1 **Parameter uncertainty quantification in an idealized GCM with**
2 **a seasonal cycle**

3 **Michael F. Howland^{1,2}, Oliver R. A. Dunbar², Tapio Schneider²**

4 ¹Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

5 ²California Institute of Technology, Pasadena, CA, USA

Corresponding author: Michael F. Howland, mhowland@mit.edu

Abstract

Climate models are generally calibrated manually by comparing selected climate statistics, such as the global top-of-atmosphere energy balance, to observations. The manual tuning only targets a limited subset of observational data and parameters. Bayesian calibration can estimate climate model parameters and their uncertainty using a larger fraction of the available data and automatically exploring the parameter space more broadly. In Bayesian learning, it is natural to exploit the seasonal cycle, which has large amplitude, compared with anthropogenic climate change, in many climate statistics. In this study, we develop methods for the calibration and uncertainty quantification (UQ) of model parameters exploiting the seasonal cycle, and we demonstrate a proof-of-concept with an idealized general circulation model (GCM). Uncertainty quantification is performed using the calibrate-emulate-sample approach, which combines stochastic optimization and machine learning emulation to speed up Bayesian learning. The methods are demonstrated in a perfect-model setting through the calibration and UQ of a convective parameterization in an idealized GCM with a seasonal cycle. Calibration and UQ based on seasonally averaged climate statistics, compared to annually averaged, reduces the calibration error by up to an order of magnitude and narrows the spread of posterior distributions by factors between two and five, depending on the variables used for UQ. The reduction in the size of the parameter posterior distributions leads to a reduction in the uncertainty of climate model predictions.

Plain Language Summary

Climate models rely on parameterizations of physical processes that cannot be resolved with available computational resources. Empirical representations of physical processes, such as turbulence and cloud physics, reduce the computational cost of simulations, but introduce new unknown parameters into the climate model. The unknown parameters contribute to uncertainties associated with climate model predictions. Historically, fixed values of the model parameters have been hand-tuned using scientific intuition and a limited amount of available data. We develop methods for the computationally efficient estimation of the unknown climate model parameters and their uncertainty systematically from data, by using gradient-free optimization and machine learning. Many processes and observable statistics of Earth's climate used to produce this data are influenced by the seasonal cycle. We demonstrate that the incorporation of seasonal information into these statistics significantly improves the resulting calibration of climate model parameters, in contrast to annually averaged information alone. We show that including seasonal information also reduces the uncertainty associated with the model parameters, which consequently reduces the uncertainty of climate model predictions.

1 Introduction

The objective of quantifying uncertainty in computational models arises in a wide range of applications, including weather and climate modeling (Schneider, Lan, et al., 2017), fluid dynamics (Duraissamy et al., 2019), and energy systems (Constantinescu et al., 2010). Often, uncertainty associated with predictions from computational models is the result of processes that cannot be resolved on the computational grid, either due to computational complexity limitations (Meneveau & Katz, 2000) or due to uncertainty associated with the process itself (Schneider, Teixeira, et al., 2017). In general circulation models (GCMs), primary uncertainties arise from the representation of subgrid-scale turbulence, convection, and cloud physics, which have a significant impact on the evolution of climate under rising greenhouse gases (Cess et al., 1989; Bony & Dufresne, 2005; Webb et al., 2013; Suzuki et al., 2013; Brient & Schneider, 2016). While clouds are associated with turbulence and convective updrafts with scales of $\mathcal{O}(10\text{ m})$, modern climate simulations have a typical horizontal resolution of $\mathcal{O}(10\text{ km})$ – $\mathcal{O}(100\text{ km})$ (Delworth et al., 2012; Hoegh-Guldberg et al., 2018). Climate simulations rely on physically motivated parameterizations that model the effects of subgrid-scale processes such as clouds and turbulence on the resolved scales (Hourdin et al., 2013). Such parameterizations come with para-

metric and structural uncertainties; quantifying these uncertainties and how they percolate into climate projections remains an outstanding challenge (Schneider, Teixeira, et al., 2017).

Physical parameterizations have historically been individually developed and calibrated using data from isolated experiments (Smagorinsky, 1963; Spalart & Allmaras, 1992; Jakob, 2010; Golaz et al., 2013; Hourdin et al., 2017). They are further adjusted so that global models that incorporate them satisfy selected large-scale observational or physical constraints, such as a closed top-of-atmosphere energy balance or reproduction of the 20th-century global-mean temperature evolution (Mauritsen et al., 2012; Hourdin et al., 2017; Schmidt et al., 2017). Model calibration is usually done manually, focusing on a small subset of model parameters and exploiting only a fraction of the available observational data.

As a step toward automating and augmenting this process, here we further develop algorithms for model calibration and uncertainty quantification (UQ) that in principle allow models to learn from large datasets and that scale to high-dimensional parameter spaces. In previous work, these algorithms have been demonstrated in simple conceptual models (Cleary et al., 2021) and in a statistically stationary idealized GCM (Dunbar et al., 2020). We take the next step and demonstrate how these algorithms can exploit seasonal variations, which for many climate statistics are large relative to the climate changes expected in the coming decades and contain exploitable information about the response of the climate system to perturbations (Knutti et al., 2006; Schneider et al., 2021).

Whereas numerical weather prediction assimilates spatiotemporally evolving trajectories of atmospheric states as initial conditions for forecasts (e.g., Kalnay (2003); Bauer et al. (2015)), in climate modeling it is preferable to assimilate time-averaged climate statistics. This focuses the learning problem on quantities of interest in climate predictions (i.e., climate statistics, including higher-order statistics such as precipitation extremes), and it avoids the need to estimate uncertain atmospheric initial conditions on which trajectories of states depend (Cleary et al., 2021). Calibration and UQ of climate models on the basis of time-averaged statistics smooths the prediction-error based objective function and enables the use of data that have different resolution than the climate simulations (Schneider, Lan, et al., 2017; Dunbar et al., 2020). Given the dual desires to avoid having to estimate atmospheric initial conditions, which are forgotten over about 2 weeks (Zhang et al., 2019), and to exploit seasonal variations, it becomes natural to choose averaging timescales between around 30 and 90 days. Such averaging timescales are the focus of this study.

We consider the calibration and UQ of convective parameters in an idealized GCM with seasonally varying insolation (Frierson et al., 2006; O’Gorman & Schneider, 2008; Bordoni & Schneider, 2008). We develop an extension of the calibrate-emulate-sample (CES) Bayesian learning methodology (Cleary et al., 2021) to enable the use of statistics computed from a non-stationary statistical state. The qualitative and quantitative impacts of the time-averaging length of the climate statistics on the parameter calibration and UQ are assessed. GCMs are generally tuned in situations which have low parameter identifiability based on available climate data (Hourdin et al., 2017). We perform numerical experiments with observable climate statistics that are informative about the convective parameters we wish to calibrate, but are not in any simple and direct way related to them. We also explore less informative climate statistics, which highlight the benefits of incorporating seasonal variations in the climate statistics that are being exploited for UQ.

The remainder of this paper is organized as follows. In section 2, the Bayesian learning methods for time-dependent problems are introduced. The numerical details of the seasonally forced GCM and UQ experiments are introduced in section 3. The calibration and UQ results are shown in section 4, and conclusions are provided in section 5.

2 Uncertainty quantification methods

The goal of this study is to estimate the probability distributions associated with model parameters θ that are used by a GCM and about which only imprecise prior information is known. We will consider a Bayesian approach to the estimation of the probability distributions of model parameters θ , where we seek $\mathbb{P}(\theta|\mathbf{y})$, the conditional probability distribution of parameters θ given observed data $\mathbf{y}$. The GCM is a computationally expensive numerical model that evolves climate states in time. From the GCM, we extract statistical information that is denoted by $\mathcal{G}(\theta)$. Here, $\mathcal{G}(\theta)$ includes the numerical integration of the GCM states and the aggregation of relevant climate statistics in time. Although the data $\mathbf{y}$ will, in general, not provide direct information about θ , we can estimate $\mathbb{P}(\theta|\mathbf{y})$ using Bayesian learning by comparing the GCM outputs $\mathcal{G}(\theta)$ with data $\mathbf{y}$.

Standard methods for Bayesian learning include Markov Chain Monte Carlo (MCMC) (Brooks et al., 2011), which typically requires $\mathcal{O}(10^5)$ forward model evaluations of $\mathcal{G}(\theta)$ to sample the posterior distribution (Geyer, 2011). Instead, we perform calibration and UQ using the recently developed CES methodology (Cleary et al., 2021), which consists of three steps: (1) Ensemble Kalman processes (Schillings & Stuart, 2017; Garbuno-Inigo et al., 2020) are used to calibrate parameters and to generate input-output pairs of the mapping $\theta \mapsto \mathcal{G}(\theta)$; (2) Gaussian process (GP) regression is used to train an emulator $\mathcal{G}_{\text{GP}}(\theta)$ of the mapping $\theta \mapsto \mathcal{G}(\theta)$ using the training points generated in the calibration step; and (3) MCMC sampling with the computationally efficient GP emulator $\mathcal{G}_{\text{GP}}(\theta)$ rather than the expensive forward model $\mathcal{G}(\theta)$ is used to estimate the posterior distribution $\mathbb{P}(\theta|\mathbf{y})$. The CES methodology has previously been used for calibration and UQ of parameters in simple model problems such as Darcy flow and Lorenz systems (Cleary et al., 2021) and for convective parameters in a statistically stationary GCM (Dunbar et al., 2020). More recent methodological developments by Lan et al. (2021) enabled the CES framework to perform simultaneous UQ on $\mathcal{O}(1000)$ parameters using deep neural network-based emulation and MCMC suited to high-dimensional spaces (Beskos et al., 2008, 2011). A schematic of the CES methodology is shown in Figure 1.

The Bayesian learning methodology used in this study is described in the following sections. In section 2.1, the inverse problem of estimating $\mathbb{P}(\theta|\mathbf{y})$ in a setting with a periodic cycle is introduced. In section 2.2, the ensemble Kalman process calibration method is outlined. Section 2.3 introduces the GP emulation in an uncorrelated transformed space, obtained by a singular value decomposition (principal component analysis) of the noise covariance matrix. Section 2.4 describes how the Bayesian posterior distribution is approximated using the GP emulator.

2.1 Seasonal GCM inverse problem

With fixed insolation (Frierson et al., 2006; O’Gorman & Schneider, 2008), GCM statistics are stationary and ergodic. With the seasonal cycle incorporated, the insolation varies as a function of the ordinal day in the simulation. The resulting GCM states are statistically cyclostationary, with a dependence on the ordinal day. The data we use are constructed accounting for the seasonally varying boundary conditions, as the time-averaged GCM statistics

$$\mathcal{G}_T(\theta, \xi, t) = \frac{1}{T} \int_t^{t+T} \mathcal{H}(t'; \theta, \xi) dt'. \quad (1)$$

Here, $\mathcal{H}(t; \theta, \xi) \in \mathbb{R}^{N_y}$ represents the output states of the GCM, depending on time t . The number of states being measured is N_y . The integration length is specified as T . The initial conditions for the GCM integration are represented by ξ , which depend on the initial time t . The time-averaged statistics are $\mathcal{G}_T(\theta, \xi, t) \in \mathbb{R}^{N_y}$.

With seasonally varying insolation, the time-averaged statistics in Eq. (1) depend on the ordinal day (t) in the GCM. With a specified integration length T , there are a corresponding

Figure 1: Schematic of the calibrate-emulate-sample (CES) methodology to estimate model parameters θ . With inputs of data y , noise covariance Σ , characteristic values of the data y_c (for normalization), and the prior $\mathbb{P}(\theta)$, the calibration stage generates input-output pairs $\{\theta_i, \mathcal{G}(\theta_i)\}_i$. The input-output mapping is emulated using Gaussian process regression in a transformed, uncorrelated space ($\tilde{\cdot}$), obtained from a truncated singular value decomposition on the noise covariance matrix Σ . The GP emulator is used for efficient sampling using MCMC to approximate the posterior distribution $\mathbb{P}(\theta|\tilde{y})$. The objective functions for calibration and sampling are denoted by $\Phi(\theta, y)$ and $\Phi_{\text{MCMC}}(\theta, \tilde{y})$, respectively.

set of time-averaged statistics

$$\mathcal{G}_{T,j}(\theta, \xi) = \frac{1}{T} \int_{t_0+(j-1)T}^{t_0+jT} \mathcal{H}(t'; \theta, \xi) dt', \quad (2)$$

where $j = 1$ through N_s is an index representing the time of year in the GCM simulation, starting from $j = 1$ at vernal equinox (time t_0). The integration windows are non-overlapping. The length of the year in the GCM is $T_{yr} = 360$ d, and the resulting number of non-overlapping time-averaged statistics are $N_s = T_{yr}/T$. In this framework, the number of batches of statistics N_s is a design parameter set by the integration timescale T . For $T = 90$ d, the data are aggregated seasonally, with $N_s = 4$. For $T = 360$ d, the statistics are averaged over the full year in the GCM, corresponding to annually averaged climate statistics, and $N_s = 1$.

The GCM data are the concatenation of the time-averaged batches over collated ordinal days

$$\mathcal{G}(\theta, \xi) = [\mathcal{G}_{T,1}(\theta, \xi_1), \mathcal{G}_{T,2}(\theta, \xi_2), \dots, \mathcal{G}_{T,N_s}(\theta, \xi_{N_s})]. \quad (3)$$

In this study, we will consider perfect-model numerical experiments, with the data y being constructed using the same GCM at target true parameters $\theta = \theta^\dagger$. The synthetic data are given by

$$y = [\mathcal{G}_{T,1}(\theta^\dagger, \xi_1), \mathcal{G}_{T,2}(\theta^\dagger, \xi_2), \dots, \mathcal{G}_{T,N_s}(\theta^\dagger, \xi_{N_s})]. \quad (4)$$

The size of y and $\mathcal{G}$ is $\mathbb{R}^N$, where $N = N_s \cdot N_y$. The resulting inverse problem relating parameters and data is

$$y = \mathcal{G}(\theta, \xi) + \eta, \quad (5)$$

where $\eta \sim N(0, \Delta)$ is a realization of normal measurement error with zero mean and covariance matrix Δ . We generate synthetic data and forward model data starting from the same initial conditions ξ in Eq. (5). Since the states of the GCM depend on the ordinal day, the averaging operation in Eq. (2) must be consistently aligned in ordinal day between the synthetic data and GCM outputs.

The states of the GCM depend on the boundary conditions, specific to the ordinal day, and the initial conditions ξ . The ensemble average of independent realizations of $\mathcal{G}(\theta, \xi_i)$ over

different initial conditions ξ_i is

$$\mathcal{G}_\infty(\boldsymbol{\theta}) = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{i=1}^M \mathcal{G}(\boldsymbol{\theta}, \xi_i). \quad (6)$$

155 In the limit of infinite realizations of the climate statistics, the dependence of the ensemble
156 average $\mathcal{G}_\infty(\boldsymbol{\theta})$ on the initial conditions trends to zero by the central limit theorem. The cen-
157 tral limit theorem applies since $\mathcal{G}(\boldsymbol{\theta}, \xi_i)$ contains a full year of data by construction (Eq. 3)
158 and is therefore a statistically stationary object with respect to the cycle. But due to compu-
159 tational limitations, only a finite number of ensemble realizations is available in practice. For
160 the synthetic data, we can similarly define $\mathbf{y}_\infty$; however, $\mathbf{y}_\infty$ is generally not accessible from
161 observations.

The atmospheric initial condition is of minimal practical value in the climate model set-
ting, and we aim to avoid its estimation. We therefore reformulate the inverse problem such
that $\mathcal{G}$ and $\mathbf{y}$ can have different initial conditions, and such that we do not require the estima-
tion of the initial conditions. To do so, we assume the forward model output $\mathcal{G}(\boldsymbol{\theta}, \xi)$, is a noisy,
finite average approximation of the infinite ensemble average of the climate statistics, that is,
 $\mathcal{G}(\boldsymbol{\theta}, \xi) = \mathcal{G}_\infty(\boldsymbol{\theta}) + N(0, \Sigma)$, where Σ is the internal variability covariance matrix of the
GCM (Cleary et al., 2021). The realizations $\mathcal{G}(\boldsymbol{\theta}, \xi)$ and $\mathbf{y}$ are both subject to the internal vari-
ability of the climate system. In this study, given the definition of $\mathbf{y}$ and $\mathcal{G}(\boldsymbol{\theta}, \xi)$, the internal
variability is the interannual variability of the GCM. As a result, the inverse problem becomes

$$\mathbf{y} = \mathcal{G}_\infty(\boldsymbol{\theta}) + \gamma, \quad (7)$$

where $\gamma \sim N(0, \Delta + \Sigma)$. In Eq. (7), the initial conditions are removed from the inference
problem. The synthetic data are collected into the matrix $Y \in \mathbb{R}^{N \times N_d}$, where N is the di-
mension of the data space and N_d is the number of data samples, or years in this study. The
internal variability covariance matrix is computed from the synthetic data matrix Y such that
 $\Sigma = \text{cov}(Y)$. We assume that Σ does not vary as a function $\boldsymbol{\theta}$ such that $\Sigma(\boldsymbol{\theta}^\dagger) \approx \Sigma(\boldsymbol{\theta})$ and
that Σ does not depend on initial conditions ξ . Invoking Gaussian error assumptions based on
the central limit theorem, the corresponding negative log-likelihood objective function then is

$$\Phi(\boldsymbol{\theta}, \mathbf{y}) = \frac{1}{2} \|\mathbf{y} - \mathcal{G}_\infty(\boldsymbol{\theta})\|_{\Delta + \Sigma}^2 \quad (8)$$

where $\|\cdot\|_A = \|A^{-1/2} \cdot\|_2$. The likelihood is (Kaipio & Somersalo, 2006)

$$\mathbb{P}(\boldsymbol{\theta}|\mathbf{y}) \propto \exp(-\Phi(\boldsymbol{\theta}, \mathbf{y})). \quad (9)$$

In this perfect-model setting, where the synthetic data and forward model are obtained
from the same GCM, there is no direct measurement error (η) and no systematic model er-
ror. To emulate measurement error, we add Gaussian noise to the GCM output, with zero mean
and covariance matrix

$$\Delta = \text{diag}(\delta_k^2). \quad (10)$$

Here, the noise standard deviation is defined such that

$$\delta_k = \min \left(C \min \left[\text{dist}(y_k + 2\sqrt{\Sigma_{kk}}, \partial\Omega_k), \text{dist}(y_k - 2\sqrt{\Sigma_{kk}}, \partial\Omega_k) \right], C_m \cdot y_k \right), \quad (11)$$

162 where $C_m = 0.1$ caps the maximum measurement error standard deviation to 10% of the mean
163 data values and $C = 0.2$ controls how close the noise-perturbed data can come within phys-
164 ical boundaries $\partial\Omega_i$ of the data (e.g., to keep relative humidities between 0 and 1 with high
165 probability).

166 2.2 Calibrate: Ensemble Kalman Inversion

167 The first stage of the CES methodology is to calibrate the parameters $\boldsymbol{\theta}$ of the model
168 based on data $\mathbf{y}$. We perform calibration with independent realizations of $\mathcal{G}(\boldsymbol{\theta}_i, \xi_i)$, viewed

as noisy approximations of $\mathcal{G}_\infty(\theta_i)$. Calibration is performed using ensemble Kalman methods, which demonstrate theoretical success, in idealized problems, and empirical success, in complex problems, to optimize parameters under such noise (Duncan et al., 2021).

The utility of the calibration stage is two-fold: (1) optimize parameters to minimize the mismatch between model output and data; and (2) provide good parameter–model output pairs $(\theta_i, \mathcal{G}(\theta_i, \xi_i))$ for training an emulator of the parameter-to-data map, with a higher density of training points near the optimal parameters. The ensemble Kalman filter (EnKF) is a Kalman filter implementation in which the covariances are approximated using Monte Carlo sampling (Evensen, 2003). The EnKF has been used widely for derivative-free state estimation in numerical weather prediction (e.g., Houtekamer and Zhang (2016)) and model-based control (e.g., Howland, Ghate, et al. (2020)). Ensemble Kalman methods for Bayesian inversion were introduced by Chen and Oliver (2012) and Emerick and Reynolds (2013). These methods provably draw samples from the posterior of linear inverse problems subject to additive Gaussian noise; however, they fail to do so more in more general, nonlinear problems. Recognizing this, the offline ensemble Kalman inversion (EKI) (Iglesias et al., 2013) algorithm was introduced for classical, optimization-based inversion. EKI generally drives the ensemble members toward consensus near the optimal solution of the inverse problem (Schillings & Stuart, 2017). The parameter update of ensemble member m at iteration step n is

$$\theta_m^{(n+1)} = \theta_m^{(n)} + C_{\theta\mathcal{G}}^{(n)} \left(\Sigma + \Delta + C_{\mathcal{G}\mathcal{G}}^{(n)} \right)^{-1} \left(\mathbf{y} - \mathcal{G}(\theta_m^{(n)}, \xi_m) \right), \quad (12)$$

where $C_{\theta\mathcal{G}}$ is the empirical cross-covariance between the parameters and the model outputs, and $C_{\mathcal{G}\mathcal{G}}$ is the empirical covariance of the model outputs.

EKI is guaranteed to find the optimal parameters in linear problems (Schillings & Stuart, 2017) and has empirical success in nonlinear problems (Dunbar et al., 2020). Several other approaches exist for estimation in nonlinear problems. Li and Reynolds (2009) and Sakov et al. (2012) developed iterative EnKF methods which improve state estimation in strongly nonlinear problems. However, the distribution of the ensemble does not converge to the posterior distribution in the limit of infinite members for nonlinear problems (Zhou et al., 2006; Annan & Hargreaves, 2007; Le Gland et al., 2009; Garbuno-Inigo et al., 2020), necessitating the emulation and sampling described in the following sections.

The number of ensemble members and EKI iterations are hyperparameters; they are set to standard values of $N_{\text{ens}} = 100$ and $N_{\text{it}} = 5$ in this study (Schillings & Stuart, 2017; Dunbar et al., 2020; Cleary et al., 2021). Each ensemble member is run through the GCM, initialized from the same initial conditions. A spin-up period of one year (360 days) is run before the statistics are computed using Eq. (2), to ensure the forward model realizations are subject to differing instantiations of internal variability in the chaotic GCM system.

2.3 Emulate: Gaussian process emulation

We emulate the mapping from parameters to model output using a machine learning method that enables the rapid execution of the mapping, compared to the computationally expensive forward model. The calibration stage (section 2.2) results in $N_t = N_{\text{ens}} \cdot N_{\text{it}}$ input-output pairs $\{\theta_i, \mathcal{G}(\theta_i, \xi_i)\}_{i=1}^{N_t}$ of model parameter to model output. Harnessing the input-output pairs as training points, Gaussian process regression is used (Rasmussen, 2003) to create an emulator, composed of a mean function and covariance function pair, where $\mathcal{G}_{\text{GP}}(\theta) \approx \mathcal{G}_\infty(\theta)$ and $\Sigma_{\text{GP}} \approx \Sigma$. Since the input-output pairs are subject to different realizations of the chaotic system, the emulator mean approximates $\mathcal{G}_\infty(\theta)$ rather than $\mathcal{G}(\theta, \xi)$ (Cleary et al., 2021; Dunbar et al., 2020).

The variables of interest in the synthetic data $Y \in \mathbb{R}^{N \times N_d}$ are correlated. The output variables from the GCM forward model $\mathcal{G}(\theta, \xi)$ are similarly correlated. The correlation between the GCM statistics results in nonzero off-diagonal covariance matrix elements. In order to maintain a diagonal covariance in the GP emulator Σ_{GP} , we transform the GCM statis-

tics into a decorrelated space using principal component analysis on $\Sigma \in \mathbb{R}^{N \times N}$ (Cleary et al., 2021). That is, we decompose the covariance matrix using the singular value decomposition

$$\Sigma = VD^2V^\top, \quad (13)$$

with a matrix of principal component vectors (or singular vectors) V and the diagonal matrix D containing the square roots of the singular values σ_i . Often, in practical data assimilation problems, $N > N_d$ (Houtekamer & Zhang, 2016). In this case, the covariance matrix is rank deficient, with $\text{rank}(\Sigma) \leq \min(N, N_d)$, with singular values $\sigma_i^2 = 0$ for $i > \text{rank}(\Sigma)$. Methods to decorrelate the data and model output with rank deficient covariance matrices are discussed in Appendix A.

The GP is trained using input-output pairs in the decorrelated space $\{\theta_i, \tilde{\mathcal{G}}(\theta_i, \xi_i)\}_{i=1}^{N_t}$ with the decorrelated space denoted by tildes, $(\tilde{\cdot})$. The resulting input-output mapping is approximated as

$$\tilde{\mathcal{G}}_\infty(\theta) \approx N(\tilde{\mathcal{G}}_{\text{GP}}(\theta), \tilde{\Sigma}_{\text{GP}}(\theta)). \quad (14)$$

The GP kernels used are a white-noise kernel added to an Automatic Relevance Determination (ARD) radial basis function kernel; further details provided in Dunbar et al. (2020); Cleary et al. (2021). The GP hyperparameters are trained using the input-output pairs. The training process results in a GP regression function $\tilde{\mathcal{G}}_{\text{GP}}(\theta)$ which takes θ as input and emulates $\tilde{\mathcal{G}}_\infty(\theta)$ in a computationally efficient fashion.

2.4 Sample: Posterior sampling using MCMC

The Bayesian posterior distribution is approximated through MCMC sampling with the trained GP emulator $\tilde{\mathcal{G}}_{\text{GP}}(\theta)$. The data $\mathbf{y}$ is normalized and transformed into the decorrelated space $\tilde{\mathbf{y}}$, as in section 2.3. The Bayesian posterior distribution is approximated as (Stuart, 2010)

$$\mathbb{P}(\theta|\tilde{\mathbf{y}}) \propto \mathbb{P}(\tilde{\mathbf{y}}|\theta)\mathbb{P}(\theta) \quad (15)$$

$$\propto \exp\left(-\frac{1}{2}\|\tilde{\mathbf{y}} - \tilde{\mathcal{G}}_{\text{GP}}(\theta)\|_{\tilde{\Sigma}_{\text{GP}}(\theta) + \tilde{\Delta}}^2 - \frac{1}{2}\log \det\left(\tilde{\Sigma}_{\text{GP}}(\theta) + \tilde{\Delta}\right)\right) \mathbb{P}(\theta). \quad (16)$$

With a normal prior distribution governed by mean $\bar{\theta}$ and covariance Γ_θ the resulting MCMC objective function is

$$\Phi_{\text{MCMC}}(\theta, \tilde{\mathbf{y}}) = \exp\left(-\frac{1}{2}\|\tilde{\mathbf{y}} - \tilde{\mathcal{G}}_{\text{GP}}(\theta)\|_{\tilde{\Sigma}_{\text{GP}}(\theta) + \tilde{\Delta}}^2 - \frac{1}{2}\log \det\left(\tilde{\Sigma}_{\text{GP}}(\theta) + \tilde{\Delta}\right) - \frac{1}{2}\|\theta - \bar{\theta}\|_{\Gamma_\theta}^2\right). \quad (17)$$

We use a random walk Metropolis MCMC algorithm. The number of MCMC samples is set to $N_{\text{MCMC}} = 200,000$ with a burn-in of 10,000.

3 Seasonal GCM uncertainty quantification

We perform numerical experiments with an idealized GCM with a seasonal cycle. The GCM simulation setup is presented in section 3.1. Various climate statistics that we extract from the GCM and use in the numerical experiments are discussed in section 3.2 and shown in section 3.3.

3.1 Seasonal simulation setup

The GCM used in this study is based on the Geophysical Fluid Dynamics Laboratory's Flexible Modeling System (Frierson et al., 2006). The GCM simulates an idealized aquaplanet with a homogeneous mixed-layer slab ocean bottom boundary condition with a depth of 1 m. The GCM has been used previously for simulations of the hydrological cycle over a range of

climates (O’Gorman & Schneider, 2008) and to characterize seasonal variability in the tropics (Bordoni & Schneider, 2008; Merlis et al., 2013a, 2013b; Kaspi & Schneider, 2011; Bischoff & Schneider, 2014, 2016; Wei & Bordoni, 2018). The GCM is axisymmetric and statistically cyclostationary. The spectral transform method is used in the horizontal directions, and finite differencing in sigma coordinates is used in the vertical direction. The horizontal resolution used is T21 with $N_\phi = 32$ discrete latitude points on the transform grid, and 20 sigma levels ($\sigma_p = p/p_s$, where p is the pressure and p_s is the local surface pressure).

A two-stream gray radiation scheme is used. The top-of-atmosphere (TOA) insolation is prescribed and varies according to a seasonal cycle (Bordoni & Schneider, 2008; Wei & Bordoni, 2018). The diurnal cycle insolation variations are neglected, with a daily average insolation applied. The longwave and shortwave optical thicknesses depend on the latitude and pressure. The radiative effects of variations of atmospheric water vapor or clouds are neglected, and therefore, water vapor and cloud feedbacks are not included in the GCM.

The convection is parameterized using a simplified quasi-equilibrium Betts-Miller (SBM) scheme (A. K. Betts, 1986; A. Betts & Miller, 1986; Frierson, 2007). Vertical profiles of temperature and humidity are used to calculate precipitation and associated temperature and humidity changes through a relaxation to moist adiabatic reference profiles (Frierson, 2007). The relaxation is included as a forcing to the temperature and humidity balances

$$\frac{\partial T}{\partial t} + \dots = S_T - f_T \frac{T - T_{\text{ref}}}{\tau} \quad (18)$$

$$\frac{\partial q}{\partial t} + \dots = S_q - f_q f_T \frac{q - q_{\text{ref}}}{\tau}, \quad (19)$$

where the dots represent the dynamical terms in the equations, and S_T and S_q represent the temperature and specific humidity forcings aside from the convection scheme. The term f_T governs the spatiotemporal activation of the convection scheme, depends on the thermodynamic state, and is a function of z . The term f_q modifies the specific humidity relaxation (O’Gorman & Schneider, 2008). The reference temperature T_{ref} is a moist adiabat, chosen so that the convection scheme conserves enthalpy integrated over vertical columns (Frierson, 2007; O’Gorman & Schneider, 2008). The reference specific humidity q_{ref} is that which corresponds to a prescribed relative humidity with respect to the moist adiabat T_{ref} . For the UQ experiments, our focus are two parameters: the prescribed reference relative humidity parameter (θ_{RH}) and the relaxation timescale parameter (θ_τ).

The GCM simulation starts from vernal equinox, and the year length is $T_{\text{yr}} = 360$ d. All climate statistics used for UQ are zonally averaged. We use the Betts-Miller convection scheme with standard reference values of the parameters of $\theta^\dagger = (\theta_{\text{RH}}^\dagger, \theta_\tau^\dagger) = (0.7, 7200 \text{ s})$ (Frierson, 2007; O’Gorman & Schneider, 2008). UQ of θ relies on a prior knowledge about the convective parameters. We use wide prior distributions of the parameters, which enforce physical constraints, such as $0 < \theta_{\text{RH}} \leq 1$ and $\theta_\tau > 0$, but are otherwise uninformative (Dunbar et al., 2020). The selected priors are $\theta_{\text{RH}} \sim \text{Logit}(N(0, 1))$ and $\theta_\tau \sim \text{Log}(N(12\text{h}, (12\text{h})^2))$, that is, normal distributions of logit- and log-transformed parameters. The parameter priors are independent, although joint prior distributions could be used in future work.

3.2 Climate statistics used for UQ

We calibrate and perform UQ on the parameters of the convection scheme using more and less informative time-averaged data from GCM simulations. The informative statistics, such as the mid-tropospheric relative humidity, are chosen because they are strongly affected by the choice of the convective parameters (e.g., O’Gorman and Schneider (2008); Schneider and O’Gorman (2008); O’Gorman et al. (2011)). For comparison, we also choose less informative statistics that are affected less by the choice of convective parameters, such as surface wind speeds. Numerical UQ experiments using the differing degrees of information in the climate statistics will illustrate the impact of incorporating higher frequencies in the climate statistics used for parameter estimation.

	Informative	Less informative
Variables	1. Relative humidity ($\sigma_p = 0.5$) 2. Precipitation rate (mm/day) 3. Probability of 90 th percentile precipitation	1. Precipitation rate (mm/day) 2. Surface wind speed (m/s)
Latitudes, N_ϕ	32	32
N_y	96	64
$N, T = 360$ d (annual)	96	64
$N, T = 90$ d (seasonal)	384	256

Table 1: Summary of the climate data used. Calibration and UQ is performed using data that are more or less informative about the convective scheme. The total size of the data used for UQ is N .

For the informative climate statistics, we choose three variables: the mid-tropospheric ($\sigma_p = 0.5$) relative humidity, the total precipitation rate, and, as a measure of precipitation intensity, the probability that a daily precipitation total exceeds the latitude-dependent 90th precipitation percentile threshold from a long control simulation with the true parameters (Dunbar et al., 2020). Since intense precipitation events are influenced by the convection scheme (O’Gorman & Schneider, 2009a, 2009b), exceedances of precipitation over a high threshold are anticipated to be informative about the convection scheme parameters. The three statistics, are evaluated at each of the $N_\phi = 32$ discrete latitudes, giving 96 total quantities of interest. For less informative statistics, we use the zonally averaged precipitation rate and surface wind speed. As with the informative statistics, $N_\phi = 32$ discrete latitudes are considered, giving 64 total quantities of interest (see Table 1 for a summary).

We nondimensionalize the GCM statistics (Eq. (A1)) with the median y_c , taken over latitude and time, of each specified data type. There is one characteristic value for each data type, e.g., one characteristic value for precipitation and a separate characteristic value for relative humidity.

3.3 Synthetic data used for UQ

As synthetic data, we used GCM output with different initial conditions, and perturbed with measurement error with covariance Δ . The synthetic states are constructed by running the seasonal GCM for 150 years at the true parameters $\theta^\dagger$. For each case of $T = 90$ d and $T = 360$ d, the GCM states are aggregated according to the method introduced in section 2.1. The informative synthetic statistics for $T = 360$ d days are shown in Figure 2a–c, which compares the $N_d = 150$ year ensemble average of the statistics with a 1 year sample subject to internal variability and perturbed with measurement noise $N(0, \Delta)$. The synthetic data for $T = 90$ d are shown in Figure 2d–f. For the data averaged over 90 d, a distinct seasonal cycle emerges with, for example, relatively high precipitation and relative humidity in the northern hemisphere summer (days 90–180). The less informative statistics of precipitation rate and surface wind speed are shown for $T = 360$ d and $T = 90$ d in Figure 3.

4 GCM calibration and UQ results

4.1 Informative statistics

The convective parameters θ_{RH} and θ_τ are calibrated using EKI with the informative statistics with $T = 90$ d and $T = 360$ d (see Table 1). The EKI calibration is performed with a synthetic data sample y (Eq. 7). Since the synthetic data is subject to internal variability, the corresponding calibration is influenced by the synthetic sample values. The synthetic data sam-

Figure 2: Informative synthetic data with (a–c) $T = 360$ d and (d–f) $T = 90$ d. The light solid lines correspond to $\mathbf{y}_\infty$. The shaded regions correspond to a 95% confidence interval around $\mathbf{y}_\infty$ with covariance $\Sigma + \Delta$. The dark solid lines with circle markers correspond to a randomly selected sample of the synthetic data $\mathbf{y}$ that has been subjected to measurement noise $N(0, \Delta)$. Day 0 corresponds to vernal equinox and the Northern hemisphere seasons are provided.

Figure 3: Less informative synthetic data with (a, b) $T = 360$ d and (c, d) $T = 90$ d. The light solid lines correspond to $\mathbf{y}_\infty$. The shaded regions correspond to a 95% confidence interval around $\mathbf{y}_\infty$ with covariance $\Sigma + \Delta$. The dark solid lines with circle markers correspond to a randomly selected sample of the synthetic data $\mathbf{y}$ that has been subjected to measurement noise $N(0, \Delta)$. Day 0 corresponds to vernal equinox and the Northern hemisphere seasons are provided.

Figure 4: Ensemble Kalman inversion convective parameter estimates using informative statistics. Mode of (a) the estimated relative humidity convective parameter θ_{RH} compared to truth value $\theta_{RH}^\dagger$ and of (b) the estimated relaxation time scale convective parameter θ_τ compared to truth value $\theta_\tau^\dagger$. The mean and standard deviation over independent instantiations with differing synthetic samples are shown by the lines and shaded regions, respectively.

ple is randomly selected from the 150 years of historical data constructed with the true convective parameters $\theta^\dagger$ (see section 3.2). Calibration is performed with 10 independent synthetic samples, and the results in this section are presented as means and standard deviations over the 10 independent realizations. The cases with $T = 90$ d and $T = 360$ d are run with the same 10 synthetic samples.

The modes of the EKI ensemble distributions are shown in Figure 4. The modes demonstrate a reduction in estimation error of the true convective parameters when the data are aggregated seasonally with $T = 90$ d. The calibrated convective parameters using annually averaged GCM statistics have higher mean bias and variability between the instantiations with different synthetic data samples. For both integration timescales, the error associated with the θ_τ estimation was larger than the estimation error for θ_{RH} . The percent error of the average mode at the last EKI iteration for θ_{RH} is 0.3% and 0.8% for $T = 90$ d and $T = 360$ d, respectively. For θ_τ , the percent errors are 1.2% and 11.5% for $T = 90$ d and $T = 360$ d, respectively. This indicates that the integration of the GCM states over an annual cycle filters information from the resulting GCM statistics that is informative, especially about the relaxation timescale.

The normalized mean square error (MSE) for θ_{RH} at EKI iteration (n) is

$$\varepsilon_{RH}^{(n)} = \frac{1}{\theta_{RH}^\dagger} \sqrt{\frac{1}{N_{ens}} \sum_{i=1}^{N_{ens}} \left(\theta_{RH,i} - \theta_{RH}^\dagger \right)^2}. \quad (20)$$

The error is computed similarly for θ_τ . For both the relative humidity and relaxation timescale parameters, $T = 90$ d reduces the parameter estimation MSE relative to $T = 360$ d by about a factor 2–3: The MSE for θ_{RH} is 0.006 and 0.015 for $T = 90$ d and $T = 360$ d, respectively; for θ_τ , the MSE is 0.06 and 0.23 for $T = 90$ d and $T = 360$ d, respectively. The seasonally aggregated data also have a smaller spread among the realizations, as visualized by the uncertainty bands around the mean values. The reductions in the MSE standard deviations were 43% and 66% for θ_{RH} and θ_τ , respectively.

The input-output pairs generated during EKI parameter calibration are used for training the GP emulator. These pairs of parameters and model evaluations are taken from four 100-member EKI iterations as well as the initial 100-member ensemble drawn from the prior, giving a total of 500 input-output pairs for GP training. The results of this study were not significantly different with more training points. The GCM outputs are mapped into a normalized and decorrelated space according to the SVD of the internal variability covariance ma-

Figure 5: Convective parameter posterior distributions computed using MCMC using a GP emulator that was trained using data aggregated with an integration timescale of $T = 360$ d and $T = 90$ d. The contours correspond to 50%, 75%, and 99% of the posterior distribution, the star is the truth, and the circle is the average of the ensemble members after the last EKI iteration. Posteriors are shown for (a) $T = 360$ d, decorrelated with full SVD, and (b) a 95%-energy truncation of the SVD. (c) Posterior for $T = 90$ d, decorrelated with 95%-energy truncated SVD. Panels (d–f) same as (a–c) with the posterior distributions shown in the physical parameter space.

Parameter space	$T = 90$ d	$T = 360$ d
Informative data	$0.5\% \pm 0.1\%$	$1.3\% \pm 0.3\%$
Less informative data	$5.5\% \pm 0.9\%$	$26.0\% \pm 7.0\%$

Table 2: The ratio (%) of the area occupied by the posterior to the area occupied by the prior. The area of the convex hull containing 75% of the posterior mass for each value of T and data type is normalized by the area containing 75% of the prior mass. The means and standard deviations over independent synthetic data realizations are provided. Ten and four realizations are used for the informative and less informative data, respectively.

trix Σ , truncated to contain 95% of the energy in the singular values. For $T = 90$ d, the truncation contains $k = 47$ singular values, while $T = 360$ d retains $k = 23$ (see Appendix B). A scalar-valued GP is trained for each of the k outputs. For comparison, we also ran the emulation and sampling steps for $T = 360$ d without SVD truncation.

The posterior distributions are approximated using MCMC sampling with the trained GP emulators. Posterior distributions for the two convective parameters for a randomly selected synthetic data sample are shown in Figure 5. In Figure 5a, the posterior distribution was computed with normalized input data, aggregated annually with $T = 360$ d, but the SVD was not truncated before emulation and sampling. The posterior resulting from the same initial ensemble and synthetic data sample but with SVD truncation is shown in Figure 5b. The truncation of the SVD smooths and inflates the posterior distribution. While the posteriors are qualitatively similar, they differ quantitatively. The true parameters are outside the region containing 99% of the posterior mass for the untruncated case but are within the region containing 75% of the posterior mass for the truncated case. Further investigation of the influence of SVD regularization on the posteriors is in Appendix B. Hereafter, all posterior results will focus on the posterior distributions estimated using emulation and sampling in the truncated SVD space (see Appendix A). The convective parameter posterior distribution using seasonally aggregated data with $T = 90$ d with the SVD truncated at 95% of the total energy is shown in Figure 5c. As in the calibration stage, the posterior mode for $T = 90$ d is closer to the true parameters than for $T = 360$ d. Compared to the posterior with $T = 360$ d, the $T = 90$ d posterior is also more compact.

We quantify the size of the posterior distributions for the two integration timescales by computing the two-dimensional area of the convex hull containing 75% of the posterior mass (middle contour level in Figure 5). The area containing 75% of the posterior mass for each value of T is normalized by the area containing 75% of the prior mass for reference. The percentages of the prior area taken up by the posteriors for $T = 90$ d and $T = 360$ d, averaged over the 10 CES instantiations, are shown in Table 2. The areas are computed in the $\log(\theta_{RH}) - \log(\theta_{\tau}/1 \text{ s})$ transformed space, in which the sampling is performed. The area of the posterior distribution resulting from the seasonal integration is approximately a factor 2 smaller than the posterior from the annually integrated data. However, compared to the wide prior, both posteriors are tight, with the posterior areas occupying around 1% of the prior area.

4.2 Less informative statistics

In many UQ applications, intuition about which quantities of interest will lead to improved parameter estimation is not available before experimentation. While the idealized GCM and convective parameterization used in this study have a rich set of investigations to provide prior knowledge about the relationship between θ and $\mathcal{G}(\theta)$, we also consider UQ of the parameters using less informative GCM statistics. To do so, we use the precipitation rate and surface wind speed, zonally averaged and integrated with $T = 360$ d and $T = 90$ d (Figure

Figure 6: Parameter calibration performed with EKI using less informative statistics. Mean square error of (a) the estimated relative humidity parameter θ_{RH} compared to truth value $\theta_{RH}^\dagger$ and of (b) the estimated relaxation time scale θ_τ compared to truth value $\theta_\tau^\dagger$. Solid lines are $T = 360$ d and dashed lines are $T = 90$ d.

3). Four instantiations of CES with differing synthetic data samples are run for each integration length of $T = 90$ d and $T = 360$ d.

The MSE of the EKI parameter calibration with the less informative GCM statistics are shown in Figure 6 for $T = 90$ d and $T = 360$ d. With precipitation rate and surface wind speed as the basis for UQ, the annually averaged data result in significantly higher parameter estimation error than seasonally aggregated data. The $T = 360$ d MSEs are 0.26 and 1.56 for θ_{RH} and θ_τ , respectively. The $T = 90$ d MSEs are 0.02 and 0.28 for θ_{RH} and θ_τ , respectively. The ensemble means and modes of the two parameters are shown in Figure 7. In Appendix C, the calibration is performed by considering the seasonal data sequentially, rather than collectively, to restrict the data size for both $T = 90$ d and $T = 360$ d to $N = 64$. The results demonstrate that the reductions in MSE are the result of the incorporation of seasonal information in the data, rather than the dimensionality of the data space.

Posterior distributions computed from CES with GP training data generated by EKI for $T = 360$ d and $T = 90$ d are shown in Figure 8. We use SVD truncation for both $T = 360$ d and $T = 90$ d (Appendix A). The posterior for $T = 360$ d is significantly larger than for $T = 90$ d, indicating that the annually averaged precipitation and surface wind speed do not provide substantive information regarding the convective parameters. As with EKI for $T = 360$ d, the posterior distribution mode has substantial error when compared to the true parameters. For $T = 90$ d, the uncertainty for θ_{RH} estimated with the less informative data has collapsed to a similar order as the uncertainty estimated with informative data. While the uncertainty is larger for θ_τ when estimated using the less informative data, the overall posterior area is reasonably collapsed, as it is of similar order to the $T = 360$ d posterior with informative data. These results suggest that seasonally averaged precipitation rate and surface wind speed are sufficient statistics to estimate convective parameters while annually averaged precipitation and wind speed introduce higher estimation error and parameter uncertainty.

As with the informative GCM statistics, we compare the sizes of the posterior distributions by computing the two-dimensional area of the convex hull containing 75% of the posterior mass. The posterior areas, normalized by the prior area, averaged over the four CES instantiations, are shown in Table 2. The posterior area for $T = 360$ d is approximately 5 times larger than the area produced with $T = 90$ d. The implications of the reduction in the size of the posterior distribution on the parametric uncertainty in the GCM are tested in section 4.2.1.

Figure 7: Parameter calibration performed with EKI using less informative statistics. Mean of (a) the estimated relative humidity parameter θ_{RH} compared to truth value $\theta_{RH}^{\dagger}$ and of (b) the estimated relaxation time scale parameter θ_{τ} compared to truth value $\theta_{\tau}^{\dagger}$. (c, d) Same as (a, b) for the modes of the ensemble members.

Figure 8: Posterior density for (a) $T = 360$ d and (b) $T = 90$ d for the less informative statistics. Contours correspond to 50%, 75%, and 99% of the posterior distribution. The location of the true parameters are indicated by the star, and the circle is the average of the ensemble members after the last EKI iteration.

Figure 9: Samples of size 100 from the posterior distributions for prediction experiments. Black are samples from the $T = 90$ d posterior and red are samples from the $T = 360$ d posterior. The true parameters are shown with dashed blue lines.

4.2.1 Prediction experiments

To demonstrate the effect of parametric uncertainty in the GCM on climate predictions, we draw samples of 100 parameter pairs from the posterior distributions resulting from the less informative statistics of precipitation rate and surface wind speed (Figure 8). The samples of parameter pairs using seasonal integration ($T = 90$ d) and annual integration ($T = 360$ d) in the UQ are shown in Figure 9. With each sample of parameters, we produce climate quantities of interest, averaged over a 20 year period. We compute ensemble statistics over the outputs from the samples (prediction uncertainty) and compare these with a 20 year simulation with the true parameters $\theta^\dagger$.

The relative humidity at $\sigma_p = 0.5$, the precipitation rate, and the intense precipitation probability (the probability of exceeding a 90th percentile threshold averaged over the 20 year simulation at $\theta^\dagger$) are shown in Figure 10. The posterior distribution was estimated using (seasonally averaged) precipitation rate data directly, but neither relative humidity nor intense precipitation probability were used. The mean absolute percent errors for the mean prediction values, averaged across latitudes, for the $T = 90$ d posterior are 1.2%, 0.6%, and 1.4%, for the relative humidity, precipitation rate, and intense precipitation probability, respectively. For the predictions with the $T = 360$ d posterior, the mean absolute percent errors are 4.7%, 1.8%, and 21.5%.

The prediction uncertainty in each climate quantity of interest is indicated by a 95% confidence interval, estimated as the values of the 2.5th and 97.5th percentiles of the variables of interest at each latitude over the convective parameter pairs (Figure 9). For both $T = 90$ d and $T = 360$ d, there is limited prediction uncertainty in the mean precipitation rate. The widths of the confidence intervals are 6% and 4% of the mean precipitation rate, averaged across latitudes, for $T = 360$ d and $T = 90$ d, respectively. The prediction uncertainties in relative humidity and intense precipitation probability are significantly higher for UQ performed with annually averaged data ($T = 360$ d) compared to UQ with seasonally averaged data ($T = 90$ d). The $T = 90$ d posterior reduces the size of the 95% confidence interval for the relative humidity by 70%, averaged across latitudes, compared to $T = 360$ d. The reduction in the size of the 95% confidence interval for the intense precipitation probability is 9%.

We also performed idealized global-warming prediction experiments. To represent global warming in the idealized GCM, the longwave opacity in the atmosphere is increased by 50%, as in O’Gorman and Schneider (2008) which results in a global-mean surface air temperature increase from 284 K in the control climate to 292 K in the warm climate. We accumulated

Figure 10: Prediction experiments for the case with no imposed warming for (a) relative humidity at $\sigma_p = 0.5$, (b) precipitation rate, and (c) intense precipitation probability. The climate statistics are averaged over 20 years. The posteriors from the less informative statistics (surface wind speed and precipitation rate) are used. The lines correspond to the mean of the predictions and the shaded regions, correspond to 95% confidence intervals. The climate statistics at the true parameters $\theta^\dagger$ are shown with a dashed line.

Figure 11: Prediction experiments for the case with imposed warming (50% increase in longwave opacity) for (a) relative humidity at $\sigma_p = 0.5$, (b) precipitation rate, and (c) intense precipitation probability. The climate statistics are averaged over 20 years. The posteriors from the less informative statistics (surface wind speed and precipitation rate) are used. The lines correspond to the mean of the predictions, and the shaded regions correspond to 95% confidence intervals. The climate statistics at the true parameters $\theta^\dagger$ are shown with a dashed line.

GCM statistics over a 20 year window after a spinup of one year, running simulations with the true GCM parameters $\theta^\dagger$ with increased longwave opacity for comparison. We use the same 100 parameter pairs drawn from the posterior distributions estimated from the less informative climate statistics.

The idealized global warming results are shown in Figure 11. As expected, the warming climate predictions with the $T = 90$ d posterior distribution are more accurate than with the $T = 360$ d posterior, especially for the relative humidity and intense precipitation probability. The mean absolute percent error for the mean prediction values, averaged across latitudes, for the $T = 90$ d posterior are 1.7%, 1.1%, and 2.2%, for the relative humidity, precipitation rate, and intense precipitation probability, respectively. For the predictions with the $T = 360$ d posterior, the mean absolute percent errors are 4.7%, 2.2%, and 10.4%. The $T = 90$ d posterior reduces the size of the 95% confidence interval for the relative humidity in a warmed climate by 75% compared to $T = 360$ d, averaged across latitudes. The reduction in the size of the 95% confidence interval for the intense precipitation probability in a warmed climate is 35%. These results demonstrate that the incorporation of the seasonal cycle in the UQ reduces the parametric uncertainty in climate predictions in this model.

5 Conclusions

We performed calibration and UQ of a convective parameterization in a seasonally forced idealized GCM. While GCMs are typically tuned using annually and globally averaged climate statistics (Hourdin et al., 2017), our results demonstrate, in an idealized setting, the qualitative, quantitative and systematic refinement of parameter distributions through the incorporation of seasonal information. Performing parameter calibration with seasonally averaged data significantly reduced the error associated with the estimated parameters compared to calibration using annually averaged data. The posterior distributions resulting from the Bayesian UQ with seasonally averaged data were reduced in size. One measure of success for Bayesian UQ, is to capture the true parameters within a high mass region of the posterior distribution (demonstrated in our results). We are also interested in the shape and size of the posterior distribution, as this provides valuable information about parameter correlations and uncertainty, with respect to observed data. We demonstrate that choosing suitable data for UQ which reduce the size of the posterior distribution can lead to significant reductions of climate prediction uncertainty.

The impact of incorporating additional frequency content is pronounced when the climate statistics used for UQ are less informative about the parameterizations. Such situations often occur in the climate modeling setting where a number of parameterizations may simultaneously and nonlinearly influence a quantity of interest and where it is not always clear which climate statistics should be used to calibrate an unknown GCM parameter.

To enable UQ in the climate model setting, we used the calibrate-emulate-sample (CES) methodology (Cleary et al., 2021), which enables the efficient and accurate estimation of Bayesian posterior distributions of parameters from noisy climate data. CES uses gradient-free optimization to calibrate parameters and generate parameter-data pairs, Gaussian process regression to emulate the parameter-to-data mapping, and Markov Chain Monte Carlo (MCMC) to sample from the posterior distribution. The emulation and sampling is performed in a decorrelated dataspace through a transformation based on the principal component analysis of the noise covariance matrix. In this study, we modified the original CES methodology by first normalizing the data to ensure all statistics are weighted equally in the UQ and second by regularizing the covariance matrix to enable the UQ of applications with ill-conditioned or rank deficient covariance matrices. We quantified the impact of the normalization and regularization, with the regularization smoothing the posterior and slightly increasing its size.

Beyond convection, Earth system models rely on a number of parameterizations of cyclostationary processes, including models of carbon accumulation and storage (Bloom et al., 2016) and of atmospheric boundary layer (ABL) turbulence which parameterize unclosed subgrid processes (Stull, 2012). Carbon storage is inherently seasonal (Rowland et al., 2014). The ABL has distinct seasonal and diurnal variations (Wyngaard, 2010; Howland, González, et al., 2020). We expect our results to be relevant to calibration and UQ for statistical variations on different timescales, beyond the seasonal cycle. For example, based on our findings presented here, we anticipate that incorporating the diurnal variation of ABL turbulence statistics may improve the UQ of subgrid scale turbulence models in GCMs. More broadly, we anticipate the findings of this study to be relevant to UQ in problems with multi-scale temporal dynamics.

The selection of the optimal aggregation timescale for climate statistics and choice of objective function remain open questions. In the context of parameter estimation, the integration timescale used to generate time-averaged statistics is a hyperparameter that dictates the trade-off between the frequency content and signal-to-noise ratio of the climate statistics. In the limiting case of $T \rightarrow 0$, the influence of the initial conditions on the state is pronounced, as in weather forecasting (Houtekamer & Zhang, 2016). Conversely, as $T \rightarrow \infty$, all frequency content is removed from the statistics, which, as this study demonstrates, can adversely impact parameter estimation. We anticipate that the selection of the filter timescale T may be problem and parameterization specific, as it relates to the question of parameter identifiability.

ity. In this study, we selected seasonal averages because of the large amplitude of the seasonal cycle in many climate statistics and indications that seasonal variations are informative about the climate change response in many climate variables (Schneider et al., 2021). Future work should develop a more generalized approach for the selection of T based on the frequency content of the quantities of interest.

Appendix A Data and model output decorrelation for rank deficient problems

To facilitate the transformation into the uncorrelated space with rank deficient covariance matrices, the singular value decomposition/principal component analysis is truncated. Since the statistical quantities of interest have a range of magnitudes, the data are first normalized, and then the singular value decomposition is truncated as a form of regularization (Hansen, 1987). A detailed investigation of the influence of the normalization and truncation on the posteriors resulting from CES is performed in Appendix B.

The data used for UQ are provided concurrently to the CES pipeline in the concatenated vector $\mathbf{y}$. In this framework, the statistical quantities in the data $\mathbf{y}$ may have a range of magnitudes, and normalization is required before regularization (e.g., Tibshirani (1996)). The data are normalized by a characteristic value associated with each individual data type $\mathbf{y}_c \in \mathbb{R}^N$. The normalized data are

$$\mathbf{y}_i^* = \mathbf{y}_i \cdot \mathbf{y}_{c,i}^{-1}, \quad (\text{A1})$$

where $(*)$ denotes normalized data. The characteristic values in $\mathbf{y}_c$ are described in the application of CES to the seasonally forced GCM in section 3. Each column of $\mathbf{Y}$ is normalized using Eq. (A1) to yield the normalized data matrix $\mathbf{Y}^*$. The normalized covariance is $\Sigma^* = \text{cov}(\mathbf{Y}^*)$. The SVD transformation is performed using covariance matrix Σ^* and data $\mathbf{y}^*$.

The SVD is truncated in order to account for rank deficient or ill-conditioned covariance matrices (Hansen, 1987). The truncated SVD is defined for covariance Σ^*

$$\Sigma^* \approx \Sigma_k^* = \mathbf{V}_k \mathbf{D}_k^{*2} \mathbf{V}_k^\top \quad \mathbf{D}_k^* = \text{diag}(\sigma_1^*, \dots, \sigma_k^*),$$

truncated at the $k \leq \text{rank}(\Sigma^*)$ singular value. The truncated singular vector matrix is $\mathbf{V}_k = [\mathbf{v}_1, \dots, \mathbf{v}_k]$ where $\mathbf{v}_i$ is the singular vector corresponding to the singular value σ_i^{*2} . The truncated SVD space is given by

$$\tilde{\Sigma}_k = \mathbf{D}_k^{-1} \mathbf{V}_k^\top \Sigma^* \mathbf{V}_k \mathbf{D}_k^{-1} \quad (\text{A2})$$

$$\tilde{\mathbf{y}}_k = \mathbf{D}_k^{-1} \mathbf{V}_k^\top \mathbf{y}^*. \quad (\text{A3})$$

In this study, the truncation location is selected as the lowest value of k such that $\sum_{i=1}^k \sigma_i^{*2} \geq 0.95 \sum_{i=1}^N \sigma_i^{*2}$, so that 95% of the variance is retained. The GP output can be mapped to the original normalized space using

$$\mathcal{G}_{\text{GP}}^* = \mathbf{V}_k \mathbf{D}_k \tilde{\mathcal{G}}_{\text{GP}} \quad (\text{A4})$$

$$\Sigma_{\text{GP}}^* = \mathbf{V}_k \mathbf{D}_k \tilde{\Sigma}_{\text{GP}} \mathbf{D}_k \mathbf{V}_k^\top. \quad (\text{A5})$$

Finally, the GP output can be transformed into the dimensional space through $\mathcal{G}_{\text{GP}} = \mathcal{G}_{\text{GP}}^* \odot \mathbf{y}_c$ and $\Sigma_{\text{GP}} = \Sigma_{\text{GP}}^* \odot \mathbf{y}_c \mathbf{y}_c^\top$, where $\odot$ denotes pointwise multiplication.

Appendix B Sensitivity of the posterior to covariance regularization

In section 2.3, the SVD is truncated as a form of regularization. Here, we detail the sensitivity of the CES posterior distributions to the SVD regularization using informative GCM statistics (see Table 1) and annually averaged data with $T = 360$ d, for which the data covariance matrix is full rank. The data are normalized, as discussed in Appendix A. The cumulative sum of the singular values is shown in Figure B1. Truncation at 95% of the singular value energy corresponds to $k = 23$. Regularization can also be performed using Tikhonov

Figure B1: Cumulative sum of the singular values from the normalized $T = 360$ d synthetic data.

Parameter space	Original	Normalized	Truncated	Tikhonov
Informative data	0.5%	0.5%	0.8%	0.9%

Table B1: The ratio (%) of the area occupied by the posterior to the area occupied by the prior. The area of the convex hull containing 75% of the posterior mass for each regularization method is normalized by the area containing 75% of the prior mass.

regularization, which inflates the diagonal elements (Hansen, 1987; Schneider, 2001). Following Hansen (1987), regularization in the form of diagonal inflation $\Lambda = \text{diag}(\lambda^2)$ is added to Σ where $\lambda = (\sigma_k^3 \sigma_{k+1})^{1/4}$. The SVD is performed such that

$$\Sigma + \Lambda = VD^2V^\top, \quad (\text{B1})$$

where D^2 is a diagonal matrix of the singular values $\sigma_i^2 + \lambda^2$. The orthonormal eigenvectors are in V . The influence of the truncation and Tikhonov regularizations on the posterior areas for $T = 360$ d are provided in Table B1. The $T = 360$ d posterior distributions are shown in Figure B2. The regularization slightly increases the posterior area and smooths the posterior roughness.

Appendix C Parameter estimation with filtering and smoothing

It is often of practical interest to estimate parameters in problems that are not statistically stationary, such as the seasonally forced GCM with time-dependent boundary conditions. In problems which are not statistically stationary (e.g., with time-evolving boundary conditions), all requisite data for parameter UQ may not be available at the initialization of the estimation. This arises in the EnKF setting for state estimation (e.g., Houtekamer and Zhang (2016)).

We can instead pose EKI as a filtering approach for parameter estimation where the synthetic data are collected sequentially in a time-dependent problem. This differs from the definition of the data in section 2.1, where the data from each season are collected together into $\mathbf{y}$ (termed smoothing). For the filtering approach, the data is observed and averaged over time $t_j \rightarrow t_j + T$ in $\mathbf{y}_{T,j}$, where j indicates the time of the year in the GCM simulation. The EKI update is performed using $\mathbf{y}_{T,j}$ and $\mathcal{G}_{T,j}(\boldsymbol{\theta}^{(k)})$ to update $\boldsymbol{\theta}^{(k+1)}$, and this process is repeated for the specified number of EKI iterations. The smoothing approach leverages the data averaged in each season concurrently ($\mathbf{y}$ and $\mathcal{G}(\boldsymbol{\theta}^{(k)})$) whereas the filtering approach uses the data averaged in each season sequentially ($\mathbf{y}_{T,j}$ and $\mathcal{G}_{T,j}(\boldsymbol{\theta}^{(k)})$), as it becomes available.

The parameter calibration results comparing filtering and smoothing are shown in Figure C1. For $T = 360$ d, filtering and smoothing are identical. For $T = 90$ d, the filtering

Figure B2: Posterior distributions using the informative statistics with $T = 360$ d. Comparison of the posteriors estimated with (a) original, (b) normalized, (c) normalized and regularized (truncation), and (d) normalized and regularized (Tikhonov diagonal inflation) covariance matrix Σ .

approach leverages the data from one season at a time to update the parameters θ . The final parameter errors for the $T = 90$ d filtering and smoothing approaches are similar, and are both significantly less than the error for the $T = 360$ d case. This result also demonstrates that the improvement in parameter estimates with $T = 90$ d, compared to $T = 360$ d, is a result of increased temporal information, rather than the dimensionality of the data space alone. Depending on the application, the filtering approach may be beneficial for computational efficiency, since it reduces the length of the forward model simulations, instead of requiring temporal integration over the timescale corresponding to an impactful low-frequency cycle.

Acknowledgments

This work was supported by the generosity of Eric and Wendy Schmidt by recommendation of the Schmidt Futures program and by the National Science Foundation (NSF, award AGS-1835860). We thank Andrew Stuart for thoughtful comments on the work and the manuscript.

Data Availability.

All computer code used in this paper is open source. The code is available at:
<http://doi.org/10.5281/zenodo.5138467>.
 An open-source Julia implementation of CES is accessible at:
<https://github.com/CliMA/CalibrateEmulateSample.jl>

References

Annan, J., & Hargreaves, J. (2007). Efficient estimation and ensemble generation in climate modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical*

Figure C1: Parameter calibration performed with EKI using less informative statistics. Mean square error of (a) the estimated relative humidity convective parameter θ_{RH} compared to truth value $\theta_{RH}^\dagger$ and of (b) the estimated relaxation time scale convective parameter θ_τ compared to truth value $\theta_\tau^\dagger$. Solid lines are $T = 360$ d and dashed lines are $T = 90$ d.

- and Engineering Sciences, 365(1857), 2077–2088.
- Bauer, P., Thorpe, A., & Brunet, G. (2015). The quiet revolution of numerical weather prediction. *Nature*, 525(7567), 47–55.
- Beskos, A., Pinski, F. J., Sanz-Serna, J. M., & Stuart, A. M. (2011). Hybrid monte carlo on hilbert spaces. *Stochastic Processes and their Applications*, 121(10), 2201–2230.
- Beskos, A., Roberts, G., Stuart, A., & Voss, J. (2008). Mcmc methods for diffusion bridges. *Stochastics and Dynamics*, 8(03), 319–350.
- Betts, A., & Miller, M. (1986). A new convective adjustment scheme. Part II: Single column tests using gate wave, bomex, atex and arctic air-mass data sets. *Quarterly Journal of the Royal Meteorological Society*, 112(473), 693–709.
- Betts, A. K. (1986). A new convective adjustment scheme. Part I: Observational and theoretical basis. *Quarterly Journal of the Royal Meteorological Society*, 112(473), 677–691.
- Bischoff, T., & Schneider, T. (2014). Energetic constraints on the position of the Intertropical Convergence Zone. *J. Climate*, 27, 4937–4951.
- Bischoff, T., & Schneider, T. (2016). The equatorial energy balance, ITCZ position, and double-ITCZ bifurcations. *J. Climate*, 29, 2997–3013.
- Bloom, A. A., Exbrayat, J.-F., Van Der Velde, I. R., Feng, L., & Williams, M. (2016). The decadal state of the terrestrial carbon cycle: Global retrievals of terrestrial carbon allocation, pools, and residence times. *Proceedings of the National Academy of Sciences*, 113(5), 1285–1290.
- Bony, S., & Dufresne, J.-L. (2005). Marine boundary layer clouds at the heart of tropical cloud feedback uncertainties in climate models. *Geophysical Research Letters*, 32(20).
- Bordoni, S., & Schneider, T. (2008). Monsoons as eddy-mediated regime transitions of the tropical overturning circulation. *Nature Geoscience*, 1(8), 515–519.
- Brient, F., & Schneider, T. (2016). Constraints on climate sensitivity from space-based measurements of low-cloud reflection. *Journal of Climate*, 29(16), 5821–5835.
- Brooks, S., Gelman, A., Jones, G., & Meng, X.-L. (2011). *Handbook of markov chain monte carlo*. CRC press.
- Cess, R. D., Potter, G., Blanchet, J., Boer, G., Ghan, S., Kiehl, J., ... others (1989). Interpretation of cloud-climate feedback as produced by 14 atmospheric general circulation models. *Science*, 245(4917), 513–516.
- Chen, Y., & Oliver, D. S. (2012). Ensemble randomized maximum likelihood method as an iterative ensemble smoother. *Mathematical Geosciences*, 44(1), 1–26.
- Cleary, E., Garbuno-Inigo, A., Lan, S., Schneider, T., & Stuart, A. M. (2021). Calibrate, emulate, sample. *Journal of Computational Physics*, 424, 109716.

- 593 Constantinescu, E. M., Zavala, V. M., Rocklin, M., Lee, S., & Anitescu, M. (2010). A
594 computational framework for uncertainty quantification and stochastic optimization in
595 unit commitment with wind power generation. *IEEE Transactions on Power Systems*,
596 26(1), 431–441.
- 597 Delworth, T. L., Rosati, A., Anderson, W., Adcroft, A. J., Balaji, V., Benson, R., ... others
598 (2012). Simulated climate and climate change in the gfdl cm2. 5 high-resolution
599 coupled climate model. *Journal of Climate*, 25(8), 2755–2781.
- 600 Dunbar, O. R., Garbuno-Inigo, A., Schneider, T., & Stuart, A. M. (2020). Calibration and un-
601 certainty quantification of convective parameters in an idealized GCM. *arXiv preprint*
602 *arXiv:2012.13262*.
- 603 Duncan, A. B., Stuart, A. M., & Wolfram, M.-T. (2021). Ensemble inference methods for
604 models with noisy and expensive likelihoods. *arXiv preprint arXiv:2104.03384*.
- 605 Duraisamy, K., Iaccarino, G., & Xiao, H. (2019). Turbulence modeling in the age of data.
606 *Annual Review of Fluid Mechanics*, 51, 357–377.
- 607 Emerick, A. A., & Reynolds, A. C. (2013). Ensemble smoother with multiple data assimi-
608 lation. *Computers & Geosciences*, 55, 3–15.
- 609 Evensen, G. (2003). The ensemble kalman filter: Theoretical formulation and practical im-
610 plementation. *Ocean dynamics*, 53(4), 343–367.
- 611 Frierson, D. M. (2007). The dynamics of idealized convection schemes and their effect on
612 the zonally averaged tropical circulation. *Journal of the Atmospheric Sciences*, 64(6),
613 1959–1976.
- 614 Frierson, D. M., Held, I. M., & Zurita-Gotor, P. (2006). A gray-radiation aquaplanet moist
615 GCM. part I: Static stability and eddy scale. *Journal of the atmospheric sciences*,
616 63(10), 2548–2566.
- 617 Garbuno-Inigo, A., Hoffmann, F., Li, W., & Stuart, A. M. (2020). Interacting langevin diffu-
618 sions: Gradient structure and ensemble kalman sampler. *SIAM Journal on Applied Dy-*
619 *nammical Systems*, 19(1), 412–441.
- 620 Geyer, C. (2011). Introduction to markov chain monte carlo. *Handbook of markov chain*
621 *monte carlo*, 20116022, 45.
- 622 Golaz, J.-C., Horowitz, L. W., & Levy, H. (2013). Cloud tuning in a coupled climate model:
623 Impact on 20th century warming. *Geophysical Research Letters*, 40(10), 2246–2251.
- 624 Hansen, P. C. (1987). The truncatedsvd as a method for regularization. *BIT Numerical Math-*
625 *ematics*, 27(4), 534–553.
- 626 Hoegh-Guldberg, O., Jacob, D., Taylor, M., Bindi, M., Brown, S., Camilloni, I., ... Zhou, G.
627 (2018). Impacts of 1.5°C global warming on natural and human systems supplemen-
628 tary material. In *Global warming of 1.5°C. An IPCC Special Report on the impacts*
629 *of global warming of 1.5°C above pre-industrial levels and related global greenhouse*
630 *gas emission pathways, in the context of strengthening the global response to the*
631 *threat of climate change, sustainable development, and efforts to eradicate poverty*.
- 632 Hourdin, F., Grandpeix, J.-Y., Rio, C., Bony, S., Jam, A., Cheruy, F., ... others (2013).
633 Lmdz5b: the atmospheric component of the ipsl climate model with revisited parame-
634 terizations for clouds and convection. *Climate Dynamics*, 40(9-10), 2193–2222.
- 635 Hourdin, F., Mauritsen, T., Gettelman, A., Golaz, J.-C., Balaji, V., Duan, Q., ... others
636 (2017). The art and science of climate model tuning. *Bulletin of the American Meteo-*
637 *rological Society*, 98(3), 589–602.
- 638 Houtekamer, P. L., & Zhang, F. (2016). Review of the ensemble kalman filter for atmo-
639 spheric data assimilation. *Monthly Weather Review*, 144(12), 4489–4532.
- 640 Howland, M. F., Ghate, A. S., Lele, S. K., & Dabiri, J. O. (2020). Optimal closed-loop wake
641 steering—part 1: Conventionally neutral atmospheric boundary layer conditions. *Wind*
642 *Energy Science*, 5(4), 1315–1338. Retrieved from [https://doi.org/10.5194/](https://doi.org/10.5194/wes-5-1315-2020)
643 [wes-5-1315-2020](https://doi.org/10.5194/wes-5-1315-2020)
- 644 Howland, M. F., González, C. M., Martínez, J. J. P., Quesada, J. B., Larranaga, F. P., Yadav,
645 N. K., ... Dabiri, J. O. (2020). Influence of atmospheric conditions on the power
646 production of utility-scale wind turbines in yaw misalignment. *Journal of Renewable*
647 *and Sustainable Energy*, 12(6), 063307.

- Iglesias, M. A., Law, K. J., & Stuart, A. M. (2013). Ensemble kalman methods for inverse problems. *Inverse Problems*, 29(4), 045001.
- Jakob, C. (2010). Accelerating progress in global atmospheric model development through improved parameterizations: Challenges, opportunities, and strategies. *Bull. Amer. Meteor. Soc.*, 91, 869–875.
- Kaipio, J., & Somersalo, E. (2006). *Statistical and computational inverse problems* (Vol. 160). Springer Science & Business Media.
- Kalnay, E. (2003). *Atmospheric modeling, data assimilation and predictability*. Cambridge, UK: Cambridge Univ. Press.
- Kaspi, Y., & Schneider, T. (2011). Winter cold of eastern continental boundaries induced by warm ocean waters. *Nature*, 471(7340), 621–624.
- Knutti, R., Meehl, G. A., Allen, M. R., & Stainforth, D. A. (2006). Constraining climate sensitivity from the seasonal cycle in surface temperature. *Journal of Climate*, 19(17), 4224–4233.
- Lan, S., Li, S., & Shahbaba, B. (2021). Scaling up bayesian uncertainty quantification for inverse problems using deep neural networks. *arXiv preprint arXiv:2101.03906*.
- Le Gland, F., Monbet, V., & Tran, V.-D. (2009). *Large sample asymptotics for the ensemble Kalman filter* (Research Report No. RR-7014). INRIA. Retrieved from <https://hal.inria.fr/inria-00409060>
- Li, G., & Reynolds, A. C. (2009). An iterative ensemble kalman filter for data assimilation. In (Vol. 14, pp. 496–505). Society of Petroleum Engineers.
- Mauritsen, T., Stevens, B., Roeckner, E., Crueger, T., Esch, M., Giorgetta, M., ... Tomassini, L. (2012). Tuning the climate of a global model. *J. Adv. Model. Earth Sys.*, 4, M00A01.
- Meneveau, C., & Katz, J. (2000). Scale-invariance and turbulence models for large-eddy simulation. *Ann. Rev. Fluid Mech.*, 32(1), 1–32.
- Merlis, T. M., Schneider, T., Bordoni, S., & Eisenman, I. (2013a). Hadley circulation response to orbital precession. Part I: Aquaplanets. *J. Climate*, 26, 740–753.
- Merlis, T. M., Schneider, T., Bordoni, S., & Eisenman, I. (2013b). The tropical precipitation response to orbital precession. *J. Climate*, 26, 2010–2021.
- O’Gorman, P. A., Lamquin, N., Schneider, T., & Singh, M. S. (2011). The relative humidity in an isentropic advection–condensation model: Limited poleward influence and properties of subtropical minima. *J. Atmos. Sci.*, 68, 3079–3093.
- O’Gorman, P. A., & Schneider, T. (2009a). The physical basis for increases in precipitation extremes in simulations of 21st-century climate change. *Proceedings of the National Academy of Sciences*, 106(35), 14773–14777.
- O’Gorman, P. A., & Schneider, T. (2009b). Scaling of precipitation extremes over a wide range of climates simulated with an idealized GCM. *Journal of Climate*, 22(21), 5676–5685.
- O’Gorman, P. A., & Schneider, T. (2008). The hydrological cycle over a wide range of climates simulated with an idealized GCM. *Journal of Climate*, 21, 3815–3832.
- Rasmussen, C. E. (2003). Gaussian processes in machine learning. In *Summer school on machine learning* (pp. 63–71).
- Rowland, L., Hill, T. C., Stahl, C., Siebicke, L., Burban, B., Zaragoza-Castells, J., ... Williams, M. (2014). Evidence for strong seasonality in the carbon storage and carbon use efficiency of an amazonian forest. *Global Change Biology*, 20(3), 979–991.
- Sakov, P., Oliver, D. S., & Bertino, L. (2012). An iterative enfk for strongly nonlinear systems. *Monthly Weather Review*, 140(6), 1988–2004.
- Schillings, C., & Stuart, A. M. (2017). Analysis of the ensemble kalman filter for inverse problems. *SIAM Journal on Numerical Analysis*, 55(3), 1264–1290.
- Schmidt, G. A., Bader, D., Donner, L. J., Elsaesser, G. S., Golaz, J.-C., Hannay, C., ... Saha, S. (2017). Practice and philosophy of climate model tuning across six us modeling centers. *Geoscientific Model Development*, 10(9), 3207–3223.
- Schneider, T. (2001). Analysis of incomplete climate data: Estimation of mean values and covariance matrices and imputation of missing values. *J. Climate*, 14, 853–871.

- 703 Schneider, T., Jeevanjee, N., & Socolow, R. (2021). Accelerating progress in climate science.
704 *Physics Today*, 74, 44–51.
- 705 Schneider, T., Lan, S., Stuart, A., & Teixeira, J. (2017). Earth system modeling 2.0: A
706 blueprint for models that learn from observations and targeted high-resolution simula-
707 tions. *Geophysical Research Letters*, 44(24), 12–396.
- 708 Schneider, T., & O’Gorman, P. A. (2008). Moist convection and the thermal stratification
709 of the extratropical troposphere. *Journal of the atmospheric sciences*, 65(11), 3571–
710 3583.
- 711 Schneider, T., Teixeira, J., Bretherton, C. S., Brient, F., Pressel, K. G., Schär, C., &
712 Siebesma, A. P. (2017). Climate goals and computing the future of clouds. *Nature*
713 *Climate Change*, 7(1), 3–5.
- 714 Smagorinsky, J. (1963). General circulation experiments with the primitive equations: I. the
715 basic experiment. *Monthly weather review*, 91(3), 99–164.
- 716 Spalart, P., & Allmaras, S. (1992). A one-equation turbulence model for aerodynamic flows.
717 In *30th aerospace sciences meeting and exhibit* (p. 439).
- 718 Stuart, A. M. (2010). Inverse problems: a bayesian perspective. *Acta numerica*, 19, 451–
719 559.
- 720 Stull, R. B. (2012). *An introduction to boundary layer meteorology* (Vol. 13). Springer Sci-
721 ence & Business Media.
- 722 Suzuki, K., Golaz, J.-C., & Stephens, G. L. (2013). Evaluating cloud tuning in a climate
723 model with satellite observations. *Geophysical Research Letters*, 40(16), 4464–4468.
- 724 Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal*
725 *Statistical Society: Series B (Methodological)*, 58(1), 267–288.
- 726 Webb, M. J., Lambert, F. H., & Gregory, J. M. (2013). Origins of differences in climate sen-
727 sitivity, forcing and feedback in climate models. *Climate Dynamics*, 40(3), 677–707.
- 728 Wei, H.-H., & Bordoni, S. (2018). Energetic constraints on the ITCZ position in idealized
729 simulations with a seasonal cycle. *Journal of Advances in Modeling Earth Systems*,
730 10(7), 1708–1725.
- 731 Wyngaard, J. C. (2010). *Turbulence in the atmosphere*. Cambridge University Press.
- 732 Zhang, F., Sun, Y. Q., Magnusson, L., Buizza, R., Lin, S.-J., Chen, J.-H., & Emanuel, K.
733 (2019). What is the predictability limit of midlatitude weather? *Journal of the At-*
734 *mospheric Sciences*, 76(4), 1077 - 1091. Retrieved from [https://journals](https://journals.ametsoc.org/view/journals/atsc/76/4/jas-d-18-0269.1.xml)
735 [.ametsoc.org/view/journals/atsc/76/4/jas-d-18-0269.1.xml](https://journals.ametsoc.org/view/journals/atsc/76/4/jas-d-18-0269.1.xml)
736 doi: 10.1175/JAS-D-18-0269.1
- 737 Zhou, Y., McLaughlin, D., & Entekhabi, D. (2006). Assessing the performance of the
738 ensemble kalman filter for land surface data assimilation. *Monthly Weather Review*,
739 134(8), 2128–2142.