

Research Article

Jim E. Griffin*, Laurențiu C. Hinoveanu and James G. Hopker

Bayesian modelling of elite sporting performance with large databases

<https://doi.org/10.1515/jqas-2021-0112>

Received December 21, 2021; accepted November 18, 2022;
published online December 8, 2022

Abstract: The availability of large databases of athletic performances offers the opportunity to understand age-related performance progression and to benchmark individual performance against the World's best. We build a flexible Bayesian model of individual performance progression whilst allowing for confounders, such as atmospheric conditions, and can be fitted using Markov chain Monte Carlo. We show how the model can be used to understand performance progression and the age of peak performance in both individuals and the population. We apply the model to both women and men in 100 m sprinting and weightlifting. In both disciplines, we find that age-related performance is skewed, that the average population performance trajectories of women and men are quite different, and that age of peak performance is substantially different between women and men. We also find that there is substantial variability in individual performance trajectories and the age of peak performance.

Keywords: Bayesian variable selection; longitudinal models; Markov chain Monte Carlo; performance monitoring; skew t distribution.

1 Introduction

The availability of large databases of elite athletic performances makes it possible to model age-related performance progression and to benchmark this progression against the World's best athletes. This information has many applications such as systematically identifying those

athletes with the potential for future success in order to prioritise funding and resource allocation (Allen, Vandenberg, and Hopkins 2014), setting performance goals by coaches and/or athletes, or guiding training programmes. These data have repeated measurements, usually taken at irregular intervals (due to the timing of different events/competitions in the international calendar) and with confounders of performance such as wind speed in sprinting, seasonality effects, and competition level. As discussed by Brander, Egan, and Yeung (2014), selection effects due to ability differences across athletes provide a further complication. For example, only exceptional younger athletes will compete at the highest levels of sport against a wider range of abilities levels in older more established counterparts. This selection bias may result in an underestimation of performance gains that an athlete may realize over their early years within senior level competition. As such, cross-sectional analytical methods including least squares or simple averages at different ages can lead to misleading results.

There are two main approaches to the statistical analysis of athletic performance over time. Firstly, a time series analysis of a small number of elite level performances can be used to understand population-level changes. For example, Stephenson and Tawn (2013) analysed world best times per year using extreme value theory to understand changing performance for men and women. Similarly, Gao, Li, and Wang (2020) analysed the finishing times in World Championship and Olympic Games semi-finals and finals using time series model to understand the effects of technology in swimming, and Kovalchik and Stefani (2013) analysed medal-winning performances at Olympic games using a regression model to understand whether the difference in performance between men and women was changing over time. Secondly, longitudinal models have been fitted to a database of performances measured over time. For example, in cricket, Boys and Philipson (2019) modelled batsmen's scores allowing for the effect of age, home advantage, different innings and quality of opposition, whereas Stevenson and Brewer (2021) modelled changing individual performance with a Gaussian process. Egidi and Gabry (2018) build a hierarchical autoregressive

*Corresponding author: Jim E. Griffin, Department of Statistical Science, University College London, London, UK,
E-mail: j.griffin@ucl.ac.uk

Laurențiu C. Hinoveanu and James G. Hopker, School of Sport and Exercise Sciences, University of Kent, Canterbury, UK,
E-mail: lh460@kent.ac.uk (L.C. Hinoveanu), J.G.Hopker@kent.ac.uk (J. Hopker)

model for player ratings in football allowing for individual, team and position effects. In ice hockey, Bradbury (2009) used a quadratic age effect and allowed for other covariates using random effects model. Wimmer et al. (2011) built a factor model for decathlon performances (across the ten component events) and used a spline functions to nonparametrically model the effect of age and the month of the year (this allows for the changes in performance levels over the athletics season).

The longitudinal approach allows for inference at the individual level (unlike the time series approach) and addresses the selection effect by accounting for different performance levels across athletes. This approach was first proposed for the analysis of sports data by Berry, Reese, and Larkey (1999) who fit Bayesian hierarchical models across different individual athletes and competition seasons. The performance for an individual in a particular season is decomposed into an individual effect, a season effect and an individual effect of age (called the individual ageing function). The use of individual ageing functions (rather than a single common ageing function for all individuals) is important since it allows for variability in the effect of ageing which is observed in data (for example, some athletes peaking earlier than others) and is a key part of longitudinal models. Berry, Reese, and Larkey (1999) model the individual ageing function as the sum of a flexible, nonparametric average ageing function and a simple (parametric) function. However, Albert (1999) suggests an alternative approach using a parametric average ageing function but a more flexible model for the departures for each individual ageing function. A range of individual and average ageing functions have been proposed for longitudinal monitoring including quadratic (Boys and Philipson 2019; Bradbury 2009), cubic (Brander, Egan, and Yeung 2014), a two parameter non-linear function (Berthelot et al., 2012; Strand, Nelson, and Grunwald 2018), and Gaussian processes (Stevenson and Brewer 2021).

In this paper, we will use a longitudinal model in a Bayesian framework (for a review of Bayesian methods for analysis of sports data the reader is referred to the work of Santos-Fernandez, Wu, and Mengersen 2019) and concentrate on centimetre-gram-second sports where performance is either measured in centimetres (e.g. distance thrown), grams (e.g. weight lifted) or seconds (e.g. time to run a specific distance). Our approach is in the spirit of Albert (1999) and uses a parametric form for the average ageing function and a flexible model for the individual ageing function. We make several contributions. Firstly, our model is more flexible than previous models in

the literature. We assume a polynomial for the average ageing effect, a spline model for the individual ageing function and a skew- t distribution for the observation errors. Secondly, we extend the Adaptively Scaled Independence proposal (Griffin, Latuszynski, and Steel 2021) in a Metropolis-Hastings sampler to allow variable selection in linear regression models with skew t distributed errors. Thirdly, we apply this model to large databases of athlete performances in both 100 m sprinting and weightlifting, and show that Markov chain Monte Carlo (MCMC) methods can be used for inference in this context. Fourthly, we discuss how the Bayesian approach can be used to draw inferences about aspects of athlete's performance such as age of peak performance at both the population and individual level.

The paper is organized as follows. Section 2 describes the data used in the study. Section 3 introduces the model to be used for the analysis and aspects of the computational methods required. Section 4 presents results for both males and females in 100 m sprinting and weightlifting. Section 5 discusses the approach and possible future work.

2 Data

Our analysis focuses on competition results from 100 m sprinting and weightlifting events obtained from the World Athletics and International Weightlifting Federation websites, respectively following Institutional ethical approval (Prop_72_2017_18). The 100 m sprint data contains results from both male and female sprinters who had at least 5 competition results between 8th January 2001 and 28th August 2021. The database contained 2834 male athletes who have a personal best below 10.5 s and 1297 female athletes who have a personal best below 11.6 s. The male data set had 95,376 observed performances, with the female data set having 48,999 observations. The ages for males athletes ranged from 12 to 45 years, whereas females ranged from 12 and 42 years.

The weightlifting data comprised competition results from 1609 male and 1212 female lifters who had at least 5 competition results between 24th April 1998 and 24th May 2021. This led to 14,557 observations in the male data set and 11,690 observations in the female data set. Weightlifters' ages ranged from 11 to 43 years old in the male data set, and between 12 and 41 years in the female data set. The two different lifts that comprise weightlifting competitions (snatch and clean and jerk) were combined into a total lift score. The scores were transformed to Sinclair totals to allow for the differences in the levels of lift scores across weight categories. This

allowed us to consider results across weight categories without adjusting for weight category in the model. The Sinclair total is calculated by multiplying the total lift score by a coefficient derived from the lifters body weight, the current World record holder's body weight in the various weight categories, and a coefficient from World record performance from the previous years.

3 Model

We assume that there is a database of measured performance for M athletes. The i th athlete has n_i performances and we define $y_{i,j}$ to be the j th performance of the i th athlete which occurs at age $t_{i,j}$ with a p -dimensional vector of additional regressors $x_{i,j}$ which are functions of measured confounders or other covariates. This allows the model to include the effects of confounders of performance such as wind speed in sprint running events. We model $y_{i,j}$ by

$$\begin{aligned} y_{i,j} &= h_i(t_{i,j}) + x_{i,j} \zeta + \epsilon_{i,j} \\ &= g(t_{i,j}) + f_i(t_{i,j}) + x_{i,j} \zeta + \sigma_i \epsilon_{i,j} \end{aligned}$$

where $h_i(t)$ is the i th athlete's individual performance trajectory (or ageing function) as a function of age, ζ are population-level regression coefficients shared by all athletes and events and $\epsilon_{i,j}$ are observation errors which are assumed to follow a standard skew t distribution (Azzalini and Capitanio 2003). The *individual performance trajectory* $h_i(t)$ is decomposed into two parts: $g(t)^t$, which represents the *population performance trajectory* (or ageing function) and $f_i(t)$ represents the difference between the i th individual performance trajectory and population performance trajectory, which we call the *excess performance trajectory*.

Our modelling approach assumes that $f_i(t)$ can be less smooth than $g(t)$ for some athletes since it represents the population average whereas individual performance trajectories can change abruptly. Therefore, we use a parametric model for the population performance trajectories and a non-linear regression model for the excess performance trajectory. This leads to the same modelling approach as the suggestion of Albert (1999). In our analyses, we assume that $g(t)$ follows a fourth order polynomial $g(t) = \sum_{k=0}^4 \eta_k (t_{i,j} - \bar{t})^k$ where \bar{t} is the mean age of a performance across all athletes. We flexibly model $f_i(t)$ by piecewise linear splines $f_i(t) = \theta_i + s_{i,j} \beta_i$ where $s_{i,j} = ((t_{i,j} - m_1)^+, \dots, (t_{i,j} - m_K)^+)$ is K -dimensional vector, θ_i is the excess performance before time m_1 for the i th athlete and controls their overall level of excess performance, β_i is a K -dimensional vector of individual-specific

regression coefficients, $(x)^+ = \max\{0, x\}$ and m_1, \dots, m_K are pre-specified knots (which are chosen to be equally spaced between the smallest and largest age observed with a performance for any athlete). Usually K is chosen to be large (100 in our examples) to allow $f_i(t)$ to have a wide-range of possible shape. Since K is large, there is potential for overfitting and so β_i must be either regularized (for example, using the Lasso) or variable selection must be performed. We use Bayesian variable selection and refer the interested reader to Denison et al. (2002) for further explanation of Bayesian methods and their application to the spline model. Other non-linear regression approaches such as Gaussian process priors (Rasmussen and Williams 2006) or B-splines (Lang and Brezger 2004) could also be used.

The choice of a skew t error distribution for the observational errors allows for both skewness and kurtosis which occur due to unusually poor performances (for example, longer times in track events or shorter throws or jumps in field events). Following Azzalini and Capitanio (2003), a random variable Y to has a standard skew t -distribution with parameters α and ν if $Y = \frac{X}{\sqrt{W}}$ where X follows a skew normal distribution with parameter α (Azzalini 1985) and $W \sim \text{Gamma}(\nu/2, \nu/2)$. This construction reduces to the Student t distribution if $\alpha = 0$. The mean and variance of Y are $\frac{\alpha\sqrt{\nu}\Gamma[0.5(\nu-1)]}{\sqrt{\alpha^2+1}\sqrt{\pi}\Gamma(0.5\nu)}$ and $\frac{\nu}{\nu-2} - \frac{\alpha^2}{\alpha^2+1} \frac{\nu}{\pi} \left\{ \frac{\Gamma[0.5(\nu-1)]}{\Gamma(0.5\nu)} \right\}^2$, respectively. The parameters have simple interpretations: α controls skewness and ν controls kurtosis. In our model, we assume that $\epsilon_{i,j}$ has shared skewness α and kurtosis ν .

We use a Bayesian analysis and assume the following priors for our model (the parametrization of all distributions is given in the Appendix A). The regression coefficients in the population-level performance trajectory are given an uninformative prior, $\pi(\eta) \propto 1$. The regression coefficients of the regressors $x_{i,j}$ are given independent horseshoe priors (Carvalho, Polson, and Scott 2010),

$$\begin{aligned} \zeta_l \mid \lambda_l^2, \tau^2, a, b &\stackrel{\text{i.i.d.}}{\sim} \text{N}\left(0, \frac{b}{a} \lambda_l^2 \tau^2\right) \text{ and} \\ \lambda_l &\stackrel{\text{i.i.d.}}{\sim} \text{C}^+ \quad \text{for } l = 1, \dots, p, \quad \tau \sim \text{C}^+ \end{aligned}$$

where C^+ is the standard half-Cauchy distribution which has density

$$p(x) = \frac{2}{\pi} (1+x)^{-1}, \quad x > 0.$$

This choices allows us to perform regularisation on ζ and so shrink out regressors which are not related to the response. The parameters of the excess performance

trajectory $f_i(t)$ follow a standard Bayesian variable selection structure by assuming that only some β_i 's are non-zero. The j th regression coefficient $\beta_{i,j}$ is non-zero (and the j th spline included in the model) if $\gamma_{i,j} = 1$ and is given the hierarchical prior

$$p(\gamma_{i,j} = 1) = \pi_i \quad \text{for } l = 1, \dots, K, \quad \pi_i \sim \text{Be}(1, (K-3)/3).$$

Let $p_i^\gamma = \sum_{j=1}^K \gamma_{i,j}$, which is the number of included splines, then prior mean of p_i^γ is $E[p_i^\gamma] = 3$ (Ley and Steel 2009) and so encourages only a small number of splines to be used for each individual. We define β_i^γ to be the regression coefficients for which $\gamma_{i,j} = 1$ and use the priors

$$\theta_i \sim N(0, \psi \sigma_i^2), \quad \beta_i^\gamma | \sigma_i^2, \gamma_i \sim N(0, \sigma_i^2 V_i^\gamma),$$

where $\gamma_i = (\gamma_{i,1}, \dots, \gamma_{i,K})$ and V_i^γ is a $(p_i^\gamma \times p_i^\gamma)$ -dimensional diagonal matrix with all diagonal terms equal to g . A hierarchical model is assumed for the individual-level observational error scale

$$\sigma_i^2 \sim \text{IG}(a, b) \quad \text{for } i = 1, \dots, M, \quad p(a) \propto 1, \quad p(b) \propto 1.$$

The remaining population-level parameters are given the following prior distributions

$$g \sim C^+, \quad \psi \sim \text{IG}(1, 1), \quad \alpha \sim N(0, 100), \quad \nu \sim \text{Ga}(2, 0.1)$$

The choice of α and ν follow the choices advocated by Frühwirth-Schnatter and Pyne (2010).

Bayesian inference can be implemented using Markov chain Monte Carlo (MCMC) with a Gibbs sampler. The sampling uses several latent variable representations to give simpler updates of the parameters. Firstly, the standard skew- t distribution can be represented as

$$\epsilon_{i,j} = \epsilon_{i,j}^* + \frac{\alpha}{\sqrt{1 + \alpha^2}} z_{i,j}$$

where $\epsilon_{i,j}^* \stackrel{\text{i.i.d.}}{\sim} N(0, \frac{1}{w_{i,j}})$, $z_{i,j} | w_{i,j}, \sigma_i^2 \stackrel{\text{i.i.d.}}{\sim} \text{TN}_{[0, \infty)}(0, \frac{1}{w_{i,j}})$, and $w_{i,j} | \nu \stackrel{\text{i.i.d.}}{\sim} \text{Ga}(\frac{\nu}{2}, \frac{\nu}{2})$, $j = 1, \dots, n_i$, $i = 1, \dots, M$. Secondly, the half-Cauchy prior on λ_l can be represented in terms of a mixture distribution on λ_l^2 (see Makalic and Schmidt 2016) giving

$$p(\lambda_l^2 | \nu_l^*) \sim \text{IG}\left(0.5, \frac{1}{\nu_l^*}\right),$$

$$p(\nu_l^*) \stackrel{\text{i.i.d.}}{\sim} \text{IG}(0.5, 1), \quad l = 1, \dots, p.$$

The same representation can be used for τ^2 ,

$$p(\tau^2 | \xi) \sim \text{IG}\left(0.5, \frac{1}{\xi}\right), \quad \pi(\xi) \sim \text{IG}(0.5, 1).$$

All steps for the Gibbs sampler are given in the Appendix A.

Bayesian variable selection is needed for spline model for each individual and we use the recently developed Adaptive Scaled Individual adaptation scheme (Griffin, Latuszynski, and Steel 2021) to update the full conditionals associated with each spline model. Griffin, Latuszynski, and Steel (2021) show that their approach can lead to faster mixing Markov chains than traditional approaches to MCMC in Bayesian variable selection in linear regression models with normal errors (see e.g. Dellaportas, Forster, and Ntzoufras 2000; Tadesse and Vannucci 2021). We extend their approach to linear regression models with skew t error distributions using a latent variable representation.

We define $d_{i,j} = (1, t_{i,j} - \bar{t}, (t_{i,j} - \bar{t})^2, (t_{i,j} - \bar{t})^3, (t_{i,j} - \bar{t})^4)$. The marginal likelihood of the observations for the i th athlete conditional on $\zeta, \eta, Z, W, a, b, g, \psi$ and γ_i is

$$m_i(\gamma_i) = g^{-p_i^\gamma/2} |\Delta_i|^{-1/2} \left(b + \frac{1}{2} \left(\sum_{j=1}^{n_i} w_{i,j} z_{i,j}^2 + \mathbf{Y}_i^{*T} \left(\mathbf{I}_{n_i} - \mathbf{S}_i^* \Delta_i^{-1} \mathbf{S}_i^{*T} \right) \mathbf{Y}_i^* \right) \right)^{-(a+n_i/2)}$$

where \mathbf{Y}_i^* is a n_i -dimensional vector whose j th entry is $\sqrt{w_{i,j}} \left[\sqrt{1 + \alpha^2} (y_{i,j} - x_{i,j} \zeta - d_{i,j} \eta) - \alpha z_{i,j} \right]$, \mathbf{S}_i^* is a $(n_i \times (p_i^\gamma + 1))$ -dimensional matrix whose j th row has entry $\sqrt{\frac{w_{i,j}}{1 + \alpha^2}}$ and subsequent entries $\sqrt{\frac{w_{i,j}}{1 + \alpha^2}} s_{i,j}$ for the j with $\gamma_{i,j} = 1$

$$\Delta_i = \mathbf{S}_i^{*T} \mathbf{S}_i^* + \text{diag} \left(\frac{1}{\psi}, \underbrace{\frac{1}{g}, \dots, \frac{1}{g}}_{p_i^\gamma} \right).$$

In our Gibbs sampler, the adaptive Metropolis–Hastings method of (Griffin, Latuszynski, and Steel 2021) can be used to update γ_i . The method uses a Metropolis–Hastings update and proposes a new state γ_i' from the proposal distribution

$$\gamma_{i,j}' | \gamma_{i,j} = 0 \sim \text{Bernoulli}(A_{i,j}),$$

$$\gamma_{i,j}' | \gamma_{i,j} = 1 \sim \text{Bernoulli}(1 - D_{i,j})$$

for $i = 1, \dots, K$. The tuning parameter $A_{i,j}$ is the probability of adding the j th variable to the model and the tuning parameter $D_{i,j}$ is the probability of deleting the j th variable from the model. The proposal value is accepted with probability

$$\min \left\{ 1, \frac{m_i(\gamma'_i)}{m_i(\gamma_i)} \right\} \times \frac{\Gamma(1 + p'_i) \Gamma((K - 3)/3 + K - p'_i) q(\gamma'_i, \gamma_i)}{\Gamma(1 + p_i) \Gamma((K - 3)/3 + K - p_i) q(\gamma_i, \gamma'_i)}$$

where

$$q(\gamma_i, \gamma'_i) = \prod_{j=1}^K \left[A_j^{(1-\gamma_{i,j})\gamma'_{i,j}} (1 - A_j)^{(1-\gamma_{i,j})(1-\gamma'_{i,j})} \times D_j^{\gamma_{i,j}(1-\gamma'_{i,j})} (1 - D_j)^{\gamma_{i,j}\gamma'_{i,j}} \right]$$

The values of these tuning parameters are adjusted during the MCMC run using the following rule

$$A_{i,j} = \kappa_i \min \left\{ 1, \frac{\hat{\pi}_{i,j}}{1 - \hat{\pi}_{i,j}} \right\}, \quad D_j = \kappa_i \min \left\{ 1, \frac{1 - \hat{\pi}_{i,j}}{\hat{\pi}_{i,j}} \right\}$$

where κ_i is a further individual-specific tuning parameter and $\hat{\pi}_{i,j}$ is an estimate of the posterior probability that $\gamma_{i,j} = 1$. A Rao–Blackwellised estimate of $\pi_{i,j}$ is used in the algorithm and κ_i is tuned using a Robbins–Munro update

$$\text{logit}_\epsilon \kappa^{(t+1)} = \text{logit}_\epsilon \kappa^{(t)} + t^{-\lambda} (a_t - 0.234),$$

where $\text{logit}_\epsilon(x) = \log(x - \epsilon) - \log(1 - x - \epsilon)$ for $0 \leq \epsilon \leq 1/2, 1/2 < \lambda \leq 1$, a_t is the Metropolis–Hastings acceptance probability at the t th iteration. Griffin, Latuszynski, and Steel (2021) shows that this scheme is optimal with $\eta = 1$ if $\gamma_{i,1}, \dots, \gamma_{i,p}$ are independent under the posterior distribution and that it leads to good performance under different levels of dependence.

4 Results and discussion

We have analyzed the two sets of data (100 m sprinting and weightlifting) using our model separately for males and females. We include the month of competition as a covariate using January as a baseline in all analyses. Wind is a confounder in sprint events and we include both linear and quadratic effects of wind in the 100 m sprint analyses. The MCMC algorithm was run for 11,000 iterations with 1000 used as a burn-in. For the MCMC results we have saved every 10th iteration and we have also started 10 chains. Altogether this leads to 10,000 MCMC recorded samples.

4.1 Ageing and athletic performance

Figures 1 and 2 show inference about the population performance trajectories for 100 m sprinting and weightlifting

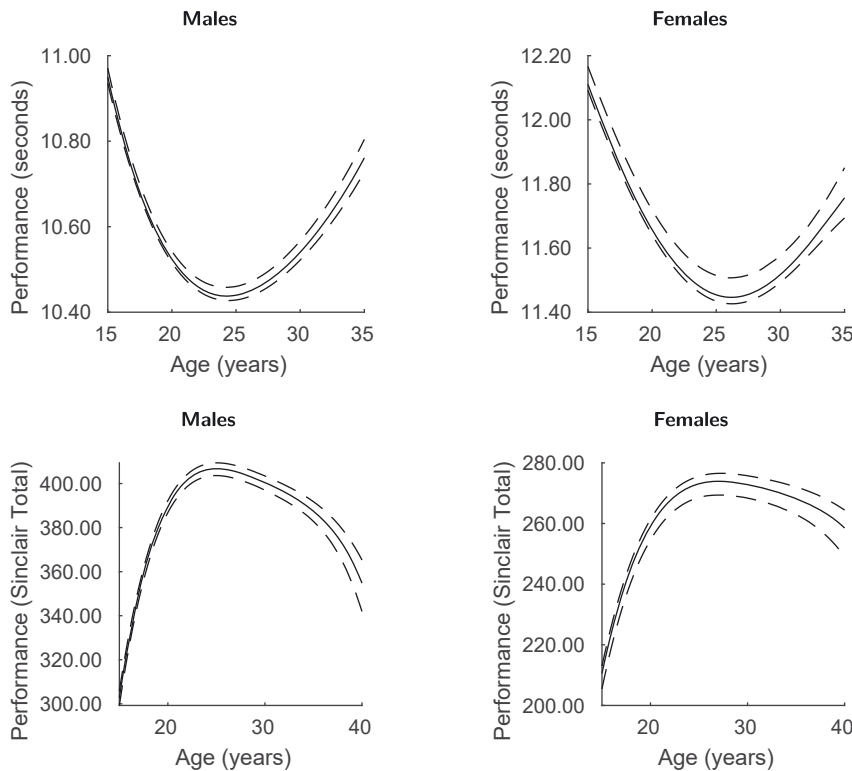


Figure 1: The population performance trajectory for 100 m sprint. The solid line is the posterior median and the dashed lines are 95% credible intervals.

Figure 2: The population performance trajectory for weightlifting. The solid line is the posterior median and the dashed lines are 95% credible intervals.

respectively. As shown in Figure 1, the population performance trajectory suggests that age of peak performance occurs about two years earlier for males than females in the 100 m. The parabolic relationship between age and performance is not symmetric with a greater rate of performance improvement (*i.e.* faster sprinting times) before peak age, with a reduced rate of performance decline subsequently (*i.e.* slower sprinting times).

This is true for both men and women. This biphasic development of athletic performance was first described by Moore (1975) and has later been characterised using quadratic and polynomial functions (Berthelot et al. 2012; Bongard et al. 2007; Stones and Kozma 1984). As shown by Figure 3, in agreement with the previous literature (Boccia et al. 2017), the rate of improvement in running performance for males appears to be much greater than females up until their late teens. However, over the early 20s improvement in female sprinting performance surpasses their male counterparts. For example, at age 20, average times for women are improving by 0.07 s per year whereas men show an improvement of 0.04 s per year.

There is no evidence of differences in the rate of deterioration of performance between the sexes at older ages. Between the age of peak performance and 30 years of age, both male and female sprint performance deteriorates at a rate of 0.03 s/year, which is similar to previous findings investigating short-duration events in swimming over a similar age range (Tanaka and Seals 1997). The similar magnitudes of decline in short duration sprinting performance with age between male and female athletes suggests that age-related decline in physiological determinants of anaerobic performance is likely to be similar between groups at the elite level. As shown in Figure 2, both male and female lifters show a rapid improvement in performance until an age of peak performance followed by

a subsequent decrease in performance. The rate of increase in performance for males is greater than females at younger ages (up to 23 years old), which is unsurprising, since functional performance capacity increases much faster for males compared to females during puberty. However, this difference in the rate of performance improvement is not evident following maturation between 23 and 30 years. After 30, average men's performance decreases at a faster rate than females. This may be due to different drop out rates between males and females once athletes pass their age of peak performance.

Interestingly, performance in weightlifting is maintained at a level close to the peak for a much longer period than seen in the 100 m data, with a clear plateau above the 0 point being evident from the mid-20's until approximately 30 in the age-related performance derivative (Figure 3). Indeed, previous research has suggested that task related duration appears to modulate the decline in physiological functional capacity (Donato et al. 2003). Moreover, there is evidence that muscle function appears to decline less rapidly in the upper limbs compared with the lower limbs (McDonagh, White, and Davies 1984), which may also play a factor in the findings. A slower rate of decline in weightlifting versus sprinting performance may therefore be attributed to a slower decline in anaerobic power of the upper body muscles.

Similar results were found between sprinting and weightlifting for age of peak performance (Figure 4) with men reaching their peak approximately 2 years earlier than women. These findings are in accordance with Haugen et al. (2018) who analysed season best results from World ranked track and field athletes. The current study observed similar results despite analysing all results from athletes competing on the World stage, not just those within the top 100. Reasons for this sex differences are unclear, but previous research has speculated that factors

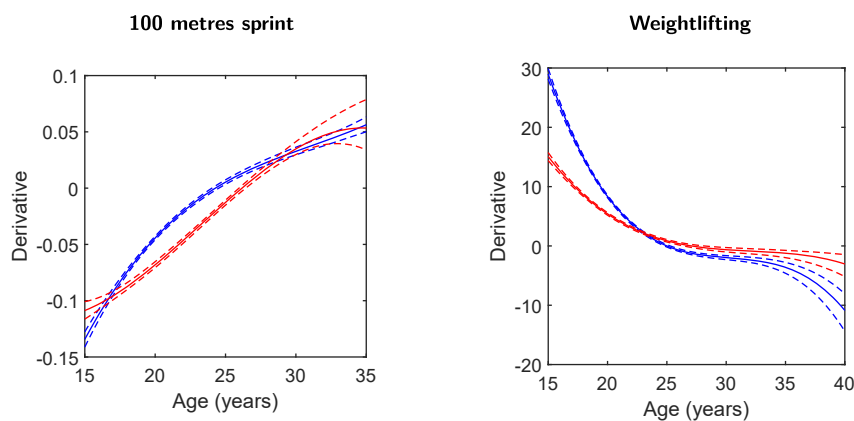


Figure 3: The derivative of the population performance trajectory for 100 m sprinting and weightlifting shown as posterior median (solid line) and 95% credible interval (dashed lines). Men and women are shown in blue and red respectively.

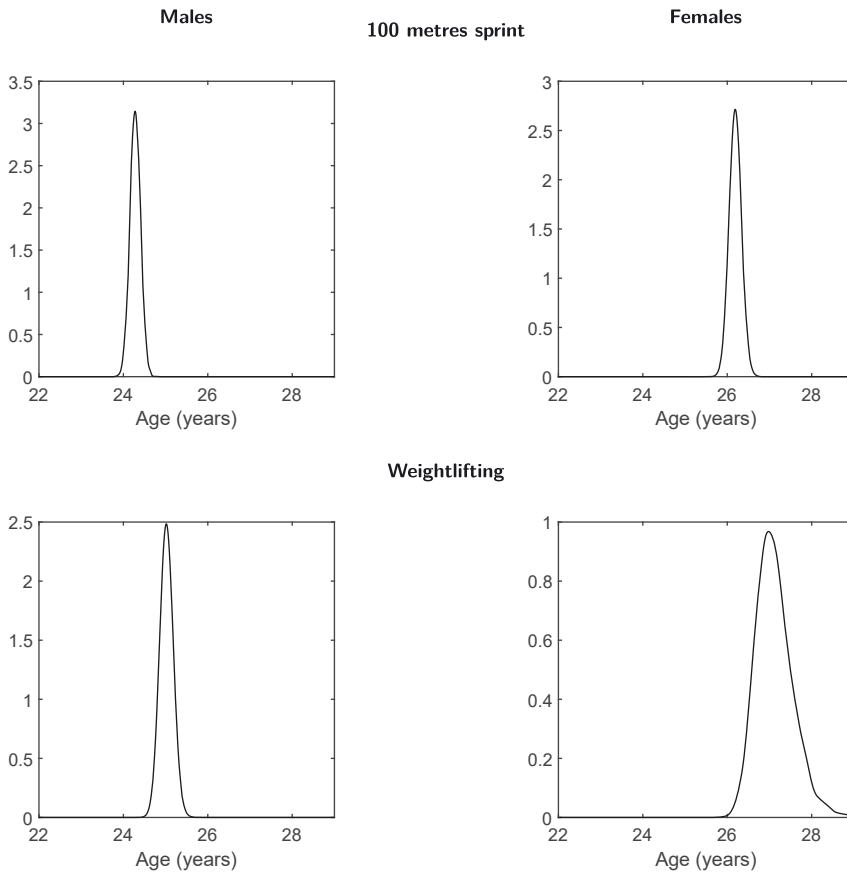


Figure 4: Posterior distribution of population peak age for 100 m sprint and weightlifting.

such as hormone-dependent changes in muscle and fat mass, access to specialized training, exposure to training and technique development, and child-bearing in females are all possible causes (Haugen et al. 2018). However, as outlined above, data suggests that women catch up some of this gender gap in their early and mid-20’s, and supports the notion that women appear to improve more quickly than men in the years immediately preceding their age of peak performance (Haugen et al. 2018).

4.2 Variability of performance in males versus females

Results of the current study demonstrate that the variance in 100 m sprinting performance is larger for females than males (Table 1). Performance for both males and females demonstrates a positively skewed error distribution with tails close to those of a normal distributions which suggests that athletes are slightly more likely to under-perform

Table 1: 100 m sprint: Parameter estimates for the model with month and wind effect (with regressors) and without either effect (without regressors).

	Without regressors			With regressors		
	$E(\sigma_i^2)$	α	ν	$E(\sigma_i^2)$	α	ν
Males	0.0293 (0.0181, 0.309)	1.29 (0.00, 1.37)	52.1 (29.9, 93.1)	0.0197 (0.0163, 0.0206)	1.21 (0.95, 1.30)	18.4 (14.6, 21.8)
Females	0.0462 (0.0279, 0.0483)	1.55 (0.52, 1.64)	49.0 (23.0, 73.2)	0.0307 (0.0215, 0.0338)	1.35 (0.53, 1.48)	19.6 (14.2, 47.5)

(i.e. run a slower time) than over-perform (i.e. run a faster time). In both males and females, the inclusion of these effects leads to a large decrease in the scale and a shorter credible intervals for α and ν . The posterior predictive distribution of the observation errors $\epsilon_{i,j}$ is shown in Figure 5.

In weightlifting, the variability in performance is much larger for males than females. This is in contrast to previous literature which suggests greater within-athlete variability in women than men (McGuigan and Kane 2004; Solberg et al. 2019). The reason for these contrasting findings is unclear, but may be due to the much larger and more representative sample of female lifters used in the current study. Moreover, the inclusion of more recent competition results may also address concerns in previous literature over level of competitiveness and participant rates of female weightlifting compared to that of their male counterparts. The error distribution is negatively skewed with a heavy tail for both males and females which, and, in a similar way to the 100 m sprinting, indicates that an under-performance is more likely than an over-performance. As shown in Figure 5, the size of skewness and kurtosis is very similar between males and females. Even though similar in nature, the error distribution is clearly more skewed and heavier tailed than that of 100 m sprinting. This may be due to the fact that sprinting involves a single performance whereas weightlifting involves a set series of lifts in which the athlete attempts to lift the largest weight possible in order to win the competition. Therefore, lifters may not lift their best possible weight due to either misjudgements of the chosen weights to lift, or competition tactics.

4.3 Effects of covariates

Previous models of athletic performance trajectories have, generally, not included adjustments for confounding

factors of performance, such as seasonality or wind effects. For 100 m sprinting, the effects of the wind and month of competition are shown in Table 1 and Figure 6. There is clear evidence that both effects are important for explaining variation in sprinting performance. The variance of the errors in the regression model are substantially reduced when regressors are included, with effects being similar for both males and females.

In agreement with the findings of previous research, the current study demonstrates that wind speed has a strong effect on performance, with greater tail winds associated with faster times (Moinat, Fabius, and Emanuel 2018). Specifically, for males a 1 m per second increase in tail wind improves sprint time by 0.05 s on average. For females, the tail wind advantage is slightly bigger, leading to an average improvement of 0.06 s. Slightly faster sprint times are also shown during the main part of the outdoor track and field racing calendar in June and July. However, interestingly, the seasonality effect is smaller than that of wind, with the sprinting time improving by only 0.04 s in June and July compared to January, suggesting that performance is remarkably stable over the year.

In weightlifting, there was evidence of a small effect across the months for both males and females. Males demonstrated an improvement of 2.5 and females of 1.8 Sinclair Total units in August compared to January. Plots of the results are shown in Figure 7. This weaker effect leads to similar estimated variability for the model with and without regressors and, as such, we only report results for the model with regressors in Table 2.

4.4 Individual performance trajectory analysis

Analysis of performance trajectories provides evidence of a typical rate of progression of successful athletes, which

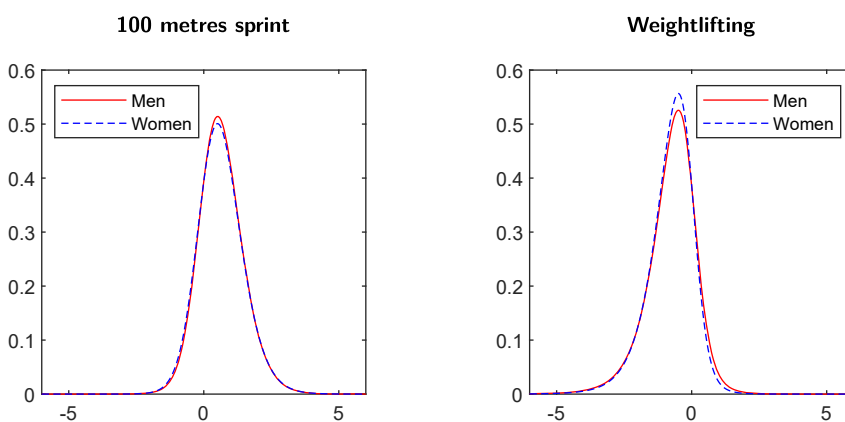


Figure 5: The posterior predictive distribution of $\epsilon_{i,j}$ for 100 m sprinting and weightlifting for men and women.

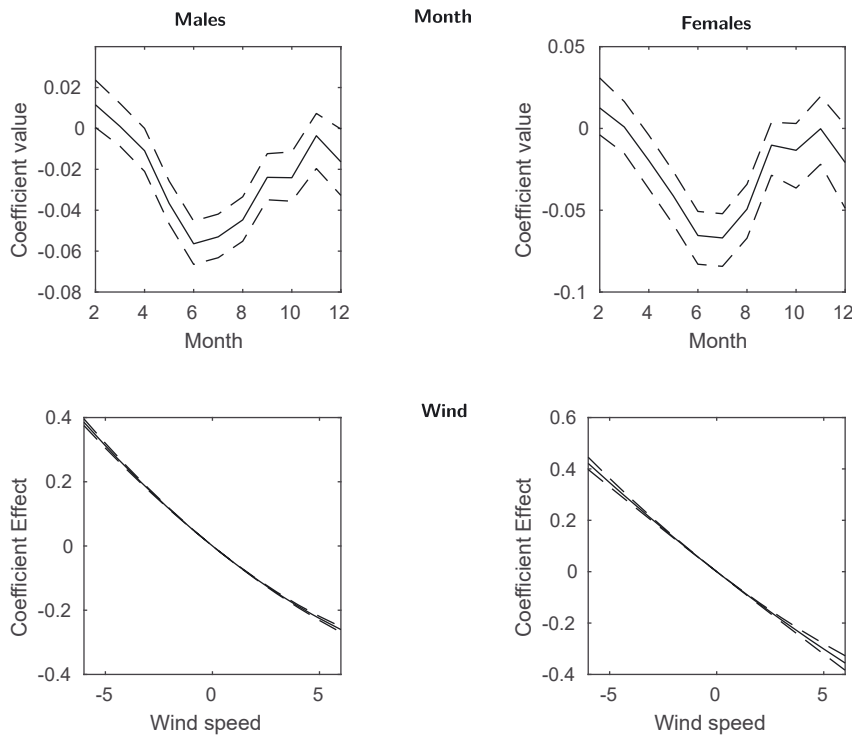


Figure 6: The effects of month and wind for 100 m sprint.




Figure 7: The effects of month for weightlifting.

Table 2: Weightlifting: Parameter estimates for the model with month effect (with regressors) and without month effect (without regressors).

	Without regressors			With regressors		
	$E(\sigma_i^2)$	α	ν	$E(\sigma_i^2)$	α	ν
Males	211 (141, 311)	-1.97 (-1.44, -2.73)	7.9 (6.3, 71.4)	206 (152, 308)	-1.91 (-2.65, -1.34)	7.8 (6.2, 42.2)
Females	134 (103, 154)	-1.73 (-1.97, -1.10)	7.3 (5.9, 9.7)	126 (100, 147)	-1.64 (-1.93, -0.60)	6.9 (5.7, 12.2)

can be used to inform benchmarks for talent identification, analysis of performance during a season to identify the impact of performance improvement strategies, and to predict performance itself. For this level of information it is important to consider performance trajectory at an

individual level in comparison to both their own historical performance level, but also that of the wider cohort of athletes. Within the current model this is provided by the measure of excess performance. In Figure 8, we present an example of individual performance trajectories (adjusted




Figure 8: Individual performance inference for Athletes 1–3: (a) shows observed performances and posterior median individual performance trajectory (solid line) with 95% credible interval (dashed line), and (b) shows adjusted excess performances and posterior median excess performance trajectory (solid line) with 95% credible interval (dashed line).

for seasonality and wind effects) for three of the World’s top 50 (based upon 2021 rankings) male and female 100 m sprinters across the course of their careers. The data points in panel (b) are adjusted by the posterior mean population performance trajectory, and month and wind effects. Figure 8 shows results for men’s 100 m sprinting. It can be seen that across his career, Athlete 1 consistently performs approximately 0.6 s faster than the rest of the population, bearing in mind the population displays performance improvement up until the mid-20’s at which point a decline in performance is seen. Similarly Athlete 3 shows a career profile that evolves at a similar rate to the population leading to a flat excess performance trajectory, but again consistently better than the wider population. By comparison, Athlete 2 was performing at the population average at age 20, but had almost caught up with Athlete 1 by age 30, showing a much

later improvement in performance and an ability to retain performance better than others. At the age of 25 years, Athlete 1 was on average 0.1 s faster than Athlete 2.

Similar patterns of performance are seen in female sprint athletes (Figure 9) such as Athlete 5 who’s performance career path largely follows the same polynomial pattern as the wider population, but with a between 0.19 and 0.42 s greater level of excess performance. Athlete 6 is an athlete who competed in the heptathlon in the earlier part of their career and specialised in sprinting disciplines later in her career. As such she realises a greater performance gain of 0.50 s into her early and mid-20’s than seen across the wider population, but as she learns the skills of sprinting the rate of her performance improvement slows to match that of the wider population by the age of 30. As such her excess performance represents a “v” shaped trajectory. Going through the American College

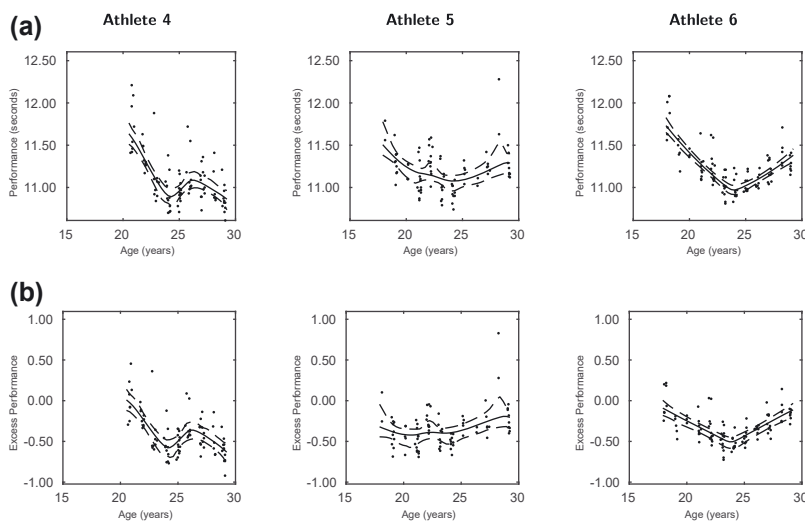


Figure 9: Individual performance inference for Athletes 4–6: (a) shows observed performances and posterior median individual performance trajectory (solid line) with 95% credible interval (dashed line), and (b) shows adjusted excess performances and posterior median excess performance trajectory (solid line) with 95% credible interval (dashed line).

athletic system, Athlete 4 shows a fast rate of improvement in her performance over her early 20s (0.57 s/year at the initial peak age of 24 years). However, recurrent injuries between the ages of 26 & 27 limits her rate of performance improvement compared to the wider population of sprinters. Following resolution of her injuries, the athlete goes on to demonstrate a second peak in her performance at 29 years of age when her rate of improvement reaches 0.63 s/year.

Figure 10 shows the posterior distribution of the age of peak performance across six 100 m sprint athletes. There are clearly differences in these distributions. There is clear evidence that Athlete 2 peaks at a later age than Athletes 1, 3 or 6. There are also some differences in shape. Particular

clear is the bi-modal posterior distribution of Athlete 4 which indicates the aforementioned two periods when this athlete was performing at their peak level.

Individual performance trajectories for 3 of the top 50 ranked male and female weightlifters (as of August 2021) are shown in Figures 11 and 12. It is evident that elite weightlifting athletes tend to compete fewer times per year than the 100 m sprinters. However, performance trajectories appear similar, with increasing (Athlete 7), polynomial (Athlete 8), and plateauing (Athlete 9) patterns observed. A key difference between the 100 m sprinting and weightlifting data sets is the skew distribution. The negative skew distribution within the weightlifting data suggests that poor performances are often a lot worse than

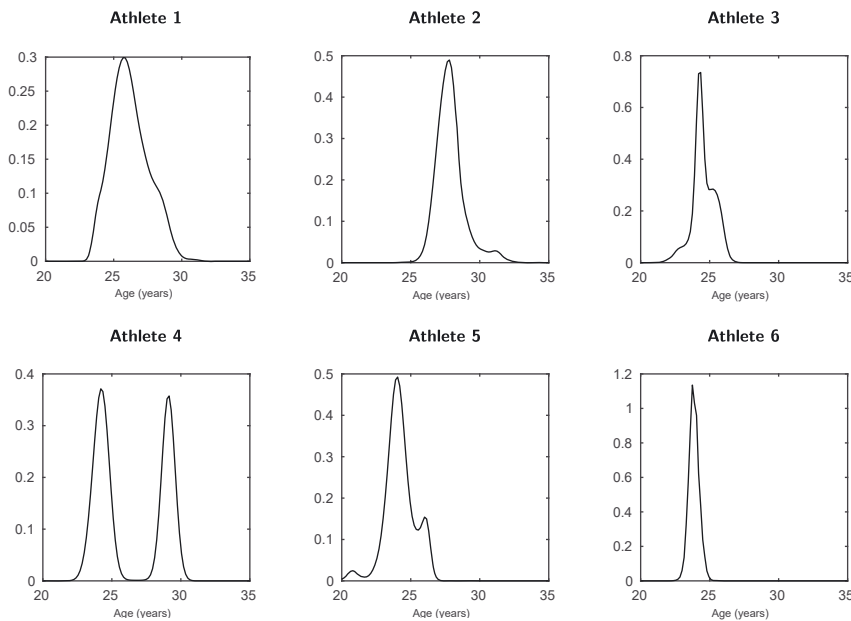


Figure 10: Posterior distributions of individual age of peak performance for Athletes 1–6.

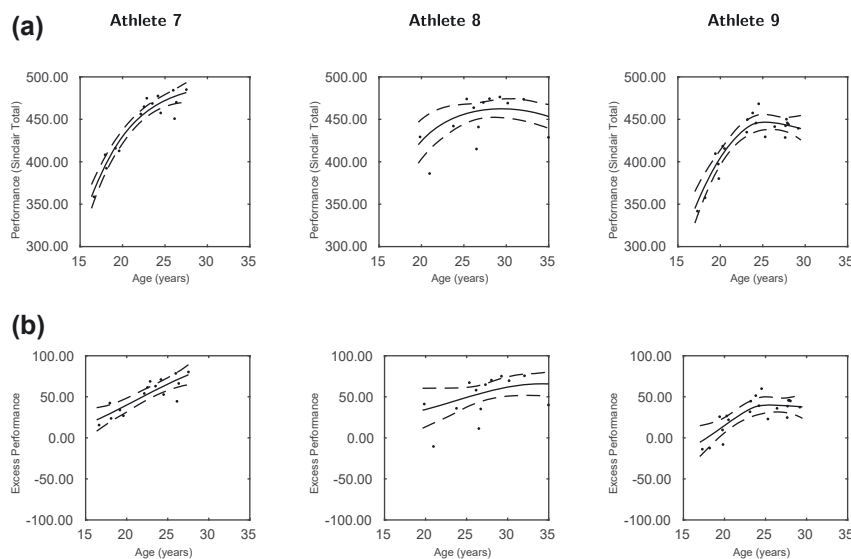


Figure 11: Individual performance inference for Athletes 7–9: (a) shows observed performances and posterior median individual performance trajectory (solid line) with 95% credible interval (dashed line), and (b) shows adjusted excess performances and posterior median excess performance trajectory (solid line) with 95% credible interval (dashed line).

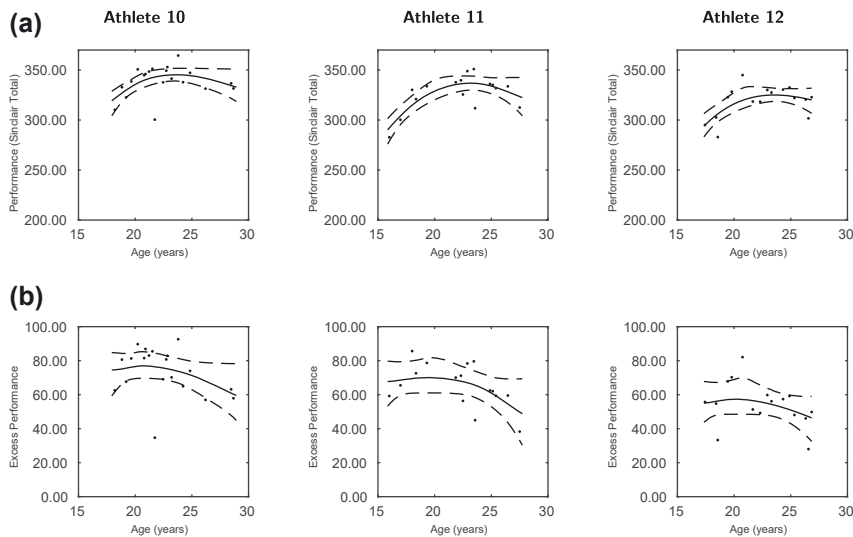


Figure 12: Individual performance inference for Athlete 10–12: (a) shows observed performances and posterior median individual performance trajectory (solid line) with 95% credible interval (dashed line), and (b) shows adjusted excess performances and posterior median excess performance trajectory (solid line) with 95% credible interval (dashed line).

exceptionally good ones. Evidence of this can be seen in the career trajectory from Athlete 8, where, despite a general upward trajectory in his level of excess performance, his worst performances lie outside the 95% confidence bounds. A similar rate of performance increase is seen in the trajectory of Athlete 7 over the early stage of his career. The rate of increase in performance shows a tendency towards plateauing towards his late 20s, but excess performance continues to increase linearly, suggesting that the lifter is still continuing to make significant gains in performance over the age matched wider population of lifters. The performance trajectory of female weightlifters (Figure 12) is much flatter than that of males, with clear evidence of a plateau in lifting performance, and excess performance from their mid-20's onwards. Notwithstanding the age polynomial, the median of excess performance suggests that these athletes are consistently performing between 50 and 100 Sinclair units above that of the wider population of female weightlifters, even in to the later stages of their careers.

5 Discussion

This study developed a Bayesian hierarchical model to investigate both population and individual level longitudinal performance trajectories over time. Performance change was modelled as a function of age using a non-parametric curve fitting approach which accounted for the estimated effects of known covariates within the regression model. The inclusion of these covariates allows us to adjust for confounders and leads to a substantially reduced variance of the observation errors in both for 100 m sprinting and weightlifting. By reducing these error variances,

the model produced tight-fitting individual performance trajectories that are appropriate for the assessment of athletic performance changes over the course of their career.

The model has potential to be applied a wide-range of sporting events within the category of centimetre-gram-second sports. For example, in swimming performance time can be modelled with pool size and depth included as confounders as both will influence performance; in timed track cycling events, the venue (altitude, humidity, track surface) has an important systematic effect on times which can be included as a confounder; finally in longer distance running events, performances can be influenced by factors including the environment and race tactics, and by including a “race” effect would allow for consistent differences in race times across a whole field. Within this manuscript we have demonstrated the model on sports measured in centimetres–grams–seconds, but it can also be applied to a wide range of others where performance is measured in rankings (e.g. Tennis), scores (e.g. Gymnastics), or where there are player metrics provided in competition (e.g. distance covered in team sports).

Several methodological limitations need to be acknowledged with this study. Career trajectories were developed giving equal importance to each performance and did not account for the strength of the competition. It is likely that athletes prioritised certain events, or performances were affected due to the relative strength of their opponents. Data were limited to results appearing on the World Athletics and International Weightlifting Federation websites, and therefore do not include lower level competitions or those from junior age-group events. We cannot discount the potential effects of doping within

the data, especially given the history of anti-doping rule violations in both 100 m sprinting and weightlifting. Future work will look at extending the model to other centimetre–gram–second sports as well as others where performance is measured by finishing position/rank or individual player metrics. This will involve adjusting for more confounders and modelling the effect of competition strength.

Acknowledgments: We are grateful to Alex Bohl and Lorenzo Gaborini who assisted us in acquiring data from the World Athletics and International Weightlifting Federation websites, respectively.

Author contributions: All the authors have accepted responsibility for the entire content of this submitted manuscript and approved submission.

Research funding: This study received funding from the Partnership for Clean Competition (Grant ID: 26), and utilised the University of Kent High Performance Computing Cluster (0000000053).

Conflict of interest statement: The authors declare no conflicts of interest regarding this article.

Appendix A: Computational method

We define $IG(a, b)$ to be an inverse gamma distribution with shape parameter a and scale parameter b which has density

$$p(x) = \frac{b^a}{\Gamma(a)} x^{-(a+1)} \exp\{-b/x\},$$

$Ga(a, b)$ to be a gamma distribution with shape parameter a and scale parameter b which has density

$$p(x) = \frac{b^a}{\Gamma(a)} x^{a-1} \exp\{-bx\},$$

$N(\mu, \Sigma)$ to be a multivariate normal with mean μ and covariance matrix Σ , $TN_A(\mu, \sigma^2)$ represents a normal distribution with mean μ and variance σ^2 truncated to the interval A , C^+ is the standard half-Cauchy distribution which has density

$$p(x) = \frac{2}{\pi} (1+x)^{-1}, \quad x > 0,$$

$GIG(p, a, b)$ is the generalized inverse Gaussian distribution with shape p and scale parameters a and b which has density

$$p(x) = \frac{(a/b)^{p/2}}{2K_p(\sqrt{ab})} x^{p-1} \exp\{-(ax + b/x)/2\}$$

If Y follows a zero-mean skew- t distribution with parameters ω^2 , α and ν can be represented as

$$Y = X + \frac{\alpha}{\sqrt{1+\alpha^2}} Z$$

where $X \sim N(0, 1/W)$ and $Z \sim TN_{[0, \infty)}(0, 1/W)$, $W \sim Ga(\nu/2, (\omega\nu)/2)$. This allows us to write the model in a convenient form for Gibbs sampling,

$$y_{i,j} = x_{i,j}\zeta + d_{i,j}\eta + s_{i,j}\beta_i^\gamma + \frac{\alpha}{\sqrt{1+\alpha^2}} z_{i,j} + \epsilon_{i,j}^*$$

where $\epsilon_{i,j}^* \sim N(0, 1/w_{i,j})$, $z_{i,j} \sim TN_{[0, \infty)}(0, 1/w_{i,j})$, $w_{i,j} \sim Ga(\nu/2, (\nu\sigma_i^2)/2)$

A.1 Gibbs sampler

We define the following notation:

\mathbf{X}_i is a n_i -dimensional vector whose j th entry is $d_{i,j}$ concatenated with $x_{i,j}$, \mathbf{Y}_i^* is a n_i -dimensional vector whose j th entry is $\sqrt{w_{i,j}} \left[\sqrt{1+\alpha^2} (y_{i,j} - x_{i,j}\zeta - d_{i,j}\eta) - \alpha z_{i,j} \right]$, \mathbf{Y}_i^{**T} is a n_i -dimensional vector whose j th entry is $y_{i,j} - \frac{\alpha}{\sqrt{1+\alpha^2}} z_{i,j}$, \mathbf{S}_i^* is a $(n_i \times (p_i^\gamma + 1))$ -dimensional matrix whose j th row has entry $\sqrt{\frac{w_{i,j}}{1+\alpha^2}}$ and subsequent entries $\sqrt{\frac{w_{i,j}}{1+\alpha^2}} s_{i,j}$ for the j with $\gamma_{i,j} = 1$ and

$$\mathbf{\Delta}_i = \mathbf{S}_i^{*T} \mathbf{S}_i^* + \text{diag} \left(\frac{1}{\psi}, \underbrace{\frac{1}{g}, \dots, \frac{1}{g}}_{p_i^\gamma \text{ times}} \right).$$

The full conditional distributions of the Gibbs sampler are as follows.

We update η and ζ jointly marginalising over $(\theta_1, \beta_1^\gamma), \dots, (\theta_M, \beta_M^\gamma)$,

$$\begin{pmatrix} \eta \\ \zeta \end{pmatrix} \sim N \left(\mathbf{\Xi}^{-1} \left[\sum_{i=1}^M \mathbf{X}_i^T \mathbf{\Sigma}_i^{-1/2} (\mathbf{I}_p - \mathbf{S}_i \mathbf{\Delta}_i^{-1} \mathbf{S}_i^T) \right. \right. \\ \left. \left. \times \mathbf{\Sigma}_i^{-1/2} \mathbf{Y}_i^{**} \right], \mathbf{\Xi}^{-1} \right),$$

where

$$\mathbf{\Xi} = \sum_{i=1}^M \mathbf{X}_i^T \mathbf{\Sigma}_i^{-1/2} (\mathbf{I}_{n_i} - \mathbf{S}_i \mathbf{\Delta}_i^{-1} \mathbf{S}_i^T) \mathbf{\Sigma}_i^{-1/2} \mathbf{X}_i + \mathbf{\Lambda}_*,$$

$$\mathbf{\Sigma}_i = \text{diag} \left(\frac{\sigma_i^2}{w_{i,1}(1+\alpha^2)}, \frac{\sigma_i^2}{w_{i,2}(1+\alpha^2)}, \dots, \frac{\sigma_i^2}{w_{i,n_i}(1+\alpha^2)} \right),$$

$$\mathbf{\Lambda}_* = \text{diag} \left(\underbrace{0, \dots, 0}_{5 \text{ times}}, \frac{a}{b\lambda_1^2 \tau^2}, \dots, \frac{a}{b\lambda_p^2 \tau^2} \right),$$

$$\lambda_l^2 \sim \text{IG}\left(1, \frac{1}{v_l^*} + \frac{\zeta_l^2 a}{2\tau^2 b}\right), \quad l = 1, \dots, p$$

$$v_l^* \sim \text{IG}\left(1, 1 + \frac{1}{\lambda_l^2}\right), \quad l = 1, \dots, p,$$

$$\tau^2 \sim \text{IG}\left(\frac{p}{2}, \frac{1}{\xi} + \frac{a}{2b} \sum_{l=2}^p \frac{\zeta_l^2}{\lambda_l^2}\right)$$

$$\xi \sim \text{IG}\left(1, 1 + \frac{1}{\tau^2}\right),$$

We update γ_i using the method in Section 3.

$$\begin{aligned} \sigma_i^2 &\sim \text{IG}\left(a + \frac{n_i}{2}, b + \frac{1}{2}\left(\sum_{j=1}^{n_i} w_{i,j} z_{i,j}^2 + \mathbf{Y}_i^{*T} (\mathbf{I}_{n_i} - \mathbf{S}_i^* \mathbf{\Delta}_i^{-1} \mathbf{S}_i^{*T}) \mathbf{Y}_i^*\right)\right), \\ \begin{pmatrix} \theta_i \\ \boldsymbol{\beta}_i^y \end{pmatrix} &\sim \text{N}\left(\mathbf{\Delta}_i^{-1} \mathbf{S}_i^{*T} \mathbf{Y}_i^*, \sigma_i^2 \mathbf{\Delta}_i^{-1}\right) \\ w_{i,j} &\sim \text{Ga}\left(\frac{\nu + 2}{2}, \frac{\nu \sigma_i^2 + \left\{\sqrt{1 + \alpha^2} [y_{i,j} - \theta_i - \mathbf{s}_{i,j}^y \boldsymbol{\beta}_i^y - x_{i,j} \zeta - d_{i,j} \eta] - \alpha z_{i,j}\right\}^2 + z_{i,j}^2}{2\sigma_i^2}\right), \\ z_{i,j} &\sim \text{TN}_{[0,+\infty)}\left(\frac{\alpha [y_{i,j} - \theta_i - \mathbf{s}_{i,j}^y \boldsymbol{\beta}_i^y - x_{i,j} \zeta - d_{i,j} \eta]}{\sqrt{1 + \alpha^2}}, \frac{\sigma_i^2}{w_{i,j}(1 + \alpha^2)}\right), \\ b &\sim \text{GIG}\left(Ma - \frac{1}{2}p + 1, 2\sum_{i=1}^M \frac{1}{\sigma_i^2}, \frac{a}{\tau^2} \sum_{l=2}^p \frac{\zeta_l^2}{\lambda_l^2}\right), \\ \pi(a) &\propto \left(\frac{b^a}{\Gamma(a)}\right)^M \left(\prod_{i=1}^M \sigma_i^2\right)^{-a} a^{p/2} \exp\left\{-\frac{a}{2\tau^2 b} \sum_{l=1}^p \frac{\zeta_l^2}{\lambda_l^2}\right\}, \end{aligned}$$

which can be updated using a Metropolis-Hastings random walk on $\log a$,

$$\begin{aligned} \psi &\sim \text{IG}\left(1 + \frac{M}{2}, 1 + \frac{1}{2} \sum_{i=1}^M \frac{(\boldsymbol{\beta}_{i,1}^y)^2}{\sigma_i^2}\right), \\ \pi(g) &\propto g^{-1/2} (1 + g)^{-1} \prod_{i=1}^M m_i(\gamma_i), \end{aligned}$$

which can be updated using a Metropolis-Hastings random walk on $\log g$,

$$\begin{aligned} p(\alpha) &\propto (1 + \alpha^2)^{\sum_{i=1}^M n_i/2} \exp\left\{-\frac{1}{2} \frac{\alpha^2}{100} - \frac{1}{2} \sum_{i=1}^M \right. \\ &\times \sum_{j=1}^{n_i} \frac{w_{i,j}}{\sigma_i^2} \left\{\sqrt{1 + \alpha^2} [y_{i,j} - \theta_i - \mathbf{x}_{i,j}^y \boldsymbol{\beta}_i^y \right. \\ &\left. \left. - x_{i,j} \zeta - d_{i,j} \eta] - \alpha z_{i,j}\right\}^2 \right\} \end{aligned}$$

which can be updated using a Metropolis-Hastings random walk on $\log \alpha$.

$$p(\nu) \propto \prod_{i=1}^M \left[\frac{\left(\frac{\nu}{2}\right)^{\nu/2}}{\Gamma\left(\frac{\nu}{2}\right)} \right]^{n_i} \left(\prod_{j=1}^{n_i} w_{i,j}\right)^{\nu/2} \exp\left[-\nu \left(\frac{\sum_{j=1}^{n_i} w_{i,j}}{2} + 0.1\right)\right]$$

which can be updated using a Metropolis-Hastings random walk on $\log \nu$.

References

- Albert, J. 1999. "Comment on "Bridging Different Eras in Sports" by Berry, Reese and Larkey." *Journal of the American Statistical Association* 94: 677–80.
- Allen, S. V., T. J. Vandenbogaerde, and W. G. Hopkins. 2014. "Career Performance Trajectories of Olympic Swimmers: Benchmarks for Talent Development." *European Journal of Sport Science* 14: 643–51.
- Azzalini, A. 1985. "A Class of Distributions Which Includes the Normal Ones." *Scandinavian Journal of Statistics* 12: 171–8.

- Azzalini, A., and A. Capitanio. 2003. "Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew T -Distribution." *Journal of the Royal Statistical Society: Series B* 65: 367–89.
- Berry, S. M., C. S. Reese, and P. D. Larkey. 1999. "Bridging Different Eras in Sports." *Journal of the American Statistical Association* 94: 661–76.
- Berthelot, G., S. Len, P. Hellard, M. Tafflet, M. Guillaume, B. Vollmer, J.-C. nad Gager, L. Quinquis, A. Marc, and J.-F. Toussaint. 2012. "Exponential Growth Combined with Exponential Decline Explains Lifetime Performance Evolution in Individual and Human Species Explains Lifetime Performance Evolution in Individuae." *Age* 34: 1001–9.
- Boccia, G., P. Moisè, A. Franceschi, F. Trova, D. Panero, A. La Torre, A. Rainoldi, F. Schena, and M. Cardinale. 2017. "Career Performance Trajectories in Track and Field Jumping Events from Youth to Senior Success: The Importance of Learning and Development." *PLoS One* 12: 1–15.
- Bongard, V., A. Y. McDermott, G. E. Dallal, and E. J. Schaefer. 2007. "Effects of Age and Gender on Physical Performance." *Age* 29: 77–85.
- Boys, R. J., and P. M. Philipson. 2019. "On the Ranking of Test Match Batsmen." *Journal of the Royal Statistical Society: Series A C* 68: 161–79.
- Bradbury, J. C. 2009. "Peak Athletic Performance and Ageing: Evidence from Baseball." *Journal of Sports Sciences* 27: 599–610.
- Brander, J. A., E. J. Egan, and L. Yeung. 2014. "Estimating the Effects of Age on NHL Player Performance." *Journal of Quantitative Analysis in Sports* 10: 241–59.
- Carvalho, C. M., N. G. Polsen, and J. G. Scott. 2010. "The Horseshoe Estimator for Sparse Signals." *Biometrika* 97: 465–80.
- Dellaportas, P., J. J. Forster, and I. Ntzoufras. 2000. "Bayesian Variable Selection Using the Gibbs Sampler." In *Generalized Linear Models: A Bayesian Perspective*, Vol. 5, edited by D. K. Dey, S. K. Ghosh, and B. K. Mallick, 273–86. Chemical Rubber Company Press. Also available at <https://eprints.soton.ac.uk/29960/>.
- Denison, D. G. T., C. C. Holmes, B. K. Mallick, and A. F. M. Smith. 2002. *Bayesian Methods for Nonlinear Classification and Regression*. Chichester: Wiley & Sons.
- Donato, A. J., K. Tench, D. H. Glueck, D. R. Seals, I. Eskurza, and H. Tanaka. 2003. "Declines in Physiological Functional Capacity with Age: A Longitudinal Study in Peak Swimming Performance." *Journal of Applied Physiology* 94: 764–9.
- Egidi, L., and J. Gabry. 2018. "Bayesian Hierarchical Models for Predicting Individual Performance in Soccer." *Journal of Quantitative Analysis in Sports* 14: 143–57.
- Frühwirth-Schnatter, S., and S. Pyne. 2010. "Bayesian Inference for Finite Mixtures of Univariate and Multivariate Skew-Normal and Skew-T Distributions." *Biostatistics* 11: 317–36.
- Gao, Z., Y. Li, and Z. Wang. 2020. "Restoring the Real World Records in Men's Swimming without High-Tech Swimsuits." *Journal of Quantitative Analysis in Sports* 16: 291–300.
- Griffin, J. E., K. Latuszynski, and M. F. J. Steel. 2021. "In Search of Lost (Mixing) Time: Adaptive Markov Chain Monte Carlo Schemes for Bayesian Variable Selection with Large P ." *Biometrika* 108: 53–69.
- Haugen, T. A., P. A. Solberg, C. Foster, R. Morán-Navarro, F. Breitschädel, and W. G. Hopkins. 2018. "Peak Age and Performance Progression in World-Class Track-And-Field Athletes." *International Journal of Sports Physiology and Performance* 13: 1122–9.
- Kovalchik, S. A., and R. Stefani. 2013. "Longitudinal Analyses of Olympic Athletics and Swimming Events Find No Gender Gap in Performance Improvement." *Journal of Quantitative Analysis in Sports* 9: 15–24.
- Lang, S., and A. Brezger. 2004. "Bayesian P-Splines." *Journal of Computational & Graphical Statistics* 13: 183–212.
- Ley, E., and M. F. J. Steel. 2009. "On the Effect of Prior Assumptions in Bayesian Model Averaging with Applications to Growth Regression." *Journal of Applied Econometrics* 24: 651–74.
- Makalic, E., and D. F. Schmidt. 2016. "A Simple Sampler for the Horseshoe Estimator." *IEEE Signal Processing Letters* 23: 179–82.
- McDonagh, M. J., M. J. White, and C. T. Davies. 1984. "Different Effects of Ageing on the Mechanical Properties of Human Arm and Leg Muscles." *Gerontology* 30: 49–54.
- McGuigan, M. R., and M. K. Kane. 2004. "Reliability of Performance of Elite Olympic Weightlifters." *The Journal of Strength & Conditioning Research* 18: 650–3.
- Moinat, M., O. Fabius, and K. S. Emanuel. 2018. "Data-driven Quantification of the Effect of Wind on Athletics Performance." *European Journal of Sport Science* 18: 1185–90.
- Rasmussen, C. E., and C. K. I. Williams. 2006. *Gaussian Processes for Machine Learning*. Cambridge: MIT Press.
- Santos-Fernandez, E., P. Wu, and K. L. Mengersen. 2019. "Bayesian Statistics Meets Sports: A Comprehensive Review." *Journal of Quantitative Analysis in Sports* 15: 289–312.
- Solberg, P. A., W. G. Hopkins, G. Paulsen, and T. A. Haugen. 2019. "Peak Age and Performance Progression in World-Class Weightlifting and Powerlifting Athletes." *International Journal of Sports Physiology and Performance* 14: 1357–63.
- Stephenson, A. G., and J. A. Tawn. 2013. "Determining the Best Track Performances of All Time Using a Conceptual Population Model for All Athletics Records." *Journal of Quantitative Analysis in Sports* 9: 67–76.
- Stevenson, O. G., and B. J. Brewer. 2021. "Finding You Feet: A Gaussian Process Model for Estimating the Abilities of Batsmen in Test Cricket." *Journal of the Royal Statistical Society: Series A C* 70: 481–506.
- Stones, M. J., and A. Kozma. 1984. "Longitudinal Trends in Track and Field Performances." *Experimental Aging Research* 10: 107–10.
- Strand, M., D. Nelson, and G. Grunwald. 2018. "Modeling Between-Subject Differences and Within-Subject Changes for Long Distance Runners by Age." *Journal of Quantitative Analysis in Sports* 14: 81–90.
- Tadesse, M. G., and M. Vannucci. 2021. *Handbook of Bayesian Variable Selection*. Boca Raton: Chapman & Hall, CRC.
- Tanaka, H., and D. R. Seals. 1997. "Age and Gender Interactions in Physiological Functional Capacity: Insight from Swimming Performance." *Journal of Applied Physiology* 82: 846–51.
- Wimmer, N., N. Fenske, P. Pyrka, and L. Fahrmeir. 2011. "Exploring Competition Performance in Decathlon Using Semi-parametric Latent Variable Models." *Journal of Quantitative Analysis in Sports* 7: 6.