Supplementary material

Hierarchical models estimation procedures: Bayesian inference

Setting up the multilevel Beverton–Holt SR model [Equation (4) in main text), we obtain the following hierarchical structure:

$$y_{it} = \alpha_{it}^{\mathrm{BH}} + \log(\beta_{it}^{\mathrm{BH}}) - \log(\beta_{it}^{\mathrm{BH}} + x_{it}) + \varepsilon_{it}\,, \tag{A.1}$$

with $\varepsilon_{it} \overset{\mathrm{iid}}{\sim} \mathrm{N}(0, \sigma_{yi}^2)$ and parameters *alpha* ($\alpha_{it}^{\mathrm{BH}}$) and *beta* ($\beta_{it}^{\mathrm{BH}}$) depending on stock-specific T time-series and, thus, being time-varying:

$$\alpha_{it}^{\mathrm{BH}} = c_{oi}^{\mathrm{BH}} + \mu_{c_{T1}}^{\mathrm{BH}} T_{it} + \mu_{c_{T2}}^{\mathrm{BH}} T_{it}^2 \tag{A.2}$$

$$\beta_{it}^{\mathrm{BH}} = d_{oi}^{\mathrm{BH}} + \mu_{d_{T1}}^{\mathrm{BH}} T_{it} + \mu_{d_{T2}}^{\mathrm{BH}} T_{it}^2\,. \tag{A.3}$$

Equation (A.1), incorporating the above relationships, is the data-level model, represented by the likelihood in the Bayesian framework:

$$y_{it} \sim \mathrm{N}(\alpha_{it}^{\mathrm{BH}} + \log(\beta_{it}^{\mathrm{BH}}) - \log(\beta_{it}^{\mathrm{BH}} + x_{it}), \sigma_{yi}^2)\,.$$

The likelihood function expresses the probability of observing the data, given the functional model and its parameters. Therefore, in terms of a joint conditional probability function for all (n) stocks, it can be written as

$$\prod_{i=1}^{n} p[y_i \mid \mu_{c_o}^{\mathrm{BH}}, \mu_{c_{T1}}^{\mathrm{BH}} T_{it}, \mu_{c_{T2}}^{\mathrm{BH}} T_{it}^2, \mu_{d_{oi}}^{\mathrm{BH}}, \mu_{d_{T1}}^{\mathrm{BH}} T_{it}, \mu_{d_{T2}}^{\mathrm{BH}} T_{it}^2]\,. \tag{A.4}$$

In the next step, stock-level models are incorporated, acting as priors for the coefficients in Equations (A.2) and (A.3). Models for the temperature-related terms allow for the possibility that *alpha* and *beta* have different degrees of sensitivity to temperature effects in the individual stocks. These across-stocks distributions are of the form

$$c_{T1i}^{\mathrm{BH}} \sim \mathrm{N}(\mu_{c_{T1}}^{\mathrm{BH}}, (\sigma_{c_{T1}}^{\mathrm{BH}})^2) = p[c_{T1i}^{\mathrm{BH}} \mid \mu_{c_{T1}}^{\mathrm{BH}}, \sigma_{c_{T1}}^{\mathrm{BH}}]\,, \tag{A.5}$$

where c_{T1i}^{BH} are assumed to be independent.

The stock-level models for the intercepts account for among-stocks differences in Beverton–Holt *alpha* and *beta* parameters arising from additional effects not included in the data-level model. For the parameter *beta*, representing K, variation is partly attributable to differences in the habitat size occupied by the individual stocks. Therefore, the intercepts in Equation (A.3) can be modelled as a function of H, and the corresponding priors become stock-specific:

$$d_{oi}^{\mathrm{BH}} \sim \mathrm{N}(\mu_{k_o} + f_{H1} H_i + f_{H2} H_i^2, (\sigma_{d_o}^{\mathrm{BH}})^2) = p[d_{oi}^{\mathrm{BH}} \mid \mu_{k_o}, f_{H1}^{\mathrm{BH}}, f_{H2}^{\mathrm{BH}}, \sigma_{d_o}^{\mathrm{BH}}]\,, \tag{A.6}$$

where d_{oi}^{BH} are assumed to be independent.

The parameters describing the prior distributions in Equations (A.5) and (A.6) are referred to as *hyperparameters* in the Bayesian framework, and the uncertainty of the latter is accounted for by the *hyperpriors*. In all cases, we use uninformative hyperpriors, imposing prior distributions N(0, 1000) on the hyperparameters.

The data- and stock-level models, together with the hyperpriors, represent the uncertainty and variability sources addressed by the hierarchical model. By expressing them as conditional probability models, they can be combined to produce the joint *posterior* distribution of the parameters in the previous levels (Berliner, 1996; Wikle, 2003):

$$\text{p[model (process and parameters) | data]} \propto \text{p [data| process, data parameters]} \\ \text{p[process | process parameters] p[parameters].} \quad \text{(A.7)}$$

The left side in the above expression is the joint posterior of all parameters. On the right side, the first conditional is the joint likelihood function (i.e. the data-level model) in Equation (A.4). The second conditional is represented by the stock-level models, such as Equations (A.5) and (A.6), and the third term corresponds to the hyperpriors.

In practice, posterior estimation in the Bayesian framework is implemented using iterative MCMC algorithms (Gilks *et al.*, 1996). The MCMC procedure produces sequences (chains) of simulations which approximate the distribution given in Equation (A.7) and, hence, capture all the levels of uncertainty in the estimation of each parameter. Also, in Bayesian inference, confidence intervals (known as *credibility intervals*) are estimated using the posterior distribution of a given term. Therefore, they have a different, more-intuitive interpretation, and describe the probability that the true value of a parameter is within a certain range.

Supplementary material references

Berliner, L. M. 1996. Hierarchical Bayesian time series models. *In* Maximum Entropy and Bayesian Methods, pp. 15–22. Ed. by K. M. Hanson, and R. N. Silver. Kluwer Academic, Dordrecht. 480 pp.

Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. 1996. Markov Chain Monte Carlo in Practice. Chapman and Hall, London. 486 pp.

Wikle, C. K. 2003. Hierarchical models in environmental science. International Statistical Review, 71: 181–199.

Table S1. Mean, maximum, and minimum K_{eq} per unit area estimates (t km^{-2}) obtained by the Ricker and BH multilevel models.

	Ricker			BH		
Cod stock	**Mean**	**Minimum**	**Maximum**	**Mean**	**Minimum**	**Maximum**
cod-2224	10.0	9.1	10.2	34.4	27.1	36.7
cod-2532	5.9	4.7	6.3	10.0	7.1	11.1
cod-347d	4.5	4.2	4.7	59.7	45.9	70.3
cod-7e-k	2.6	2.2	3.2	5.7	4.2	9.5
cod-arct	2.6	2.3	2.7	14.4	9.9	15.2
cod-coas	0.8	0.5	1.2	1.7	1.3	2.2
cod-farp	9.5	7.5	10.6	17.4	12.1	21.7
cod-iceg	22.1	20.4	22.8	72.9	54.2	81.1
cod-kat	15.3	13.5	15.8	26.0	20.8	27.9
cod2j3kl	5.5	3.2	7.3	12.5	5.4	21.3
cod3m	26.2	24.8	26.6	65.3	59.5	67.0
cod3no	2.4	0.6	3.0	4.3	2.0	5.4
cod3pn4rs	4.0	2.0	6.4	5.8	3.5	11.0
cod3ps	2.4	2.1	2.8	5.2	4.1	6.7
cod4tvn	7.4	7.0	7.6	13.7	11.9	14.5
cod4vsw	4.4	4.0	4.4	15.5	12.0	16.1
cod4x	3.7	3.6	3.7	6.9	6.7	7.0
codgb	5.0	3.7	5.6	11.6	7.2	14.4
codgom	6.6	5.9	6.7	20.3	16.2	21.0
codvia	4.6	4.2	5.3	12.4	10.1	17.2
codviia	6.7	4.4	7.4	14.6	7.9	18.0

Table S2. Mean, maximum, and minimum K_{max} per unit area estimates (t km^{-2}) obtained by the Ricker and BH multilevel models. The column labelled “Change” refers to the change in mean K_{max} induced by an increase in current mean T by 3°C.

Cod stock	Ricker			BH			Change
	Mean	**Minimum**	**Maximum**	**Mean**	**Minimum**	**Maximum**	
cod-2224	21.8	17.0	23.2	33.9	26.8	36.1	17.49%
cod-2532	7.6	5.3	8.6	9.9	7.0	11.0	26.93%
cod-347d	33.7	25.7	40.0	59.4	45.7	70.1	45.40%
cod-7e-k	3.3	2.5	5.1	5.7	4.2	9.5	44.27%
cod-arct	9.4	6.0	10.0	14.3	9.9	15.2	32.64%
cod-coas	1.0	0.7	1.3	1.7	1.3	2.2	53.71%
cod-farp	14.1	9.1	17.9	16.5	11.3	20.5	53.79%
cod-iceg	63.8	49.4	70.0	72.2	53.7	80.3	26.08%
cod-kat	19.9	16.0	21.2	25.6	20.4	27.5	20.42%
cod2j3kl	11.2	3.5	23.1	12.5	5.4	21.2	235.68%
cod3m	52.8	46.1	54.4	64.2	58.4	66.0	7.47%
cod3no	2.5	1.0	3.2	4.3	1.9	5.3	67.67%
cod3pn4rs	4.1	2.4	7.7	5.8	3.5	10.9	147.47%
cod3ps	3.3	2.6	4.4	5.1	4.0	6.5	75.63%
cod4tvn	10.4	9.2	10.9	13.5	11.8	14.3	16.20%
cod4vsw	8.3	6.7	8.5	15.3	11.8	15.9	1.40%
cod4x	4.1	3.9	4.2	6.8	6.6	6.9	8.59%
codgb	7.1	4.2	9.0	11.5	7.1	14.3	40.35%
codgom	11.7	9.0	12.1	20.1	16.0	20.7	15.17%
codvia	7.0	5.7	9.6	12.3	10.0	17.0	56.89%
codviia	9.5	5.0	11.8	14.3	7.7	17.6	39.69%

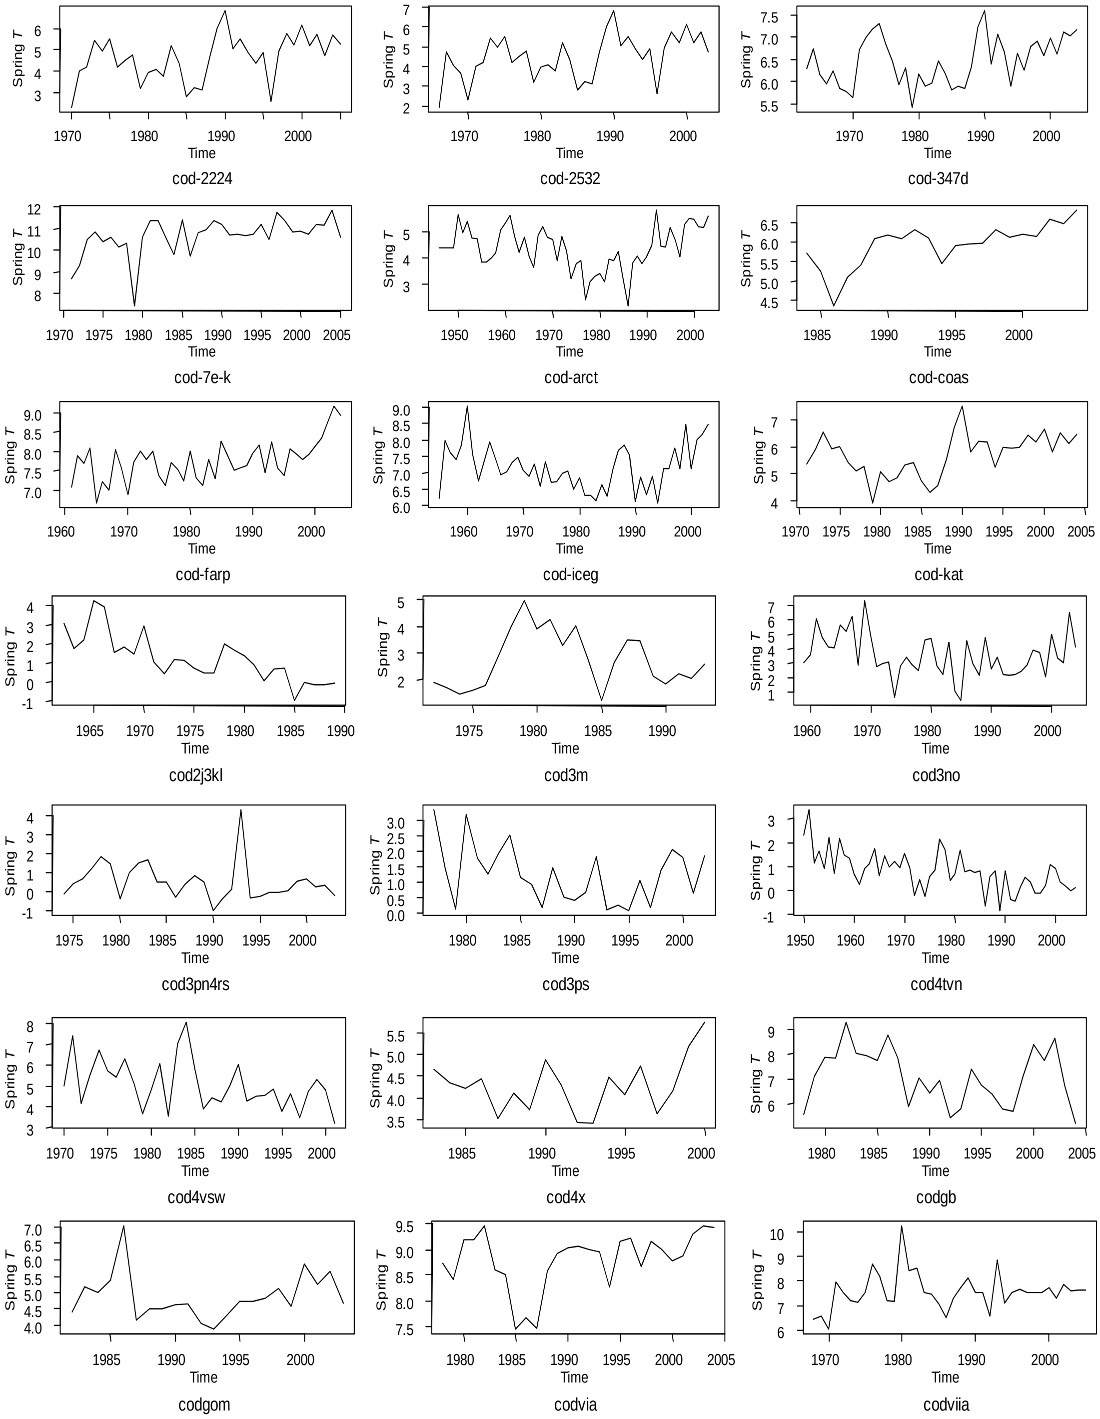

Figure S1. Time-series of the stock-specific spring (March–May) temperature (at 0–100 m) used in the model. See Table 1 in main text for the stock codes.

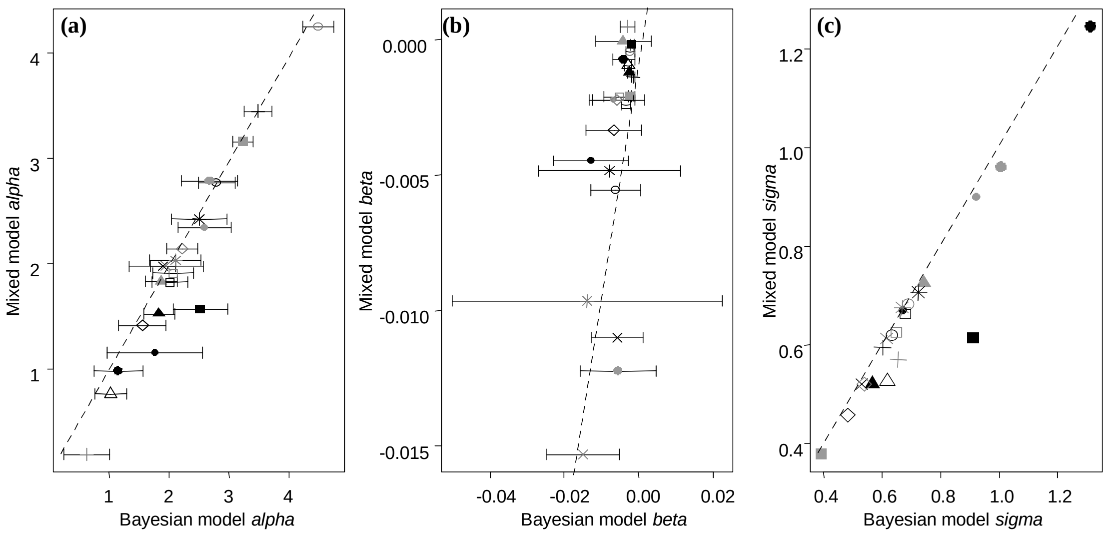

Figure S2. Comparison of the stock-specific Ricker SR model parameters (a) *alpha*, (b) $\beta_i^{\mathrm{RIC}} = -1/beta$, and (c) residual standard errors (sigma), as estimated by the Bayesian and the mixed models. The simple Ricker models with stock-specific errors, without the effects of temperature and habitat size, were fitted (MR.3 model structure in Table 2 in the main text). The horizontal solid lines correspond to the 95% credibility intervals of the Bayesian model parameters. The dashed lines are the 1 : 1 lines. See Figure 2 in the main text for stock symbol codes.

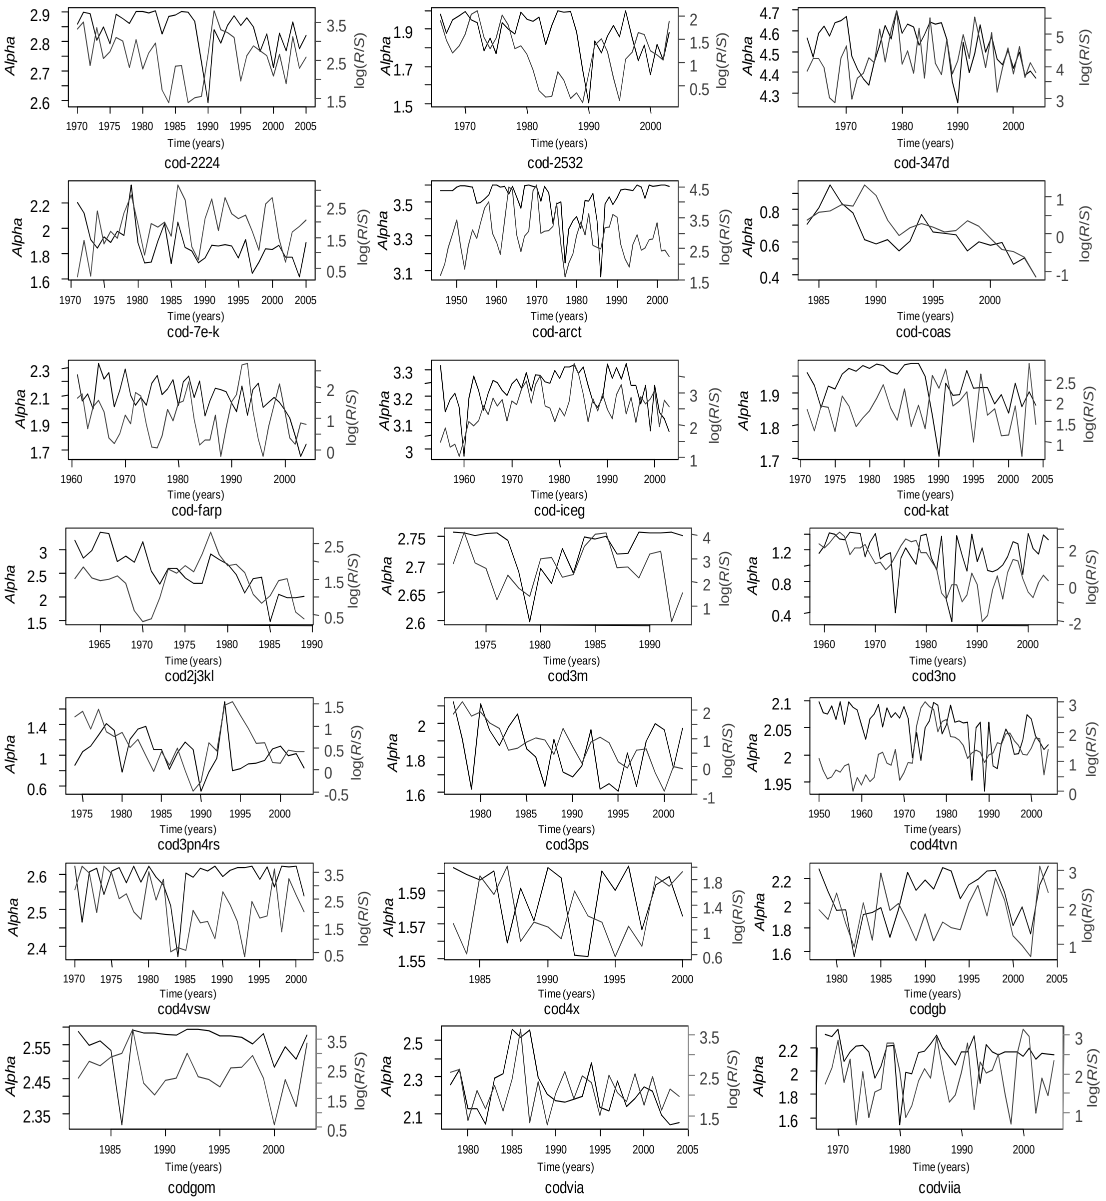

Figure S3. Fluctuations of stock-specific *alpha* values, as predicted by the Bayesian Ricker SR model (black lines, left vertical axis) and of the recruitment survival (log(*R*/*S*); grey lines, right vertical axis) in time. See Table 1 in main text for stock codes.

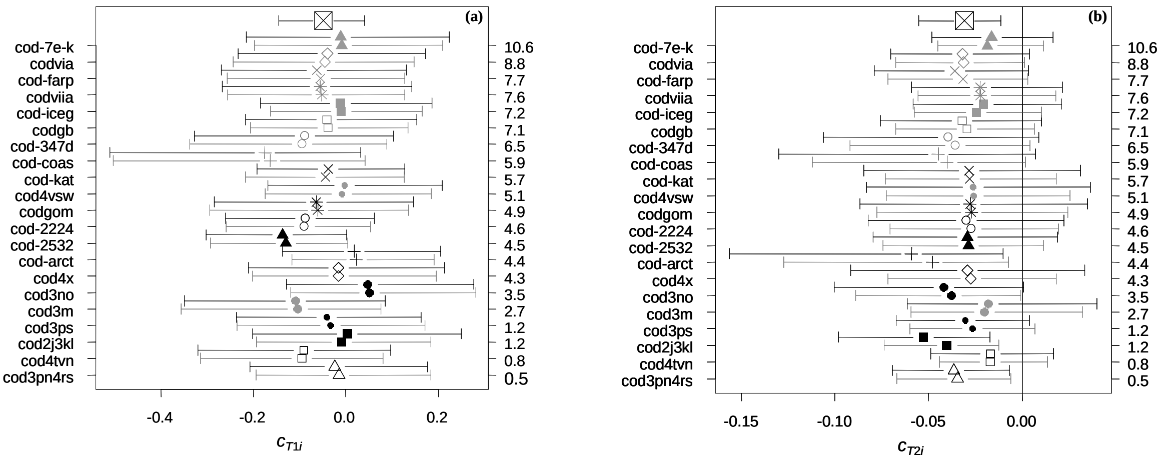

Figure S4. The *T*-related slopes (a) c_{T1i} , and (b) c_{T2i} and 95% credibility intervals obtained from the Ricker (black bars) and BH (grey bars) Bayesian SR models. See Figure 2 in main text for stock symbol codes.

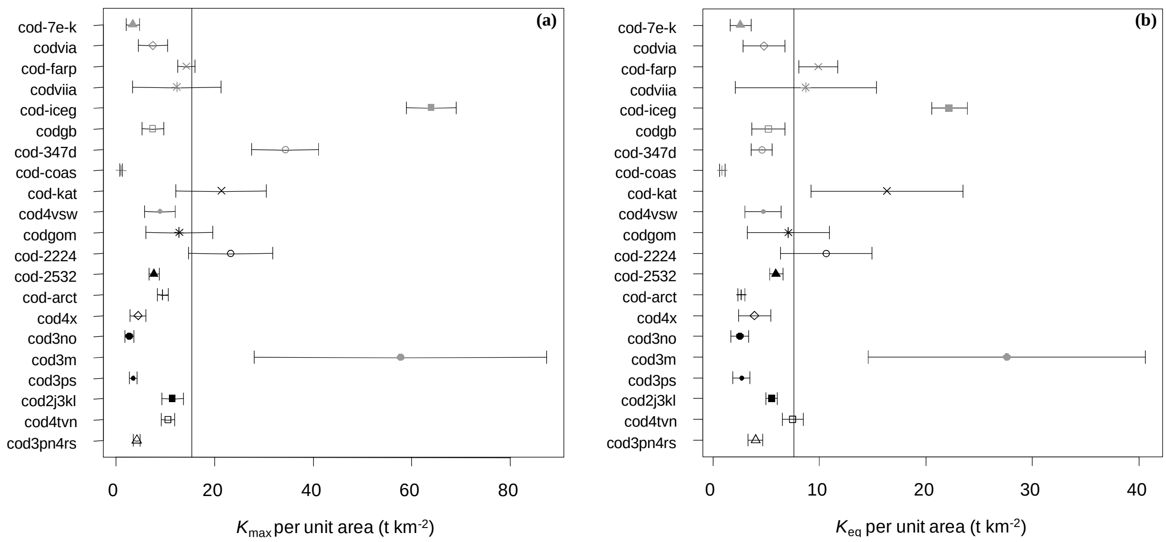

Figure S5. The stock-specific (a) K_{max} (±s.e.) and (b) K_{eq} (±s.e.), estimated as t km^{-2}, obtained from the Bayesian Ricker model ordered by increasing mean temperature (right vertical axis). See Figure 2 in main text for stock symbol codes.