Bettors' reaction to match dynamics — Evidence from in-game betting

Rouven Michels[*][†], Marius Ötting[†], Roland Langrock[†]

Abstract

It is still largely unclear to what extent bettors update their prior assumptions about the strength and form of competing teams considering the dynamics during the match. This is of interest not only from the psychological perspective, but also as the pricing of live odds ideally should be driven both by the (objective) outcome probabilities and also the bettors' behaviour. Using state-space models (SSMs) to account for the dynamically evolving latent sentiment of the betting market, we analyse a unique high-frequency data set on stakes placed during the match. We find that stakes in the live-betting market are driven both by perceived pre-game strength and by in-game strength, the latter as measured by the Valuing Actions by Estimating Probabilities (VAEP) approach. Both effects vary over the course of the match.

1 Introduction

Given the economic relevance of betting markets — the gross gaming revenue was reported as 41.7 billion Euro in 2020 in Europe (European Gaming & Betting Association, 2020) — it is of much interest to understand the behaviour of the market's participants, i.e. the bettors. From the bookmakers' perspective, it is first of all crucial to avoid inefficiencies in the pricing of

---

[*]corresponding author: r.michels@uni-bielefeld.de
[†]Bielefeld University, Universitätsstraße 25, 33615 Bielefeld, Germany

odds. Furthermore, for profit maximisation it may be beneficial to exploit particular patterns in bettors' actions, such as a favourite-longshot bias (or the reverse bias). Such patterns and potential biases are also interesting from the psychological perspective, and several studies have indeed already focused on betting behaviour, e.g. investigating the 'gambler's fallacy' (Tversky and Kahneman, 1974; Clotfelter and Cook, 1993), the reaction to a perceived 'hot hand' (Sundali and Croson, 2006; Paul et al., 2014), but also the drivers of demand for sports bets more generally (Humphreys et al., 2013; Deutscher et al., 2019).

Sports betting takes place both in the pre-game (bets placed before kick-off) and in the in-game (bets placed during games) market. Although only 45% of the betting volume in Europe (European Gaming & Betting Association, 2020) is placed in the pre-game market, empirical research to date has largely focused on this market, investigating market (in-)efficiencies, over- and underreaction of bettors, drivers of betting volume and forecasting of match outcomes, to name but a few (see, e.g. Thaler and Ziemba, 1988; Vergin, 2001; Dixon and Pope, 2004; Feddersen et al., 2017; Deutscher et al., 2018; Brown and Reade, 2019; Butler et al., 2021; Durand et al., 2021). In contrast, for the in-game market, the existing literature to date has mostly focused on inefficiencies (see, e.g., Debnath et al., 2003; Choi and Hui, 2014; Croxson and Reade, 2014), such that in particular the betting behaviour in this important and highly dynamic market is not yet well understood. The principle question motivating this contribution thus is the following: What are the drivers of bet placements in the in-game market? And in particular: To what extent are these related to match dynamics?

For the pre-game betting market, building a forecast based on information such as the (perceived) teams' strength is comparable, conceptually, to a *fundamental* analysis in the financial market (Abarbanell and Bushee, 1997). In the latter, the incentive to buy stocks — the counterpart of a sports bet placement — is that people feel confident that the 'true' value of a company is larger than the actual stock price. Market analysts estimate the former by means of, e.g., the debt-equity-ratio and the earnings before interest, taxes, depreciation, and amortisation (EBITDA; Mukherji et al., 1997, Quirin et al., 2000, Baresa et al., 2013). Thus, investments based on a fundamental analysis correspond to an expectation of net profits in the long run. In contrast, exploiting short-time fluctuations in the financial market can be attempted based on a *technical* analysis (see, e.g. Brown and Jennings, 1989; Edwards et al., 2018). In that case, the investor examines several quantitative indicators such as moving averages to build a short-term forecast (Zhu and Zhou, 2009). While we do not want to stretch the analogy too much, we argue that in-game bettors can, in principle, follow not only a strategy similar to a fundamental analysis, essentially considering the current score

2

relative to fixed variables such as the team strength, but also a strategy more similar to a technical analysis, incorporating information from short-term measures related for example to ball possession, tackles, goal-scoring opportunities or passes completed. To what extent either of these two strategy types drives the placement of bets in the in-game market is the focus of this work.

We investigate the effects of both fixed team information and in-game dynamics on the stakes placed in the in-game market based on two large and high-resolution data sets from the 2017/18 season of the German Bundesliga. The first data set covers detailed betting data, specifically in-game betting odds and volumes for all 306 matches played in the 2017/18 season. This unique data set was provided by a large European bookmaker and allows to investigate not only in-game odds (as has been done in the existing literature) but also stakes placed during the match. The second data set, provided by Pappalardo et al. (2019), comprises WyScout event data, consisting of information on events such as shots on goal, passes, 1-on-1 situations, and set pieces for all 306 matches. The combination of these two large data sets enables us to investigate the potential effects of in-game dynamics on betting behaviour (rather than on the match outcome; such dependence has been studied, e.g., in Weimar and Wicker, 2017, Brechot and Flepp, 2020). Moreover, we investigate how these effects may vary over the course of a match — the expectation being that the importance of in-game (pre-game) information will increase (decrease) over time.

2 Data

We use in-game stakes placed during all 306 matches of the 2017/2018 German Bundesliga season to investigate the investment behaviour of bettors during football matches. In particular, the corresponding data considered cover bets on the match outcome, i.e. home win or away win (we exclude bets on a draw from our analysis). These data, which were provided to us by one of the largest bookmakers in Europe (with most of its customers located in Germany), have a 1 Hz resolution. This temporal resolution is finer than necessary with respect to our research objective, such that to simplify the modelling we aggregate the second-by-second stakes into intervals of one minute. To compare stakes across teams and matches, we calculate *relativestake* per team and minute, where for each interval we divide the amount of stakes placed on each team by the total amount placed on either team. The two teams' relative stakes thus always sum to 1, such that it is sufficient to analyse the relative stakes from the point of view of one team only. Therefore, without loss of generality, in the following

we will consider the relative stakes placed on the home team only. The processed data set then comprises $N = 306$ time series, $\{y_{n,t}\}$, $n = 1, \ldots, 306$, with t, $t = 1, \ldots, T$, indicating the minute of the match. Towards the end of the injury time, much less stakes are placed, such that we truncate all time series at minute $t = 85$ to avoid the need to deal with the much increased variation. Overall, $26\,010$ observations of relative stakes are considered.

To investigate the drivers of betting behaviour, as represented here by the relative stakes being placed, we consider both static (pre-game) as well as dynamic (in-game) covariate information. For the former, we use the pre-game odds as a proxy for the (relative) pre-game strength of a team, as betting odds in general are accurate probabilistic forecasts for match outcomes (see, e.g., Spann and Skiera, 2009), and hence constitute reliable measures for the strength of teams (relative to their opponents). To simplify interpretation, we consider the implied probabilities, i.e. the inverse of the decimal odds, adjusted for the overround (or 'vig'). To account for both teams' strength, we consider the difference of their pre-game winning probabilities (*prewindiff*).

To additionally investigate the effect of in-game actions on the stakes placed, an additional comprehensive data set on in-game match events is considered. These data, again collected at 1 Hz, were provided by the company WyScout and were made publicly available by Pappalardo et al. (2019). They contain information on all relevant actions during the match, together with a time stamp as well as the associated location on the pitch, indicated by x– and y– coordinates. From the WyScout in-game data, we extract covariate information related to the match action, aiming to measure how the in-game team strength evolves dynamically throughout the match. To this end, several simple summary statistics could be used, e.g. the number of shots, the number of passes or the average distance of actions to the opposing goal. However, such simple metrics have been shown to be relatively poor predictors for the match outcome in football (Mackenzie and Cushion, 2012; Carling et al., 2014).

To capture the in-game strength of a team appropriately we thus consider a more refined metric for measuring the *value* of actions on the pitch, namely the so-called Valuing Actions by Estimating Probabilities (VAEP) approach (Decroos et al., 2019). The idea of the VAEP is to measure the value of any action, e.g. a pass or a tackle, with respect to both the probability of scoring and the probability of conceding a goal. For illustration, Figure 1 shows an example sequence of actions and their associated VAEP values, obtained using predictive machine learning methods, in particular gradient-boosted trees — see the Appendix for more details. From the action-level VAEP values, we build the covariate *vaepdiff*, where we consider the differences between the teams' VAEP values aggregated over 1-minute intervals. The higher

Figure 1: The attack which led to a goal scored by Amine Harit in the match between Borussia Dortmund and FC Schalke 04 (4:4) on November 25, 2017. The numbers indicate the VAEP values of the corresponding actions.

the value of *vaepdiff*, the more the momentum of the match is with the team for which the relative stakes are modelled. The top panel in Figure 2 shows an example of how *vaepdiff* evolves over time along with our response variable, the relative stakes, for a scoreless draw. A sequence of positive VAEP differences between minutes 15 and 25 is followed by a shift towards increased betting on Hamburger SV, whereas around minute 40 a decrease in the relative stakes may be caused by negative VAEP differences, i.e. Werder Bremen gaining some momentum. The bottom panel of Figure 2 shows a second example match, one in which goals were scored. The first goal in minute 29 here leads to a shift towards bets on FSV Mainz 05. However, the relative stakes are mostly smaller than 0.5, likely caused by RB Leipzig being the favourite in the match, as their pre-game winning probability was larger than for FSV Mainz 05.

To explore the relationship between the VAEP and the relative stakes in our data, we consider their serial cross-correlations. Figure 3 shows the first 20 cross-correlation lags for all matches in our data (truncated at the time the first goal was scored, to allow for a meaningful

Figure 2: Time series of the relative stakes placed on a win of Hamburger SV (playing against Werder Bremen) and FSV Mainz 05 (playing against RB Leipzig). The vertical bars denote the differences between the competing teams' VAEP values from the perspective of Hamburg and Mainz, respectively, and dashed lines indicate goals. In the bottom panel, the 2-0 and 3-0 were scored in minutes 85 and 90, respectively — we truncate all time series in minute 85.

comparison), indicating a small positive correlation at lower lags. Table 1 displays summary statistics on the VAEP and on all remaining variables considered in our analysis.

Table 1: Descriptive statistics of the variable analysed, *relativestake*, as well as the covariates *prewindiff* and *vaepdiff*. Since the covariate *prewindiff* is constant across matches, the standard deviation was calculated on the basis of one value per match.

	n	mean	st. dev.	min	max
relativestake	26,010	0.493	0.313	0.000	1.000
prewindiff	26,010	0.139	0.317	-0.740	0.851
vaepdiff	26,010	0.004	0.161	-1.091	1.167

Figure 3: Cross-correlations between the VAEP and the relative stakes, with each grey line referring to one of the 306 time series considered, and the red line showing the overall mean for each lag.

3   Modelling relative stakes

The time series of relative stakes shows strong positive serial correlation (cf. Figure 2). This correlation does not result from a potential direct effect of the stakes placed in any given minute on the stakes placed in subsequent one-minute intervals, as individual bettors (for the most part) act independently of each other. Instead, the correlation is induced by the market progressing through different phases: for example, in the top panel of Figure 2, the relative stakes placed on the home team are relatively low between minutes 10–15. In contrast, between minutes 25 and 65, bets are placed predominantly on the home team. Such different phases can be formalised as a latent variable within a state-space model (SSM), which therefore constitutes the natural modelling approach for our data: stakes are driven by the current underlying market phase, more intuitively to be understood as the market sentiment (e.g. leaning towards bets being placed on the home team), and the unobserved market sentiment evolves probabilistically over time, exhibiting serial correlation and hence inducing serial correlation also in the observed time series of stakes.

More specifically, such an SSM comprises two processes, (i) an unobserved, serially corre-lated state process $\{g_{n,t}\}$, in our application to be seen as a proxy of the market sentiment, with higher values implying a higher tendency towards bet placement on the home team, and (ii) an observed state-dependent process $\{y_{n,t}\}$, which is driven by $\{g_{n,t}\}$ and in our case is the time series of relative stakes placed. For simplicity, the subscript n, indicating the match considered, is omitted in the following. The state variables are assumed to be first-order Markovian, i.e.

$$f(g_t|g_1,\ldots,g_{t-1}) = f(g_t|g_{t-1}),$$

and the observations are assumed to be conditionally independent of each other and of previous states, given the current state:

$$f(y_t|g_1,\ldots,g_t,y_1,\ldots,y_{t-1}) = f(y_t|g_t).$$

In the next section, we develop the specific model formulation, i.e. the precise form of the conditional distributions $f(g_t|g_{t-1})$ and $f(y_t|g_t)$, used to model the time series of relative stakes. Details on the implementation of the maximum likelihood estimation of the model parameters are provided in the Appendix.

3.1 Baseline state-space model

Our response variable y_t, the relative amount of stakes placed on the home team, is continuous-valued with support $[0,1]$, rendering the beta distribution a natural choice for modelling purposes. The support of the regular beta distribution is $(0,1)$, such that we use the beta-inflated distribution (BEINF) to account for the fact that in some intervals stakes are placed on one team only (in which case $y_t = 0$ or $y_t = 1$). We follow the parametrisation proposed by Rigby et al. (2019), such that

$$y_t \sim \text{BEINF}(\mu_t, \sigma, p, q), \quad \text{with} \quad f(y_t) = \begin{cases} p, & \text{for } y_t = 0; \\ (1-p-q)h(y_t), & \text{for } 0 < y_t < 1; \\ q, & \text{for } y_t = 1, \end{cases}$$

for $0 \leq y_t \leq 1$. Here $h(y_t)$ is the density function of the regular beta distribution, i.e.

$$h(y_t, a, b) = \frac{y_t^{a-1}(1-y_t)^{b-1}}{\text{B}(a,b)},$$

with the beta function $B(a, b)$. The shape parameters a and b are not directly amenable to regression modelling, such that we consider a reparametrisation in terms of the beta distribution's mean μ_t and its standard deviation σ (Rigby et al., 2019). From μ_t and σ the associated shape parameters are obtained as $a = \mu_t(1 - \sigma^2)/\sigma^2$ and $b = (1 - \mu_t)(1 - \sigma^2)/\sigma^2$. Figure 4 displays the exact dependence structure of our SSM as a directed acyclic graph. Specifically, to account for the dynamic nature of the relative stakes within matches, the mean μ_t is assumed to be time-varying and is modelled as follows:

$$\mu_t = \text{logit}^{-1}(\alpha_0 + \alpha\, prewindiff + g_t). \tag{1}$$

The rationale of the latter two components of the linear predictor is as follows. We include the difference of the teams' pre-game win probabilities (*prewindiff*) to address systematic effects such as potential favourite-longshot biases, or their reverse, throughout matches (see, e.g., Cain et al., 2000, Cain et al., 2003). In addition, information on the current market sentiment is included via g_t, thereby accounting for phases with stronger preferences in bet placement on either team. If a team shows positive actions on the pitch, then the market sentiment towards placing bets on that team should improve, which would be reflected by an increase in g_t.

The unobserved variable corresponding to the market sentiment, g_t, is modelled as an autoregressive process of order 1, with additional covariate dependence:

$$g_t = \phi g_{t-1} + \beta\, vaepdiff_{t-1} + \omega \eta_t, \tag{2}$$

with $\eta_t \overset{\text{iid}}{\sim} \mathcal{N}(0, 1)$, $\omega > 0$ and $\phi \in (-1, 1)$. The in-game covariate *vaepdiff* is included in the state variable g_t, as it seems natural to assume that potential effects of positive actions by a team will not necessarily be instantaneous — i.e. affecting the mean of the relative stakes only in the very minute the action took place — but rather accumulate and persist over some time. For example, from the bettor's perspective, a single one-minute interval with high *vaepdiff* values is not as likely to affect his or her betting decision as a positive spell of say 15 minutes of the team considered, with overall elevated *vaepdiff* values. This is accounted for by allowing positive contributions of *vaepdiff* to accumulate in g_t, which for $\phi > 0$ is persistent and hence to some extent memorises these contributions.

9

Figure 4: Dependence structure of the SSM used to model the relative stakes y_t, driven by the state variable g_t corresponding to the market sentiment. Additional covariate dependence is assumed in both state-dependent and state process, for the former considering static covariate information (pre-game winning probabilities are constant within matches), for the latter considering dynamic covariate information built from event data.

3.2 Varying-coefficient state-space model

The model described in the previous section is to be regarded as a baseline model, which attempts to formalise a) how the relative stakes evolve during a match depending on the market sentiment, and b) how the covariates of primary interest, i.e. *prewindiff* and *vaepdiff*, are most naturally incorporated into this model. We will now extend this baseline model to improve its realism by allowing the effects of these covariates to vary over time and by introducing additional control variables.

Varying coefficients

Due to the dynamic nature of football matches, it may very well be the case that the effects of *prewindiff* and *vaepdiff* vary over time. Specifically, it seems intuitively plausible that the strengths of the teams, as measured by the pre-game win probabilities, are a strong

10

predictor for betting activity at the beginning of a match when little additional information is available. Vice versa, towards the end of a match the teams' performances during the match might become increasingly important for understanding betting dynamics. To allow the effects of *prewindiff* and *vaepdiff* to vary over time, we replace α and β in Eqs. (1) and (2) by time-varying parameters α_t and β_t, respectively. To avoid a priori assumptions on the functional forms of α_t and β_t, we model these functions nonparametrically using B-splines.

Since the seminal paper by Eilers and Marx (1996), the class of B-splines has rapidly gained popularity in nonparametric statistical modelling, and in recent years, B-spline-based modelling of functional effects has been embedded also in various types of SSMs (see, e.g. de Souza and Heckman, 2014; Hambuckers et al., 2018; Mews et al., 2022). In our setting, α_t and β_t, the time-varying effects of *prewindiff* (on the mean relative stakes) and *vaepdiff* (on the state variable indicating the market sentiment), respectively, are modelled as linear combinations of a finite number of section-wise defined basis functions,

$$\alpha_t = \sum_{k=1}^{K} \nu_k^\alpha B_k(t), \quad \beta_t = \sum_{k=1}^{K} \nu_k^\beta B_k(t), \tag{3}$$

for $t = 1, \ldots, 85$, where $B_1, \ldots, B_k$, $k = 1, \ldots, K$, are fixed, equidistant B-spline basis functions of order three. We use cubic polynomial B-splines to obtain a twice continuously differentiable function, thus leading to smooth density estimates (Langrock et al., 2017). To prevent overfitting we add a roughness penalty term, thus considering so-called penalised B-splines, i.e. P-splines (Eilers and Marx, 1996). Specifically, we penalise high values of the second-order differences of adjacent coefficients in the linear combinations above. The sum of these second-order differences corresponds to an approximation of the integrated squared curvature of the functional estimate. The resulting penalised log-likelihood function is then given as follows (cf. Langrock et al., 2017):

$$\ell_p = \log\left(\mathcal{L}_{\text{approx}}\right) - \frac{\lambda_\alpha}{2} \sum_{k=3}^{K} (\Delta^2 \nu_k^\alpha)^2 - \frac{\lambda_\beta}{2} \sum_{k=3}^{K} (\Delta^2 \nu_k^\beta)^2,$$

with the unpenalized likelihood function $\mathcal{L}_{\text{approx}}$ (see Eq. (8) in the Appendix), the second-order differences $\Delta^2 \nu_k = \nu_k - 2\nu_{k-1} + \nu_{k-2}$, and smoothing parameters λ_α and λ_β to control the bias-variance trade-off. For $\lambda_\alpha, \lambda_\beta \to \infty$ the varying coefficients α_t and β_t simplify to a linear effect (Eilers and Marx, 1996). In other words, this nonparametric approach can capture complex time-varying effects if necessary and otherwise will typically collapse to simple linear

modelling due to the penalisation of non-zero curvature.

Following Fahrmeir et al. (2013), the tuning parameters λ_α and λ_β are chosen via the Akaike Information Criterion (AIC), $\text{AIC} = -2 \cdot l + 2 \cdot \widehat{df}$, where l is the unpenalised likelihood under the fitted model and $\widehat{df}$ is an estimate of the degrees of freedom. The latter is obtained as the trace of the product of the Fisher information matrix for the unpenalized likelihood and the inverse Fisher information matrix for the penalized likelihood (see Gray, 1992). We consider the following two-dimensional grid from which the smoothing parameters are chosen:

$$\Lambda_\alpha \times \Lambda_\beta = \{0.05, 0.25, 1, 5, 25, 100, 500\} \times \{0.05, 0.25, 1, 5, 25, 100, 500\}. \tag{4}$$

For each combination of $\lambda_\alpha \in \Lambda_\alpha$ and $\lambda_\beta \in \Lambda_\beta$, the model is fitted and the AIC calculated, then selecting the combination of λ_α and λ_β that yields the lowest AIC value. A more detailed discussion of the implementation of P-splines can be found in Langrock et al. (2017).

Additional control variables

As indicated by the bottom panel of Figure 2, the relative stakes placed on the home team might also be affected by the current score. We thus now also include the difference in the current score at minute t ($scorediff_t$), calculated from the home team's point of view, such that positive values correspond to a lead of the home team. Since considering the difference in the current score alone does not fully reflect a team's winning chances, we further include the in-game win probability in minute t, which is derived by the betting odds in minute t ($winprobteam_t$). As for the pre-game win probabilities, we again consider the inverse of the odds adjusted for the bookmaker's vig.

We add these further covariates to the predictor for the mean of the BEINF distribution:

$$\mu_t = \text{logit}^{-1}\Big(\alpha_0 + \sum_{k=1}^{K} \nu_k^\alpha B_k(t) \; prewindiff + \zeta_1 scorediff_t + \zeta_2 winprobteam_t + g_t\Big). \tag{5}$$

4 Results

4.1 Baseline state-space model

For the baseline SSM specified by (1) and (2), with the effects of *prewindiff* and *vaepdiff* assumed to be constant over time, the parameter estimates are given in Table 2. The persistence in the state process was estimated to be fairly strong ($\hat{\phi} = 0.968$), corresponding to a

positive correlation in the proportional allocation of stakes. In other words, if bets are placed predominantly on either of the two teams, this pattern tends to persist for some time. The estimate of the intercept α_0 is negative, indicating that overall higher stakes are placed on the away team (see also Levitt, 2004, for similar results) when the pre-game win probabilities of both teams are identical (i.e. *prewindiff* $= 0$). A possible explanation of this effect is that bettors may be underestimating the home advantage. The effects of *prewindiff* and *vaepdiff* were both estimated to be positive, confirming the intuition that team strength, both prior to the match and as manifested during the match itself, is valued by bettors. According to the AIC, this model is preferred over the model excluding *prewindiff* (ΔAIC $= 339.17$) as well as over the model excluding *vaepdiff* (ΔAIC $= 522.04$). This is corroborated also by the 95% confidence intervals given in Table 2.

Table 2: Parameter estimates with 95% confidence intervals for the baseline SSM.

parameter	estimate	95% CI
ϕ	0.968	$[0.964; 0.971]$
ω	0.249	$[0.238; 0.261]$
σ	0.300	$[0.296; 0.303]$
α_0	-0.195	$[-0.278; -0.113]$
α $\left(prewindiff\right)$	2.395	$[2.151; 2.640]$
β $\left(vaepdiff\right)$	0.600	$[0.550; 0.651]$

4.2 Varying-coefficient state-space model

The varying-coefficient SSM, which allows the effects of *prewindiff* and *vaepdiff* to change as the match progresses, was estimated using $K = 10$ basis functions to build the functional effects according to (3). The model was fitted for all combinations of tuning parameters from $\Lambda_\alpha \times \Lambda_\beta$ as specified in (4), with the optimal choice $(\lambda_\alpha, \lambda_\beta) = (1, 5)$ according to the AIC. The model including time-varying effects of *prewindiff* and *vaepdiff* is clearly favoured over the simpler model reported in Section 4.1 (ΔAIC $= 270.62$)[1].

Regarding the estimated parametric effects, the results again confirm serial correlation in the state process ($\hat{\phi} = 0.963$). When the market sentiment towards the home team is improved (as induced for example by a large *vaepdiff* value), such that g_t takes higher values, this process tends to persist in this phase for the next few minutes. For the effect of the

---

[1]To select the tuning parameters λ_α and λ_β, we also considered the BIC and the Hannan-Quinn criterion, which led to very similar results.

Figure 5: Time varying P-spline effects of the pre-game winning probabilities ($\lambda_\alpha = 1$) as well as the VAEP variable ($\lambda_\beta = 5$). The white area indicates the intervals data was considered for. The grey area indicates where outer knots are set.

difference in the current score and the in-game win probability, ζ_1 and ζ_2 are both estimated to be positive ($\hat{\zeta}_1 = 0.234$, $\hat{\zeta}_2 = 0.336$) — teams having the lead and those with a higher in-game win probability are thus preferred by bettors.

The nonparametrically estimated time-varying effects of *prewindiff* and *vaepdiff* are shown in Figure 5. The effect of *prewindiff* on stake placement is estimated to be positive throughout the match, but with the effect size decreasing as the match progresses. This matches the intuition that the effect of *prewindiff* — a variable measuring the overall strength of a team but not taking into account the actions on the pitch — should be largest when new information is limited, i.e. at the very beginning of matches. The estimated effect size of *prewindiff* decreases approximately linearly over time. In contrast, for *vaepdiff* we find a highly non-linear functional form of the effect size over time. The effect of *vaepdiff* — a variable measuring the in-game strength of a team — is estimated to be slightly positive throughout the first half of a match, followed by a much more rapid increase in the second half. Bettors thus do value actions on the pitch, with positive spells of a team leading to a shift in the market's sentiment and eventually an increase in the relative stakes placed on that team. Perhaps most

interestingly, the importance of in-game actions as drivers of betting volumes very rapidly increases towards the end of matches.

The results show that bettors incorporate information both on the perceived (pre-game) team quality as well as in-game dynamics when devising their betting strategy. This is in accordance with evidence found for financial markets, where more than 85% of investors rely on both fundamental as well as technical analysis (Lui and Mole, 1998), thus incorporating both the valuation of a company as well as the stock's more short-term momentum.

4.3   Implications for bookmakers

Figure 6: Time series of the win probabilities of Hamburger SV for one example match from the data set (Hamburger SV vs. Werder Bremen). The vertical bars denote the differences between both teams' VAEP values from the point of view of Hamburger SV. Despite phases of dominance seeming to alternate, especially towards the end of the match between minutes 60 and 80, the corresponding win probability — derived by the in-game odds — does not fluctuate accordingly. The changes in the win probability are almost exclusively resulting from the time remaining.

The estimated effect of in-game actions on stake placement indicates that bettors' stake placement is driven by in-game dynamics. In contrast, for the example match Hamburger SV vs. Werder Bremen, Figure 6 illustrates that the betting odds and hence the implied winning probabilities appear to be largely unaffected by in-game actions as proxied by the

15

VAEP differences. As bookmakers thus do not seem to incorporate such information into their betting odds, this raises the question of whether bettors can exploit such potential inefficiencies. To that end, we evaluate an example simple betting strategy: we consider all matches that are tied after halftime and evaluate the outcome of a strategy placing a bet of 1 euro on the home team whenever the *vaepdiff* exceeds certain thresholds. Table 3 displays the returns when following this strategy for our data set, distinguishing several match periods. While positive returns could have been obtained for relatively large *vaepdiff* values in the final phase of a match, it should be noted here that this combination involves only a few data points. The betting strategy would in fact have led to negative returns in the majority of settings, with some of the negative returns substantially larger than the usual vig of about 5%.

Table 3: Possible returns for bettors using simple in-game betting strategies on final match outcomes based on different thresholds for *vaepdiff*. Only matches tied in the second half are considered.

	$vaepdiff > 0.02$	$vaepdiff > 0.03$	$vaepdiff > 0.05$
minute 45-60	-0.27	-0.20	-0.14
minute 60-75	-0.08	-0.17	-0.28
minute 75 until end	-0.01	0.04	0.32

The models developed here could also be used by bookmakers, for example when setting odds or developing automatic fraud detection systems. In particular, the models can be used to study the sensitivity of stake placement with respect to changes in the odds, thereby potentially identifying opportunities for profit maximisation. Furthermore, model-based forecasts of relative stakes can be used to inspect unusual betting behaviour, i.e. to identify stake placements that are not well-explained by the model.

To obtain real-time predictions of relative stakes, depending on the odds, we can calculate one-step-ahead forecasts under the fitted SSM. For this, we use the same discretisation strategy as for the parameter estimation (cf. Appendix), such that the state distribution in minute $t + 1$ is predicted, based on which we obtain the distribution of relative stakes according to Eq. (5). To illustrate such one-step-ahead prediction, Figure 7 displays the time series of relative stakes placed within the example matches already shown in Figure 2, together with the model-based one-step-ahead forecasts during the second half. In addition to showcasing how the model can be used in particular for outlier detection, the figure also indicates that our model provides adequate predictions of the relative stakes for the two example time series.

Figure 7: Observed relative stakes (grey line) and one-step-ahead predictions of relative stakes under our model and conditional on the first half. To illustrate the uncertainty, the light blue points are simulated from the one-step-ahead predicted BEINF distribution at time $t + 1$ (500 simulated observations for each time point). The blue line indicates the predicted mean.

5 Discussion

Our results show that bettors react to in-game dynamics. With this effect being strongest in the second half, it seems that bettors try to exploit the information provided by in-game dynamics. However, evaluating a corresponding simple betting strategy where bets are placed on teams with positive actions, we found that the potential returns to bettors are negative — thus indicating an overreaction by bettors. Such an overreaction to positive events has in fact already been reported in sports betting (Durand et al., 2021; Ötting et al., 2021) and is also known to occur in financial markets (see, e.g., Ma et al., 2005; Piccoli et al., 2017). In particular, Huang et al. (2007) showed that price-earning ratios in the financial market deviate from expectations based on fundamentals, in other words that stock prices can be subject to an overreaction. For the in-game betting market, our results indicate a similar momentum-induced overreaction by bettors.

From a bookmaker's point of view, the overreaction of bettors described above can improve profits. Our analysis does however not reveal whether the bookmaker's current pricing strategy is profit-maximising. A framework for maximising bookmaker revenue by taking into account in-game information was developed by Lorig et al. (2021).

Moreover, for bookmakers, the model-based prediction of relative stakes can be beneficial in detecting unusual betting behaviour. In particular, if the observed relative stakes are substantially higher than the predicted relative stakes — e.g. outside the 99% quantile of the associated forecast distribution — this could indicate insider information or fraud. However, a comprehensive fraud detection system should additionally take into account the absolute (rather than only the relative) stakes (Ötting et al., 2018).

In more general terms, while our study focused on a specific aspect of live betting — namely to what extent in-game dynamics affect bettor behaviour — it also illustrates the immense potential of sports data, the availability of which has improved thanks to data providers such as WyScout or StatsBomb. The analysis of such complex data requires sophisticated statistical modelling, and SSMs as applied in the present paper constitute a versatile framework to accommodate the time series nature of most sports data (cf. Koopmeiners, 2012; Green and Zwiebel, 2018; Ötting et al., 2020; Mews and Ötting, 2022). We thus anticipate an uptake of this type of modern statistical modelling tools in future research in particular into the dynamics of live-betting markets, but also in other sports settings.

Acknowledgements

References

Abarbanell, J. S. and Bushee, B. J. (1997). Fundamental analysis, future earnings, and stock prices. *Journal of Accounting Research*, 35(1):1–24.

Baresa, S., Bogdan, S., and Ivanovic, Z. (2013). Strategy of stock valuation by fundamental analysis. *UTMS Journal of Economics*, 4(1):45–51.

Bartolucci, F. and De Luca, G. (2001). Maximum likelihood estimation of a latent variable time-series model. *Applied Stochastic Models in Business and Industry*, 17(1):5–17.

Brechot, M. and Flepp, R. (2020). Dealing with randomness in match outcomes: how to rethink performance evaluation in European club football using expected goals. *Journal of Sports Economics*, 21(4):335–362.

Brown, A. and Reade, J. J. (2019). The wisdom of amateur crowds: Evidence from an online community of sports tipsters. *European Journal of Operational Research*, 272(3):1073–1081.

Brown, D. P. and Jennings, R. H. (1989). On technical analysis. *The Review of Financial Studies*, 2(4):527–551.

Butler, D., Butler, R., and Eakins, J. (2021). Expert performance and crowd wisdom: Evidence from English Premier League predictions. *European Journal of Operational Research*, 288(1):170–182.

Cain, M., Law, D., and Peel, D. (2000). The favourite-longshot bias and market efficiency in UK football betting. *Scottish Journal of Political Economy*, 47(1):25–36.

Cain, M., Law, D., and Peel, D. (2003). The favourite-longshot bias, bookmaker margins and insider trading in a variety of betting markets. *Bulletin of Economic Research*, 55(3):263–273.

Carling, C., Wright, C., Nelson, L. J., and Bradley, P. S. (2014). Comment on 'Performance analysis in football: A critical review and implications for future research'. *Journal of Sports Sciences*, 32(1):2–7.

Choi, D. and Hui, S. K. (2014). The role of surprise: Understanding overreaction and underreaction to unanticipated events using in-play soccer betting market. *Journal of Economic Behavior & Organization*, 107:614–629.

Clotfelter, C. T. and Cook, P. J. (1993). The "gambler's fallacy" in lottery play. *Management Science*, 39(12):1521–1525.

Croxson, K. and Reade, J. (2014). Information and efficiency: Goal arrival in soccer betting. *The Economic Journal*, 124(575):62–91.

de Souza, C. P. and Heckman, N. E. (2014). Switching nonparametric regression models. *Journal of Nonparametric Statistics*, 26(4):617–637.

Debnath, S., Pennock, D. M., Giles, C. L., and Lawrence, S. (2003). Information incorporation in online in-game sports betting markets. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 258–259.

Decroos, T., Bransen, L., Van Haaren, J., and Davis, J. (2019). Actions speak louder than goals: Valuing player actions in soccer. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1851–1861.

Deutscher, C., Frick, B., and Ötting, M. (2018). Betting market inefficiencies are short-lived in German professional football. *Applied Economics*, 50(30):3240–3246.

Deutscher, C., Ötting, M., Schneemann, S., and Scholten, H. (2019). The demand for English Premier League soccer betting. *Journal of Sports Economics*, 20(4):556–579.

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

Durand, R., Patterson, F., and Shank, C. (2021). Behavioral biases in the NFL gambling market: Overreaction to news and the recency bias. *Journal of Behavioral and Experimental Finance*, 31:100522.

Edwards, R. D., Magee, J., and Bassetti, W. C. (2018). *Technical Analysis of Stock Trends*. CRC Press.

Eilers, P. H. C. and Marx, B. D. (1996). Flexible smoothing with B-splines and penalties. *Statistical Science*, 11(2):89 – 121.

European Gaming & Betting Association (2020). European Online Gambling 2020 Edition.

Fahrmeir, L., Kneib, T., Lang, S., and Marx, B. (2013). *Regression: Models, Methods and Applications*. Springer-Verlag, Berlin.

Feddersen, A., Humphreys, B. R., and Soebbing, B. P. (2017). Sentiment bias and asset prices: Evidence from sports betting markets and social media. *Economic Inquiry*, 55(2):1119–1129.

Gray, R. J. (1992). Flexible methods for analyzing survival data using splines, with applications to breast cancer prognosis. *Journal of the American Statistical Association*, 87(420):942–951.

Green, B. and Zwiebel, J. (2018). The hot-hand fallacy: Cognitive mistakes or equilibrium adjustments? evidence from Major League Baseball. *Management Science*, 64(11):5315–5348.

Hambuckers, J., Kneib, T., Langrock, R., and Silbersdorff, A. (2018). A Markov-switching generalized additive model for compound Poisson processes, with applications to operational loss models. *Quantitative Finance*, 18(10):1679–1698.

Huang, Y., Tsai, C.-H., and Chen, C. R. (2007). Expected p/e, residual p/e, and stock return reversal: Time-varying fundamentals or investor overreaction? *International Journal of Business and Economics*, 6(1):11.

Humphreys, B. R., Paul, R. J., and Weinbach, A. P. (2013). Consumption benefits and gambling: Evidence from the NCAA basketball betting market. *Journal of Economic Psychology*, 39:376–386.

Kitagawa, G. (1987). Non-Gaussian state-space modeling of nonstationary time series. *Journal of the American Statistical Association*, 82(400):1032–1041.

Koopmeiners, J. S. (2012). A comparison of the autocorrelation and variance of NFL team strengths over time using a Bayesian state-space model. *Journal of Quantitative Analysis in Sports*, 8(3).

Langrock, R. (2011). Some applications of nonlinear and non-gaussian state–space modelling by means of hidden markov models. *Journal of Applied Statistics*, 38(12):2955–2970.

Langrock, R., Kneib, T., Glennie, R., and Michelot, T. (2017). Markov-switching generalized additive models. *Statistics and Computing*, 27(1):259–270.

Levitt, S. D. (2004). Why are Gambling Markets Organised so Differently from Financial Markets? *The Economic Journal*, 114(495):223–246.

Lorig, M., Zhou, Z., and Zou, B. (2021). Optimal bookmaking. *European Journal of Operational Research*, 295(2):560–574.

Lui, Y.-H. and Mole, D. (1998). The use of fundamental and technical analyses by foreign exchange dealers: Hong Kong evidence. *Journal of International Money and Finance*, 17(3):535–545.

Ma, Y., Tang, A. P., and Hasan, T. (2005). The stock price overreaction effect: Evidence on Nasdaq stocks. *Quarterly Journal of Business and Economics*, 44(3/4):113–127.

Mackenzie, R. and Cushion, C. (2012). Performance analysis in football: A critical review and implications for future research. *Journal of Sports Sciences*, 31(3):639–676.

Mews, S., Langrock, R., Ötting, M., Yaqine, H., and Reinecke, J. (2022). Maximum approximate likelihood estimation of general continuous-time state-space models. *Statistical Modelling*, –(–):in press.

Mews, S. and Ötting, M. (2022). Continuous-time state-space modelling of the hot hand in basketball. *AStA Advances in Statistical Analysis*, –(–):in press.

Mukherji, S., Dhatt, M. S., and Kim, Y. H. (1997). A fundamental analysis of Korean stock returns. *Financial Analysts Journal*, 53(3):75–80.

Ötting, M., Langrock, R., and Deutscher, C. (2018). Integrating multiple data sources in match-fixing warning systems. *Statistical Modelling*, 18(5-6):483–504.

Ötting, M., Langrock, R., Deutscher, C., and Leos-Barajas, V. (2020). The hot hand in professional darts. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 183(2):565–580.

Ötting, M., Michels, R., Langrock, R., and Deutscher, C. (2021). The reaction to news in live betting. *Unpublished results*.

[dataset] Pappalardo, L., Cintia, P., Rossi, A., Massucco, E., Ferragina, P., Pedreschi, D., and Giannotti, F. (2019). A public data set of spatio-temporal match events in soccer competitions. *Scientific Data*, 6(1):1–15.

Paul, R. J., Weinbach, A. P., and Humphreys, B. (2014). Bettor belief in the "hot hand" evidence from detailed betting data on the NFL. *Journal of Sports Economics*, 15(6):636–649.

Piccoli, P., Chaudhury, M., Souza, A., and Silva, W. (2017). Stock overreaction to extreme market events. *The North American Journal of Economics and Finance*, 41:97–111.

Quirin, J. J., Berry, K. T., and O'Brien, D. (2000). A fundamental analysis approach to oil and gas firm valuation. *Journal of Business Finance & Accounting*, 27(7-8):785–820.

Rigby, R. A., Stasinopoulos, M. D., Heller, G. Z., and De Bastiani, F. (2019). *Distributions for Modeling Location, Scale, and Shape: Using GAMLSS in R*. CRC Press.

Spann, M. and Skiera, B. (2009). Sports forecasting: a comparison of the forecast accuracy of prediction markets, betting odds and tipsters. *Journal of Forecasting*, 28(1):55–72.

Sundali, J. and Croson, R. (2006). Biases in casino betting: The hot hand and the gambler's fallacy. *Judgement and Decision Making*, 1(1):1.

Thaler, R. H. and Ziemba, W. T. (1988). Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic perspectives*, 2(2):161–174.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.

Vergin, R. C. (2001). Overreaction in the NFL point spread market. *Applied Financial Economics*, 11(5):497–509.

Weimar, D. and Wicker, P. (2017). Moneyball revisited: Effort and team performance in professional soccer. *Journal of Sports Economics*, 18(2):140–161.

Zhu, Y. and Zhou, G. (2009). Technical analysis: An asset allocation perspective on the use of moving averages. *Journal of Financial Economics*, 92(3):519–544.

Zucchini, W., MacDonald, I. L., and Langrock, R. (2016). *Hidden Markov Models for Time Series: An Introduction Using R*. Chapman and Hall/CRC.

Appendix

A.1   Computation of the VAEP values

The VAEP values are obtained using predictive machine learning as proposed in Decroos et al. (2019), fitting gradient-boosted trees to the data on in-game actions across five different seasons as well as one World Cup and one European Championship (data provided by Pappalardo et al., 2019). We obtain positive but small values for successful passes and dribblings in midfield, whereas corresponding actions yield larger VAEP values when occurring closer to the goal.

A.2   Calculating and optimising the SSM likelihood

The likelihood of our non-linear and non-Gaussian SSM involves multiple integrals, which we evaluate numerically by finely discretising the state process as first suggested by Kitagawa (1987). This discretisation corresponds to a reformulation of the continuous-state SSM as a discrete-state hidden Markov model (HMM) with a large state space (Zucchini et al., 2016). Using the Markov property assumed for the state process and the conditional independence of the observations, given the states, the likelihood can be written as

$$\mathcal{L}(\theta) = f(y_1, \ldots, y_T) = \int \ldots \int f(g_1)f(y_1|g_1) \prod_{t=2}^{T} f(g_t|g_{t-1}f(y_t|g_t) \; dg_T \ldots g_1, \qquad (6)$$

where the vector θ contains all model parameters. While the integrand has a simple form, the multiple integration makes this expression analytically intractable. Our approach to (approximately) evaluating the right side of Eq. (6) is based on very finely discretising the state space, thus replacing the integrals by approximating sums. More specifically, for the state process g_t we consider a possible range $[c_0, c_m]$, $c_0, c_m \in \mathbb{R}$, chosen sufficiently large as to cover virtually all possible values the process may take. We split this range into m equidistant intervals $C_i = (c_{i-1}, c_i)$, $i = 1, \ldots, m$, of length $h = (c_m - c_0)/m$, and denote the midpoint of the i–th interval by $c_i^\star$. Using quadrature with a simple midpoint rule over each of the intervals $C_1, \ldots, C_m$, expression (6) is approximated by

$$\mathcal{L}_{\text{approx}} = h^T \sum_{i_1=1}^{m} \ldots \sum_{i_T=1}^{m} f(c_{i_1}^\star)f(y_1|g_1 = c_{i_1}^\star) \prod_{t=2}^{T} f(c_{i_t}^\star|g_{t-1} = c_{i_{t-1}}^\star)f(y_t|g_t = c_{i_t}^\star). \qquad (7)$$

While the calculation of this expression is not computationally feasible given the large num-

ber of summands (m^T), this reformulation allows us to apply recursive techniques from the HMM toolbox, in particular for calculating the approximate likelihood (for a single match). Specifically, the approximated likelihood in (7) is *precisely* the likelihood of a particular m-state HMM. The initial state distribution of this HMM is given by the m-dimensional vector $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_m)$ with $\delta_i = hf(c_i^\star)$. Similarly, the $m \times m$ transition probability matrix (t.p.m.) Γ of the HMM has entries $\gamma_{ij} = hf(c_j^\star | c_i^\star)$ — the approximate probability of the state process moving from interval C_i to interval C_j. In our model, Γ additionally depends on covariates, such that we use the notation $\Gamma^{(t)}$ to make it explicit that the t.p.m. varies over time. Finally, the m-state HMM involves the state-dependent densities $f(y_t | c_i^\star)$, $i = 1, \ldots, m$, representing the approximate density of the observation y_t, given that the state process g_t is in the subinterval C_i at time t. Having recognised (7) as the likelihood of this particular HMM, we can use the corresponding powerful tools for inference, first and foremost the forward algorithm for efficiently evaluating (7), yielding the matrix product expression

$$\mathcal{L}_{\text{approx}} = \boldsymbol{\delta}\mathbf{P}(y_1)\boldsymbol{\Gamma}^{(1)}\mathbf{P}(y_2)\boldsymbol{\Gamma}^{(2)}\mathbf{P}(y_3)\ldots\boldsymbol{\Gamma}^{(T-2)}\mathbf{P}(y_{T-1})\boldsymbol{\Gamma}^{(T-1)}\mathbf{P}(y_T)\mathbf{1}, \qquad (8)$$

with $\mathbf{1} = (1, \ldots, 1) \in \mathbb{R}^m$. This discretisation trick allows us to calculate an arbitrarily accurate approximation of the SSM likelihood at computational cost $\mathcal{O}(m^2 T)$ only (see also, for example, Bartolucci and De Luca, 2001; Langrock, 2011; Zucchini et al., 2016; Mews et al., 2022).

To evaluate the likelihood for the complete data set, independence of stakes placed across matches is assumed, such that the joint likelihood for all 306 matches is simply the product of the likelihoods of the individual matches. Parameter estimation is then carried out numerically by optimising the likelihood — or, in case of the varying-coefficient model, the penalised likelihood — using a Newton-Raphson-type scheme.