Climate
of the Past

Bayesian parameter estimation and interpretation for an intermediate model of tree-ring width

S. E. Tolwinski-Ward[1], K. J. Anchukaitis[2], and M. N. Evans[3]

[1]Institute for Mathematics Applied to Geosciences, National Center for Atmospheric Research, Boulder, CO, USA
[2]Woods Hole Oceanographic Institution, Woods Hole, MA, USA
[3]Department of Geology and Earth System Science Interdisciplinary Center, University of Maryland, College Park, MD, USA

Correspondence to: S. E. Tolwinski-Ward (tolwinsk@ucar.edu)

Abstract. We present a Bayesian model for estimating the parameters of the VS-Lite forward model of tree-ring width for a particular chronology and its local climatology. The scheme also provides information about the uncertainty of the parameter estimates, as well as the model error in representing the observed proxy time series. By inferring VS-Lite's parameters independently for synthetically generated ring-width series at several hundred sites across the United States, we show that the algorithm is skillful. We also infer optimal parameter values for modeling observed ring-width data at the same network of sites. The estimated parameter values covary in physical space, and their locations in multidimensional parameter space provide insight into the dominant climatic controls on modeled tree-ring growth at each site as well as the stability of those controls. The estimation procedure is useful for forward and inverse modeling studies using VS-Lite to quantify the full range of model uncertainty stemming from its parameterization.

1   Introduction

Forward models of the physical or biological processes by which climate variability is imprinted on natural archives provide important tools for understanding such "proxies" as recorders of climate (Evans et al., 2013). The VS-Lite model (Tolwinski-Ward et al., 2011) provides one such forward model for the climate controls on tree-ring width chronologies. Under this model, just four parameters determine a simulated chronology's response to local mean monthly air temperature and monthly model-simulated soil moisture. These

parameters connect the local climatology to the modeled controls on growth and the climatic signal contained in the simulated chronology. Thus, in order to use VS-Lite to study the relationship between climate and proxies in the real world, an objective method for choosing the model parameters for any particular site or region is necessary.

Ideally, parameterization should be based on a first-principles understanding of the science represented by the model. The growth response parameters in the VS-Lite model are loosely interpretable as temperature and soil moisture thresholds above which growth begins or is no longer sensitive to climatic fluctuations, respectively. This interpretation is consistent with the well-established biological Law of the Minimum (Taylor, 1934), expressed in dendrochronology by the idea that tree growth is determined by the most limiting factor (Fritts, 2001). However, it is debatable whether each of these VS-Lite parameters has directly measurable analogs in the natural world. Recent physiological studies have advanced the scientific community's understanding of climatic thresholds for xylogenesis (e.g., Rossi et al., 2007; Deslauriers et al., 2008), but the results are not well-enough developed or cataloged across tree species or climate regimes to be generalizable to all forward model simulations. Even if they were, it is not clear how measurable quantities in the real world are related to vastly simplified model quantities. In particular, VS-Lite operates using monthly data, while the cellular-level processes it is intended to mimic vary at daily and shorter timescales. Minimum and optimum parameters therefore represent a simplification that may be inferred from, but may not be strictly interpretable as, biophysical limits on cambial activity itself. Given the

limited ability of direct observations to constrain the model parameters, it is necessary to estimate VS-Lite's parameters numerically using monthly climate inputs, observed ring-width series, and partial knowledge of the model error structure. At sites where VS-Lite is believed to provide a reasonable intermediate complexity proxy system model for tree-ring width variations, such an estimation procedure can also be used to estimate the parameters of the error model arising from the model's incomplete representation of the proxy system.

Three general approaches to objective numerical parameter estimation of forward models of tree-ring growth have been explored in the literature. The first is presented by iterative local schemes that optimize the fit of simulated model quantities to their observed counterparts under changing parameter combinations. In the iterative scheme of Fritts et al. (1999), for example, one model parameter is changed at a time in a search over a continuous region of parameter of space for the set of parameters producing optimal model fit. Modern computing power makes such schemes possible to run at many locations, but the approach does not account for potential interactions between parameters. The scheme may also locate a local optimum in parameter space but miss information, including other optima, provided by a more global search. A second class of approaches attempts to avoid this last pitfall by running growth models at a preselected, discretized set of parameter combinations believed to cover the entire physically plausible regions of parameter space, and the single combination that produces the best match of modeled outputs to its observed counterparts is deemed optimal (e.g., Misson et al., 2004; Misson, 2004; Tolwinski-Ward et al., 2011). This parameterization method is only practical where the number of parameters to be constrained by the available data is relatively small. Both local and global previous parameterization schemes have the shortcoming that they provide only point estimates of optimal parameters, and their results do not include any information about model sensitivity to the parameter choices. In addition, while the ranges of parameters included in the search space are generally chosen through consideration of physically plausible bounds on their values, the search algorithms lack any more sophisticated use of science-based understanding of where the most likely parameter values lie. Bayesian modeling and inference represents a third approach in which conditional probability distributions of the unknown parameters are inferred given the information in model input climate data and observed series of tree-ring width data. Because the product of such an analysis is a probability distribution, the Bayesian approach also has the advantage of automatically providing uncertainty and sensitivity estimates along with parameter estimates. The approach also automatically safeguards against over-tuning when the calibration data series are short and/or noisy. Such an approach was explored by Gaucherel et al. (2008) to infer the parameters of a complex, biology-resolving model of tree-ring growth. Although not a forward model of tree-ring width given climate, work by Boreux et al. (2009) also employed a Bayesian approach to estimate a latent regional tree-ring width signal and the parameters of an accompanying hierarchical model to interpret several neighboring tree-ring width chronologies.

Here we present and test a Bayesian statistical model to infer parameter estimates for VS-Lite for simulating a particular chronology from co-located climatic inputs. Our scheme is efficient enough to run in a matter of minutes in a modern laptop computing environment, and returns probabilistic posteriors that can be analyzed to infer uncertainty and sensitivity information as described above. In addition, the scheme takes particular advantage of the capability of the Bayesian framework to assimilate expert scientific prior knowledge into the inference. Version 2.3 code for the scheme is freely available with the VS-Lite v2.3 model at the National Oceanic and Atmospheric Administration (NOAA) Paleoclimatology software library (http://www.ncdc.noaa.gov/paleo/softlib/softlib.html). We test the skill of the parameterization approach in several hundred independent idealized experiments using synthetically generated tree-ring width data. As an application of the method to observed ring-width chronologies, we also independently estimate parameters for several hundred observed chronologies across the continental United States and present a graphical method for interpreting the fitted model parameters in terms of climatic controls on tree-ring growth at each site.

2 Model, data, and methods

While climate is a spatiotemporal phenomenon, and networks of tree-ring width can be used to make inferences about the past space–time variations in climatic fields, trees only experience and respond to their local environments. This is reflected in the structure of VS-Lite, which models tree-ring width variations as a function of local climate. Here we take the view that the model provides a first-order representation of the general mechanisms by which climatic variability is imprinted on all climatically sensitive trees. Within this framework, VS-Lite's growth response parameters and the parameters of an accompanying error model adjust the response of VS-Lite to fit better the observed growth responses that vary from site to site depending on species, specific local environmental conditions, and the varying sensitivity of the signal in individual tree-ring width series that compose a given chronology.

2.1 Summary of VS-Lite and parameters

A complete description of the VS-Lite model is given by Tolwinski-Ward et al. (2011). However, we briefly summarize the basic structure of the model and its parameterization here. VS-Lite is a substantially simplified, monthly time-step version of the full Vaganov–Shashkin model of tree-ring

growth (Vaganov et al., 2006, 2011). At its core, VS-Lite is a parsimonious representation of the Principle of Limiting Factors with respect to local monthly temperature and soil moisture, and with growth modulated by local insolation. In its current version (version 2.2), insolation is determined from site latitude, and soil moisture is determined from local monthly temperature and precipitation via a simple bucket model (Huang et al., 1996). Non-dimensional scaled growth responses $g_T(m, y)$ and $g_M(m, y)$ to monthly time-step temperature and soil moisture content, respectively, are key to determining the extent of simulated growth in each modeled month m and year y. These responses have the piecewise linear forms

$$
g_T(m, y) = \begin{cases} 0 & T(m, y) \leq T_1; \\ \frac{T(m,y)-T_1}{T_2-T_1} & T_1 \leq T(m, y) \leq T_2; \\ 1 & T_2 \leq T(m, y) \end{cases} \quad (1)
$$

and

$$
g_M(m, y) = \begin{cases} 0 & M(m, y) \leq M_1; \\ \frac{M(m,y)-M_1}{M_2-M_1} & M_1 \leq M(m, y) \leq M_2; \\ 1 & M_2 \leq M(m, y) \end{cases} \quad (2)
$$

The parameters T_1 and M_1 thus represent thresholds in temperature and soil moisture content below which growth cannot occur, while T_2 and M_2 are thresholds above which growth is insensitive to climatic variability. The overall monthly growth rate is given by $g(m, y) = \min\{g_T(m, y), g_M(m, y)\}$, to mimic the Principle of Limiting Factors (Fritts, 2001) that the more limiting environmental variable controls growth. The simulated annual-resolution ring-width series result from taking an inner product of these overall growth rates with estimates of mean relative monthly insolation derived from trigonometric functions of latitude. Thus, the climatic variable that tends to produce lesser values in its growth response function controls the modeled climate signal contained in the simulated proxy series. The relationship of the parameter values T_1, T_2, M_1, and M_2 to the model's climate inputs is therefore critical in determining which variable gets "recorded" by the synthetic trees.

We denote by $\hat{\mathbf{W}}_{\text{VSL}}$, the deterministic VS-Lite estimate of a tree-ring width series given fixed temperature and precipitation inputs and fixed VS-Lite growth response parameters $\boldsymbol{\theta}_{\text{VSL}} = (T_1, T_2, M_1, M_2)'$. Formally, we model the zero-mean, unit-variance time series of ring-width data $\mathbf{W}$ at a location of interest by a scaling of the VS-Lite output plus stochastic noise to account for non-climatic noise in the data and for the effects on the time series of processes that VS-Lite does not resolve:

$$
\mathbf{W} = \sqrt{1 - \sigma_W^2}\, \hat{\mathbf{W}}_{\text{VSL}} + \mathbf{e} \quad (3)
$$
$$
\text{with } \sigma_W^2 = \text{Var}(e_t) \quad .
$$

We allow for the possibility that the time series of errors $\mathbf{e}$ may follow either a white noise model, so that $e_t \sim$

$N(0, \sigma_W^2)$, or else an AR(1) model, $e_t \sim N(\phi_1 e_{t-1}, \tau^2)$, with the conditions for stationarity for the noise process imposing $\sigma_W^2 = \frac{\tau^2}{1-\phi_1^2}$ in the latter case (Shumway and Stoffer, 2006). Since both the time series of data and of VS-Lite output are standardized to have zero mean and unit variance, the coefficient $\sqrt{1 - \sigma_W^2}$ on the estimate from VS-Lite gives the proportion of the observed data's standard deviation that can be explained by VS-Lite as "signal", and so the signal-to-noise ratio is $\text{SNR} = \sqrt{\frac{1-\sigma_W^2}{\sigma_W^2}}$. We write the parameters of the error model as $\boldsymbol{\theta}_e = \sigma_W^2$ for white noise and $\boldsymbol{\theta}_e = (\phi_1, \tau^2)^T$ for AR(1) errors. We treat these as additional model parameters to be estimated at any given site, whose values provide information about the degree to which VS-Lite can be used to explain the observed variations in the data.

2.2 Approach to model parameter estimation

We follow a Bayesian approach to estimating the model parameters at a particular location. Let $\boldsymbol{\theta}$ denote the vector of model parameters we would like to estimate, and $\mathbf{W}(\mathbf{T}, \mathbf{P})$ the vector of observed ring-width data, which depends on vectors of monthly temperature and precipitation data covering the same interval in time as the ring width data. The Bayesian paradigm allows inference on the posterior distribution $\pi(\boldsymbol{\theta}|\mathbf{W}(\mathbf{T}, \mathbf{P}))$ of the parameters given the climate and ring width data in terms of the likelihood $f(\mathbf{W}(\mathbf{T}, \mathbf{P})|\boldsymbol{\theta})$ of the ring-width data given the climate and the parameters, as well as a prior distribution $\pi(\boldsymbol{\theta})$ on the parameters via Bayes' law:

$$
\pi(\boldsymbol{\theta}|\mathbf{W}(\mathbf{T}, \mathbf{P})) \propto f(\mathbf{W}(\mathbf{T}, \mathbf{P})|\boldsymbol{\theta})\pi(\boldsymbol{\theta}). \quad (4)
$$

Given the likelihood and prior parameter models, Markov chain Monte Carlo techniques produce an ensemble of draws from the posterior distribution (Gilks et al., 1996), from which estimates of the parameters and their associated uncertainties can be made.

In the present setting, the parameters $\boldsymbol{\theta}$ consist of the set $\boldsymbol{\theta}_{\text{VSL}} = (T_1, T_2, M_1, M_2)^T$ used to compute the deterministic response of VS-Lite to input climate data, and the set $\boldsymbol{\theta}_e$ that describe the VS-Lite model error structure, so that $\boldsymbol{\theta} = (T_1, T_2, M_1, M_2, \boldsymbol{\theta}_e)^T$. Given the data-level model in Eq. (3), the likelihood in Eq. (4) may be thus be written as

$$
f(\mathbf{W}(\mathbf{T}, \mathbf{P})|\boldsymbol{\theta}) \propto \frac{1}{|\Sigma_e|^{1/2}}\exp \quad (5)
$$
$$
\left(-\frac{1}{2}(W - \sqrt{1 - \sigma_W^2}\,\hat{\mathbf{W}}_{\text{VSL}})^T \Sigma_e^{-1} (W - \sqrt{1 - \sigma_W^2}\,\hat{\mathbf{W}}_{\text{VSL}})\right).
$$

Note that the dependence on $\boldsymbol{\theta}_{\text{VSL}}$ is implicit in the estimate $\hat{\mathbf{W}}_{\text{VSL}}$, which is computed for a particular set of growth response parameters. The dependence on the error model parameters $\boldsymbol{\theta}_e$ is through the covariance matrix Σ_e and the noise process variance σ_W^2. The specific form of these depends on

whether a white noise or AR(1) noise model is appropriate at a particular site; these details are discussed in the supplementary material. Note that in its current version, VS-Lite also requires several parameters of the Leaky Bucket model of soil moisture (Huang et al., 1996). We do not estimate those here, as the soil moisture model may be viewed as an ancillary component of VS-Lite that may be replaced by a more sophisticated hydrological model or direct measurements of soil moisture. In effect, our current approach transfers the uncertainty associated with these parameters to uncertainty in the soil moisture response parameters M_1 and M_2.

In modeling the prior distribution of the VS-Lite growth response parameters, we first make the assumption that each parameter is independent of the others. This assumption allows us to model their joint prior distribution as the product of individual prior models for each. We put relatively broad but informative priors on the growth response parameters, with shapes and supports consistent with current scientific understanding of tree growth responses to temperature and moisture. While the original full VS model was built around conifer physiology, VS-Lite is a generalized implementation of the principle of limiting factors and piecewise growth functions, and should be a valid representation of growth in both gymnosperms and angiosperms. Although it is conceivable that different priors might be constructed for different tree species, it is difficult to disentangle the influence of species on growth response from the influences of specific site characteristics, non-climatic influences, and random effects, and no comprehensive identification of the differences in growth response exists for the dozens of species used for dendrochronology. We thus construct species-independent priors using general information about growth across all climate-sensitive trees.

Of the four growth response parameters, the literature provides the most information about T_1, the threshold temperature for growth to begin. The physiological experiments of Körner and Hoch (2007) at a montane site in Switzerland indicated that mean seasonal soil temperatures below 6–7 °C would not permit growth. An assessment of root and air temperatures at a few dozen tree-line sites by Körner and Paulsen (2004) gave a value of 6.7 ± 0.8 °C for this growth threshold, and histological measurements and analyses of Rossi et al. (2007) and Deslauriers et al. (2008) for conifers in the Alps gave a range of 5.8–8.5 °C. Hoch and Körner (2009) found that two montane conifer species maintained cambial activity even when grown at 6 °C. Körner (2012) inferred a global mean tree-line isotherm near 6 °C and cessation of growth at 5 °C (Körner, 2008); 0 °C is the theoretical limit below which plant tissue formation cannot occur (Körner, 2012). We thus chose to model the temperature threshold for growth by $T_1 \sim \beta(9, 5, 0, 9)$, a four-parameter beta distribution with shape parameters $\alpha = 9$ and $\beta = 5$ supported on the interval [0, 9]. This choice puts the mode of the probability density function at 6 °C, assigns zero probability below freezing (0 °C) or above 9 °C, and places 90 % of the total

probability in the interval (3.8 °C, 7.5 °C) (see blue curve in Fig. 1a and e).

The biologically based information available about T_2, the threshold above which growth is no longer sensitive to temperature variations, is more uncertain. Vaganov et al. (2006) give a default value of 18 °C for the full Vaganov–Shashkin model based on a few intensive case studies at a limited number of Russian tree-ring sites, but use a value of 15 °C in an example model run, demonstrating the range of uncertainty associated with this parameter. The analogous parameter in the TreeRing2000 model has a default value of 23 °C (Fritts et al., 1999). Data shown by Williams et al. (2011) suggest a broad plateau where ring width in Alaskan *Picea glauca* ceased increasing with June and July mean temperatures between approximately 10 and 13 °C, depending on site hydrology. On the other hand, Garfinkel and Brubaker (1980) showed no change in the regression of ring width on temperature in the same species even at temperatures approaching 15 °C. Carrer et al. (1998) inferred a lower optimal summer temperature threshold of 13 °C for *Picea abies* and 16 °C for *Larix decidua*. Although this information sheds some light on the threshold for sensitivity to temperature, the majority of these studies are based on empirical data at monthly to seasonal timescales, as opposed to direct studies of cambial activity in response to temperature. To reflect the uncertainty inherent in the wide range of these estimates, as well as uncertainty in their direct applicability as parameter estimates in VS-Lite, we model the prior by $T_2 \sim \beta(3.5, 3.5, 10, 24)$. This choice limits probability mass to the interval (10 °C, 24 °C), and distributes probability symmetrically about a mean of 17 °C with a standard deviation of 2.5 °C, and 90 % of the total probability in the interval (12.9 °C, 21.1 °C) (blue curve, Fig. 1b and f).

Very little biologically based information is available to constrain either moisture parameter. We use the default parameters developed by Vaganov et al. (2006) to define broad priors on M_1 and M_2. Default values for M_1, interpretable as soil wilting point, are 0.02 v/v (Vaganov et al., 2006) and 0.01 v/v (Fritts et al., 1999). The latter source also sets a moisture optimum at 0.109 v/v, so the value of M_1 should certainly fall well below this value. We set the prior mean at 0.035 with standard deviation of 0.02 v/v, with no probability mass outside of (0 v/v, 0.1 v/v), by letting $M_1 \sim \beta(1.5, 2.8, 0, 0.1)$, with the interval (0.006 v/v, 0.073 v/v) containing 90 % of the prior probability (blue curve, Fig. 1c and g). The default for M_2 is 0.8 of typical soil saturation levels, and the Leaky Bucket model of soil moisture employed by VS-Lite never allows soil to be saturated to a value more than 0.75 v/v (Huang et al., 1996). We set $M_2 \sim \beta(1.5, 2.5, 0.1, 0.5)$. This gives the prior a mean of 0.25 v/v, a standard deviation of 0.1 v/v, nonzero probability on (0.1 v/v, 0.5 v/v), and 90 % of the prior probability mass in (0.125 v/v, 0.406 v/v) (blue curve, Fig. 1d and h).

In assigning priors to the parameters $\boldsymbol{\theta}_e$, we enforce stationarity on the time series of errors, require that the

Prior (blue) and posterior (red) densities of VS–Lite parameters, Pseudoproxy Experiments
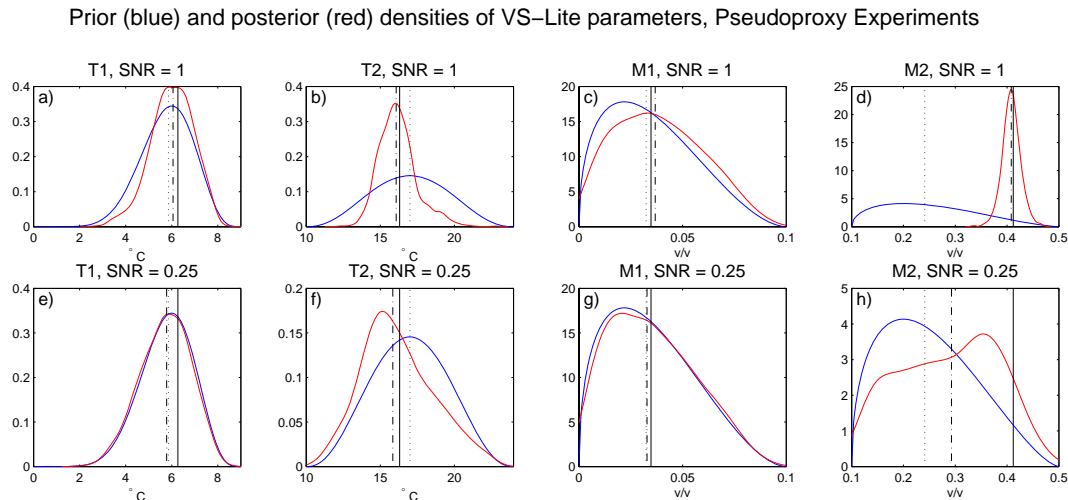

Fig. 1. (a–d) Prior (blue) and estimated posterior (red) densities for the parameters of VS-Lite at the Sipsey Wilderness site in Alabama, for the pseudoproxy experiment with SNR = 1. Plot of posterior density given by a kernel smoothing of the frequency distribution of ensemble members. Solid black vertical lines give the target values of the growth response parameters, dotted black lines prior medians, and dash-dot black lines the posterior medians. **(e–h)** As in **(a–d)**, but for PPE with SNR = 0.25.

magnitude of the process be no greater than the unit-variance observations, and require that the lag-1 autocorrelation be non-negative, but do not supply the priors with any other information. For the white-noise model, these criteria result in a prior $\sigma_W^2 \sim U(0, 1)$. For the AR(1) model, they imply a normalized indicator function on the set $\{\phi_1, \tau^2 : \phi_1 \in (0, 1), \tau^2 < (1 - \phi_1^2)\}$ (see, e.g., Shumway and Stoffer, 2006). At each site the parameter estimation procedure is first run assuming white errors. The residuals of simulated ring-width index using the posterior median growth response parameters are then fit with a white noise error model and an AR(1) model. If the AR(1) model has a lesser value of Schwarz's Bayesian criterion (Robert, 2007, Sect. 7.2.3), then the parameter calibration is re-run for that site under the assumption of AR(1) model errors.

The posterior distribution (Eq. 4) is sampled using a Metropolis–Hastings algorithm embedded within a Gibbs sampler, which is a standard Markov chain Monte Carlo (MCMC) approach (Chib and Greenberg, 1996). To check for convergence, we run three chains with 30 000 iterations each after a burn-in period of 300 iterations. In the rare case that the R-hat statistic (Gelman and Rubin, 1992) indicates the MCMC has not converged, we re-run the sampler with a greater number of iterations until the sampler has converged. The autocorrelation functions of the MCMC chains indicate that, at most sites, autocorrelation in the parameter sampling chains is no longer significant past a lag of 20. We conservatively subsample every 50th value of each of the 3 chains to ensure independence of samples, resulting in a collection of 1800 samples for each parameter value at each site.

The ensemble output may be used in several different ways. First, point estimates from the posterior ensemble,

such as the posterior median or the maximum likelihood parameter set, may be used as calibrated parameter values that optimize the fit between model-simulated ring width data and a target ring width series, given fixed input climate data. However, the posteriors contain additional information beyond point estimates. Their spread indicates the uncertainty in the parameter estimates, as well as the degree to which the climate and target ring-width data inform the parameter values. Hence a measure of the model sensitivity to each parameter may also be gleaned from the posterior spread. The Monte Carlo ensemble of parameter values may also be used to run modeling studies where accounting for the effect of parameter uncertainty is important for interpretation of the results.

2.3 Experimental design

We perform our study using estimates of monthly temperature and precipitation from the 4 km × 4 km resolved gridded Parameter-elevation Regressions on Independent Slopes Model (PRISM) data product (Daly et al., 2008). We use the mean of the monthly maximum and minimum temperature fields as well as the accumulated precipitation field, and neglect the inherent PRISM measurement and model error. The climate product is used at the grid cells co-located with 277 sites associated with observed tree-ring width chronologies across the continental United States. These sites form the subset of chronologies used in a multi-proxy hemispheric temperature reconstruction by Mann et al. (2008) that also overlap with the PRISM data from 1895 to 1984. The choice of this 90 yr interval represents a balance between the availability of proxy observations and climate data. All the associated chronologies are freely available online on the NOAA

Paleoclimate Reconstructions Network/Proxy Data webpage (http://www.ncdc.noaa.gov/paleo/pubs/pcn/pcn-proxy.html).

2.3.1 Pseudoproxy experiment (PPE)

To evaluate the skill of the Bayesian parameter estimates, we perform two so-called "pseudoproxy experiments" (PPEs) (Smerdon, 2012) using synthetically generated tree-ring width data. At each site, we first perform a preliminary run of the Bayesian scheme described above using the observed chronologies and PRISM-derived climate data for the interval 1895–1984. The sampling scheme is run to convergence with a white noise model at each site, but at sites where the residuals between VS-Lite estimates and the observed data display temporal autocorrelation, the sampling scheme is re-run with the AR(1) noise model. The posterior medians and the estimated noise models from this step parameterize regionally realistic tree growth responses to climate, and comprise a set of known PPE parameter targets that we try to recover to test our methodology. We run the VS-Lite model over the same interval using this target parameter set and the PRISM climate estimates to produce 277 synthetic ring-width signals. We scale the simulated signal and add noise according to Eq. (3) such that the signal-to-noise ratio (SNR) is 1 in the first PPE, and is SNR = 0.25 in the second PPE. These values represent optimistic and pessimistic estimates of the SNRs of real-world proxies, respectively (Smerdon, 2012), and the resulting 277 time series constitute the synthetic data set we use to infer the known parameter values using the Bayesian scheme. In this inference step of the PPEs, we condition on the known climate data and pseudoproxy ring-width data over the entire interval to estimate the growth response parameters and error model parameters. The form of the error model (white or AR(1)) at each site is also assumed known, rather than inferred. The PPEs are designed as a test of our parameter calibration scheme over a realistic range of tree responses to climate, but in an idealized model world where both the signal-formation and noise processes are perfectly known and represented except for the value of their parameters. The ability of the parameter inference procedure to recover the known target parameters in the PPEs thus provides an upper bound on the skill for the procedure in real-world scenarios.

Comparing the growth parameter posterior distributions to the known targets allows us to quantify the skill of the estimation scheme. The numerics return $N = 1800$ draws from the posterior distribution, so we compute Monte Carlo estimates of the root-mean-square error and bias in the pseudoproxy context using the "true" target parameter values θ:

$$\text{RMSE} \approx \sqrt{\frac{1}{N}\sum_{n=1}^{N}(\hat{\theta}_n - \theta)^2}; \tag{6}$$

$$\text{Bias} \approx \frac{1}{N}\sum_{n=1}^{N}(\hat{\theta}_n - \theta). \tag{7}$$

The extent to which the prior and posterior distributions differ indicates the degree to which the climate, ring-width data, and the VS-Lite model structure constrain the value of each parameter. To quantify this "Bayesian learning" at each site, we examine the ratio of posterior to prior variance. Parameters whose posteriors are well-constrained by the data will have much smaller posterior variance than prior variance, while parameters that are not well-constrained will have posteriors that resemble their priors, and hence variance ratios close to one. Finally, we compute the estimated signal-to-noise ratios by the posterior median of $\sqrt{\frac{1-\sigma_{\hat{W}}^2}{\sigma_{\hat{W}}^2}}$, and compare to the known SNR used to construct the synthetic data across the 277 sites.

2.3.2 Observed proxy model calibration (OPMC)

We also calibrate the VS-Lite model parameters independently at each of the 277 network sites to perform an observed proxy model calibration (referred to hereafter by the acronym OPMC). The estimates of the growth response parameters are conditioned on the PRISM-derived climate series and observed ring-width index series for the 45 yr interval 1940–1984, and the parameters of the error model are estimated using the data in the complementary 45 yr interval 1895–1939. Split calibration/validation intervals help prevent calibration of artificial skill when selecting parameters (Cook et al., 1999; Evans et al., 2002). We also perform the analysis with the growth parameter and error parameter calibration intervals reversed to gauge the dependence of the calibration on the choice of calibration interval.

We take the posterior medians as point estimates of each parameter, and look at their spatial distribution across the experimental domain. As in the PPEs, we also look at the ratio of posterior to prior variance to assess Bayesian learning at each site. We also examine the estimated signal-to-noise ratio of the chronology at each site, and compare both the reduction of error (RE) statistic and correlation with observations for VS-Lite simulations computed with calibrated parameters and simulations computed with parameters held at their prior medians. In the case of parameter calibration using real data, the parameters are unknown, and so we cannot compute RMSE or bias. Instead, we seek to interpret the fitted parameters in terms of the climate controls on growth at each site. We first classify each site as having growth that is either temperature-limited, moisture-limited, or as having mixed climatic controls. To do so, we run the VS-Lite model at each site with the parameters' posterior medians and examine the growth response functions during June, July, and August, when insolation is at its peak and allows the bulk of modeled growth to occur. We compute the proportion of summer

months over the entire simulation in which the growth response to soil moisture (temperature) is strictly less than the growth response to temperature (soil moisture). If the modeled proportion is significantly more than the null hypothesis of half, then the site is classified as M-limited (T-limited). Sites for which the proportion cannot be statistically distinguished from 0.5 are classified as mixed-control sites. We then examine the structure of the parameter point estimates in multi-dimensional parameter space for each class of sites.

To evaluate the sensitivity of the parameter estimation scheme to the choice of prior distributions, the observed proxy model calibration described above is also performed with uniform prior distributions with the same supports as those for the literature-informed four-parameter beta priors described in Sect. 2.2. The uniform prior is a standard non-informative choice against which to check the sensitivity of posterior results to more complicated priors (see Gelman et al., 2003, Sect. 6.8). The posteriors derived under the four-parameter beta priors informed by the literature are compared with those derived using the uniform priors.

3 Results

For a single site, the parameter estimation procedure described here can be run on a MacBook Pro with 2.7 GHz processor in under 3.5 min for three chains of length 30 300 in series when the white noise model is used, and in close to 9 min when the AR(1) error model is used.

For the experiments performed here, trace plots of the MCMC chains and R^2 statistics close to 1 both indicated that the samples had converged adequately (see Supplement Fig. 1 for trace plots at a representative site). The procedure returns 106 out of 277 sites with a white noise error model, and 171 sites with an AR(1) error model when run using the observed data. At two representative sites, one using the white noise error model and one using the AR(1) error model, the reproducibility of the point estimates was checked by running the estimation procedure twice for the PPE with SNR = 1 and for the OPMC. In both cases, the difference between estimates of the two temperature thresholds T_1 and T_2 from two runs of the algorithm was less than a tenth of 1 °C; moisture thresholds M_1 and M_2 were reproduced to within 0.004 v / v (so the percentage change in the value of the point estimates of the growth parameters is less than 1 % except for M_1, which is reproduced to within 6 %). The ratio of posterior to prior variance was reproduced in all cases to within 0.05 (percent change of less than 7 % except for M_1, to within 11 %), and the estimated signal-to-noise ratios were within 0.02 of one another (reproduced to less than 0.2 %) for the two runs (further details are in the Supplementary Tables 1 and 2).
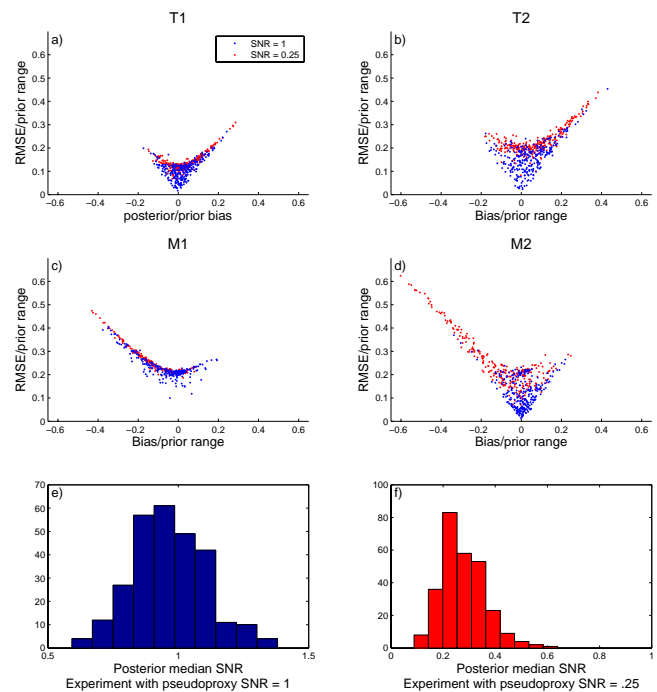

Fig. 2. (a–d) root-mean-squared error versus bias of parameter estimates in the pseudoproxy experiment, with both statistics shown relative to the length of each prior's support for the PPE with SNR = 0.25 (red) and SNR = 1 (blue). Note that the structure in the scatter plots is a result of the fact that by definition $\mathrm{RMSE}(X) = \sqrt{\mathrm{Var}(X) + \mathrm{Bias}^2(X)}$. **(e)** Histogram of median posterior estimated signal-to-noise ratio (SNR) across all 277 sites, for the pseudoproxy experiment where the true SNR = 1. **(f)** As in **(e)** but for the experiment with SNR = 0.25.

3.1 Pseudoproxy experiment (PPE) results

A plot of prior and smoothed posterior distributions of the four growth response parameters for one representative site in the network is shown in Fig. 1 for both PPEs. The variation in learning between parameters is evident in this figure, as conditioning the data tends to inform values of M_2 and T_2 to a greater extent at this particular site than the lower thresholds T_1 and M_1. Comparing the shape of the posteriors between pseudoproxy experiments also shows that the data inform the estimation to a greater (lesser) degree when the data have higher (lower) signal-to-noise ratio. A set of such posterior distributions exists for every site in the experimental network, and we compute statistics on these distributions to assess the skill of the parameter calibration method.

Both posterior bias and root-mean-squared error tend to be on the order of 20 % or less of the length of the prior interval for estimates of the parameters (Fig. 2a, b, d). The negative bias for the parameter M_2 in the experiment with the lower signal-to-noise ratio is an exception, and is on the order of 60 % of the length of the prior interval for the sites with the worst bias. The decrease in bias for the higher
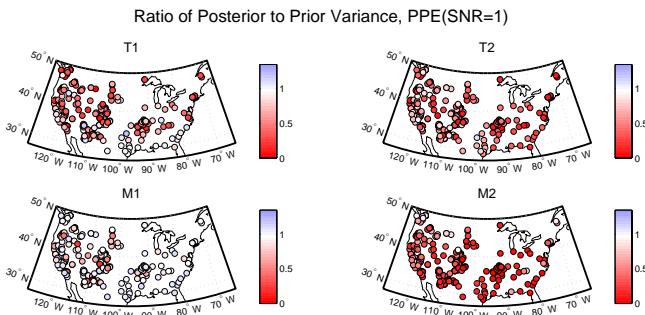
Fig. 3. Ratio of posterior to prior variance for the four growth response parameters as a measure of Bayesian learning in the pseudoproxy experiment with SNR = 1. The color scale is calibrated so that sites with values of the ratio greater than/equal to one/less than one have blue/white/red coloration.
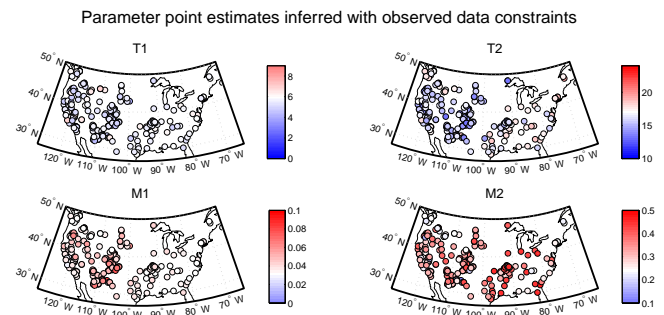

Fig. 4. Posterior medians of VS-Lite growth parameters. Note that the color scale for each parameter ranges over the interval on which the prior is supported, and is calibrated so that white indicates the prior median.

SNR experiment suggests that the bias is an artifact of the choice of prior distribution being centered at lower values than those that seem to fit the data best. Figure 2e and f show histograms of the posterior median values of the estimated signal-to-noise ratios in both experiments across the 277 sites. The individual estimates of SNR vary from site to site because each posterior is conditioned on only a single realization of the noisy synthetic ring-width data. Across the independent posteriors at all sites, however, the distributions of the estimates are centered on the correct SNR values for the respective experiments.

The ratio of posterior to prior variance is shown in Fig. 3 for the PPE with unit SNR. In general, "Bayesian learning" tends to be high for the upper thresholds T_2 and M_2, as evidenced by ratios less than one at most sites across the network. This result is indicative of high model sensitivity to the value of these parameters. By contrast, more sites have variance ratios close to one for the parameter T_1, especially in the southern United States and along the southeastern seaboard of the United States, and most sites across the network have variance ratios near one for the parameter M_1. Thus the marginal posteriors of T_1 and M_1 at these sites are similar in spread to the priors, indicating model insensitivity to their values. The variance ratios are higher across all variables and sites for the PPE with lower signal-to-noise ratio (Supplement Fig. 2), which is consistent with the general expectation that noisier data cannot constrain the solution as well as less noisy data.

3.2 Observed proxy model calibration (OPMC) results

The point estimates of the VS-Lite parameters in the observed proxy model calibration experiments, given by the posterior medians, show some spatial structure (Fig. 4). Although the varying network density makes it difficult to distinguish geographical patterns, parameters do tend to take on values close to those of their nearest neighbors. In particular, the estimated values of T_2 tend to occupy the lower end of
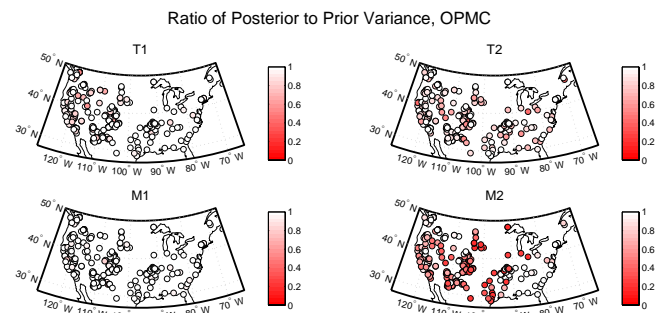
Fig. 5. Ratio of posterior to prior variance for the four growth response parameters as a measure of Bayesian learning in the observed proxy experiment. The color scale is calibrated so that sites with smaller (larger) values of the ratio, indicating greater (lesser) Bayesian learning, have darker (lighter) coloration.

the prior support in the west of the United States, with more variation between sites east of the Continental Divide. Values of M_1 also tend to fall near the upper end of the prior for arid and semi-arid sites in the west. Point estimates of M_2 are close to the upper end of the prior range at most sites across space, and preferred values of T_1 tend to be close to the prior median but show some variation for sites within close range of one another.

The spatial distribution of Bayesian learning, parameterized by the ratio of posterior to prior variance, is highest for the parameter M_2 away from the eastern seaboard and northernmost Pacific Northwest sites (Fig. 5). This result indicates that simulations of the data are sensitive to this parameter at all but those sites. Sites with high and low Bayesian learning on the parameter T_2 are interspersed throughout the domain. The parameters T_1 and especially M_1 appear to have very little influence on the data, as conditioning on the data constrains their posterior distributions little if at all.

Sites along the eastern seaboard tend to have low SNR, while those along the west coast south of 40° N tend to have high SNR, and the SNR ratio elsewhere is mixed (Fig. 6, top panel), indicating variation across sites in the degree to
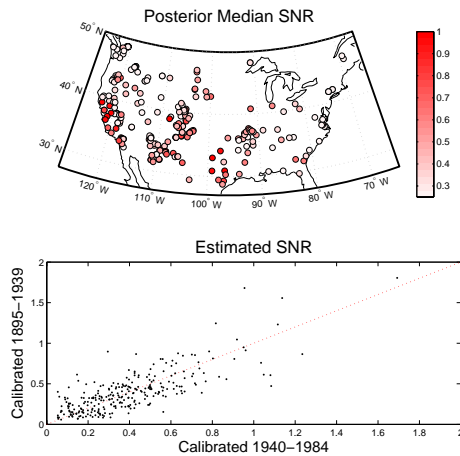
Fig. 6. Top: map of posterior median signal-to-noise ratio (SNR) across sites in network, estimated by comparing simulated and observed data during 1895–1939. Bottom: comparison of posterior median SNR estimated during 1895–1939 and 1940–1984 for values of growth response parameters estimated using data in the complementary intervals.

Fig. 7. Left: correlation of VS-Lite simulated ring width and observed ring width during the complement of the growth response parameter calibration interval for simulations run with posterior median parameters and prior median parameters. Calibrated (uncalibrated) simulations at 160 (87) sites have positive significant correlations with observations at $p < 0.05$. Right: RE statistic for VS-Lite simulations in the complement of the growth response parameter calibration interval ring-width index using posterior median parameters and prior median parameters. RE is greater than zero and thus indicates skillful simulations at 74 (44) of calibrated (uncalibrated) sites.

which VS-Lite can be used to explain the observed ring-width chronologies. The posterior median signal-to-noise ratios are similar when estimated during the first and second halves of the 90 yr interval (bottom panel), indicating robustness in the model performance to the choice of calibration interval, and the posterior median SNR estimated between 1940 and 1984 is within 1, 2, and 3 posterior standard deviations of the posterior median SNR estimated between 1895 and 1939 for 55, 87, and 95 % of sites respectively.

Calibration of the parameters improves the skill of the VS-Lite simulations in predicting the observed data. The correlation of the simulated ring width with observations in the complement of the interval used to estimate the growth response parameters increases uniformly across the sites (Fig. 7, left panel). At the Andrew Johnson Woods site in Ohio, for example, the correlation of calibrated (uncalibrated) simulated ring width with observations is typical of this metric of skill across chronologies simulated in this study, at $\rho = 0.30$ ($\rho = -0.13$) with the observed ring-width time series. In fact, calibrating the parameters at this representative site is crucial in determining the simulated climate controls on growth, with calibrated (uncalibrated) simulated growth limited by moisture (temperature). Time-series plots of these simulations, as well as for another site with more dramatic improvement resulting from calibration, are shown in Supplement Figs. 8 and 9. Calibration also increases the number of sites with positive RE statistic to 74 sites relative to 44 out of 277 for simulations run with prior median parameter values (Fig. 7, right panel).

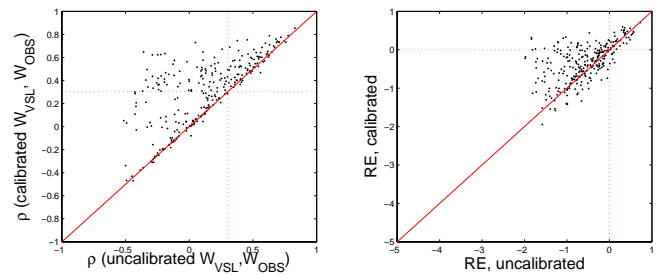Although we treat the simulation at each site as independent of the others, there is some spatial covariance of the residuals of VS-Lite simulations run with calibrated parameters, due to spatial covariance in the underlying climate fields used as model inputs. However, both the range and strength of the residual spatial structure not accounted for by VS-Lite are decreased by running simulations with the calibrated parameters, rather than the prior median parameters at every site (see Supplement Fig. 7 and accompanying description for further details). These results suggest that the parameter calibration improved the ability of VS-Lite to explain variance in the observed tree-ring width chronologies.

The point estimates of the parameters cluster in anomaly-parameter space according to each site's classification as temperature-limited, moisture-limited, or mixed-control sites (Fig. 8). The estimated values of T_1 fall below the mean local JJA temperature (not shown). In other words, the mean summer temperatures all fall above the threshold for nonzero growth at every site, which is consistent with the fact that the chronologies we used at these locations were in fact sampled from living trees. Data at sites classified as temperature-sensitive constrain all estimates of T_2 to values above the mean local summer temperature. Summer temperature variations therefore influence modeled growth at these sites. Sites classified as moisture-limited or as having mixed controls tend to have values of T_2 that fall below local summer mean temperatures; thus temperature variability will have less of an effect on modeled growth. The results for the moisture parameters are similar. All sites have calibrated values of M_1 falling below the climatological mean soil moisture content, so that there is enough moisture for modeled growth to occur across the experimental domain in summer. Calibrated values of M_2 are greater than local climatological mean soil moisture for all sites classified as moisture-limited, but mixed-control and temperature-limited sites tend to have values of M_2 that fall below the climatological summer mean.
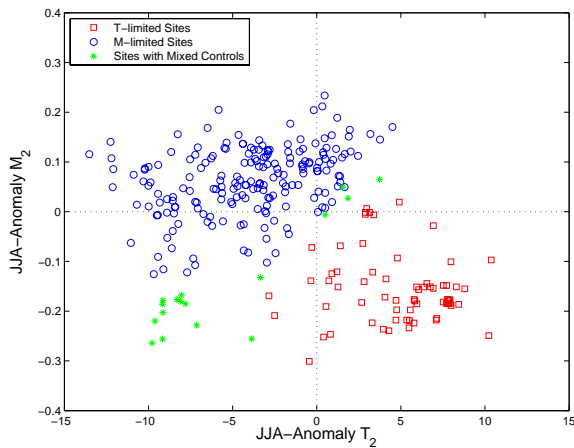
Fig. 8. Plot of estimated M_2 versus estimated T_2 at each site, with both parameters measured as anomalies relative to the local climatological mean summer temperature. The color of points denotes the classification of the controls on modeled growth at each site. $x = 0$ and $y = 0$ define the mean local summer temperature and soil moisture content, respectively.

Given that the parameter estimates of the lower thresholds T_1 and M_1 fall below mean summer climatological values at all sites, the distribution of the anomaly point estimates in $T_2 \times M_2$ space contains the most information about the modeled climate controls on growth (Fig. 8). In the second quadrant of the plot (defined by anomaly $T_2 < 0$ and anomaly $M_2 > 0$ one would expect sites where moisture generally limits summer growth, since climatological temperatures tend to fall above the optimal temperature growth limit, but soil moisture tends to fall below its optimal growth limit. All of the sites whose parameterizations end up in this quadrant are in fact classified as moisture-limited by our classification scheme. The fourth quadrant (anomaly $T_2 > 0$, anomaly $M_2 < 0$) would seem to define a region of parameter space describing temperature-limited growth, and indeed the sites whose estimated parameters are in this quadrant are nearly all classified this way. The sites that fall into quadrant III and far from the quadrant boundaries are mixed-control sites, as one would expect for locations where trees are sensitive to both variations in summer moisture and temperature variability.

The modeled climate controls on growth break the continental United States into roughly three regions. In both the pseudoproxy and observed proxy experiments, the northwest contains mainly temperature-controlled sites (red markers in quadrant IV of Fig. 8), while moisture-controlled sites fill the west and midwest (blue markers in quadrant II of Fig. 8), and mixed-control sites are most common in the southeast and along the eastern seaboard (green markers in quadrant III Fig. 8). This pattern is generally consistent with our knowledge of the climate sensitivity of the North American tree-ring network (e.g., Meko et al., 1993).

The point estimates of T_2 and M_2 derived from the scientifically based priors described here are generally similar in value to those derived from uniform priors (Supplement Fig. 3). This result is consistent with the fact that the ratios of posterior to prior variance for these parameters tend to be appreciably less than one at most sites, thus indicating that the posteriors are dominated by information from the data. By contrast, estimates of T_1 and M_1 tend to differ more appreciably, but the associated posterior to prior variance ratios for these parameters are close to one. In general, the difference in the point estimates tends to be greatest where the data inform the posteriors least, while sites with high learning (and hence posterior-to-prior variance ratios close to zero) exhibit little difference in posteriors derived from either prior (Supplement Fig. 4).

4 Discussion

The Bayesian inference scheme is skillful in recovering the known parameters used to create pseudoproxy ring-width series. Although real-world analogs to the pseudoproxy target model parameters may not be known, the skill in the pseudoproxy context supports the notion that the approach will estimate parameters that optimize VS-Lite's fit to observed tree-ring width chronologies. The results of the PPE and OPMC are similar in terms of their spatial distributions of the modeled controls on growth, as well as the model sensitivities to the growth response parameters. These correspondences further support the applicability of pseudoproxy experiment results to studies using observed proxy data. The skill of observed ring-width index simulations is also shown to improve after calibrating the parameters, as compared to simulations run with parameter values set to the prior medians.

In addition to point estimates of the parameters, the spread of the posterior distributions also provides measures of the estimation uncertainty and the model sensitivity to the parameters. We find that the VS-Lite model is generally least (most) sensitive to the value of M_1 (M_2), as the ratio of posterior to prior variance is very close to one (zero) at all sites in both pseudoproxy and observed proxy experiments (Figs. 3 and 5). The model sensitivity to the temperature thresholds T_1 and T_2 depends on the particular site.

At sites where little Bayesian learning occurs due to a low fraction of variance explained by VS-Lite, the posterior inference is determined almost entirely by the prior model. This result can be seen by comparison of the posterior-to-prior variance ratios of the PPE performed with SNR = 1 and SNR = 0.25. Much less "learning" occurs in the latter experiment, and so the posterior-to-prior variance ratios tend to be closer to one than in the former experiment. Defaulting to the prior when the data contain little information is a standard feature of Bayesian analysis and underscores the importance of careful prior elicitation based on all available scientific evidence. This feature is ideal to the degree that one has

faith in the representation of the underlying science and its inherent uncertainty reflected in the prior distributions. Our publicly available code includes flexible options for users to define their own priors, should new information from future field studies of tree growth render the default set of priors described here obsolete.

The choice of finite support or range is another component of prior specification that may heavily influence the posterior inference, as in the case of the parameter M_2 at most sites in this study. The posterior distributions of this parameter at most sites show high probability mass toward the upper bound of the compact prior support (see Fig. 1 and bottom right panel of Fig. 4), indicating that the data alone imply values of this parameter above the region allowed by the prior. However, the upper limit of the prior represents a physical constraint on biological thresholds for optimal moisture conditions for plant growth, as excessive soil moisture values may become detrimental to plant growth (Kozlowski, 1984). Given that the modeling of soil moisture within VS-Lite is known to be simplistic (Tolwinski-Ward et al., 2011), we believe the posteriors here represent an objective compromise between the data and prior knowledge of the parameter, given the uncertainty of VS-Lite.

The location of parameter point estimates relative to local climatological means in multidimensional anomaly parameter space presents a graphical tool for understanding the climate controls on the modeled ring width signal (Fig. 8). At sites where VS-Lite reasonably represents growth, this type of plot could help identify and predict changes in the climate–proxy relationship that result from climatic nonstationarities driving mean environmental conditions across biological thresholds. In such cases of "divergence" (D'Arrigo et al., 2004; Carrer and Urbinati, 2006), one would expect the point representing the optimal set of parameter choices to cross from one quadrant into another after the climatic shift. At any site where the VS-Lite model fits the associated observed chronology well, the proximity of the parameters fitted by the present methodology to the boundaries between quadrants in Fig. 8 might be used to diagnose the potential for calibration–interval relationships between ring-width and climate to degrade under past shifts in climatology, and thus provide an important diagnostic before a given ring-width series is used for climate reconstruction. The spread of parameter ensemble members produced by our estimation technique could even be used in a Monte Carlo sense to quantify the probability that, for instance, a chronology characterized by temperature-limited growth in the calibration interval "crosses the line" into a moisture-limited, mixed-control, or complacent growth response regime for a shift in climatology of a given number of degrees and volumetric soil moisture content.

This approach may be particularly valuable for reconstructions that combine information contained in ring-width series with that from other proxy sources. Currently, the parameter estimation scheme is in use for probabilistic recon-

structions of paleoclimate from tree-ring width data and co-located isotopic dendrochronologies; preliminary results are sketched by Evans et al. (2013), and details are in preparation by Tolwinski-Ward et al. (2013). At the southwestern American site in question, the ring-width variations display strong moisture sensitivity during the calibration interval; however, the isotopic data suggest an anomalously wet interval in the pre-instrumental period. The reconstruction employs a Bayesian hierarchical model in which VS-Lite is used to link the variations in the proxy data to past variations in temperature and moisture. In that setting, the parameter-estimation procedure explored here is used to derive an ensemble of parameters consistent with the VS-Lite model and observed tree-ring width indices during a calibration interval. As described above in discussion of Fig. 8, the draws of the parameter values from the posterior allow for quantification of the probability that ring-width growth was insensitive to hydroclimatic variability during the reconstruction interval due to a general increase in available moisture resources indicated by the isotopic data.

In future studies, uncertainty in the parameter posteriors might be reduced by modeling the fields of parameters as functions of space, elevation, species, or other parameterizations of site environment. This is the classic principle of "borrowing strength" across spatially covarying sites (Tukey, 1986). Sites that are close in space seem to take on similar optimal parameter values (Fig. 4). One possible approach is to model the parameters as a transformation of one or more spatial processes (see for example Chiles and Delfiner, 1999, Sect. 6.2 on anamorphosis). The spatial modeling could then proceed by using a standard additive model with latitude, longitude, and elevation as covariates giving fixed effects, and random effects given by a Gaussian field with Matérn covariance structure whose parameters could be inferred within the Bayesian framework. Since the parameter values are interpretable in terms of the climate controls on growth, any modeled spatial structure of the parameter fields likely holds information about the spatial distribution of climate controls that can be linked to mechanistic causes, such as orography, regional drought patterns, the timing of snowmelt, regional climatological means and variances, or teleconnections to larger global patterns of climatic variability. The scheme also could be extended to develop parameter sets that depend on species or climatic regions by fitting a single set of parameters to a set of trees of a common species or regional location. From a physiological point of view, at parameter–site combinations where a high degree of Bayesian learning occurs, parameter estimates resulting from the estimation procedure could also be viewed as hypotheses for field studies aiming to provide observed physiological evidence either supporting or refuting the inferred values.

5 Conclusions

The Bayesian calibration scheme presented here skillfully recovers parameter estimates near the values used to create synthetic tree-ring width data. The spread of the posterior distributions shows that the model fit to data is generally sensitive (insensitive) to the value of the moisture threshold M_2 (M_1), and may or may not be sensitive to the temperature threshold parameters depending on location. Estimates of the VS-Lite model's uncertainty provided by the scheme appear to be robust outside of the interval used for calibration. The location of estimated parameters relative to local climatology in multidimensional parameter space provides insight into the climate controls on modeled tree-ring growth, and may also provide information about the potential for instability in the calibration–interval climate–ring-width relationships before the instrumental record. The output of the estimation procedure enables users of VS-Lite to represent fully the range of model error stemming from uncertainty in its parameterization in both forward simulations of tree-ring width and inverse paleoclimate estimation settings.

Supplementary material related to this article is available online at: http://www.clim-past.net/9/1481/2013/cp-9-1481-2013-supplement.zip.

References

Boreux, J.-J., Naveau, P., Guin, O., Perreault, L., and Bernier, J.: Extracting a common high frequency signal from Northern Quebec black spruce tree-rings with a Bayesian hierarchical model, Clim. Past, 5, 607–613, doi:10.5194/cp-5-607-2009, 2009.

Carrer, M. and Urbinati, C.: Long-term change in the sensitivity of tree-ring growth to climate forcing in *Larix decidua*, New Phytol., 170, 861–872, 2006.

Carrer, M., Anfodillo, T., Urbinati, C., and Carraro, V.: High altitude forest trees sensitivity to global warming: results from long-term and short-term analyses in the Italian Eastern Alps, in: The impacts of Climate variability on Forests. Lecture Notes in Earth Sciences, 74, edited by: Beniston, M. and Innes, J. L., 171–189, Springer, Berlin, Germany, 1998.

Chib, S. and Greenberg, E.: Markov Chain Monte Carlo Simulation Methods in Econometrics, Economet. Theor., 12, 409–431, 1996.

Chiles, J. and Delfiner, P.: Geostatistics: Modeling Spatial Uncertainty, Wiley Series in Probability and Statistics, New York, NY, 1999.

Cook, E., Meko, D., Stahle, D., and Cleaveland, M.: Drought reconstructions for the continental United States, J. Climate, 12, 1145–1162, 1999.

Daly, C., Halbleib, M., Smith, J., Gibson, W., Doggett, M., Taylor, G., Curtis, J., and Pasteris, P.: Physiographically sensitive mapping of climatological temperature and precipitation across the coterminous United States, Int. J. Climatol., 28, 2031–2064, doi:10.1002/joc.1688, 2008.

D'Arrigo, R. D., Kaufmann, R. K., Davi, N., Jacoby, G. C., Laskowski, C., Myneni, R. B., and Cherubini, P.: Thresholds for warming-induced growth decline at elevational tree line in the Yukon Territory, Canada, Global Biogeochem. Cy., 18, GB3021, doi:10.1029/2004GB002249, 2004.

Deslauriers, A., Rossi, S., Anfodillo, T., and Saracino, A.: Cambial phenology, wood formation and temperature thresholds in two contrasting years at high altitude in southern Italy, Tree Physiol., (Oxford, UK), 28, 863–871, 2008.

Evans, M., Kaplan, A., and Cane, M.: Pacific sea surface temperature field reconstruction from coral δ^{18}O data using reduced space objective analysis, Paleoceanography, 17, 1, doi:10.1029/2000PA000590, 2002.

Evans, M. N., Tolwinski-Ward, S. E., Thompson, D. M., and Anchukaitis, K. J.: Applications of proxy system modeling in high resolution paleoclimatology, Quaternary Sci. Rev., 76, 16–28, doi:10.1016/j.quascirev.2013.05.024, 2013.

Fritts, H. C.: Tree Rings and Climate, The Blackburn Press, New York, 2001.

Fritts, H. C., Shashkin, A., and Downes, G.: A Simulation Model of Conifer ring growth and cell structure, in: Tree Ring Analysis, edited by: Wimmer, R. and Vetter, R., Chap. 1, 3–32, Cambridge University Press, Cambridge, UK, 1999.

Garfinkel, H. and Brubaker, L.: Modern climate–tree-growth relationships and climatic reconstruction in sub-Arctic Alaska, Nature, 286, 872–874, doi:10.1038/286872a0, 1980.

Gaucherel, C., Campillo, F., Misson, L., Guiot, J., and Boreux, J.: Comparison of a process-based tree-growth model: Comparison of optimization, MCMC, and Particle Filtering algorithms, Environ. Modell. Softw., 23, 1280–1288, 2008.

Gelman, A. and Rubin, D.: Inference from Iterative Simulation Using Multiple Sequences, Statist. Sci., 7, 457–472, doi:10.1214/ss/1177011136, 1992.

Gelman, A., Carlin, J., Stern, H., and Rubin, D.: Bayesian Data Analysis, Chapman and Hall, Boca Raton, Florida, 2003.

Gilks, W., Richardson, S., and Spiegelhalter, D. (Eds.): Markov Chain Monte Carlo in Practice, Chapman and Hall, Boca Raton, FL, 1996.

Hoch, G. and Körner, C.: Growth and carbon relations of tree line forming conifers at constant vs. variable low temperatures, J.

Ecol., 97, 57–66, 2009.

Huang, J., van den Dool, H. M., and Georgakakos, K. P.: Analysis of model-calculated soil moisture over the United States (1931–1993) and applications to long-range temperature forecasts, J. Climate, 9, 1350–1362, 1996.

Körner, C.: Winter crop growth at low temperature may hold the answer for alpine treeline formation, Plant Ecol. Div., 1, 3–11, 2008.

Körner, C.: Alpine treelines: functional ecology of the global high elevation tree limits, Switzerland: Springer Basel, 2012.

Körner, C. and Hoch, G.: A Test of treeline theory on a montane permafrost island, Arct. Antarct. Alp. Re., 38, 113–119, doi:10.1657/1523-0430(2006)038[0113:ATOTTO]2.0.CO;2, 2007.

Körner, C. and Paulsen, J.: A world-wide study of high altitude treeline temperatures, J. Biogeogr., 31, 713–732, doi:10.1111/j.1365-2699.2003.01043.x, 2004.

Kozlowski, T.: Responses of woody plants to flooding, in: Flooding and Plant Growth, edited by: Kozlowski, T., 129–163, Academic Press, New York, 1984.

Mann, M. E., Zhang, Z., Hughes, M. K., Bradley, R. S., Miller, S. K., Rutherford, S., and Ni, F.: Proxy-based reconstructions of hemispheric and global surface temperature variations over the past two millennia, Proc. Nat. Acad. Sci. USA, 105, 13252–13257, doi:10.1073/pnas.0805721105, 2008.

Meko, D., Cook, E. R., Stahle, D. W., Stockton, C. W., and Hughes, M. K.: Spatial Patterns Of Tree-Growth Anomalies In The United States And Southeastern Canada, J. Climate, 6, 1773–1786, 1993.

Misson, L.: MAIDEN: a model for analyzing ecosystem processes in dendroecology, Can. J. For. Res., 34, 874–887, doi:10.1139/x03-252, 2004.

Misson, L., Rathgeber, C., and Guiot, J.: Dendroecological analysis of climatic effects on Quercus petraea and Pinus halepensis radial growth using the process-based MAIDEN model, Can. J. For. Res., 34, 888–898, doi:10.1139/x03-253, 2004.

Robert, C.: The Bayesian Choice: From Decision-Theoretic Foundations fo Computational Implementation, Springer Texts in Statistics, Paris, 2007.

Rossi, S., Deslauriers, A., Anfodillo, T., and Carraro, V.: Evidence of threshold temperatures for xylogenesis in conifers at high altitudes, Oecologia, 152, 1432–1939, doi:10.1007/s00442-006-0625-7, 2007.

Shumway, R. and Stoffer, D.: Time Series Analysis and Its Applications, Springer Texts in Statistics, New York, NY, 2006.

Smerdon, J.: Climate models as a test bed for climate reconstruction methods: pseudoproxy experiments, WIREs Clim. Change, 3, 63–77, doi:10.1002/wcc.149, 2012.

Taylor, W. P.: Significance of Extreme or Intermittent Conditions in Distribution of Species and Management of Natural Resources, with a Restatement of Liebig's Law of Minimum, Ecology, 15, 374–379, 1934.

Tolwinski-Ward, S. E., Evans, M. N., Hughes, M. K., and Anchukaitis, K. J.: An efficient forward model of the climate controls on interannual variation in tree-ring width, Clim. Dynam., 36, 2419–2439, doi:10.1007/s00382-010-0945-5, 2011.

Tolwinski-Ward, S. E., Tingley, M. P., Nychka, D. W., Evans, M. N., Hughes, M. K., and Leavitt, S. W.: Probabilistic reconstructions of local temperature and soil moisture from potentially nonstationary tree-ring width data, in preparation, 2013.

Tukey, J.: The Collected Works of John W. Tukey, Volune IV Philosophy and Principles of Data Analysis: 1965–1986, Wadsworth & Brooks/Cole Advanced Books, Statistics/Probability Series, Murray Hill, New Jersey, 1986.

Vaganov, E. A., Hughes, M. K., and Shashkin, A. V.: Growth dynamics of conifer tree rings: Images of past and future environments, Ecol. Studies 183, Springer-Verlag, Berlin/Heidelberg/New York, 2006.

Vaganov, E. A., Anchukaitis, K. J., and Evans, M. N.: How well understood are the processes that create dendroclimatic records? A mechanistic model of climatic control on conifer tree-ring growth dynamics, in: Dendroclimatology: Progress and Prospects, edited by: Hughes, M., Swetnam, T., and Diaz, H., Developments in Paleoecological Research, Chap. 3, Springer-Verlag, Dordrecht, Netherlands, 2011.

Williams, A., Xu, C., and McDowell, N.: Who is the new sheriff in town regulating boreal forest growth?, Environ. Res. Lett., 6, 041004, doi:10.1088/1748-9326/6/4/041004, 2011.