Taking error into account when fitting models using approximate Bayesian computation

Article

Supplemental Material

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

1. Randomly select d of the accepted (θ^i, X^i).
2. For each of these d pairs (θ^0, X^0):
    a. Do ABC taking $D = X^0$ using all remaining model runs as the simulations:
        i. Accept the m θ^i values that minimise $\rho(X^i, X^0)$.
    b. For each parameter j:
        i. Calculate p_j^0, the proportion of accepted $\theta_j^i \leq \theta_j^0$
3. Plot the distribution of all p_j^0, and check for uniformity.

Algorithm S1. Original coverage algorithm used in [1].

Table S1. Model parameters. Literature values were derived in [1].

Symbol	Definition	Literature Value	Unit
B_o	taxon-specific normalisation constant	967	kJ / day
E	activation energy	0.25	eV
E_c	energy cost of tissue	3.6	kJ / g
E_f	energy from food	10.6	kJ / g
E_s	energy cost of synthesis	3.6	kJ / g
h	half saturation coefficient	3.5	g / 0.01 g
IG_m	maximum ingestion rate	0.70	g / day / g
M_b	mass at birth	0.011	g
M_c	mass of cocoon	0.015	g
M_m	maximum asymptotic mass	0.5	g
M_p	mass at sexual maturity	0.25	g
r_B	growth constant	0.177	day^{-1}
r_m	maximum energy to reproduction	0.182	kJ / g / day
s	movement speed	0.004	m / day

Textbox 1. Exact Bayesian posteriors for the polynomial example.

Our Bayesian regression setting can be expressed as:
$$y|\beta \sim N(X\beta, s^2 I_{n \times n}), \qquad \beta \sim N(0, t^2 I_{p+1,p+1}).$$

Here y is a vector of n observations, β is a vector of $p + 1$ unknown parameters, s and t are known quantities and X is a $n \times p + 1$ design matrix. For polynomial regression this has ijth entry v_i^{j-1} for some underlying values $v_1, v_2, \ldots, v_n$.

Using standard results on conditional normal distributions it follows that:
$$\begin{pmatrix} \beta \\ y \end{pmatrix} \sim N\left(0, \begin{pmatrix} t^2 I_{p+1,p+1} & t^2 X^T \\ t^2 X & t^2 XX^T + s^2 I_{n \times n} \end{pmatrix}\right),$$
and therefore:
$$\beta|y \sim N(t^2 X^T[t^2 XX^T + s^2 I_{n \times n}]^{-1} y, \ t^2 I_{p+1,p+1} - t^4 X^T[t^2 XX^T + s^2 I_{n \times n}]^{-1} X).$$
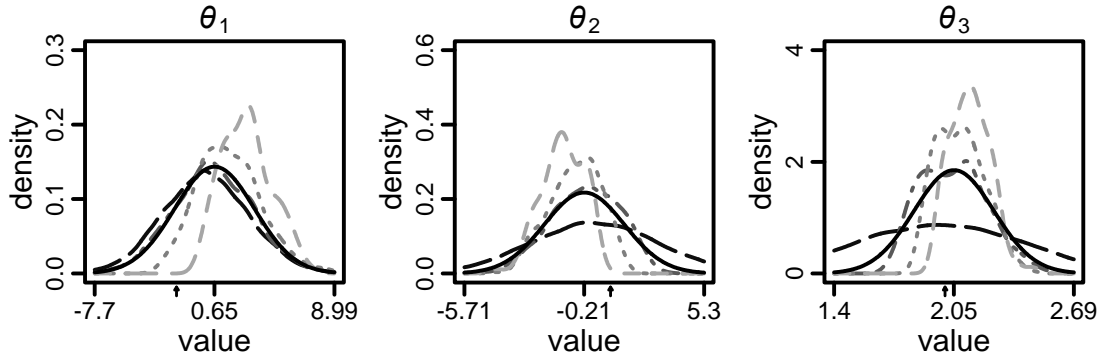
Figure S1. Posterior distributions for the polynomial model. True values for θ_1, θ_2 and θ_3 were -2, 1 and 2, marked by arrows on the x-axes, respectively, and priors were $N(0, 3)$. Solid black lines indicate the posteriors produced by exact Bayesian regression; these are centred differently because of the added noise and the priors that were used. On the horizontal axes, ticks are placed at the mean of the exact posterior density and three standard deviations above and below. Grey lines are 'rejection ABC', with m = 100, 330, 1000 and 10000 acceptances, indicated by light to dark grey respectively. The posteriors obtained with m = 1000 did not significantly differ from those obtained by exact Bayesian regression, but except for b_2 with m = 10000 all others did (Kolmogorov-Smirnov, $p < 0.01$, sample sizes standardised by taking 100 samples from each posterior).
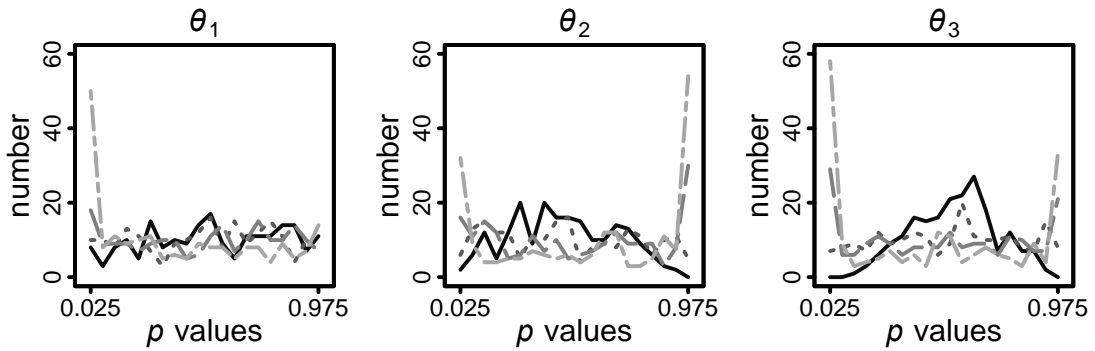

Figure S2. Coverage tests for the polynomial example using 'rejection ABC'. Plots show relative frequency of p, which is the proportion of accepted parameter values that were less than the true value in 200 ABC analyses using "pseudo-data". This "pseudo-data" was generated by adding $N(0, \lambda)$ noise estimated from the 'best-fitting run' to all simulation results and then selecting the 200 runs closest to the data. Relative frequencies for m = 100, 275, 1000 and 10000 accepted runs are indicated by dark to light grey lines respectively. Coverage was indistinguishable from uniformity for m = 1000 only (Kolmogorov-Smirnov, $p > 0.01$).
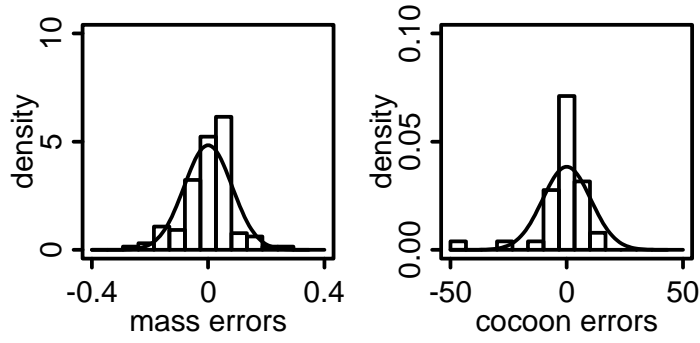
Figure S3. Error distributions for the earthworm IBM. The bars represent the distribution of the errors between the 'best-fitting run' (Figure 3) and the empirical data, separately for the earthworm masses and cocoon counts. The black lines illustrate normal distributions with mean 0 and the estimated standard deviations of 0.08 and 10.4 respectively. Neither distribution was significantly different from normal (Kolmogorov-Smirnov, $p > 0.01$).
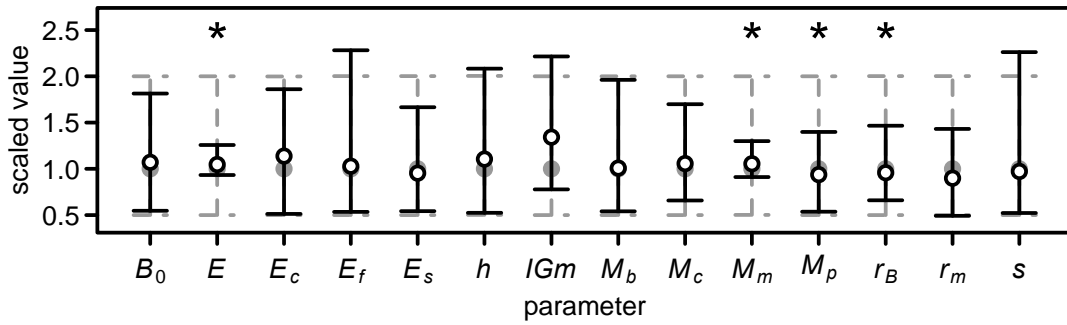

Figure S4. Posterior distributions for the earthworm IBM, 'error-calibrated ABC'. Grey lines show the priors; black lines the posteriors. Circles represent medians, whiskers 95% credible intervals. Asterixes mark significant narrowing (Levene's test, $p < 0.01$). All parameter values were scaled by dividing by the corresponding literature values [2].
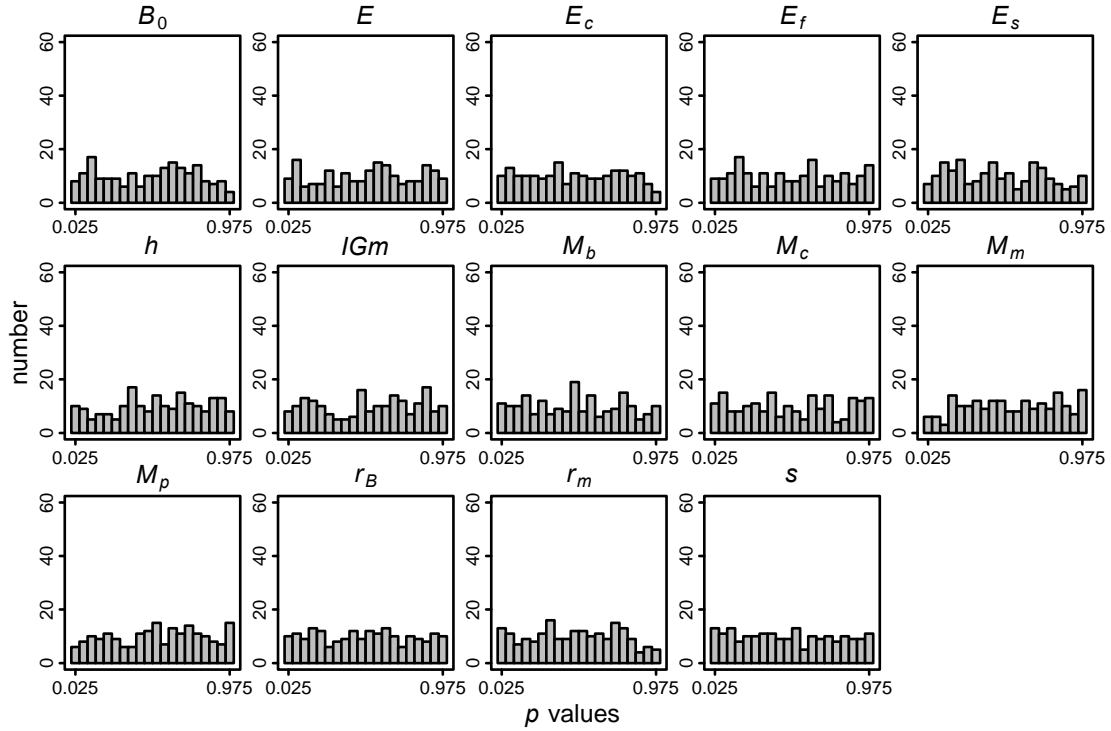
Figure S5. Coverage tests for the earthworm IBM, 'error-calibrated ABC'. Plots show relative frequency of *p*, which is the proportion of accepted parameter values that were less than the true value in 200 ABC analyses using "pseudo-data". This "pseudo-data" was generated by adding $N(0, \lambda)$ noise estimated from the 'best-fitting run' to all simulation results and then selecting the 200 runs closest to the data. All coverage tests were indistinguishable from uniformity (Kolmogorov-Smirnov; $p > 0.01$).
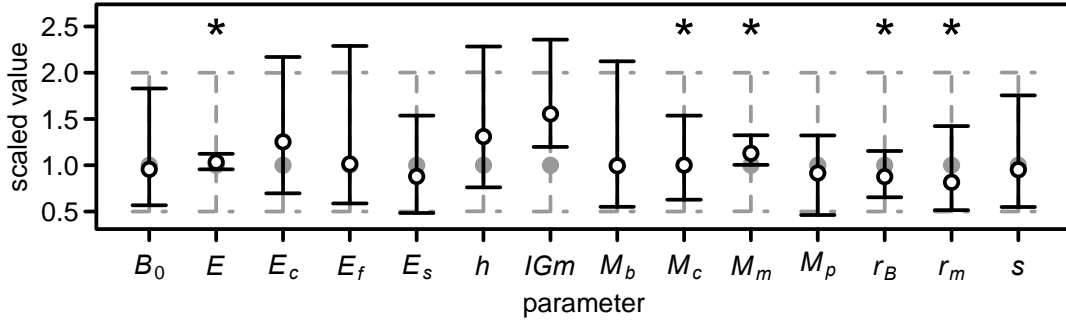

Figure S6. Posterior distributions for the earthworm IBM, 'rejection ABC', 100 acceptances. Grey lines show the priors; black lines the posteriors. Circles represent medians, whiskers 95% credible intervals. Asterixes mark significant narrowing (Levene's test, $p < 0.01$). All parameter values were scaled by dividing by the corresponding literature values [2].
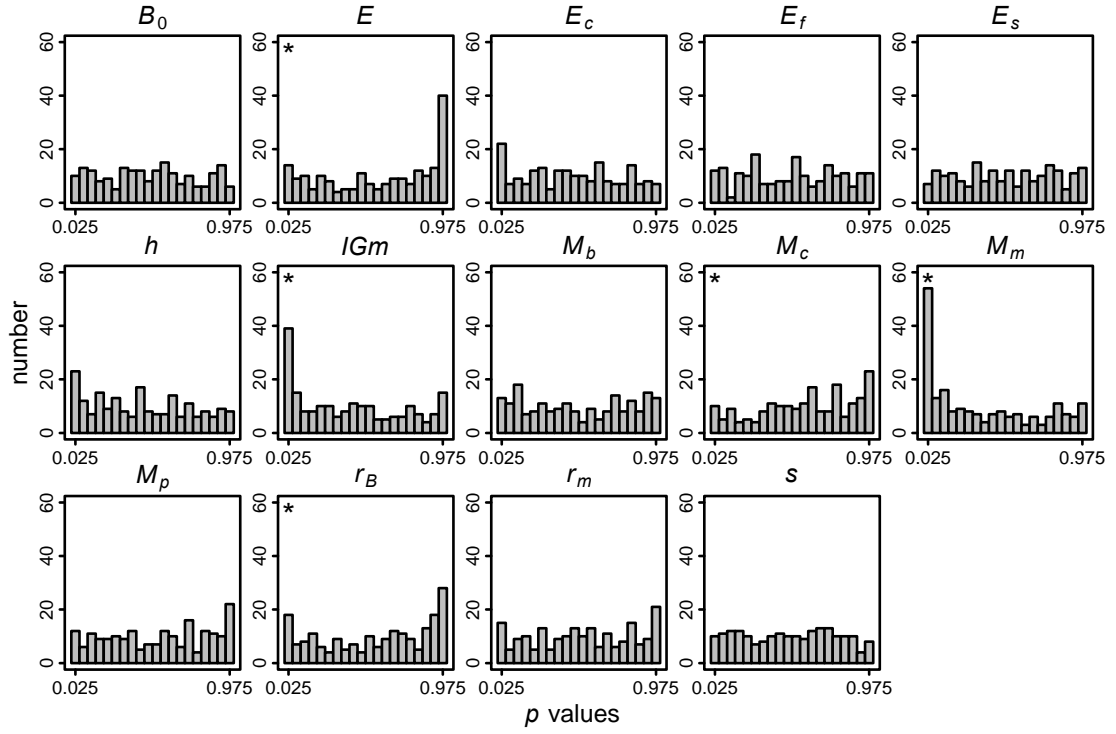
Figure S7. Coverage tests for the earthworm IBM: 'rejection ABC', 100 acceptances. Plots show relative frequency of *p*, which is the proportion of accepted parameter values that were less than the true value in 200 ABC analyses using "pseudo-data". This "pseudo-data" was generated by adding $N(0, \lambda)$ noise estimated from the 'best-fitting run' to all simulation results and then selecting the 200 runs closest to the data. Coverage tests marked by asterixes were significantly different from uniform (Kolmogorov-Smirnov; $p < 0.01$).
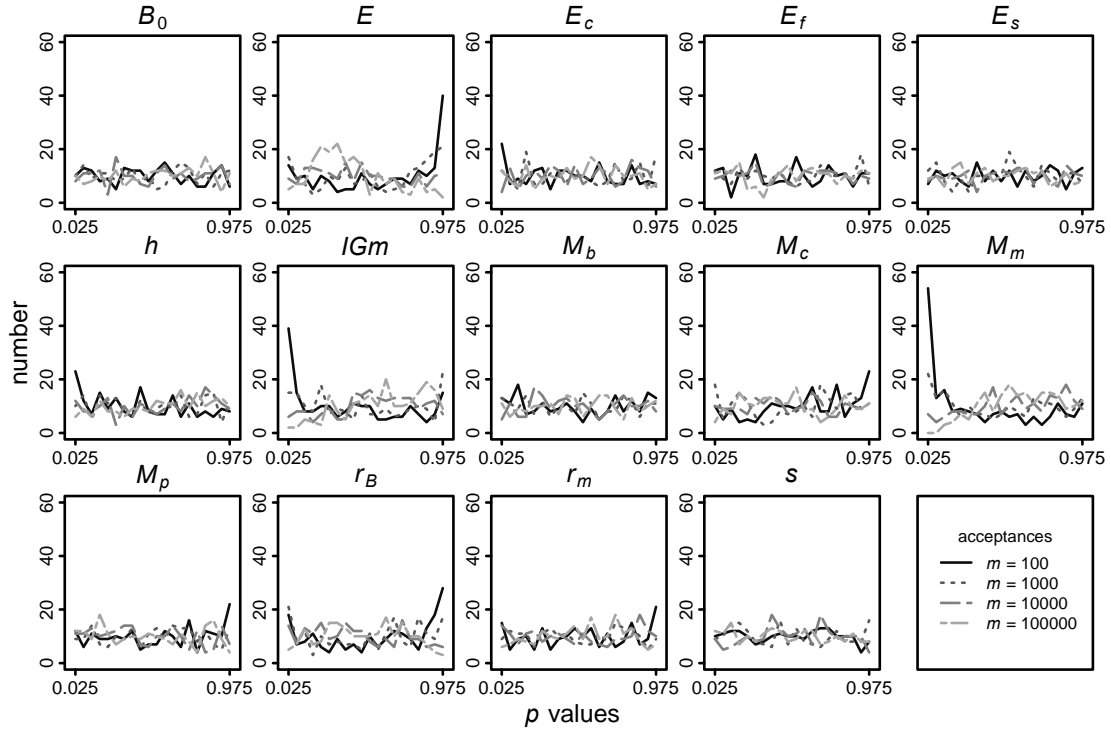
Figure S8. Coverage tests for the earthworm IBM: 'rejection ABC', varying number of acceptances *m*. Plots show relative frequency of *p*, which is the proportion of accepted parameter values that were less than the true value in 200 ABC analyses using "pseudo-data". Relative frequency *p* of accepted parameter values that were less than the true value in 200 ABC analyses using "pseudo-data". This "pseudo-data" was generated by adding $N(0, \lambda)$ noise estimated from the 'best-fitting run' to all simulation results and then selecting the 200 runs closest to the empirical data. At least one parameter was significantly different from uniform (Kolmogorov-Smirnov, $p < 0.01$) for all different values of *m*.
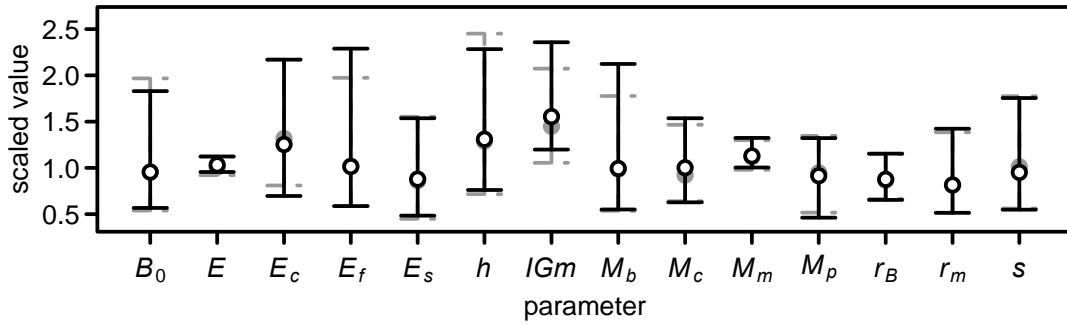

Figure S9. Posterior distributions for the earthworm IBM, 'rejection ABC', 100 acceptances. Grey lines show the posteriors after using the full 160 data points; black lines after using every 6th data point only. Circles represent medians, whiskers 95% credible intervals. No posteriors differed significantly between the two methods (Kolmogorov-Smirnov, $p > 0.01$). All parameter values were scaled by dividing by the corresponding literature values [2].

References

[1] van der Vaart, E., Beaumont, M.A., Johnston, A. & Sibly, R.M. 2015 Calibration and evaluation of individual-based models using Approximate Bayesian Computation. *Ecological Modelling* **312**, 182-190.

[2] Johnston, A.S.A., Hodson, M.E., Thorbek, P., Alvarez, T. & Sibly, R.M. 2014 An energy budget agent-based model of earthworm populations and its application to study the effects of pesticides. *Ecological Modelling* **280**, 5-17.