

Mews, Sina; Ötting, Marius

Article — Published Version

Continuous-time state-space modelling of the hot hand in basketball

AStA Advances in Statistical Analysis

Provided in Cooperation with:

Springer Nature

Suggested Citation: Mews, Sina; Ötting, Marius (2021) : Continuous-time state-space modelling of the hot hand in basketball, AStA Advances in Statistical Analysis, ISSN 1863-818X, Springer, Berlin, Heidelberg, Vol. 107, Iss. 1-2, pp. 313-326,
<https://doi.org/10.1007/s10182-021-00410-y>

This Version is available at:

<https://hdl.handle.net/10419/286852>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>



Continuous-time state-space modelling of the hot hand in basketball

Sina Mews¹ · Marius Ötting¹

Received: 30 October 2020 / Accepted: 26 May 2021 / Published online: 29 June 2021
© The Author(s) 2021

Abstract

We investigate the hot hand phenomenon using data on 110,513 free throws taken in the National Basketball Association. As free throws occur at unevenly spaced time points within a game, we consider a state-space model formulated in continuous time to investigate serial dependence in players' success probabilities. In particular, the underlying state process can be interpreted as a player's (latent) varying form and is modelled using the Ornstein-Uhlenbeck process. Our results support the existence of the hot hand, but the magnitude of the estimated effect is rather small as the underlying success probabilities are elevated by only a few percentage points.

Keywords Free throws · Hot hand · Irregularly sampled data · Ornstein-Uhlenbeck process · Sports analytics · State-space model

1 Introduction

In several areas of society, it remains an open question whether a “hot hand” effect exists, according to which humans may temporarily enter a state during which they perform better than on average. While this concept may occur in different fields, such as among hedge fund managers and artists (Jagannathan et al. 2010; Liu et al. 2018), it is most prominent in sports. Sports commentators and fans — especially in basketball — often refer to players as having a “hot hand”, and being “on fire” or “in the zone” when they show a (successful) streak in performance. In the academic literature, the hot hand has gained great interest since the seminal paper by Gilovich et al. (1985), who investigated a potential hot hand effect in basketball. They found no evidence for its existence and attributed the hot hand to a cognitive illusion, much to the disapproval of many athletes and fans. Still, the results provided by Gilovich et al. (1985) have often been used as a primary example for showing that humans

✉ Sina Mews
sina.mews@uni-bielefeld.de

¹ Department of Business Administration and Economics, Bielefeld University, Bielefeld, Germany

over-interpret patterns of success and failure in random sequences (see, e.g. Thaler and Sunstein 2009; Kahneman 2011).

During the last decades, many studies have attempted to replicate or refute the results of Gilovich et al. (1985), analysing sports such as volleyball, baseball, golf, and especially basketball. Bar-Eli et al. (2006) provide an overview of 25 studies on the hot hand, 11 of which were in favour of the hot hand phenomenon, while 13 studies provided evidence against the hot hand, and one study remained inconclusive. Several more recent studies, often based on large data sets, support the existence of the hot hand (see, e.g. Raab et al. 2012; Green and Zwiebel 2017; Miller and Sanjurjo 2018; Chang 2019). Notably, Miller and Sanjurjo (2018) show that the original study from Gilovich et al. (1985) suffers from a selection bias. Using the same data as in the original study by Gilovich et al. (1985), Miller and Sanjurjo (2018) account for that bias, and their results do reveal a hot hand effect. However, there are also recent studies which provide mixed results (see, e.g. Wetzels et al. 2016) or which do not find evidence for the hot hand, such as Morgulev et al. (2020). Thus, more than 30 years after the study of Gilovich et al. (1985), the existence of the hot hand remains highly disputed.

Moreover, the literature does not provide a universally accepted statistical modelling approach for the hot hand effect. While some studies regard it as serial correlation in *outcomes* (see, e.g. Gilovich et al. 1985; Dorsey-Palmateer and Smith 2004; Miller and Sanjurjo 2018), others consider it as serial correlation in *success probabilities* (see, e.g. Albert 1993; Wetzels et al. 2016; Ötting et al. 2020). The latter modelling approach translates into a latent (state) process underlying the observed performance — intuitively speaking, a measure for a player's form — which can be elevated without the player necessarily being successful in every attempt. In our analysis, we follow this approach and hence consider state-space models (SSMs) to investigate the hot hand effect in basketball. Specifically, we analyse free throws from more than 9000 games played in the National Basketball Association (NBA), totalling in 110,513 observations. In contrast, Gilovich et al. (1985) use data on 2600 attempts in their controlled shooting experiment.

Free throws in basketball, or similar events in sports with game clocks, occur at unevenly spaced time points. These varying time lengths between consecutive attempts may affect inference on the hot hand effect if the model formulation does not account for the temporal irregularity of the observations. As an illustrative example, consider an irregular sequence of throws with intervals ranging from, say, two seconds to 15 minutes. For intervals between attempts that are fairly short (such as a few seconds), players will most likely be able to retain their form from the last shot. On the other hand, if several minutes elapse before players take their next shot, it becomes less likely that they are able to retain their form from the last attempt. However, we found that existing studies on the hot hand do not account for different interval lengths between attempts. In particular, studies investigating serial correlation in success probabilities usually consider discrete-time models that require the data to follow a regular sampling scheme and thus, cannot (directly) be applied to irregularly sampled data. In our contribution, we overcome this limitation by formulating our model in continuous time to explicitly account for irregular time intervals between free throws in basketball. Specifically, we consider a stochastic

differential equation (SDE) as latent state process, namely the Ornstein-Uhlenbeck (OU) process, which represents the underlying form of players fluctuating continuously around their average performance.

In the following, Sect. 2 presents our data set and covers some descriptive statistics. Subsequently, in Sect. 3, the continuous-time SSM formulation for the analysis of the hot hand effect is introduced, while its results are presented in Sect. 4. We conclude our paper with a discussion in Sect. 5.

2 Data

We extracted data on all basketball games played in the NBA between October 2012 and June 2019 from <https://www.basketball-reference.com/>, covering both regular seasons and playoff games. For our analysis, we consider data only on free throw attempts as these constitute a highly standardised setting without any interaction between players, which is usually hard to account for when modelling field goals in basketball. Specifically, when considering field goals in basketball, several additional factors may affect the outcome of a shot, such as the position of the player, the position of the opposing team’s players, and the effort of the defence. If this information on a shot is ignored, a corresponding hot hand analysis might suffer from an omitted-variable bias.

We included all players who took at least 2000 free throws in the period considered, totalling in 110,513 free throws from 44 players. For each player, we included only those games in which he attempted at least four free throws to ensure that throws did not only follow successively (as players receive up to three free throws if they are fouled). In our analysis of the hot hand effect, we are interested in within-game variations in a player’s form. A single sequence of free throw attempts thus consists of all throws taken by one player in a given game, totalling in 15,075 throwing sequences with a median number of 6 free throws per game (min: 4; max: 39).

As free throws occur irregularly within a basketball game, the information on whether an attempt was successful needs to be supplemented by its time point $t_k, k = 1, \dots, T$, where $0 \leq t_1 \leq t_2 \leq \dots \leq t_T$, corresponding to the time already played (in minutes) as indicated by the game clock. For each player p in his n -th game, we thus consider an irregular sequence of binary variables $Y_{t_1}^{p,n}, Y_{t_2}^{p,n}, \dots, Y_{t_{T,p,n}}^{p,n}$, with

$$Y_k^{p,n} = \begin{cases} 1 & \text{if free throw attempt at time } t_k \text{ is successful;} \\ 0 & \text{otherwise.} \end{cases}$$

In our sample, the proportion of successful free throw attempts is obtained as 0.784. However, there is considerable heterogeneity in the players’ throwing success as the corresponding empirical proportions range from 0.451 (Andre Drummond) to 0.906 (Stephen Curry). Players can receive up to three free throws (depending on the foul) in the NBA, which are then thrown in quick succession, and the proportion of successful free throws differs substantially between the three attempts, with 0.769, 0.8, and 0.883 obtained for the first, second, and third free throw, respectively. To

account for the position of the throw in a player's set of (at most) three free throws, we hence include the dummy variables *ft2* and *ft3* in our analysis. In our sample, 54.5% of all free throws correspond to the first, 43.7% to the second, and only 1.8% to the third attempt in a set (cf. Table 1). Furthermore, as the outcome of a free throw is likely affected by intermediate information on the game — such as a close game leading to pressure situations — we consider several further covariates, which were also used in previous studies (see, e.g. Toma 2017; Morgulev et al. 2020). Specifically, we consider the current score difference (*scorediff*), a home dummy variable (*home*), and a dummy variable indicating whether the free throw occurred in the last 30 seconds of the quarter (*last30*). Corresponding summary statistics are shown in Table 1.

In Table 2, example throwing sequences used in our analysis are shown for free throws taken by LeBron James in five NBA games. These throwing sequences illustrate that free throw attempts often appear in clusters of two or three attempts at the same time (depending on the foul), followed by a time period without any free throws. Therefore, it is important to take into account the different lengths of the time intervals between consecutive attempts as the time elapsed between free throws affects a player's underlying form.

3 Continuous-time modelling of the hot hand

3.1 State-space model specification

Following the idea that the throwing success depends on a player's current (latent) form (see, e.g. Albert 1993; Wetzels et al. 2016; Ötting et al. 2020), we model the observed free throw attempts using a state-space model formulation as represented in Fig. 1. The observation process corresponds to the binary sequence of a player's throwing success, while the state process can be interpreted as a player's underlying form (or “hotness”). We further include the covariates introduced in Sect. 2 in the model, which possibly affect a player's throwing success. In particular, we model the binary response of throwing success $Y_{i_k}^{p,n}$ using a Bernoulli distribution with the associated success probability $\pi_{i_k}^{p,n}$ being a function of the player's current state $S_{i_k}^{p,n}$ and the covariates. Dropping the superscripts p and n for notational simplicity from now on, we thus have

Table 1 Descriptive statistics of the covariates

	Mean	St. dev.	Min.	Max.
<i>Scorediff</i>	0.576	9.860	− 45	49
<i>Home</i>	0.514	–	0	1
<i>Last30</i>	0.093	–	0	1
<i>ft2</i>	0.437	–	0	1
<i>ft3</i>	0.018	–	0	1

Table 2 Throwing sequences of LeBron James

Miami Heat @ Houston Rockets, November 12, 2012												
$y_{t_k}^{\text{James},1}$	0	1	0	1	1	0	1	1	1	1	1	1
t_k (in min.)	11.01	11.01	23.48	23.48	46.68	46.68	47.09	47.09	47.09	47.09	47.09	47.09
Miami Heat @ Los Angeles Clippers, November 14, 2012												
$y_{t_k}^{\text{James},2}$	0	1	0	0	1	1	1	1	1	1	1	1
t_k (in min.)	9.47	23.08	23.08	24.73	36.95	36.95	41	41	41	42.77	42.77	42.77
Milwaukee Bucks @ Miami Heat, November 21, 2012												
$y_{t_k}^{\text{James},3}$	0	1	0	1	0	1	1	1	1	1	1	1
t_k (in min.)	7.5	7.5	10.02	10.02	35.62	35.62	42.68	42.68	42.68	42.68	42.68	42.68
Cleveland Cavaliers @ Miami Heat, November 24, 2012												
$y_{t_k}^{\text{James},4}$	1	1	1	0	1	1	1	1	1	1	1	1
t_k (in min.)	11.62	21.07	21.07	21.38	21.38	31.95	31.95	31.95	32.68	32.68	32.68	32.68
Los Angeles Lakers @ Miami Heat, January 23, 2014												
$y_{t_k}^{\text{James},5}$	0	1	0	1	1	0	0	0	1	1	0	1
t_k (in min.)	9.28	9.28	20.53	25.97	25.97	31.57	31.57	31.57	33.43	33.43	42.1	47.62




Fig. 1 Dependence structure of our SSM: the throwing success Y_{t_k} is assumed to be driven by the underlying (latent) form of a player. To explicitly account for the irregular time intervals between observations, we formulate our model in continuous time

$$Y_{t_k} \sim \text{Bern}(\pi_{t_k}), \quad \text{logit}(\pi_{t_k}) = S_{t_k} + \beta_{0,p} + \beta_1 \text{home} + \beta_2 \text{scorediff} + \beta_3 \text{last30} + \beta_4 \text{ft2} + \beta_5 \text{ft3}, \tag{1}$$

where $\beta_{0,p}$ is a player-specific intercept to account for differences between players' average throwing success. To address the temporal irregularity of the free throw attempts, we formulate the stochastic process $\{S_t\}_{t \geq 0}$ in continuous time. Furthermore, we require the state process to be continuous-valued to allow for gradual changes in a player's form, rather than assuming a finite number of discrete states (e.g. three states interpreted as cold vs. normal vs. hot; cf. Wetzels et al. 2016; Green and Zwiebel 2017). In addition, the state process ought to be stationary such that in the long-run a player returns to his average form. A natural candidate for a corresponding stationary, continuous-valued and continuous-time process is the OU process, which is described by the following SDE:

$$dS_t = \theta(\mu - S_t)dt + \sigma dB_t, \quad S_0 = s_0, \tag{2}$$

where $\theta > 0$ is the reversion parameter indicating the strength of reversion to the long-term mean $\mu \in \mathbb{R}$, while $\sigma > 0$ is the diffusion parameter controlling the strength of fluctuations, and B_t denotes the Brownian motion. We further specify $\mu = 0$ to ensure that our model is identifiable, leading to the state process fluctuating around a player's average form, given the current covariate values. In particular, positive values of the state process indicate higher success probabilities, whereas negative values indicate decreased throwing success, given the player's average ability and the current game characteristics.

As shown in Fig. 1, we model the hot hand effect as serial correlation in success probabilities as induced by the state process. Specifically, the observed free throw attempts are conditionally independent, given the underlying states, while the unobserved state process induces correlation in the observation process. Regarding the hot hand effect, the reversion parameter θ of the OU process is thus of main interest as it governs the speed of reversion (to the average form). The smaller θ , the longer it takes for the OU process to return to its mean and hence the higher the serial correlation. To assess whether a model including serial dependence (i.e. an SSM) is actually needed to describe the structure in the data, we additionally fit a benchmark model without the underlying state variable S_{t_k} in Eq. (1). Consequently, the

benchmark model corresponds to the absence of any hot hand effect, i.e. a standard logistic regression model. We compare the fit of both models to the data using a cross-validated likelihood score as suggested by Celeux and Durand (2008). While standard information criteria such as AIC and BIC are often considered for this purpose, it may very well be the case that they underpenalise the flexible parameters of the OU process in our SSM — a similar issue is well-known for nonparametric and mixed models (see, e.g. Hurvich et al. 1998; Müller et al. 2013). For the SSM considered here, it is anything but clear to what extent the latent states should be counted as parameters and contribute to the penalisation term (see, e.g. Auger-Méthé et al. in press). We circumvent these difficulties by instead using a cross-validated likelihood score to compare our models. We split the data set into 20 random samples, where 90% of each individual’s throwing sequences serve as training data, while the remaining 10% serve as test data. Fitting each model to the training data, we subsequently consider the out-of-sample log-likelihood on the remaining 10% of the throwing sequences to assess the models.

3.2 Statistical inference

The likelihood of the continuous-time SSM given by Eqs. (1) and (2) involves integration over all possible realisations of the continuous-valued state S_{t_k} , at each observation time t_1, t_2, \dots, t_T . For simplicity of notation, let the integer $\tau = 1, 2, \dots, T$ denote the *index* of the observation in the time series, such that Y_{t_τ} shortens to Y_τ and S_{t_τ} shortens to S_τ . Further, t_τ represents the *time* at which the observation τ was collected. Then the likelihood of a single throwing sequence y_1, \dots, y_T is given by

$$\begin{aligned} \mathcal{L}_T &= \int \dots \int p(y_1, \dots, y_T, s_1, \dots, s_T) ds_T \dots ds_1 \\ &= \int \dots \int p(s_1) \Pr(y_1 | s_1) \prod_{\tau=2}^T p(s_\tau | s_{\tau-1}) \Pr(y_\tau | s_\tau) ds_T \dots ds_1, \end{aligned} \tag{3}$$

where we assume that each player starts a game in his stationary distribution $S_1 \sim \mathcal{N}\left(0, \frac{\sigma^2}{2\theta}\right)$, i.e. the stationary distribution of the OU process. Further, we assume Y_τ to be Bernoulli distributed with corresponding state-dependent probabilities $\Pr(y_\tau | s_\tau)$ additionally depending on the current covariate values (cf. Eq. (1)), while the transition density of the state process $p(s_\tau | s_{\tau-1})$ is normally distributed as determined by the conditional distribution of the OU process:

$$S_\tau | S_{\tau-1} = s \sim \mathcal{N}\left(e^{-\theta\Delta_\tau} s, \frac{\sigma^2}{2\theta} (1 - e^{-2\theta\Delta_\tau})\right), \tag{4}$$

where $\Delta_\tau = t_\tau - t_{\tau-1}$ denotes the time difference between consecutive observations. Note that some of the observations actually occur at the exact same time according to the game clock, namely if fouled players are assigned two or three successive free throws. In such cases, the conditional distribution in Eq. (4) is given by

$S_\tau | S_{\tau-1} = s \sim \mathcal{N}(s, 0)$, for $\Delta_\tau = 0$, such that effectively, we assume a player's underlying form to be the same for two or three successive attempts.

Due to the T integrals in Eq. (3), the likelihood calculation is intractable. To render its evaluation feasible, we approximate the multiple integral by finely discretising the continuous-valued state space as first suggested by Kitagawa (1987). The discretisation of the state space can effectively be seen as a reframing of the model as a continuous-time hidden Markov model (HMM) with a large but finite number of states, enabling us to apply the corresponding efficient machinery. In particular, we use the forward algorithm to calculate the likelihood, defining the possible range of state values as $[-2, 2]$, which we divide into 100 intervals. We chose the interval $[-2, 2]$ such that the essential range of possible values of the state process is covered. The latter can be examined by looking at the (estimated) stationary distribution of the OU process, which we also monitored while fitting the model. For the OU process, a possible choice of the interval range would be $[-3\sigma/\sqrt{2\theta}, 3\sigma/\sqrt{2\theta}]$ (corresponding to three times the standard deviation in either direction as suggested by Fridman and Harris 1998). For further details on the approximation of the likelihood via state discretisation, see Ötting et al. (2020) for discrete-time and Mews et al. (2020) for continuous-time SSMs.

As we are interested in within-game variations of a player's form, we assume single throwing sequences of players to be mutually independent, conditional on the model parameters, such that the likelihood over all games and players is simply calculated as the product of the individual likelihoods. The model parameters, i.e. the reversion parameter and the diffusion coefficient of the OU process as well as the regression coefficients, are then estimated by numerically maximising the (approximate) joint likelihood. The resulting parameter estimators are unbiased and consistent — corresponding simulation experiments are shown in Mews et al. (2020).

4 Results

According to the cross-validated likelihood scores, the continuous-time model formulation including a potential hot hand effect is preferred over the benchmark model without any underlying state process. On average, the difference in the out-of-sample log-likelihood between the models is 3.8 in favour of the continuous-time SSM, with the latter scoring higher likelihoods in 17 out of 20 cross-validation sets. The parameter estimates of the OU process, which represents the underlying form of a player, as well as the estimated regression coefficients are shown in Table 3. In particular, the estimate for the reversion parameter θ of the OU process is fairly small, thus indicating serial correlation in the state process over time. To be more specific, the autocorrelation function (ACF) of the OU process is proportional to $\exp(-\theta\Delta)$, where Δ denotes the time interval, based on which we can deduce that $\exp(-0.042 \cdot 71) \approx 0.05$. Consequently, the ACF decreases by 95% over a time interval of 71 minutes, indicating that the hot hand phenomenon occurs over a whole game. Given that an NBA game lasts 48 minutes, our results indicate that the correlation between a player's form at the beginning and end of a game is still around 0.13 (cf. Fig. 2).

Table 3 Parameter estimates with 95% confidence intervals

Parameter	Estimate	95% CI	Odds ratios
θ (Reversion parameter)	0.042	[0.016; 0.109]	–
σ (Diffusion)	0.101	[0.055; 0.185]	–
β_1 (Home)	0.023	[– 0.009; 0.055]	1.023
β_2 (Scorediff)	0.030	[0.011; 0.048]	1.030
β_3 (Last30)	0.003	[– 0.051; 0.058]	1.003
β_4 (ft2)	0.223	[0.192; 0.254]	1.250
β_5 (ft3)	0.421	[0.279; 0.563]	1.523




Fig. 2 Continuous-time autocorrelation function of the OU process based on the estimated reversion parameter provided in Table 3. The shaded area corresponds to the 95% CI

Based on the estimated parameters of the OU process, the corresponding stationary distribution is estimated as $\mathcal{N}(0, 0.348^2)$, indicating a rather small range of the state process. The latter becomes apparent also when simulating state trajectories based on the parameter estimates of the OU process (cf. Fig. 3). Still, the associated success probabilities, given that all covariate values are fixed to zero, vary considerably during the time of an NBA game (cf. right y-axis of Fig. 3). More specifically, the state process fluctuates between $- 0.68$ and 0.68 around 95% of the time, implying that for a player with median throwing success the corresponding success probabilities lie between 0.69 and 0.9, conditional on the covariates being equal to zero. In this specific case, a player’s underlying form thus changes his throwing success by up to $- 12$ and $+ 8$ percentage points (the change in percentage points is not symmetrical around the average throwing success as we model the predictor on the logit scale, cf. Eq. (1)). While the state process and hence the resulting success probabilities slowly fluctuate around the average throwing success (given the covariates), the simulated state trajectories reflect the temporal persistence of the players’ underlying form. Thus, our results

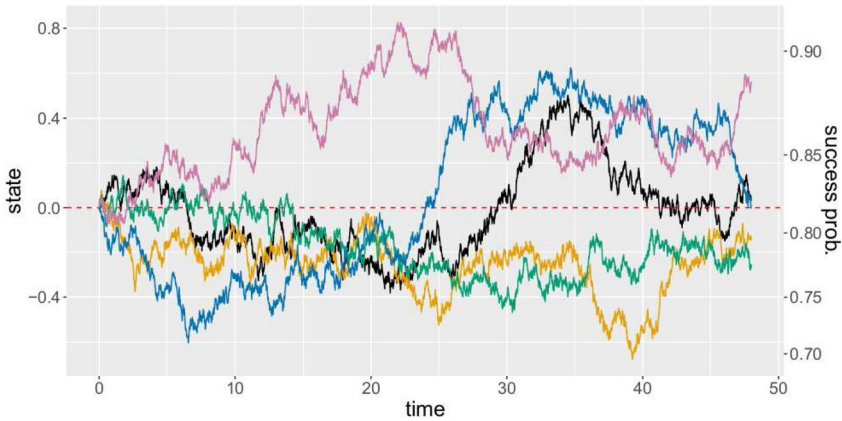


Fig. 3 Simulation of possible state trajectories for the length of an NBA game based on the estimated parameters of the OU process. The red-dashed line indicates the intercept (here: the median throwing success over all players), around which the processes fluctuate. The right y-axis shows the success probabilities resulting from the current state (left y-axis), given that the explanatory variables equal 0. The graphs were obtained by application of the Euler-Maruyama scheme with initial value 0 and step length 0.01

suggest that players can temporarily enter a state in which their success probability is considerably higher than their average performance, which provides evidence for a hot hand effect.

Regarding the estimated regression coefficients, the player-specific intercepts $\hat{\beta}_{0,p}$ range from -0.311 to 2.192 (on the logit scale), reflecting the heterogeneity in players' throwing success. The estimates for β_1 to β_5 are displayed in Table 3 together with their 95% confidence intervals and the resulting odds ratios. The chance in making a free throw is slightly increased if the game is played at home (β_1) or if a free throw occurs in the last 30 seconds of a quarter (β_3), but both corresponding confidence intervals include the zero. In contrast, the confidence interval for the score difference (β_2) does not include the zero and its effect is positive, indicating that the odds to make a free throw increase by 3% per score the team is leading. The position of the throw, i.e. whether it is the first, second (β_4), or third (β_5) attempt in a row, has the largest effect of all covariates considered: compared to the first free throw, the odds of a hit increase considerably if it is the second or, in particular, the third attempt, i.e. an increase in odds of 25% and 52.3%, respectively, which was already indicated by the descriptive analysis presented in Sect. 2. However, this strong effect on the success probabilities is probably caused by the fact that three successive free throws are only awarded if a player is fouled while shooting a three-point field goal, which, in turn, is more often attempted by players who regularly perform well at free throws.

To further investigate how the hot hand may evolve during a game, we compute the most likely state sequences, corresponding to the underlying form of a player. Specifically, we seek

$$(s_1^*, \dots, s_T^*) = \operatorname{argmax}_{s_1, \dots, s_T} \Pr(s_1, \dots, s_T | y_1, \dots, y_T),$$

where s_1^*, \dots, s_T^* denotes the most likely state sequence given the observations. As we transferred our continuous-time SSM to an HMM framework by finely discretising the state space (cf. Sect. 3.2), we can use the Viterbi algorithm to calculate such sequences at low computational cost (Zucchini et al. 2016). Figure 4 shows the most likely states underlying the throwing sequences presented in Table 2. While the decoded state processes fluctuate around zero (i.e. the player’s average throwing success), the state values vary slightly over the time of an NBA game. Over all players and games, the decoded states range from -0.42 (in a game of Josh Smith) to 0.46 (Andre Drummond), leading to changes in the success probabilities of about -10 and $+11.4$ percentage points, respectively, given the corresponding current game characteristics.

The decoded state sequences in Fig. 4 further allow to illustrate the advantages and the main idea of our continuous-time modelling approach. For example, consider the throwing sequence in the second match shown, where LeBron James only made a single free throw of his first four attempts. The decoded state at throw number 3 is -0.092 (cf. Fig. 4) and the time passed between throw number 3 and 4 is 1.65 minutes (cf. Table 2). Thus, the value of the state process at throw number 4 is drawn from a normal distribution, given the decoded state of the previous attempt, with mean $e^{-0.042 \cdot 1.65}(-0.092) = -0.086$ and variance $\frac{0.101^2}{2 \cdot 0.042} (1 - e^{-2 \cdot 0.042 \cdot 1.65}) = 0.016$ (cf. Eq. (4)). Accordingly, the value of the state process for throw number 5 is drawn

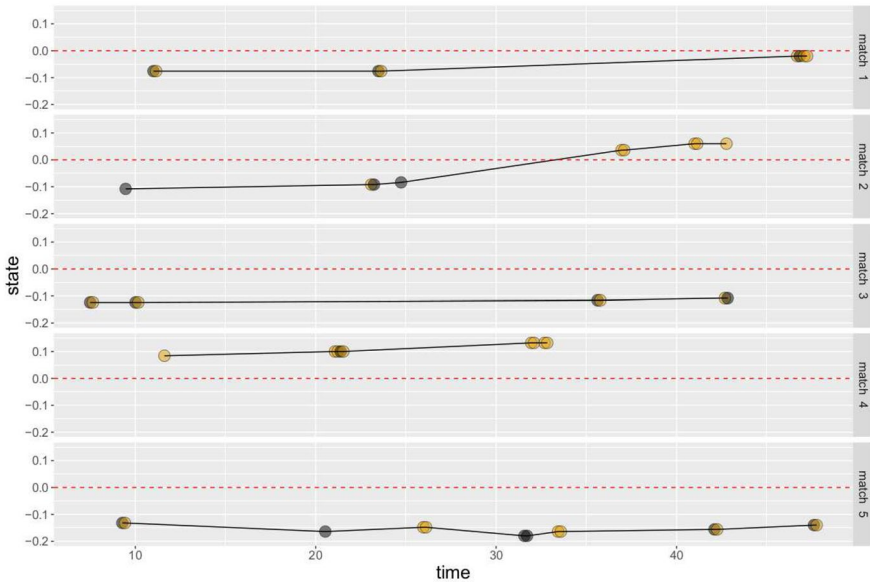


Fig. 4 Decoded states underlying the throwing sequences of LeBron James shown in Table 2. Successful free throws are shown in yellow, missed shots in black

from a normal distribution with mean -0.050 and variance 0.078 , conditional on the decoded state of -0.084 at throw number 4 and a relatively long time interval of 12.22 minutes. As highlighted by these example calculations, the conditional distribution of the state process takes into account the interval length between consecutive attempts: the more time elapses, the higher the variance in the state process and hence, the less likely is a player to retain his form, with a tendency to return to his average performance.

5 Discussion

In our analysis of the hot hand, we used SSMs formulated in continuous time to model throwing success in basketball. Focusing on free throws taken in the NBA, our results provide evidence for a hot hand effect as the underlying state process exhibits some persistence over time. In particular, the model including a hot hand effect is preferred over the benchmark model without any underlying state process based on cross-validated likelihood scores. Although we provide evidence for the existence of a hot hand, the magnitude of the hot hand effect is rather small as the underlying success probabilities are elevated by only a few percentage points (cf. Figs. 3 and 4).

A clear advantage of the continuous-time modelling framework considered in this contribution is that it enables us to explicitly account for the unevenly spaced time points of free throws in basketball. As argued in Sect. 1, it seems intuitively plausible that players are more likely to retain their form the less time elapses between two consecutive attempts. Our modelling approach thus naturally mirrors these considerations by using the OU process to model the evolution of a player's latent form over time, while accommodating varying interval lengths between observations. In contrast, SSMs operating in discrete time require that time intervals are of equal length and therefore, are not (directly) applicable to irregularly spaced observations in time. Although it is possible to temporally aggregate irregularly spaced data to generate regular (i.e. equidistant) time intervals between observations and thus, to render discrete-time models applicable, this approach comes with several drawbacks. For instance, the temporal aggregation of the data introduces subjectivity regarding the choice of the discrete-time modelling resolution and discards information on the exact observation times. Moreover, it has been shown that applying discrete-time models to data with irregularly spaced events can lead to biased estimates (see, e.g. Delsing et al. 2005; Barbour et al. 2013; de Haan-Rietdijk et al. 2017). Therefore, using continuous-time models to analyse irregularly spaced observations in time avoids the pitfalls mentioned above and, in addition, is conceptually appealing as the model's interpretation does not depend on the time resolution of the data at hand.

A minor drawback of the analysis arises from the fact that there is no universally accepted statistical modelling approach for the hot hand. While we evaluate the magnitude of the hot hand effect based on the ACF (see Sect. 4), it is not clear how strong the autocorrelation in the state process should be to reflect a distinct hot hand pattern, rendering definite conclusions almost impossible. Based on the ACF,

however, the magnitude of the estimated hot hand effect can easily be compared to other studies which consider (discrete- or continuous-time) SSMs to investigate autocorrelation in players' underlying success probabilities. A further limitation of our approach is related to the general problem of disentangling the hot hand effect from possible model misspecifications. Although we include several covariates in our model, which were also used in previous studies on free throw shooting in basketball, we cannot rule out an omitted-variable bias. In addition, possible heterogeneity across players regarding covariate effects has not been considered in our analysis. Consequently, if player-specific effects or covariates which affect the outcome of a free throw are missing in our model, this may lead to biased estimates of the OU process and hence, to biased results on the hot hand effect.

In general, the modelling framework considered provides great flexibility with regard to distributional assumptions. In particular, the response variable is not restricted to be Bernoulli distributed (or Gaussian, as is often the case when making inference on continuous-time SSMs), such that other types of response variables used in hot hand analyses (e.g. Poisson) can be implemented by changing just a few lines of code. Our continuous-time SSM can thus easily be applied to other sports, and the measure for success does not have to be binary as considered here. For readers interested in adopting our code to fit their own hot hand model, the electronic supplement of this article provides the data and code used for the analysis.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10182-021-00410-y>.

Acknowledgements We thank Roland Langrock, Christian Deutscher, and Houda Yaqine for stimulating discussions and helpful comments. We are also grateful to the Associate Editor, Dominik Liebl, and two anonymous reviewers for their insightful and very constructive feedback that helped to improve this article. This research was supported by Deutsche Forschungsgemeinschaft (Grant 431536450).

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Albert, J.: A statistical analysis of hitting streaks in baseball: comment. *J. Am. Stat. Assoc.* **88**(424), 1184–1188 (1993)

- Auger-Méthé, M., Newman, K., Cole, D., Empacher, F., Gryba, R., King, A.A., Leos-Barajas, V., Fleming, J.M., Nielsen, A., Petris, G., Thomas, L.: A guide to state-space modeling of ecological time series. *Ecol. Monogr.* (in press)
- Bar-Eli, M., Avugos, S., Raab, M.: Twenty years of hot hand research: review and critique. *Psychol. Sport Exer.* **7**(6), 525–553 (2006)
- Barbour, A.B., Ponciano, J.M., Lorenzen, K.: Apparent survival estimation from continuous mark-recapture/resighting data. *Method Ecol. Evol.* **4**(9), 846–853 (2013)
- Celeux, G., Durand, J.B.: Selecting hidden Markov model state number with cross-validated likelihood. *Comput. Stat.* **23**(4), 541–564 (2008)
- Chang, J.C.: Predictive Bayesian selection of multistep Markov chains, applied to the detection of the hot hand and other statistical dependencies in free throws. *R. Soc. Open Sci.* **6**(3), 841 (2019)
- Delsing, M.J.M.H., Oud, J.H.L., De Bruyn, E.E.J.: Assessment of bidirectional influences between family relationships and adolescent problem behavior: discrete vs. continuous time analysis. *Eur. J. Psychol. Assess.* **21**(4), 226–231 (2005)
- Dorsey-Palmateer, R., Smith, G.: Bowlers' hot hands. *Am. Stat.* **58**(1), 38–45 (2004)
- Fridman, M., Harris, L.: A maximum likelihood approach for non-Gaussian stochastic volatility models. *J. Bus. Econ. Stat.* **16**(3), 284–291 (1998)
- Gilovich, T., Vallone, R., Tversky, A.: The hot hand in basketball: on the misperception of random sequences. *Cognit. Psychol.* **17**(3), 295–314 (1985)
- Green, B., Zwiebel, J.: The hot-hand fallacy: cognitive mistakes or equilibrium adjustments? Evidence from Major League Baseball. *Manage. Sci.* **64**(11), 4967–5460 (2017)
- de Haan-Rietdijk, S., Voelkle, M.C., Keijsers, L., Hamaker, E.L.: Discrete- vs. continuous-time modeling of unequally spaced experience sampling method data. *Front. Psychol.* **8**, 1849 (2017)
- Hurvich, C.M., Simonoff, J.S., Tsai, C.L.: Smoothing parameter selection in nonparametric regression using an improved Akaike information criterion. *J. R. Stat. Soc.: Ser. B* **60**(2), 271–293 (1998)
- Jagannathan, R., Malakhov, A., Novikov, D.: Do hot hands exist among hedge fund managers? An empirical evaluation. *J. Financ.* **65**(1), 217–255 (2010)
- Kahneman, D.: *Thinking, Fast and Slow*. Farrar Straus and Giroux, New York (2011)
- Kitagawa, G.: Non-Gaussian state-space modeling of nonstationary time series. *J. Am. Stat. Assoc.* **82**(400), 1032–1041 (1987)
- Liu, L., Wang, Y., Sinatra, R., Giles, C.L., Song, C., Wang, D.: Hot streaks in artistic, cultural, and scientific careers. *Nature* **559**(7714), 396–399 (2018)
- Mews, S., Langrock, R., Ötting, M., Yaşın, H., Reinecke, J.: Maximum approximate likelihood estimation of general continuous-time state-space models. *arXiv preprint arXiv:201014883* (2020)
- Miller, J.B., Sanjurjo, A.: Surprised by the hot hand fallacy? A truth in the law of small numbers. *Econometrica* **86**(6), 2019–2047 (2018)
- Morgulev, E., Azar, O.H., Bar-Eli, M.: Searching for momentum in NBA triplets of free throws. *J. Sports Sci.* **38**(4), 390–398 (2020)
- Müller, S., Scealy, J.L., Welsh, A.H.: Model selection in linear mixed models. *Stat. Sci.* **28**(2), 135–167 (2013)
- Ötting, M., Langrock, R., Deutscher, C., Leos-Barajas, V.: The hot hand in professional darts. *J. R. Stat. Soc. (Ser. A)* **183**(2), 565–580 (2020)
- Raab, M., Gula, B., Gigerenzer, G.: The hot hand exists in volleyball and is used for allocation decisions. *J. Exper. Psychol.: Appl.* **18**(1), 81–94 (2012)
- Thaler, R.H., Sunstein, C.R.: *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Penguin, London (2009)
- Toma, M.: Missed shots at the free-throw line: analyzing the determinants of choking under pressure. *J. Sports Econ.* **18**(6), 539–559 (2017)
- Wetzels, R., Tutschkow, D., Dolan, C., van der Sluis, S., Dutilh, G., Wagenmakers, E.J.: A Bayesian test for the hot hand phenomenon. *J. Math. Psychol.* **72**, 200–209 (2016)
- Zucchini, W., MacDonald, I.L., Langrock, R.: *Hidden Markov Models for Time Series: An Introduction Using R*. Chapman & Hall/CRC, Boca Raton (2016)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.