent time horizon is one way in which Figure 4 could be improved upon. The timely communication of gamblers' past losses could also help bookmakers who plan to increase their customer base by competing on price. By comparison, some British bookmakers appear to be currently competing to create new bets with high betting odds, which consumers observe, but high hidden house edges (Newall, 2018).

There are many aspects of any new warning label that should first be empirically tested. Any new warning label should ideally be trialled in both the laboratory and in the field before being introduced into the population in order to mitigate the risk of any behavioural backfiring (Stibe & Cugelman, 2016). For example, qualitative research suggests that some young drinkers use ABV% information to find cost-effective ways of maximising their alcohol consumption (Jones & Gregory, 2009; Maynard *et al.*, 2017). While we think that few gamblers will want to purposively maximise their gambling losses, there are other issues from other public health domains of potential relevance to gambling risk communication. Any warning label should ideally be as attention grabbing as other product features. One eye-tracking study of alcohol warning labels showed that participants spent only 7% of their time looking at current UK alcohol warning labels (Kersbergen & Field, 2017). The tobacco literature demonstrates that larger warning labels are the most effective (Hammond, 2011), which is relevant to any new warning label and to present return-to-player warning labels. Previous gambling research showed that many regular electronic machine gamblers had not even seen return-to-player warning labels before (Collins *et al.*, 2014), and that many gamblers find return-to-player warning labels confusing (Collins *et al.*, 2014; Harrigan *et al.*, 2017; Rowe *et al.*, 2017). For example, the behavioural risk communication literature suggests that visual aids could boost gamblers' risk understanding (Garcia-Retamero & Cokely, 2017).

Another approach is to educate gamblers not about statistical risks, but about observable features that are correlated with these risks. For example, with alcohol, there are many salient product features such as alcohol type and bottle size that are correlated with ABV%. This paper found that gamblers' losses increased as the potential payoff implied by the betting odds increased, and this occurred both within and between bet types. Bets on longshots have higher average losses than bets on favourites. This is unlike European roulette, where the potential payoff and event probability are perfectly negatively correlated to create a constant average loss of 2.7%. A public health awareness campaign could therefore seek to warn gamblers

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

about the above-average risks of longshot soccer bets. There is some evidence that sports bettors do learn to improve their returns over time, partly by avoiding longshots (Feess *et al.*, 2014). This natural learning process could be enhanced with more explicit advice.

It could be the case that warning labels and education campaigns are an insufficient intervention to debias consumers about the complex statistical information underlying modern soccer bets (Weiss-Cohen *et al.*, 2018) and given the messages in gambling marketing prompting more frequent and riskier gambling (Hing *et al.*, 2018, 2019). This may particularly be the case in soccer betting, given evidence that sports fans are prone to wishful thinking (Babad, 1987) and are overconfident even when the financial stakes are high (Simmons & Massey, 2012). If this is the case, than less freedom-preserving public health interventions should also be considered (Nuffield Council on Bioethics, 2007). For example, a stronger intervention is to 'restrict choice', which is already being used for UK electronic gambling machines, where the maximum bet size will be reduced from £100 to £2 from April 2019 (Casey, 2018). Tax incentives are also being used in UK alcohol policy, with stronger forms of alcohol being subject to different rates of tax duty than lower-strength alcohol. Any public health intervention should be informed by the fundamental data on product risk that this paper attempts to uncover for soccer betting.

Gambling-related harm increases as gamblers' total losses increase (Markham *et al.*, 2014, 2016). However, gamblers' average losses must be combined with betting volume to determine gamblers' total losses. But the betting volume of human sports bettors constitutes data that we did not have access to for this paper. Data on human betting volume could plausibly highlight a different set of high-loss soccer bets that are important for the promotion of responsible gambling. We hope that gambling operators will share more data on gamblers' actual behaviour (Cassidy *et al.*, 2013) in order to determine the extent to which these machine learning strategies reflect actual differences in human betting skill and to enable further research on responsible gambling in humans.

Questions might also be asked as to what machine learning can do to help inform soccer bettors. Although better predictive models than those used in the current paper could always be designed, this should not disqualify our inferences on different betting strategies' *relative* risks. Machine learning has progressed recently in relatively constrained skill-based decision-making domains. Chess champion Garry Kasparov lost to Deep Blue in 1997. Go champion Lee Sedol lost to AlphaGo in 2016. And various top poker professionals have lost special matches to machines (Brown & Sandholm, 2017; Moravčík *et al.*, 2017), with poker being a skill-based gambling form (Potter van Loon *et al.*, 2015). Although machine play is still imperfect for all but

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

the simplest games, machines have attained a high enough level that experts in chess (Kasparov, 2017), go (Metz, 2016) and poker (Newall, 2013) now routinely train with machines. This suggests that machine intelligence should be good enough to inform non-expert soccer bettors as well. Using machine intelligence to inform responsible gambling strategies in humans could well become more internationally relevant over time as the USA moves towards widespread legalised sports betting (Purdum, 2018).

Acknowledgements

Thanks to Tadgh Hegarty and Lukasz Walasek for their useful feedback.

Funding

The authors received no funding for this research.

References

Andersson, P., J. Edman and M. Ekman (2005), 'Predicting the World Cup 2002 in soccer: Performance and confidence of experts and non-experts', *International Journal of Forecasting*, **21**(3): 565–576.

Andersson, P., D. Memmert and E. Popowicz (2009), 'Forecasting outcomes of the World Cup 2006 in football: Performance and confidence of bettors and laypeople', *Psychology of Sport and Exercise*, **10**(1): 116–123.

Angelini, G. and L. De Angelis (2017), 'PARX model for football match predictions', *Journal of Forecasting*, **36**(7): 795–807.

Ayton, P. (1997), 'How to Be Incoherent and Seductive: Bookmakers' Odds and Support Theory', *Organizational Behavior and Human Decision Processes*, **72**(1): 99–115.

Babad, E. (1987), 'Wishful thinking and objectivity among sports fans', *Social Behaviour*, **2**(4): 231–240.

Baio, G. and M. Blangiardo (2010), 'Bayesian hierarchical model for the prediction of football results', *Journal of Applied Statistics*, **37**(2): 253–264.

Bajgrowicz, P. and O. Scaillet (2012), 'Technical trading revisited: False discoveries, persistence tests, and transaction costs', *Journal of Financial Economics*, **106**(3): 473–491.

Brown, N. and T. Sandholm (2017), 'Superhuman AI for heads-up no-limit poker: Libratus beats top professionals', *Science*, **359**(6374): 418–424.

Buhagiar, R., D. Cortis and P. W. S. Newall (2018), 'Why do some soccer bettors lose more money than others?', *Journal of Behavioral and Experimental Finance*, **18**(2018): 85–93.

Casey, J. (2018), *UK Government Confirms FOBT Change Now Set for April 2019*. Available: https://www.gambling.com/news/uk-government-confirms-fobt-change-now-set-for-april-2019-1674400.

Cassidy, R., C. Loussouarn and A. Pisac (2013), *Fair Game: Producing Gambling Research – The Goldsmiths Report*. London: Goldsmiths, University of London.

Chagas, B. T. and J. F. Gomes (2017), 'Internet Gambling: A Critical Review of Behavioural Tracking Research', *Journal of Gambling Issues*, **36**.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

Chen, B. and J. Ankenman (2006), *The Mathematics of Poker*, Pittsburgh, Pennsylvania: ConJelCo LLC.

Collins, D., S. Green, J. d'Ardenne, H. Wardle and S. Williams (2014), *Understanding of Return to Player Messages: Findings from User Testing*, London: NatCen Social Research.

Constantinou, A. C. and N. E. Fenton (2013), 'Profiting from arbitrage and odds biases of the European football gambling market', *The Journal of Gambling Business and Economics*, **7**(2): 41–70.

Constantinou, A., C. & Fenton, N. (2017), 'Towards Smart-Data: Improving predictive accuracy in long-term football team performance', *Knowledge-Based Systems*, **124**: 93–104.

Cortis, D. (2015), 'Expected values and variances in bookmaker payouts: A theoretical approach towards setting limits on odds', *The Journal of Prediction Markets*, **9**(1): 1–14.

Dixon, M. J. and P. F. Pope (2004), 'The value of statistical forecasts in the UK association football betting market', *International Journal of Forecasting*, **20**(4): 697–711.

Dixon, M. and M. Robinson (1998), 'A birth process model for association football matches', *Journal of the Royal Statistical Society: Series D (The Statistician)*, **47**(3): 523–538.

Feess, E., H. Müller and C. Schumacher (2014), 'The favorite-longshot bias and the impact of experience', *Business Research*, **7**(2): 217–234.

Forrest, D. (2008), "Soccer betting in Britain', in *Handbook of Sports and Lottery Markets*, eds. D. B. Hausch and W. T. Ziemba, Elsevier, 421–446.

Gambleaware (2017), *Remote Gambling Research: Interim report on Phase II*. Available: https://about.gambleaware.org/media/1549/gamble-aware_remote-gambling-research_phase-2_pwc-report_august-2017-final.pdf.

Garcia-Retamero, R. and E. T. Cokely (2017), 'Designing visual aids that promote risk literacy: a systematic review of health research and evidence-based design heuristics', *Human Factors*, **59**(4): 582–627.

Goddard, J. (2005), 'Regression models for forecasting goals and match results in association football', *International Journal of Forecasting*, **21**(2): 331–340.

Goddard, J. and I. Asimakopoulos (2004), 'Forecasting football results and the efficiency of fixed-odds betting', *Journal of Forecasting*, **23**(1): 51–66.

Hammond, D. (2011), 'Health warning messages on tobacco products: a review', *Tobacco Control*, **20**(5): 327–337.

Harik, G., E. Cantú-Paz, D. E. Goldberg and B. L. Miller (1999), 'The gambler's ruin problem, genetic algorithms, and the sizing of populations', *Evolutionary computation*, **7**(3): 231–253.

Harrigan, K., D. Brown and K. Barton, R. (2017), *Classification of slot machines in Ontario: Providing relevant information to players*. Guelph, Ontario, Canada: Gambling Research Exchange Ontario.

Harvey, C. R. and Y. Liu (2015), 'Backtesting', *Journal of Portfolio Management*, **42**(1): 13–28.

Hertwig, R. (2017), 'When to consider boosting: some rules for policy-makers', *Behavioural Public Policy*, **1**(2): 143–161.

Hill, I. (1974), 'Association football and statistical inference', *Applied Statistics*, **23**(2): 203–208.

Hing, N., A. M. Russell, E. Li and P. Vitartas (2018), 'Does the uptake of wagering inducements predict impulse betting on sport?', *Journal of Behavioral Addictions*, **7**(1): 146–157.

Hing, N., A. M. Russell, A. Thomas and R. Jenkinson (2019), 'Wagering Advertisements and Inducements: Exposure and Perceived Influence on Betting Behaviour', *Journal of Gambling Studies*.

Hsieh, F. Y., D. A. Bloch and M. D. Larsen (1998), 'A simple method of sample size calculation for linear and logistic regression', *Statistics in Medicine*, **17**(14): 1623–1634.

Hvattum, L. M. and H. Arntzen (2010), 'Using ELO ratings for match result prediction in association football', *International Journal of Forecasting*, **26**(3): 460–470.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

Jones, S. C. and P. Gregory (2009), 'The impact of more visible standard drink labelling on youth alcohol consumption: Helping young people drink (ir) responsibly?', *Drug and Alcohol Review*, **28**(3): 230–234.
Kahneman, D. and A. Tversky (1979), 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica*, **47**(2): 263–292.
Kasparov, G. (2017), *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins*, London, UK: John Murray.
Kaunitz, L., S. Zhong and J. Kreiner (2017), 'Beating the bookies with their own numbers-and how the online sports betting market is rigged', *arXiv preprint arXiv:1710.02824*.
Kersbergen, I. and M. Field (2017), 'Alcohol consumers' attention to warning labels and brand information on alcohol packaging: Findings from cross-sectional and experimental studies', *BMC Public Health*, **17**(1): 123.
Khazaal, Y., A. Chatton, J. Billieux, L. Bizzini, G. Monney, E. Fresard, G. Thorens, G. Bondolfi, N. El-Guebaly and D. Zullino (2012), 'Effects of expertise on football betting', *Substance Abuse Treatment, Prevention, and Policy*, **7**: 18.
Koning, R. H. (2000), 'Balance in competition in Dutch soccer', *Journal of the Royal Statistical Society: Series D (The Statistician)*, **49**(3): 419–431.
Korn, D. A. and H. J. Shaffer (1999), 'Gambling and the health of the public: Adopting a public health perspective', *Journal of Gambling Studies*, **15**(4): 289–365.
Kuypers, T. (2000), 'Information and efficiency: an empirical study of a fixed odds betting market', *Applied Economics*, **32**(11): 1353–1363.
Leitner, C., A. Zeileis and K. Hornik (2010), 'Forecasting sports tournaments by ratings of (prob) abilities: A comparison for the EURO 2008', *International Journal of Forecasting*, **26**(3): 471–481.
Lessmann, S., M. Sung, J. E. Johnson and T. Ma (2012), 'A new methodology for generating and combining statistical forecasting models to enhance competitive event prediction', *European Journal of Operational Research*, **218**(1): 163–174.
Livingstone, C. and R. Woolley (2007), 'Risky business: A few provocations on the regulation of electronic gaming machines', *International Gambling Studies*, **7**(3): 361–376.
Lopez-Gonzalez, H., A. Estévez and M. Griffiths (2017), 'Marketing and advertising online sports betting: a problem gambling perspective', *Journal of Sport and Social Issues*, **41**(3): 256–272.
Markham, F. and M. Young (2015), '"Big gambling": the rise of the global industry-state gambling complex', *Addiction Research & Theory*, **23**(1): 1–4.
Markham, F., M. Young and B. Doran (2016), 'The relationship between player losses and gambling-related harm: evidence from nationally representative cross-sectional surveys in four countries', *Addiction*, **111**(2): 320–330.
Markham, F., M. Young and B. Doran (2014), 'Gambling expenditure predicts harm: evidence from a venue-level study', *Addiction*, **109**(9): 1509–1516.
Maynard, O. M., T. Langfield, A. S. Attwood, E. Allen, I. Drew, A. Votier and M. R. Munafò (2017), 'No impact of calorie or unit information on ad libitum alcohol consumption', *Alcohol and Alcoholism*, **53**(1): 12–19.
McFadden, D. (1974), 'Conditional Logit Analysis of Qualitative Choice Behavior', in *Frontiers in Econometrics*, ed. P. Zarembka, New York: Academic Press, 105–142.
McFadden, D. and K. Train (2000), 'Mixed MNL models for discrete response', *Journal of Applied Econometrics*, **15**(5): 447–470.
Metz, C. (2016), The sadness and beauty of watching google's ai play go. Available: https://www.wired.com/2016/03/sadness-beauty-watching-googles-ai-play-go/.
Mohan, C. (1955), 'The gambler's ruin problem with correlation', *Biometrika*, **42**(3/4): 486–493.
Moravčík, M., M. Schmid, N. Burch, V. Lisý, D. Morrill, N. Bard, T. Davis, K. Waugh, M. Johanson and M. Bowling (2017), 'DeepStack: Expert-level artificial intelligence in heads-up no-limit poker', *Science*, **356**(6337): 508–513.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

Newall, P. (2013), *Further Limit Hold 'Em: Exploring the Model Poker Game*, Las Vegas, NV: Two Plus Two Publishing.
Newall, P. W. S., A. Thobhani, L. Walasek and C. Meyer (2018), *"Impulsiveness and urgency:" Gambling advertising and the 2018 soccer World Cup*. Available: https://psyarxiv.com/3uc9s.
Newall, P. W. S. (2015), 'How bookies make your money', *Judgment and Decision Making*, **10**(3): 225–231.
Newall, P. W. S. (2018), 'Dark nudges in gambling', *Addiction Research & Theory*, **27**(2): 65–67.
Newall, P. W. S. (2017), 'Behavioral complexity of British gambling advertising', *Addiction Research & Theory*, **25**(6): 505–511.
Nuffield Council on Bioethics (2007), *Public Health: Ethical Issues*, UK: Cambridge Publishers.
Oberstone, J. (2009), 'Differentiating the top English premier league football clubs from the rest of the pack: Identifying the keys to success', *Journal of Quantitative Analysis in Sports*, **5**(3).
Orford, J. (2010), *An Unsafe Bet? The Dangerous Rise of Gambling and the Debate We Should Be Having*, Singapore: John Wiley & Sons.
Parke, A., A. Harris, J. Parke, J. Rigbye and A. Blaszczynski (2015), 'Responsible marketing and advertising in gambling: a critical review', *The Journal of Gambling Business and Economics*, **8**(3): 21–35.
Parke, J., A. Parke and A. Blaszczynski (2016), *Key Issues in Product-Based Harm Minimisation: Examining Theory, Evidence and Policy Issues Relevant in Great Britain*, London: Responsible Gambling Trust.
Peters, E. (2017), 'Educating good decisions', *Behavioural Public Policy*, **1**(2): 162–176.
Pleskac, T. J. and R. Hertwig (2014), 'Ecologically rational choice and the structure of the environment', *Journal of Experimental Psychology: General*, **143**(5): 2000–2019.
Portman Group (2017), *A Review of Research into the impacts of Alcohol Warning Labels on attitudes and behaviour*. Available: http://www.portmangroup.org.uk/docs/default-source/alcohol-health-toolkit/final-for-publication-08-sept-17.pdf?sfvrsn=2.
van Loon Potter, R. J., M. J. van den Assem and D. van Dolder (2015), 'Beyond chance? The persistence of performance in online poker', *PloS One*, **10**(3): e0115479.
Purdum, D. (2018), *Inside how sports betting went mainstream*. Available: http://www.espn.com/chalk/story/_/id/24310393/gambling-how-media-daily-fantasy-new-thinking-us-pro-sports-commissioners-helped-sports-betting-become-accepted.
Rowe, B., D. De Ionno, J. Holland, C. Flude, L. Brodrick, A. Vinay and M. Moutos (2017), *Responsible Gambling: Collaborative Innovation Identifying Good Practice and Inspiring Change*, Revealing Reality.
Simmons, J. P. and C. Massey (2012), 'Is optimism real?', *Journal of Experimental Psychology: General*, **141**(4): 630–634.
Smith, M. A., D. Paton and L. Vaughan Williams (2009), 'Do bookmakers possess superior skills to bettors in predicting outcomes?', *Journal of Economic Behavior & Organization*, **71**(2): 539–549.
Stark, D. and D. Cortis (2017), 'Balancing the book: Is it necessary and sufficient?', *The Journal of Gambling Business and Economics*, **11**(1): 1–6.
Stibe, A. and B. Cugelman (2016), 'Persuasive backfiring: when behavior change interventions trigger unintended negative outcomes', in *Persuasive Technology*, eds. A. Meschtscherjakov, B. De Ruyter, V. Fuchsberger, M. Murer and M. Tscheligi, Springer, 65–77.
Štrumbelj, E. (2014), 'On determining probability forecasts from betting odds', *International Journal of Forecasting*, **30**(4): 934–943.
Štrumbelj, E. and M. R. Šikonja (2010), 'Online bookmakers' odds as forecasts: The case of European soccer leagues', *International Journal of Forecasting*, **26**(3): 482–488.
Thaler, R. H. and C. R. Sunstein (2008), *Nudge: Improving Decisions about Health, Wealth, and Happiness*, New Haven, CT: Yale University Press.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press

Thorp, E. O. (1998), 'The invention of the first wearable computer', *Wearable Computers, 1998. Digest of Papers. Second International Symposium on Wearable Computers*. IEEE, 4.

Titman, A., D. Costain, P. Ridall and K. Gregory (2015), 'Joint modelling of goals and bookings in association football', *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **178**(3): 659–683.

Vaughan Williams, L. (1999), 'Information efficiency in betting markets: A survey', *Bulletin of Economic Research*, **51**(1): 1–39.

Weiss-Cohen, L., E. Konstantinidis, M. Speekenbrink and N. Harvey (2018), 'Task complexity moderates the influence of descriptions in decisions from experience', *Cognition*, **170**: 209–227.

Willoughby, K. A. (2002), 'Winning games in Canadian football: A logistic regression analysis', *The College Mathematics Journal*, **33**(3): 215–220.

Woolley, R., C. Livingstone, K. Harrigan and A. Rintoul (2013), 'House edge: hold percentage and the cost of EGM gambling', *International Gambling Studies*, **13**(3): 388–402.

Yarkoni, T. and J. Westfall (2017), 'Choosing prediction over explanation in psychology: Lessons from machine learning', *Perspectives on Psychological Science*, **12**(6): 1100–1122.

https://doi.org/10.1017/bpp.2019.9 Published online by Cambridge University Press