

An efficient Bayesian approach to learning droplet collision kernels: Proof of concept using “Cloudy”, a new n-moment bulk microphysics scheme

Melanie Bieli¹, Oliver Dunbar², Emily K De Jong¹, Anna Jaruga¹, Tapio Schneider², and Tobias Bischoff¹

¹Caltech

²California Institute of Technology

November 26, 2022

Abstract

The small-scale microphysical processes governing the formation of precipitation particles cannot be resolved explicitly by cloud resolving and climate models. Instead, they are represented by microphysics schemes that are based on a combination of theoretical knowledge, statistical assumptions, and fitting to data (“tuning”). Historically, tuning was done in an ad-hoc fashion, leading to parameter choices that are not explainable or repeatable. Recent work has treated it as an inverse problem that can be solved by Bayesian inference. The posterior distribution of the parameters given the data—the solution of Bayesian inference—is found through computationally expensive sampling methods, which require over $O(10^5)$ evaluations of the forward model; this is prohibitive for many models. We present a proof-of-concept of Bayesian learning applied to a new bulk microphysics scheme named “Cloudy”, using the recently developed Calibrate-Emulate-Sample (CES) algorithm. Cloudy models collision-coalescence and collisional breakup of cloud droplets with an adjustable number of prognostic moments and with easily modifiable assumptions for the cloud droplet mass distribution and the collision kernel. The CES algorithm uses machine learning tools to accelerate Bayesian inference by reducing the number of forward evaluations needed to $O(10^2)$. It also exhibits a smoothing effect when forward evaluations are polluted by noise. In a suite of perfect-model experiments, we show that CES enables computationally efficient Bayesian inference of parameters in Cloudy from noisy observations of moments of the droplet mass distribution. In an additional imperfect-model experiment, a collision kernel parameter is successfully learned from output generated by a Lagrangian particle-based microphysics model.

1 **An efficient Bayesian approach to learning droplet**
2 **collision kernels: Proof of concept using “Cloudy”, a**
3 **new n -moment bulk microphysics scheme**

4 **Melanie Bieli¹, Oliver R. A. Dunbar¹, Emily K. de Jong², Anna Jaruga¹,**
5 **Tapio Schneider¹, Tobias Bischoff¹**

6 ¹Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena, CA, USA

7 ²Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA, USA

8 **Key Points:**

- 9 • Historically, microphysics schemes were tuned to data in an ad-hoc way, result-
10 ing in parameter values that are not repeatable or explainable
- 11 • Bayesian inference puts uncertainty quantification and parameter learning on solid
12 mathematical grounds, but is computationally expensive
- 13 • We present a proof-of-concept of computationally efficient Bayesian learning ap-
14 plied to a new bulk microphysics scheme called “Cloudy”

Corresponding author: Melanie Bieli, melanie@caltech.edu

15 Abstract

16 The small-scale microphysical processes governing the formation of precipitation
 17 particles cannot be resolved explicitly by cloud resolving and climate models. Instead,
 18 they are represented by microphysics schemes that are based on a combination of the-
 19 oretical knowledge, statistical assumptions, and fitting to data (“tuning”). Historically,
 20 tuning was done in an ad-hoc fashion, leading to parameter choices that are not explain-
 21 able or repeatable. Recent work has treated it as an inverse problem that can be solved
 22 by Bayesian inference. The posterior distribution of the parameters given the data—the
 23 solution of Bayesian inference—is found through computationally expensive sampling
 24 methods, which require over $\mathcal{O}(10^5)$ evaluations of the forward model; this is prohibitive
 25 for many models.

26 We present a proof-of-concept of Bayesian learning applied to a new bulk micro-
 27 physics scheme named “Cloudy,” using the recently developed Calibrate-Emulate-Sample
 28 (CES) algorithm. Cloudy models collision-coalescence and collisional breakup of cloud
 29 droplets with an adjustable number of prognostic moments and with easily modifiable
 30 assumptions for the cloud droplet mass distribution and the collision kernel. The CES
 31 algorithm uses machine learning tools to accelerate Bayesian inference by reducing the
 32 number of forward evaluations needed to $\mathcal{O}(10^2)$. It also exhibits a smoothing effect when
 33 forward evaluations are polluted by noise. In a suite of perfect-model experiments, we
 34 show that CES enables computationally efficient Bayesian inference of parameters in Cloudy
 35 from noisy observations of moments of the droplet mass distribution. In an additional
 36 imperfect-model experiment, a collision kernel parameter is successfully learned from out-
 37 put generated by a Lagrangian particle-based microphysics model.

38 Plain Language Summary

39 Clouds contain gazillions of cloud droplets, which grow by colliding and sticking
 40 together with each other, and eventually they become big enough to fall out as rain. Keep-
 41 ing track of every one of these droplet in weather and climate models is impossible, so
 42 the formation of rain has to be represented by simplified models, so-called “microphysics
 43 schemes”. These schemes have become a bit like black boxes, with baked-in statistical
 44 assumptions and some empirical parameters whose values are somewhat obscure and not
 45 explainable. We show that we can use a method called Bayesian inference to determine
 46 the values of these parameters in a way that is both mathematically sound and reason-
 47 ably fast. The idea of Bayesian inference is to come up with a first guess about the pos-
 48 sible values of the parameters, and then to systematically refine that guess using observed
 49 data. We apply this method to a new microphysics scheme that we developed and named
 50 “Cloudy”. To be honest, the data we use for our experiments are not real observations
 51 from, say, satellites, but are generated by Cloudy itself. With real observations, the prob-
 52 lem becomes hairier, so what we do here is only a proof-of-concept — but hey, it’s a start!

53 1 Introduction

54 Cloud microphysics comprises all processes controlling the formation and growth
 55 of cloud droplets and ice crystals and their fallout as precipitation. These processes play
 56 a key role in the climate system, affecting surface precipitation, latent heating and cool-
 57 ing, cloud radiative properties, and cloud chemistry. Due to the small scales on which
 58 they occur (sub-microns to centimeters), explicitly simulating the growth of individual
 59 cloud particles in a turbulent cloud requires model resolutions at least as small as the
 60 Kolmogorov scale, which is about 1 mm in the Earth’s atmosphere. With horizontal grid
 61 spacings of about 10-50 km, state-of-the-art climate models are (and will remain) orders
 62 of magnitude too coarse to resolve the vast number of hydrometeors—typically about
 63 10^8 in 1 m^3 of cloudy air—on a global scale.

64 Instead, global climate models (GCMs) represent microphysical processes by means
 65 of statistical parameterizations, which are developed based on a combination of phys-
 66 ical understanding, statistical assumptions, heuristics, and tuning to observations. GCMs
 67 typically employ so-called bulk schemes, which assume some functional form of the cloud
 68 droplet size distribution (DSD) and step one or more statistical moments of the distri-
 69 bution forward in time. More computationally expensive representations of microphysics
 70 include bin schemes, which partition the DSD into discrete size bins and model the pro-
 71 cesses affecting the particles in each bin. More recent developments include Lagrangian
 72 particle-based schemes (e.g., Andrejczuk et al., 2010; Riechelmann et al., 2012; Shima
 73 et al., 2009, 2019), which simulate an ensemble of computational particles (“super-particles”
 74 or “super-droplets”), each representing a large number of real cloud and precipitation
 75 particles. The closest approximations to “ab initio” calculations of microphysical pro-
 76 cesses are performed by direct numerical simulations (DNS), which track the motion and
 77 growth of each individual cloud particle. To do that, they need to resolve the smallest
 78 scales of turbulence (millimeter and submillimeter scales), which limits typical domain
 79 sizes to less than 1 m^3 .

80 Morrison, van Lier-Walqui, Fridlind, et al. (2020) provide an overview of the dif-
 81 ferent approaches to the numerical modeling of microphysics and argue that the Lagrangian
 82 particle-based methods overcome several shortcomings of traditional bulk and bin schemes.
 83 A particularly attractive property of Lagrangian schemes is that in the limiting case where
 84 each computational particle represents a single real particle and the model resolution ap-
 85 proaches that of a DNS, the Lagrangian scheme converges to the particle-by-particle DNS
 86 (Dziekan & Pawlowska, 2017). Due to their computational cost, though, Lagrangian particle-
 87 based methods will likely not replace the bulk schemes in global models within the next
 88 1–2 decades. However, Lagrangian particle-based schemes can be used for cloud mod-
 89 eling and, as shown in this study, they can provide a benchmark for testing bulk schemes.

90 All three types of microphysics schemes (bulk, bin, and Lagrangian) rely on em-
 91 pirical parameters to compute process rates. Ultimately, this is a consequence of the fact
 92 that there is no known complete set of equations governing these microscopic processes,
 93 i.e., there is no microphysics analog to the Navier-Stokes equations. Given the limited
 94 theoretical knowledge, data and observations play a crucial role in constraining the val-
 95 ues of empirical parameters. Historically, the process of determining (“tuning”) these
 96 values has not been approached in a systematic and transparent way. However, a num-
 97 ber of recent studies demonstrate the use of Bayesian techniques to parameter estima-
 98 tion in bulk microphysics schemes (e.g., Posselt & Vukicevic, 2010; Posselt, 2016; van
 99 Lier-Walqui et al., 2014). In a two-part paper, Morrison, van Lier-Walqui, Kumjian, and
 100 Prat (2020) introduce the Bayesian Observationally-constrained Statistical-physical Scheme
 101 (BOSS), a framework for the bulk parameterization of microphysics, which is designed
 102 to learn microphysical parameter distributions from data by means of Bayesian infer-
 103 ence. The second part (van Lier-Walqui et al., 2020) gives a demonstration in the form
 104 of a perfect-model experiment, which shows that BOSS can be used in conjunction with
 105 a Markov chain Monte Carlo (MCMC) sampling algorithm to estimate parameters from
 106 synthetically generated rain observations. A key feature of BOSS is its adjustable com-
 107 plexity: While traditional schemes have fixed numbers of prognostic moments, BOSS al-
 108 lows the number of prognostic moments to be chosen flexibly, depending on the appli-
 109 cation and the observations available.

110 In a similar vein, to model parametric uncertainty with a strong mathematical foun-
 111 dation, we use a Bayesian framework, where model parameters are described by random
 112 variables. We propose a prior form of the distributions and refine them systematically
 113 with observed data, using a process known as Bayesian inference, Bayesian calibration,
 114 or uncertainty quantification. As in Schneider et al. (2017), we use the word “data” for
 115 any information source that is used as a ground truth against which a microphysics scheme
 116 is compared and calibrated, including both observations of natural clouds (e.g., satel-

117 lite products or in-situ airborne observations) and output of higher-resolution or more
 118 physical model simulations. There is also much to be gained from using a perfect-model
 119 setting in a proof-of-concept, where the same model is used for both generating data and
 120 for inversion; however, the resulting parameter learning results will necessarily be op-
 121 timistic. The standard tool of Bayesian inference is MCMC sampling, which represents
 122 this data-refined distribution empirically by providing a large set of samples drawn from
 123 it. The main drawback of MCMC methods is their computational cost: They require large
 124 numbers of model evaluations (typically about 10^5 to 10^6), which is not feasible for ex-
 125 pensive models such as GCMs.

126 Recent work by Cleary et al. (2021) presents a method to perform an approximate
 127 Bayesian inversion of computationally expensive models for which derivatives are not read-
 128 ily available and whose evaluations may be polluted by noise, for example, from chaotic
 129 internal variability. This three-step method called Calibrate-Emulate-Sample (CES) has
 130 been shown to be effective for inferring parameter distributions in a convection scheme
 131 of a GCM exhibiting these properties (Dunbar et al., 2021; Howland et al., 2021). The
 132 calibrate step of the algorithm consists of ensemble Kalman inversion or variants such
 133 as the ensemble Kalman sampler (EKS; Garbuno-Inigo, Hoffmann, et al., 2020), which
 134 are used to find pairs of parameters and their respective model outputs, automatically
 135 focusing on a region of the parameter space that is likely to have produced the observed
 136 data. In the emulate step, a Gaussian process (Rasmussen & Williams, 2006) is trained
 137 on these parameter-output pairs and serves as a surrogate (emulator) of the original (ex-
 138 pensive) forward model. In the sampling step, the fast-to-evaluate emulator is used in
 139 the likelihood of an MCMC algorithm to sample the posterior distribution in a compu-
 140 tationally efficient manner.

141 The goal of this study is twofold: First, we introduce a new bulk microphysics frame-
 142 work that was designed for consistent representation of microphysical processes across
 143 models with different resolutions and physics. The model, called “Cloudy” (available at
 144 <https://github.com/CliMA/Cloudy.jl>), is broadly similar to BOSS (Morrison, van
 145 Lier-Walqui, Kumjian, & Prat, 2020; van Lier-Walqui et al., 2020), e.g., in that the num-
 146 ber of prognostic moments is modifiable and that it can learn from data; however, there
 147 are also a few important differences, e.g., in that Cloudy allows for separate learning of
 148 collision kernels and DSDs, hence facilitating the finding of physically realizable solu-
 149 tions. Cloudy currently simulates collision-coalescence and collisional breakup of cloud
 150 droplets (with future development plans including an extension to other warm-rain pro-
 151 cesses and ice microphysics), in a way that the governing equations for the moments of
 152 the DSD can easily be related to the specific properties of collision kernels.

153 The second goal is to demonstrate that parameters in Cloudy can be learned from
 154 data in a computationally efficient way, through the approximate Bayesian inversion per-
 155 formed by CES. We present a suite of perfect-model experiments (where Cloudy itself
 156 is used to generate the data used for Bayesian inversion), as well as an experiment us-
 157 ing data from simulations generated by PySDM (Bartman et al., 2021), a high-performance
 158 Python implementation of the super-droplet method (SDM) for representing liquid mi-
 159 crophysics (package available at <https://github.com/atmos-cloud-sim-uj/PySDM>),
 160 with an additional process added to represent droplet breakup. In the spirit of a proof-
 161 of-concept, both Cloudy and PySDM are run in a computationally cheap zero-dimensional
 162 “box” framework.

163 This paper is organized as follows: Section 2 describes the underlying concepts and
 164 equations of Cloudy. In Section 3, we give a brief introduction to the Bayesian approach
 165 to solving inverse problems, together with an overview of the CES method. Section 4
 166 explains the model experiments, including a summary of the PySDM model. The results
 167 of the experiments are shown in Section 5. The paper concludes with a summary of the
 168 findings in Section 6.

2 Cloudy Model Description

Cloudy is a flexible bulk microphysics model that simulates collision-coalescence and collisional breakup of cloud droplets. By “flexible,” we mean the following:

- The number of prognostic moments can be adjusted to the requirements of the cloud droplet mass distribution function and availability of data for calibration.
- A modular design facilitates experimenting with different collision kernels and cloud droplet mass distribution functions.
- It is set up for Bayesian inference, i.e., parameters of collision kernels do not have to be fixed but instead can be learned from data.

The three main inputs required to run the model are: an initial droplet mass distribution function, a collision kernel specifying the rate of collisions between particles, and a coalescence efficiency defining the fraction of collisions that result in coalescence of the particles into one larger drop, as opposed to collisions that result in breakup of the particles into smaller fragments. Cloudy then simulates how the distribution (characterized by a set of n prognostic moments) evolves over time as a result of the droplet interactions defined by the given collision kernel and coalescence efficiency. The number of prognostic moments is determined ab initio by the type of the cloud droplet mass distribution to be simulated (e.g., a Gamma distribution or a mixture of Gamma distributions). Assuming a fixed distribution type is the central closure assumption made in Cloudy; the number of prognostic moments has to be chosen such that the distribution parameters can be computed from the prognostic moments. Note that not all distributions can be inferred uniquely from their moments; for example, a lognormal distribution is not uniquely defined by its moments because its moment generating function does not converge.

The mathematical core of the model consists of two equations: the stochastic collection equation (SCE; Smoluchowski, 1916) and the stochastic breakup equation (SBE; e.g., Pruppacher & Klett, 1978), both expressed in terms of the DSD moments. The SCE describes the time rate of change of $f = f(m, t)$, the mass distribution function of liquid water droplets, due to the process of collision and coalescence. The distribution f depends on droplet mass m and time t ; generally, it will also depend on position in space, but we neglect this dependence in our zero-dimensional setting here. The mass distribution function is defined such that $f(m, t) dm$ denotes the number of droplets with masses in the interval $[m, m + dm]$ per unit volume at time t . We will mostly refer to $f(m, t)$ by the term “particle mass distribution” (PMD) rather than by “droplet size distribution,” even though the two expressions could be used interchangeably for spherical water droplets (the only type of droplet considered in this study), where there is a one-to-one map between droplet size and droplet mass. By deviating from the traditional terminology, we account for the possibility that a future version of Cloudy may include non-spherical and non-liquid particles.

Following Beheng (2010), the SCE can be written as

$$\left. \frac{\partial f(m, t)}{\partial t} \right|_{\text{coal}} = \frac{1}{2} \int_{m'=0}^{\infty} f(m', t) f(m-m', t) C(m', m-m') dm' - f(m, t) \int_{m'=0}^{\infty} f(m', t) C(m, m') dm'. \quad (1)$$

The collection kernel $C(m', m'') = K(m', m'') E_c(m', m'')$ (units: cubic centimeters per particle per second) describes the rate at which two droplets of masses m' and m'' come into contact and coalesce. It is the product of the collision kernel $K(m', m'')$ and the dimensionless coalescence efficiency $E_c(m', m'')$, which denotes the fraction of droplets that coalesce into a drop of mass $m' + m''$ upon colliding with each other. Throughout this paper, we will assume E_c to have a constant value, even though in reality it depends on the kinetic energy of the two colliding droplets (e.g., Low & List, 1982; Beard & Ochs, 1995). The first term on the right-hand side of Eq. (1) describes the rate of increase of

216 the number of drops having a mass m due to collision and coalescence of drops of masses
 217 m' and $m-m'$ (where the factor $\frac{1}{2}$ avoids double counting); the second term describes
 218 the rate of reduction of drops of mass m due to collision and coalescence of drops hav-
 219 ing a mass m with other drops.

The SBE describes the time evolution of the PMD due to collision-induced breakup and is given by

$$\begin{aligned} \left. \frac{\partial f(m, t)}{\partial t} \right|_{\text{breakup}} &= \frac{1}{2} \int_{m'=0}^{\infty} f(m', t) dm' \int_{m''=0}^{\infty} f(m'', t) B(m', m'') P(m; m', m'') dm'' \\ &\quad - f(m, t) \int_{m''=0}^{\infty} \frac{f(m'', t) B(m', m'')}{m + m''} dm'' \int_{m'=0}^{m+m''} m' P(m'; m, m'') dm'. \end{aligned} \quad (2)$$

The breakup kernel $B(m', m'') = K(m', m'')(1 - E_c(m', m''))$ (units: cubic centimeters per particle per second) defines the rate at which two droplets of masses m' and m'' come into contact and break apart. The function $P(m, m', m'')$ is the mass distribution function of the fragments m produced by collisional breakup of two droplets of masses m' and m'' , with $P(m, m', m'') dm$ giving the number of drops in the mass interval between m and $m+dm$ resulting from the breakup. We use the exponential fragment distribution introduced by Feingold et al. (1988),

$$P(m; m', m'') = \nu^2 (m' + m'') \exp(-\nu m), \quad (3)$$

220 where $\nu = (qM_0/M_1)$; M_0 (cm^{-3}) is the initial value of the zeroth PMD moment (i.e.,
 221 the initial number of droplets), M_1 (g cm^{-3}) is the first PMD moment (i.e., the water
 222 content), and q is a positive integer characterizing the fragment concentration.

Conservation of mass dictates that the mass of the sum of all fragments must equal the mass of the two colliding drops, i.e.,

$$\int_0^{m'+m''} m P(m; m', m'') dm = m' + m''. \quad (4)$$

223 If the fragment distribution conserves mass exactly, the last integral in Eq. (2) evalu-
 224 ates to $m+m''$ and, after division by the denominator of the previous term, results in
 225 a multiplication by 1. However, the fragment distribution by Feingold et al. (1988) used
 226 in this study does not fulfill conservation of mass exactly, so the last integral in Eq. (2)
 227 cannot be omitted.

We will rewrite Eqs. (1) and (2) in terms of the moments M_k of $f(m, t)$, which are the prognostic microphysical variables in Cloudy. They are defined by

$$M_k = \int_0^{\infty} m^k f(m, t) dm. \quad (5)$$

The time rate of change of the k -th moment of f is obtained by multiplying Eqs. (1) and (2) by m^k and integrating over the droplet mass, which yields

$$\begin{aligned} \left. \frac{\partial M_k}{\partial t} \right|_{\text{coal}} &= \frac{1}{2} \int_0^{\infty} \int_0^{\infty} \left((m + m')^k - m^k - m'^k \right) C(m, m') f(m, t) f(m', t) dm dm' \\ &=: I_{k, \text{coal}}, \\ \left. \frac{\partial M_k}{\partial t} \right|_{\text{breakup}} &= \frac{1}{2} \int_0^{m'+m''} m^k P(m; m', m'') dm \int_0^{\infty} f(m', t) dm' \int_0^{\infty} f(m'', t) B(m', m'') dm'' \\ &\quad - \int_0^{\infty} m^k f(m, t) dm \int_0^{\infty} \frac{f(m'', t) B(m, m'')}{m + m''} dm'' \int_0^{m+m''} m' P(m'; m, m'') dm' \\ &=: I_{k, \text{breakup}}. \end{aligned} \quad (6)$$

Thus, the time evolution of the k th moment M_k due to collision-coalescence and collisional breakup is given by

$$\frac{\partial M_k}{\partial t} = I_{k, \text{ coal}} + I_{k, \text{ breakup}} \quad (7)$$

228 To step the prognostic moments forward in time, Cloudy computes the right-hand
 229 side of Eq. (7) using Monte Carlo integration, a technique for numerical integration using
 230 random numbers. Compared with other numerical integration methods, Monte Carlo
 231 integration has the following advantages: It converges in any dimension, regardless of
 232 the smoothness of the integrand, albeit only at a rate of $\mathcal{O}(\sqrt{N})$ (where N is the number
 233 of random samples used to compute the integral). It is also conceptually simple and
 234 parallelizable. A further desirable feature is the fact that it adds a stochastic element
 235 to the simulations. The resulting internal model variability provides a straightforward
 236 way of estimating the covariance of the observational noise needed as input to the CES
 237 algorithm in a perfect-model setting (see section 3.2). With a deterministic integration
 238 method such as a numerical quadrature, there would not be any randomness in the simulated
 239 data, and artificial noise would have to be added to mimic observational noise.

240 The drawback of Monte Carlo integration is its computational cost: Many samples—
 241 thousands or even millions—may be required to obtain results of acceptable accuracy.
 242 Numerous techniques have been developed to reduce the variance of the Monte Carlo
 243 estimator, and hence the number of samples needed (e.g., Kleijnen & Rubinstein, 2013).
 244 For our specific application, 200 samples turned out to produce sufficiently accurate results
 245 (see Appendix A for implementation details of the Monte Carlo integration). Monte
 246 Carlo integration can easily be extended to higher dimensions and is therefore typically
 247 applied to compute high-dimensional integrals, but we use it here to demonstrate that
 248 CES works even with models that produce noisy output, as well as for its general-purpose
 249 functionality and robustness with respect to the integrand. For the simple model setup
 250 presented here, the computational cost of the Monte Carlo integration is easily affordable.
 251 The algorithm can even be applied in larger-scale settings as its embarrassingly parallel
 252 nature can be exploited very efficiently by GPUs (e.g., Kanzaki, 2011; Borowka et
 253 al., 2019).

254 Evaluating the integrals on the right-hand side of Eq. (7) requires knowledge of the
 255 PMD function $f(m, t)$. The initial distribution $f(m, t = 0)$ is specified by the user. There
 256 is a priori no reason to assume that the PMD would retain its functional form as particles
 257 are colliding, are forming new drops, and are breaking apart. However, to uniquely
 258 identify the distribution f at each time step, one would have to keep track of infinitely
 259 many moments of the PMD, which is obviously not practicable. The truncation of this
 260 infinite system is the moment closure problem, which all moment-based microphysics schemes
 261 have to address in some form. In Cloudy, as in most bulk microphysics schemes, the closure
 262 is achieved by assuming an analytic functional form for the PMD, and allowing for the
 263 parameters of the PMD to change over time while the type of distribution itself is
 264 kept fixed. Thus, a Cloudy simulation consists of the following steps, which are summarized
 265 in Fig. 1:

- 266 1. The user specifies the collision kernel $K(x, y)$, initial PMD $f(m, t = 0; \xi_0)$ with
 267 distribution parameters ξ_0 , the coalescence efficiency E_c (assumed to be constant,
 268 as mentioned above), and the end time t_{end} of the simulation. In addition, a map
 269 $h_{\xi \rightarrow M}$ and a map $h_{M \rightarrow \xi}^{-1}$ have to be supplied. The former defines how to compute
 270 the prognostic PMD moments $M = \{M_k\}_{k=1, \dots, n}$ from the parameters ξ of the
 271 PMD, and the latter defines the inverse map from the prognostic moments to the
 272 parameters. Note that except for simple distributions, $h_{M \rightarrow \xi}^{-1}$ does not have a closed-
 273 form representation, and the PMD parameters have to be determined by solving
 274 an optimization problem.

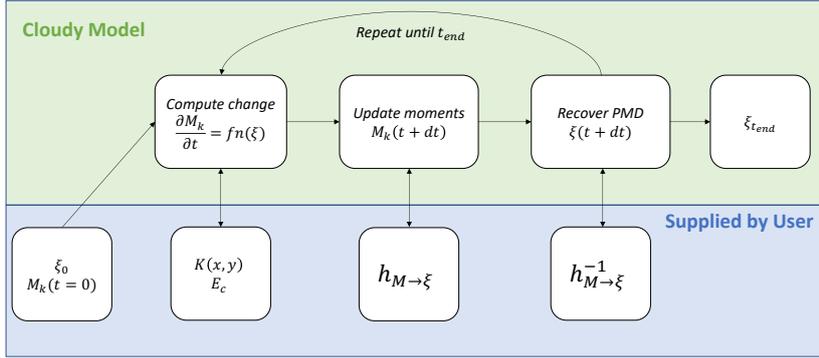


Figure 1. Summary of the computational steps performed by Cloudy and the required user-defined input. See text for notation.

- 275 2. The contributions of collision-coalescence and collisional breakup to the time evolution of each prognostic moment (right-hand side of Eq. (7); denoted $fn(\xi)$ in Fig. 1) is computed.
- 276
- 277
- 278 3. The prognostic moment are stepped forward in time.
- 279
- 280 4. The new parameters of the PMD are computed from the updated moments using $h_{M \rightarrow \xi}^{-1}$.
- 281 5. Steps 2–4 are repeated until t_{end} is reached.

282 One of the guiding principles in developing Cloudy was to make the scheme amenable to learning from data. Its modular design makes it easy to experiment with different kernels and PMDs, and the number of prognostic moments is determined by the user-provided parameter-to-moment map $h_{\xi \rightarrow M}$. We focus here on learning parameters of the collision kernel $K(x, y)$, though alternatively (and with slight modifications of the setup), Cloudy can be used to learn parameters of the PMD instead or in addition.

283

284

285

286

287

288 The approach to the closure problem is a notable difference between Cloudy and BOSS. In contrast to Cloudy, BOSS does not assume a functional form for the DSD. Instead, the diagnostic moments are expanded as multivariate generalized power series of the set of prognostic moments, and the DSD normalization method of Morrison et al. (2019) is used to relate the moments to one another statistically. While the assumption of a DSD form results in a loss of generality, it also allows for a clear separation of the parameters associated with the physics of collision-coalescence and breakup (as defined by the parameters of the collision kernel) from those associated with the droplet population (as defined by the parameters of the DSD). This improves the interpretability of the scheme and helps ensure physical realizability of the simulations.

289

290

291

292

293

294

295

296

297

298 3 Methods

299 3.1 Bayesian parameter estimation

The estimation of model parameters such as the coefficients of a collision-coalescence kernel can be formulated as a Bayesian inverse problem, whose solution—the posterior

distribution of the unknown parameters given the observed data—is given by Bayes’ rule:

$$f_{\theta|y}(\theta|y) = \frac{f_{y|\theta}(y|\theta)f_{\theta}(\theta)}{f_y(y)} \propto f_{y|\theta}(y|\theta)f_{\theta}(\theta). \quad (8)$$

300 Here, $f_{\theta|y}(\theta|y)$ is the posterior probability density function (PDF) of the parameters θ
 301 given the data y , $f_{y|\theta}(y|\theta)$ is the likelihood function of the data given the parameters,
 302 $f_{\theta}(\theta)$ is the prior PDF of the parameters, and the normalization factor $f_y(y)$ is the marginal
 303 PDF of the data. In a Bayesian framework, the unknown parameters are thus modeled
 304 as random variables, and the posterior distribution can be written in terms of contribu-
 305 tions from both prior information about the parameters and the likelihood of the observed
 306 data. Note that sampling the posterior distribution using MCMC methods does not re-
 307 quire knowledge of the normalization factor $f_y(y)$.

We assume that the data y are linked to the parameter vector θ according to the additive relationship

$$\begin{aligned} y &= \mathcal{G}(\theta) + \eta_s \\ &= \bar{\mathcal{G}}(\theta) + \eta_y + \eta_s \\ &= \bar{\mathcal{G}}(\theta) + \eta. \end{aligned} \quad (9)$$

308 Here, the forward map $\mathcal{G} : \mathbb{R}^p \rightarrow \mathbb{R}^d$ maps a parameter vector $\theta \in \mathbb{R}^p$ to a d -dimensional
 309 output space; the error term η_s denotes structural error, which arises from a model’s in-
 310 ability to accurately represent its target due to deficiencies in its mathematical struc-
 311 ture. Because of the randomness introduced by the Monte Carlo integration (see section 2),
 312 the output of the forward map \mathcal{G} is polluted by noise. We can think of $\mathcal{G}(\theta)$ as noisy ob-
 313 servations of an underlying “true”, deterministic forward map $\bar{\mathcal{G}}: \mathcal{G}(\theta) = \bar{\mathcal{G}}(\theta) + \eta_y$,
 314 where η_y is observational noise. The total error $\eta \in \mathbb{R}^d$ is thus the sum of two terms,
 315 observational noise η_y and structural error η_s , which are assumed to be independent. The
 316 choices we make for η_y and η_s will encode our assumptions about structure and origin
 317 of the error in a given scenario.

The parameter-to-data map \mathcal{G} consists of two components. The main component is the map $\Psi : \Omega \rightarrow \mathbb{R}^d$, involving a forward simulation of Cloudy using parameters sampled from a physical parameter domain Ω . The map Ψ can be thought of as a dynamical model, whose output depends on p model parameters that we wish to learn (here, the model parameters to be learned are parameters of collision kernels). So that computational methods interface only with unbounded distributions, we choose to work always with unbounded parameter distributions θ in tandem with an invertible transformation function $\mathcal{T}: \Omega \rightarrow \mathbb{R}^p$. The combined map \mathcal{G} from parameters to data thus takes the form

$$\mathcal{G} = \Psi \circ \mathcal{T}^{-1}. \quad (10)$$

Sampling the posterior distribution (Eq. 8) requires the repeated evaluation of the data likelihood, whose distribution is given by

$$f_{y|\theta}(y|\theta) = f_{\eta}(y - \bar{\mathcal{G}}(\theta)), \quad (11)$$

318 where f_{η} is the PDF of the noise η . Initially, only the noisy map \mathcal{G} is available, such that
 319 the likelihood given by Eq. (11) cannot be evaluated. The underlying true model $\bar{\mathcal{G}}(\theta)$
 320 only becomes accessible through the training of a Gaussian process emulator in the mid-
 321 dle stage of CES (see section 3.2 and Dunbar et al., 2021)—this emulator learns a smooth
 322 function which is a surrogate for $\bar{\mathcal{G}}(\theta)$.

323 We work with two different model-data scenarios. In a first scenario, we consider
 324 a perfect-model experiment where we assume that there is no structural uncertainty, $\eta_s =$
 325 0. We also choose η_y to be realizations of random noise due to measurement error. When
 326 using data y that are statistical aggregates such as moments of the PMD or other av-
 327 erages, and assuming a perfect (unbiased) model, it is reasonable to use a central-limit

Table 1. Overview of collision kernels, kernel parameters, their prior distributions and constraints. For any element X of a parameter θ living in the unconstrained space, the mapping to the corresponding parameter value x in the constrained space with a uniform prior is given approximately by the transformation $\mathcal{T}^{-1}(X) = (x_{\text{upper}} \exp(X) + x_{\text{lower}})/(\exp(X) + 1)$.

Kernel	Name	Unconstrained Parameter	Prior	Constrained Parameter	Constraint
$K(x, y) = b(x + y)$	Sum-of-masses	$\theta = [B]$	$\theta \sim \mathcal{N}(0, 1)$	$\phi = [b]$	$[b_{\text{lower}}, b_{\text{upper}}] = [10^2, 10^4] \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$
$K(x, y) = a + b(x + y)$	Sum-of-masses plus constant	$\theta = [A, B]$	$\theta \sim \mathcal{N}(0, 1)$	$\phi = [a, b]$	$[a_{\text{lower}}, a_{\text{upper}}] = [10^{-7}, 10^{-5}] \text{ cm}^3 \text{ s}^{-1}$ $[b_{\text{lower}}, b_{\text{upper}}] = [10^2, 10^4] \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$
$K(x, y) = \begin{cases} b(x + y) & x < c \text{ or } y < c \\ a & \text{otherwise} \end{cases}$	Piecewise	$\theta = [A, B, C]$	$\theta \sim \mathcal{N}(0, 1)$	$\phi = [a, b, c]$	$[a_{\text{lower}}, a_{\text{upper}}] = [10^{-7}, 10^{-5}] \text{ cm}^3 \text{ s}^{-1}$ $[b_{\text{lower}}, b_{\text{upper}}] = [10^2, 10^4] \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$ $[c_{\text{lower}}, c_{\text{upper}}] = [5.0 \times 10^{-11}, 5.0 \times 10^{-9}] \text{ g}$

328 theorem, so that η_y is a draw from a (multivariate) normal random variable with zero
329 mean and covariance matrix Γ_y .

330 For each experiment in this scenario, we approximate the observational covariance
331 by running the forward model 10 times with true parameter values and estimating Γ_y
332 as the sample covariance matrix of the resulting data, which are 3-dimensional vectors
333 of PMD moments (note that in order to give a non-singular estimate of the covariance
334 matrix, the number of samples needs to be greater than the dimensionality of the data).
335 The data y are taken to be the sample mean of these 10 vectors of PMD moments.

336 In a second scenario, the data y are generated by PySDM (see section 4) instead
337 of Cloudy. Structural uncertainties arise from the differences between the modeling ap-
338 proaches of Cloudy and PySDM (bulk vs. Lagrangian particle-based), with PySDM ar-
339 guably simulating droplet-droplet interactions in a more realistic way (e.g., due to the
340 lack of a closure assumption).

Here, we opt for a simple representation of the model discrepancy η_s as a (multi-
variate) normal random variable with constant mean $m_s \in \mathbb{R}^d$ and covariance Γ_s . Ex-
plicitly, we write this as

$$y = \mathcal{G}(\theta) + m_s + \tilde{\eta}, \quad \text{where } \tilde{\eta} \sim N(0, \Gamma_y + \Gamma_s) \quad (12)$$

341 As both PySDM and Cloudy have the same model parameters, we can estimate
342 the bias m_s by running each model ten times with at the true parameter value and tak-
343 ing the difference of the sample means of the resulting output. Apart from the addition
344 of a bias term, Eq. (12) also differs from the perfect-model version (Eq. 9) in that the
345 covariance of the noise $\tilde{\eta}$ is the sum of the Cloudy and PySDM noise covariance matri-
346 ces.

347 In general, assessing a model's adequacy to reproduce the given data (even when
348 the model is perfectly calibrated) is a difficult task (e.g., Brockwell & Davis, 1996; Kennedy
349 & O'Hagan, 2001; Weisberg, 2014). Our choice can be seen as a special case of the ap-
350 proach followed by Kennedy and O'Hagan (2001), where η_s is modeled with a Gaussian
351 process.

352 3.2 Calibrate-Emulate-Sample

353 The Calibrate-Emulate-Sample (CES) method (Cleary et al., 2021) is designed for
354 Bayesian inversion in settings where the forward model is too computationally expen-
355 sive and/or noisy for direct sampling of the posterior using, e.g., MCMC. We give a con-
356 ceptual overview of CES and refer to Cleary et al. (2021) and Dunbar et al. (2021) for
357 more detailed descriptions. The method accelerates Bayesian learning by substantially
358 reducing the number of forward model evaluations required from the $\mathcal{O}(10^5)$ – $\mathcal{O}(10^6)$ eval-
359 uations typically needed for MCMC to $\mathcal{O}(10^2)$ evaluations. It consists of three stages:

- 360 • The **calibration** stage uses ensemble Kalman inversion (EKI; Iglesias et al., 2013)
361 or variants thereof such as ensemble Kalman sampling (EKS; Garbuno-Inigo, Hoff-
362 mann, et al., 2020; Garbuno-Inigo, Nüsken, & Reich, 2020) to solve the experi-
363 mental design problem of choosing good training points for the subsequent em-
364 ulation stage. EKI and EKS are derivative-free methods that place training points
365 of the parameter-to-data map in the vicinity of where the Bayesian posterior dis-
366 tribution of the parameters is concentrated. They are highly parallelizable, scale
367 well to high-dimensional problems, and are well suited to dealing with noisy for-
368 ward model evaluations (Duncan et al., 2021). In this study, we use EKS, whose
369 ensemble approximates the Bayesian posterior.
- 370 • In the **emulation** stage, the samples from the calibration stage are used to train
371 a Gaussian process regression model, which serves as an emulator that approx-
372 imates the original parameter-to-data map but is much cheaper to evaluate.
- 373 • The **sampling** stage uses MCMC methods to sample the posterior distribution
374 of the parameters, using the cheap emulator instead of the original (expensive) for-
375 ward model.

376 Cloudy is very cheap to run in the prototype setting of this paper; thus, CES is not
377 necessary from the perspective of computational cost here. But we apply the method
378 with an eye toward larger-scale applications in GCMs later, and for the smoothing prop-
379 erty of the Gaussian process emulator, which increases robustness with respect to the
380 noise induced by the Monte Carlo integration (see section 2 and Dunbar et al. (2021)).
381 A proof-of-concept with a cheap model also has the advantage that results can be com-
382 pared against traditional techniques of Bayesian inversion (namely, direct MCMC sam-
383 pling without an emulator), which would not be feasible with a more computationally
384 expensive model.

385 4 Experimental Setup

386 The goal of the experiments is to demonstrate that kernel parameters in Cloudy
387 can be learned from data using the CES method. For this purpose, parameter learning
388 is performed in a perfect-model setting, where the data are generated by running Cloudy
389 with the “true” parameter values that are then to be learned by Bayesian inversion. This
390 setup tests if the true parameters are identifiable in the absence of model uncertainty,
391 i.e., in a scenario where the constrained model, with the correct parameters, is able to
392 reproduce the data to within noise. In all experiments, the data are values of the zeroth,
393 first, and second moment of the PMD at the end time t_{end} of the simulation. The ze-
394 roth moment of the PMD is equal to the total number concentration, while the second
395 and third moments are proportional to the mass mixing ratio and radar reflectivity fac-
396 tor, respectively. Thus, the first three PMD moments are directly related to quantities
397 that can be in principle obtained from remote sensing systems.

398 We also present an experiment where the data are generated by PySDM (Bartman
399 et al., 2021), a Lagrangian super-droplet scheme based on the Monte Carlo algorithm
400 by Shima et al. (2009), which models collisional growth of cloud droplets without using
401 the stochastic collection equation. Instead, it represents the cloud droplet population by
402 a number of computational super-droplets, each corresponding to some multitude of real
403 droplets with identical properties (including size and position). The collision and coa-
404 lesence of these super-droplets is modeled stochastically: Within each time step, only
405 a discrete sample of super-droplet pairs is considered. This is done to reduce the com-
406 putational cost from $\mathcal{O}(N_s^2)$, which would result from considering all pairs, to $\mathcal{O}(N_s)$,
407 where N_s is the number of super-droplets. Each of these candidate pairs then collides
408 with a probability that depends on the multiplicities of the two colliding super-droplets,
409 i.e., on the numbers of real droplets they represent. A comprehensive description of the

410 method is given in Shima et al. (2009), and its implementation in PySDM is detailed in
 411 Bartman et al. (2021).

For the purpose of this paper, a simple breakup implementation was added to PySDM. In this implementation, described in Appendix B, a breakup results in exactly two fragments, each of which carries half of the sum of the masses of the two colliding drops. When learning from data generated by PySDM, the same breakup implementation is used in Cloudy, by defining a fragment distribution

$$P(m, m', m'') = 2\delta\left(m - \frac{m' + m''}{2}\right), \quad (13)$$

412 where δ is the Dirac delta function.

413 We will present results of the following experiments (see also Table 1):

- 414 • Collision kernel of the form $K(x, y) = b(x+y)$: The parameter b is learned in a
 415 perfect-model setting, for $b = 2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$, $4000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$, $6000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$,
 416 and $8000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$. In addition to CES, two of these four experiments ($b =$
 417 $2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$ and $4000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$) are also carried out with a brute-force
 418 MCMC sampling, to compare results and performance of the two methods. As the
 419 name suggests, brute-force MCMC sampling involves repeatedly (10^5 times) evalu-
 420 ating Cloudy itself rather than using the predictions of an emulator.
- 421 • Collision kernel of the form $K(x, y) = b(x+y)$: The parameter $b = 2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$
 422 is learned from data generated by PySDM, using CES.
- 423 • Collision kernel of the form $K(x, y) = a + b(x+y)$: The parameters a and b are
 424 learned in a perfect-model setting using CES, for $a = 4.0 \times 10^{-6} \text{ cm}^3 \text{ s}^{-1}$ and
 425 $b = 3000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$.
- Collision kernel of the form

$$K(x, y) = \begin{cases} b(x+y) & x < c \text{ or } y < c, \\ a & \text{otherwise.} \end{cases} \quad (14)$$

426 The parameters a , b , and c are learned in a perfect-model setting using CES, for
 427 $a = 2.0 \times 10^{-6} \text{ cm}^3 \text{ s}^{-1}$, $b = 3000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$, $c = 1.0 \times 10^{-9} \text{ g}$.

428 The kernels chosen for this suite of experiments represent a sequence of increasingly dif-
 429 ficult learning tasks (from learning one parameter to learning three parameters) that are
 430 used to assess the ability of the CES method to retrieve parameters and provide uncer-
 431 tainty quantification. How realistically these kernels represent droplet-droplet interac-
 432 tions from a physical perspective is of lesser concern for this purpose, but they are nev-
 433 ertheless inspired by established kernels: The sum-of-masses kernel $K(x, y) = b(x+y)$
 434 is known as a Golovin kernel (Golovin, 1963), while the piecewise defined kernel is a sim-
 435 pler variant of a Long (1974) kernel, which is quadratic for small droplets and linear for
 436 large droplets¹.

The PMD is assumed to be a Gamma distribution parameterized by $\xi = [N_0, \alpha, \beta]$, where N_0 is a scaling constant (corresponding to the total number of droplets), α is the shape parameter, and β is the rate parameter:

$$f(m, t) = \frac{N_0 \beta^\alpha}{\Gamma(\alpha)} m^{\alpha-1} \exp(-\beta m). \quad (15)$$

¹To be precise, both the Golovin kernel and the Long kernel are *collection* kernels, i.e., they represent the product of a collision kernel and the coalescence efficiency. This study is concerned with learning *collision* kernels, but since the coalescence efficiency is simply assumed to be constant, the resulting collection kernel is proportional to the underlying collision kernel.

The parameter vector ξ changes over time, as the shape of the distribution evolves. For Gamma mass distribution functions, specifying the first three moments is sufficient to uniquely determine the distribution parameters ξ , hence Cloudy solves Eq. (6) for $k = 0, 1, 2$ (but the number of prognostic moments can be adjusted to the requirements of any given closure distribution). The map $h_{\xi \rightarrow m}$ from the distribution parameters ξ to the PMD moments $M = [M_0, M_1, M_2]$ and its inverse $h_{M \rightarrow \xi}^{-1}$ are given by

$$\begin{aligned} h_{\xi \rightarrow M}(N_0, \alpha, \beta) &= \left[N_0, \frac{N_0 \alpha}{\beta}, \frac{N_0 \alpha (\alpha + 1)}{\beta^2} \right] = [M_0, M_1, M_2] = M, \\ h_{M \rightarrow \xi}^{-1}(M_0, M_1, M_2) &= \left[M_0, \frac{1}{\frac{M_0 M_2}{M_1^2} - 1}, \frac{M_1 M_0}{M_0 M_2 - M_1^2} \right] = [N_0, \alpha, \beta] = \xi. \end{aligned} \quad (16)$$

437 The definition of $h_{M \rightarrow \xi}^{-1}$ shows that α goes to zero when the product of M_0 and M_2 in-
 438 creases over time (M_1 is approximately constant and does not cause much variation in
 439 the value of α). Small values of α lead to instabilities in Cloudy and eventually cause
 440 it to crash. The reason is that in the regime of $\alpha \approx 0$, the small changes in the prog-
 441 nostic moments that the adaptive time stepping produces correspond to large changes
 442 in the underlying distribution parameters. In addition, sampling from Gamma distri-
 443 butions (which is done in the Monte Carlo approximation of the coalescence and breakup
 444 integrals described in Appendix A) becomes inaccurate and inefficient when the shape
 445 parameter is small (e.g., Best, 1983; Liu et al., 2017). Collision-coalescence decreases M_0
 446 and tends to increase M_2 , while collisional breakup increases M_0 and tends to decrease
 447 M_2 . The combined effect can be such that M_0 grows faster than M_2 declines (or vice
 448 visa), resulting in an increasing product of the two moments. For the experiments pre-
 449 sented in this paper, we circumvented this problem by choosing the settings (constraints
 450 for the kernel parameters, duration of the simulations, initial condition, value of the co-
 451 alescence efficiency) such that the resulting simulations were stable.

452 The simulations are initialized with 10^4 particles with a mean mass and standard
 453 deviation 0.33×10^{-9} g, corresponding to initial moments $M_0 = 10^4 \text{ cm}^{-3}$, $M_1 = 3.30 \times$
 454 $10^{-6} \text{ g cm}^{-3}$, and $M_2 = 2.18 \times 10^{-15} \text{ g}^2 \text{ cm}^{-3}$. They are run for a simulation time pe-
 455 riod of 60 s.

456 5 Results

457 5.1 Evolution of the PMD moments

458 Unless the coalescence efficiency is set to 0 or 1, the time evolution of the PMD is
 459 the result of two competing processes: Collision-coalescence reduces the number of droplets
 460 (decreasing M_0) and creates larger drops (increasing M_2), while collisional breakup gen-
 461 erates more but smaller droplets (increasing M_0 and decreasing M_2). Both processes con-
 462 serve the liquid water mass (M_1). To illustrate the effects of differences in the relative
 463 strength of these two processes, Fig. 2 shows the time evolution of M_0 , M_1 , and M_2 for
 464 coalescence efficiencies of 0.8, 0.9, and 1.0 (with a coalescence efficiency of 1.0, there is
 465 no breakup, i.e., all collisions result in coalescence). For each value of the coalescence
 466 efficiency in Fig. 2a, 10 simulations with identical collision kernels are run; the result-
 467 ing spread in the moment evolution is due to the randomness inherent in the Monte Carlo
 468 integration used to compute Eq. (7). As is to be expected, the number of droplets in the
 469 simulation without breakup decreases monotonically, moving toward its theoretical limit
 470 of a single drop containing all the liquid water mass. With a coalescence efficiency of 0.9,
 471 the droplet number decreases more slowly over the time period shown, and with a co-
 472 alescence efficiency of 0.8, breakup produces more droplets than coalescence removes. Mass
 473 is conserved most exactly for the no-breakup simulation, where determining the time rate
 474 of change of the moments does not involve the computation of the breakup term (Eq. 6).
 475 For coalescence efficiencies of 0.8 and 0.9, mass is conserved to within 6% of the initial
 476 mass. These deviations from mass conservation are due to the Monte Carlo approxima-

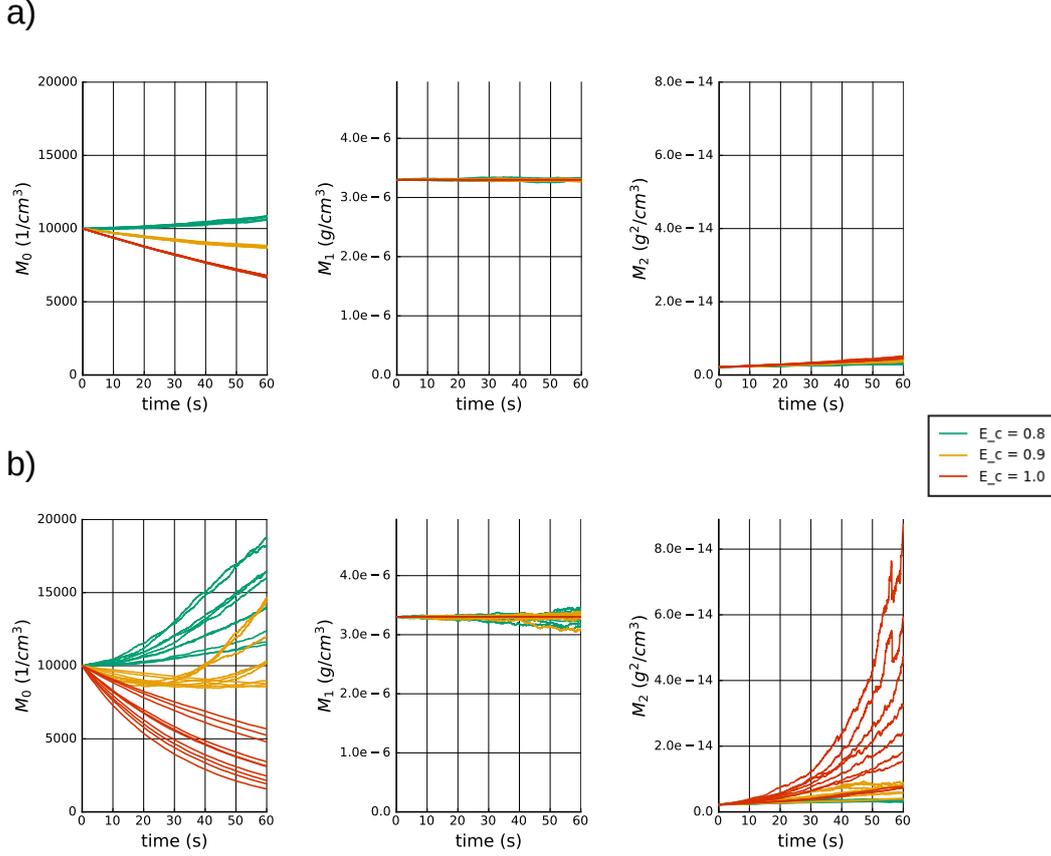


Figure 2. Evolution of the (left) zeroth, (middle) first, and (right) second moment of the PMD, for different values of the coalescence efficiency E_c (0.8, 0.9, and 1.0). All simulations use a kernel of the form $K(x, y) = b(x + y)$. (a) Simulations for each value of E_c are repeated 10 times with the same kernel ($b = 2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$). (b) Simulations for each value of E_c are also repeated 10 times, but each time with a kernel parameter that is randomly drawn from its prior distribution.

477 tion to the coalescence and breakup integrals (Eq. 6) and due to numerical time step-
 478 ping errors.

479 In Fig. 2b, each simulation is performed with a new kernel parameter b drawn from
 480 its prior distribution, such that the observed spread in the moment evolution is due to
 481 the combined randomness of sampling the kernel parameter and of the Monte Carlo in-
 482 tegration. The former clearly accounts for a greater share of the variability, as can be
 483 seen by comparing Figs. 2b and Fig. 2a (note that both figures use the same y axes). Fig-
 484 ure 2b also shows that the bounds on the kernel parameters (Table 1) are large enough
 485 for the M_0 evolution of different coalescence efficiencies to overlap in some cases.

486 5.2 “Sum-of-Masses” Kernel

487 Figure 3a shows the posterior distributions generated by the CES method for four
 488 different values of the parameter b in a sum-of-masses kernel $K(x, y) = b(x + y)$. The
 489 results are shown in the transformed, “unconstrained” space where the CES algorithm
 490 takes place and where the prior distribution of the parameter vector θ is defined. In Fig. 3b,

491 the same results are shown in the constrained space where the model input lives. Note
 492 that our discussion of the posterior distributions is based on Fig. 3a, and all following
 493 results will be displayed only in the unconstrained space.

494 In all four experiments, the maximum a posteriori estimate is a good estimate of
 495 the true parameter value. The narrowest of the four distributions and hence the most
 496 certain parameter estimate is obtained for the smallest parameter value ($b = 2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$),
 497 while the largest of the four parameter values ($b = 8000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$) results in the
 498 distribution with the largest spread. The distribution spread reflects the underlying noise
 499 in the data, which varies with the magnitude of the kernel parameter: the larger its value,
 500 the larger (in absolute value) the error in the Monte Carlo estimate of the coalescence
 501 and breakup integrals describing the time rates of change of the distribution moments
 502 (Eq. 6), and hence the larger the resulting variance in the data. This effect gets multi-
 503 plied because the adaptive time stepper uses smaller step sizes when the solution is chang-
 504 ing fast, leading to more evaluations of the coalescence and breakup integrals over the
 505 course of a simulation (about 10 times more evaluations for $b = 8000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$ than
 506 for $b = 2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$).

507 While the Gaussian approximation obtained from the ensemble mean and covari-
 508 ance of the last EKS iteration is a good approximation of the posterior distributions for
 509 $b = 4000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$ and $b = 6000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$, it underestimates the mass in the
 510 tails of the posterior for $b = 2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$ and does not capture the more cusp-
 511 like shape of the posterior for $b = 8000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$. In all practical applications, the
 512 shape of the posterior is (by definition) unknown a priori and may differ substantially
 513 from the Gaussian approximation obtained in the calibration stage. Accurate uncertainty
 514 quantification thus requires sampling the posterior.

515 The results of these four experiments show that the CES method is able to retrieve
 516 the optimal parameter and provide uncertainty quantification in a perfect-model setting,
 517 for a one-parameter kernel. For comparison, Fig. 3 also shows the posterior distributions
 518 obtained from brute-force MCMC sampling without the calibration and emulation stages,
 519 which is about 10^3 times slower than CES. The similar shapes of the posterior distribu-
 520 tions from these two methods confirm that CES produces a high-quality approximation
 521 to the true solution of this problem. Brute-force sampling was only possible for the two
 522 smaller parameter values; the higher noise levels in the simulations with the larger pa-
 523 rameters caused the MCMC algorithm to get stuck in local maxima of the objective func-
 524 tion. While there are advanced Monte Carlo methods such as simulated annealing (Kirkpatrick
 525 et al., 1983) that are less susceptible to local trapping, CES has the advantage of per-
 526 forming well even with simple MCMC implementations, thanks to the smoothing prop-
 527 erty of the GP emulator.

528 Learning the parameter of a sum-of-masses kernel from data generated by PySDM
 529 results in the posterior distribution shown in Fig. 4. As described in section 3.1, this ex-
 530 periment differs from the perfect-model experiments in that its underlying equation in-
 531 cludes a bias term representing the model discrepancy, and an inflated noise represent-
 532 ing the combined stochasticity of Cloudy and PySDM (Eq. 12). CES is able to provide
 533 uncertainty quantification and a good estimate of the true parameter value in this mod-
 534 ified setup. Note however that in the absence of the “gold standard” posterior distribu-
 535 tion obtained from brute-force MCMC sampling, the quality of the uncertainty quan-
 536 tification is not easily measured. This underlines the importance of testing any approach
 537 to uncertainty quantification in a setting that allows for comparison of the resulting pos-
 538 terior distribution with that obtained from a method such as MCMC, which provably
 539 converges to the desired posterior distribution (e.g., Robert & Casella, 2005).

540

5.3 “Sum-of-Masses Plus Constant” Kernel

541

542

543

We will visualize the output of each of the three stages of the CES algorithm using the example of the “sum-of-masses plus constant” kernel $K(x, y) = a + b(x + y)$ with $a = 4 \times 10^{-6} \text{ cm}^3 \text{ s}^{-1}$ and $b = 3000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$.

544

545

546

547

548

549

550

551

552

553

554

In the first stage (Fig. 5), the EKS algorithm transforms an initial ensemble of $J = 50$ members through successive updates into approximate samples of the posterior distribution. The initial ensemble, randomly drawn from the relatively uninformative prior distributions of the parameters, is spread broadly over the parameter space. Over the course of subsequent iterations (each of which requires 50 model evaluations), the ensemble becomes concentrated near the true parameter values, with the sample mean and covariance of the ensemble sampled from a Gaussian approximation of the posterior distribution. This is a difference to EKI (Iglesias et al., 2013), a closely related optimization method whose iterative updates result in a collapse of the ensemble onto the optimal parameter. EKS produces better training points for the emulator but in its present forms usually requires more iterations.

555

556

557

558

559

560

561

The calibration stage generates $N_{it} \times J = 500$ parameter-data pairs, which are used to train the emulator. The Gaussian process emulator predicts the mean and the variance at any data point in its input space, conditional on the training data (Fig. 6). Thanks to the well chosen training points, which are concentrated around the mean or mode of the posterior distribution, the predictions are most confident (have smallest variance) near the optimal parameter (around $[A, B] = [-0.43, -0.88]$), i.e., near the minimum of the objective function.

562

563

564

565

566

567

568

569

570

571

The MCMC algorithm samples the posterior distribution using the predictions of the emulator instead of actual model evaluations. Figure 7 shows kernel density estimates of the MCMC results, with contours containing 5%, 10%, 50%, 75%, 90%, and 99% of the posterior mass. The true value (blue dot) is captured within the 5% contour of the posterior density, showing that the maximum a posteriori estimate of the parameters obtained by the CES method is a good approximation of the true optimum. Both the mean and shape of the MCMC sampled distribution differ from the distribution of the last EKS ensemble (red dots in Figure 5). Since EKS relies on a Gaussian assumption for the posterior distribution, its output may diverge from the true posterior when that assumption does not hold.

572

5.4 Piecewise Kernel

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

The piecewise kernel represents a scenario where the collision rates of droplets smaller than some (not precisely known) threshold differ from those of droplets exceeding that threshold. Figure 8 shows the posterior distributions for the three parameters of this kernel, together with their true values and prior distributions. The difference between the prior and posterior distributions reflects the information about the kernel parameters learned from the data. In this example, the largest information gain was achieved for the linear rate b , whose posterior shows the smallest uncertainty. However, the information gain from the prior to the posterior of the other two parameters is smaller than that of previous examples, especially for the mass threshold (parameter c). This reflects the difficulty of finding data that provide the necessary constraints on all three parameters of the piecewise kernel, and that make the inverse problem sufficiently well posed. For example, if the mass threshold c is not exceeded during a model run, the resulting output will only be sensitive to the linear rate b . Similarly, the information content of the model output is limited when the mass threshold is too small, when the effect of a is much bigger than that of b (or vice versa), etc. Which parameter can best be retrieved depends on the choice of true parameter values and their constraints, which has to ensure that the model output produced in the calibration stage is sensitive to all under-

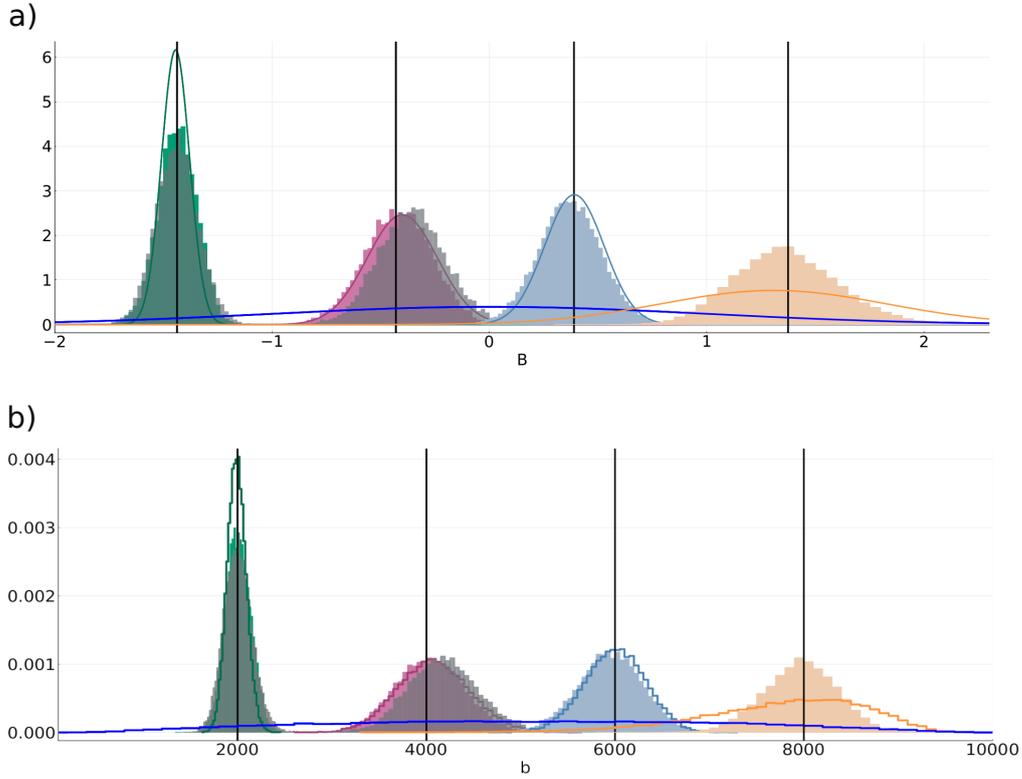


Figure 3. Histograms of MCMC samples of the posterior distributions obtained by the CES algorithm, for the inverse problem of finding the parameter b of a “sum-of-masses” kernel $K(x, y) = b(x + y)$. The results in (a) are shown in the unbounded space where the CES algorithm is performed; in (b), the same results are shown in the bounded space where the model input lives. Different colors correspond to different values of the true parameter b , each of which is marked by a vertical black line (from left to right, in the bounded space: $b = 2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$, $b = 4000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$, $b = 6000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$, $b = 8000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$). The lines show Gaussian approximations to the posterior distributions, which are specified by the ensemble mean and standard deviation of the parameters in the last EKS iteration. All four experiments have the same prior parameter distribution shown as the dark blue line. The two additional histograms in grey (for $b = 2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$ and $b = 4000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$) show the results of brute-force MCMC sampling.

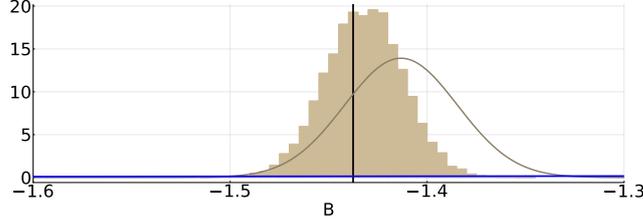


Figure 4. Histogram of MCMC samples produced by CES, showing the posterior distribution of the parameter b of a “sum-of-masses” kernel $K(x, y) = b(x + y)$, given data y generated by PySDM. The Gaussian approximation to the posterior distributions shown as a line is specified by the ensemble mean and standard deviation of the parameters in the last EKS iteration. The prior distribution of the parameter is shown in blue, and the true parameter value ($b = 2000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$ in the bounded space where the model input lives) is marked by a vertical black line. The plot is shown in the unbounded parameter space.

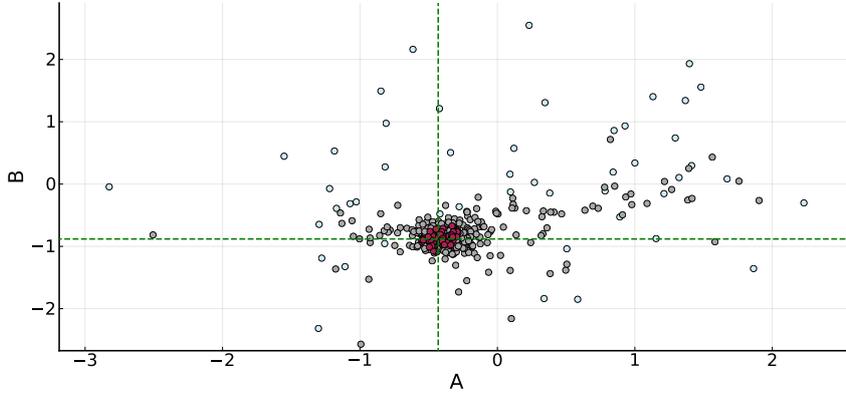


Figure 5. Evolution of the ensemble over subsequent EKS iterations, for the “sum-of-masses plus constant” kernel $K(x, y) = a + b(x + y)$ with $a = 4 \times 10^{-6} \text{ cm}^3 \text{ s}^{-1}$ and $b = 3000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$. The initial ensemble (iteration 0) is highlighted in light blue; the subsequent eight iterations are colored in gray; and the final ensemble (iteration 9) is highlighted in red. The intersection of the dashed green lines represents the true parameter values used to generate observational data. The plot is shown in the unbounded parameter space.

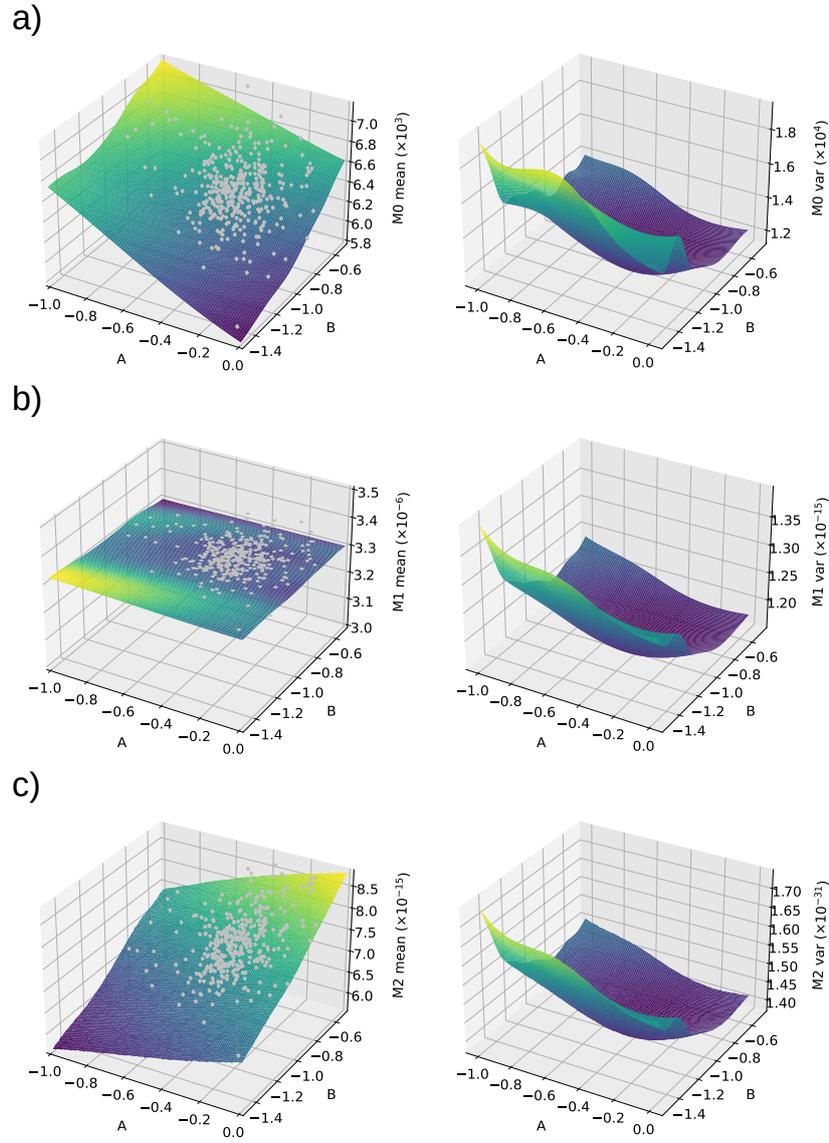


Figure 6. Predictions of the Gaussian process emulator for the “sum-of-masses plus constant” kernel $K(x, y) = a + b(x + y)$ with $a = 4 \times 10^{-6} \text{ cm}^3 \text{ s}^{-1}$ and $b = 3000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$: (left) predicted mean and (right) predicted variance, for (a) M_0 , (b) M_1 , and (c) M_2 . The grey dots represent training points generated during the calibration stage of CES (there are additional ones that fall outside the plotting domain). The predictions are shown in the unbounded parameter space.

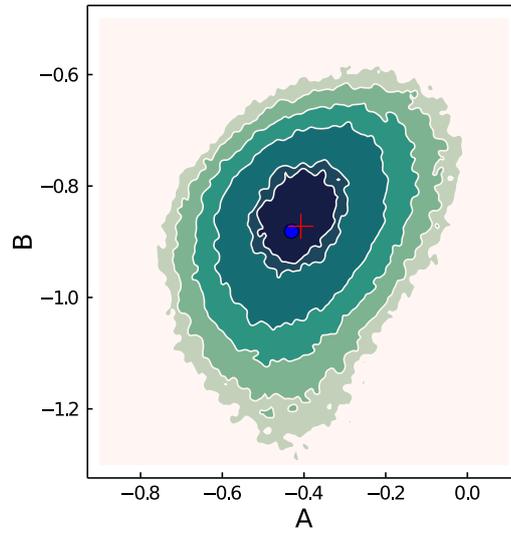


Figure 7. Density plot of MCMC samples of the posterior distribution, for the “sum-of-masses plus constant” kernel $K(x, y) = a + b(x + y)$ with $a = 4 \times 10^{-6} \text{ cm}^3 \text{ s}^{-1}$ and $b = 3000 \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-1}$. The contours contain 5%, 10%, 50%, 75%, 90% and 95% of the sampled distribution. The blue dot marks the true parameters, and the red cross is the average across ensemble members in the last EKS iteration. The posterior densities are shown in the unbounded parameter space.

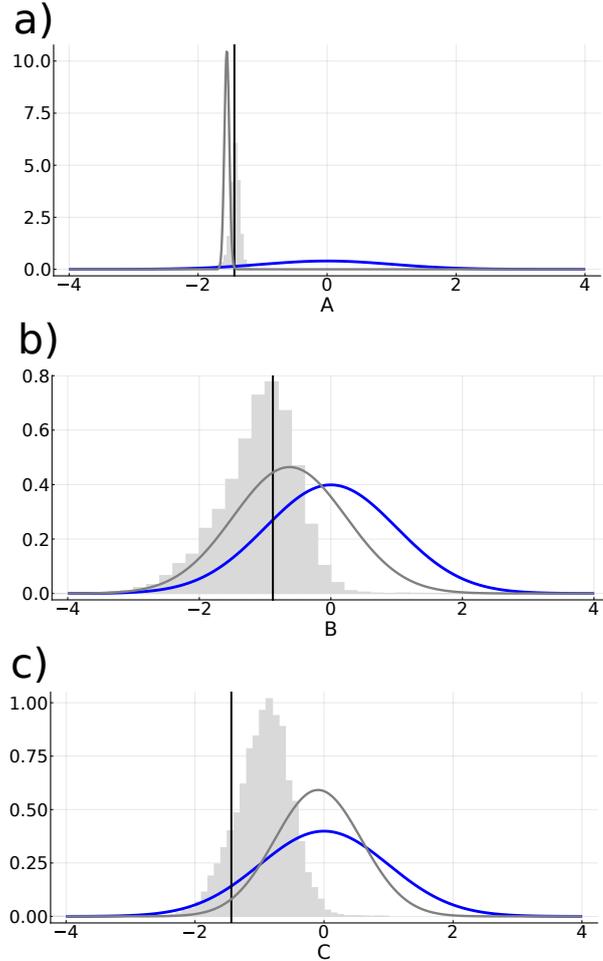


Figure 8. Histograms of MCMC samples of the posterior distributions obtained by the CES algorithm, for the parameters a , b , and c of a piecewise kernel (Eq. 14), given the zeroth, first, and second moment of the PMD at time $t_{\text{end}} = 60$ s. The prior distribution of the parameters is shown in blue, and the true parameter values are marked by vertical black lines. The Gaussian approximations to the posterior distributions shown as grey lines are specified by the ensemble mean and standard deviation of the parameters in the last EKS iteration. The plot is shown in the unbounded parameter space.

590 lying parameters. Choosing prior parameter ranges that give rise to informative data while
 591 not being overly narrow required some exploration of the parameter space.

592 The problem of finding data that are informative enough to constrain microphys-
 593 ical processes is not limited to this proof-of-concept study, where the only data available
 594 are moments of the PMD. In fact, the space- and ground-based observations available
 595 today generally remain incomplete for directly constraining individual microphysical pro-
 596 cess rates in schemes (Morrison, van Lier-Walqui, Fridlind, et al., 2020), and choosing
 597 and combining data for use in Bayesian inversion will arguably be the ultimate challenge
 598 in developing data-informed microphysics schemes.

6 Summary and Conclusions

This paper introduces Cloudy, a flexible microphysics scheme that simulates collision-coalescence and collisional breakup of cloud droplets. We have shown how parameters of the collision kernels describing these droplet-droplet interactions can be learned from data through a computationally efficient Bayesian inversion.

The main points of this study can be summarized as follows:

- Cloudy is a bulk scheme for the collision-coalescence and collisional breakup of cloud droplets. By virtue of its flexible and modular design, the number of prognostic moments can be adjusted to the requirements of the particle mass distribution (PMD), and both PMD and collision kernel can easily be changed. Cloudy is broadly similar to BOSS, the scheme introduced by Morrison, van Lier-Walqui, Kumjian, and Prat (2020), with important differences in how the closure problem is formulated.
- We have looked at microphysics parameterizations through the lens of Bayesian inverse problems and have configured Cloudy to learn parameters of collision kernels from data using Calibrate-Emulate-Sample (CES; Cleary et al., 2021).
- CES is a three-stage approach to Bayesian inversion that is about a factor 1000 faster than traditional techniques. It makes estimation and uncertainty quantification of unknown parameters possible for computationally expensive and/or noisy models.
- CES is able to retrieve posterior parameter distributions in a suite of perfect-model experiments where Cloudy itself generates the data used to constrain the scheme. Results of experiments with different collision kernels show that most posterior distributions capture the true parameter values within 5% of the posterior mass.
- Moving beyond perfect-model experiments, we have learned collision kernel parameters from output generated by PySDM (Bartman et al., 2021), a Lagrangian particle-based microphysics model. In this experiment, we represent model error resulting from the closure assumption in Cloudy (an assumption that PySDM does not need to make) as a simple bias term. This modification in the setup of the inverse problem allows CES to retrieve the posterior distribution of the “true” parameter, not of that which minimizes the mismatch with the PySDM data.

Taken together, they constitute a proof of concept that informing microphysics schemes with data through Bayesian learning is possible in a computationally efficient way. This makes data-informed but physics-based approaches to modeling microphysics a more attractive and viable avenue for future parameterization development. Such approaches have been gaining traction in recent years (e.g., van Lier-Walqui et al., 2020; Morrison, van Lier-Walqui, Fridlind, et al., 2020) as they incorporate the existing physical understanding of microphysical processes while taking advantage of data and statistical tools to bridge knowledge gaps. Bayesian methods are particularly well suited to this task because they allow for continuous updates as the physical theory and understanding of these processes evolve. However, these strengths can only be brought to bear in combination with microphysics schemes that can be constrained in a rigorous and transparent way.

Ultimately, the goal is to inform these schemes by natural observations of clouds and precipitation. This will require methods that account for structural uncertainty, which is uncertainty resulting from the inadequacy of a model to reproduce a given set of data even with the “correct” parameter values, e.g., because it lacks some processes that have been present in producing the data. Neglecting to account for structural uncertainty results in parameter estimates that do not necessarily represent the true physics but that minimize the mismatch between the model output and the given data and hence maximize the predictive accuracy of the emulator (e.g., Kennedy & O’Hagan, 2001). Methods for quantifying structural uncertainty are less well developed than those for parametric uncertainty, but an established approach is to model the structural error as a Gaus-

650 sian process at the interface of model and data (Kennedy & O’Hagan, 2000). An alter-
 651 native is to use Gaussian processes or other machine learning techniques—for example,
 652 neural networks or learning from a dictionary of candidate terms (Brunton et al., 2016;
 653 Schneider et al., 2021)—directly where structural model errors actually occur, for exam-
 654 ple, in the collision kernel. In our example, the direct correspondence of the collision and
 655 breakup kernels between Cloudy and PySDM allowed us to instead use a simple addi-
 656 tive bias term. However, incorporating structural uncertainty in a rigorous way will be
 657 a crucial element to fully exploit the potential of Bayesian inversion in constraining mi-
 658 crophysics models.

659 Appendix A Monte Carlo Integration

660 Monte Carlo integration is a numerical technique that uses random numbers to ap-
 661 proximate integrals. The core idea is to estimate the integral to be calculated by the sam-
 662 ple mean of a sequence of random numbers, whose expected value is the exact value of
 663 the integral. There exist many variants and modifications of Monte Carlo integration that
 664 aim to reduce the variance of the estimator and hence the number of samples needed to
 665 achieve the desired accuracy.

666 Our goal here is to provide some implementation details of the Monte Carlo inte-
 667 gration that is used to compute the time rate of change of the PMD moments in Cloudy,
 668 i.e., the right-hand side of Eq. 7. For a comprehensive treatment of Monte Carlo inte-
 669 gration, the reader is referred to, e.g., Robert and Casella (2005).

Suppose that the integrand $h(x)$ (where $x \in \mathbb{R}^d$) can be written as a product of
 a function $T(x)$ and a probability density $p(x)$. We want to estimate the value of the in-
 tegral

$$I = \int h(x) \, dx = \int T(x)p(x) \, dx = \mathbb{E}_p [T(x)] \quad (\text{A1})$$

Monte Carlo integration consists of generating samples $\{X_1, \dots, X_N\}$ from the density
 p and approximating the integral (A1) by

$$\hat{I}_N = \frac{1}{N} \sum_{i=1}^N T(X_i) \quad (\text{A2})$$

670 Due to the strong law of large numbers, \hat{I}_N converges to I with probability 1 as $N \rightarrow$
 671 ∞ . If the variance σ^2 of $T(X_i)$ is finite, then the standard error, σ_x^2 , is given by

$$\sigma_x^2 := \mathbb{E} \left[(\hat{I}_N - I)^2 \right] = \frac{\sigma}{N} \quad (\text{A3})$$

672 The error is thus independent of the dimension d of the integral.

673 To compute the coalescence integral (see Eq. 7; repeated here for convenience),

$$I_{k,\text{coal}} = \frac{1}{2} \int_0^\infty \int_0^\infty \left((m + m')^k - m^k - m'^k \right) C(m, m') f(m) f(m') dm dm', \quad (\text{A4})$$

674 we define

$$p1(x, y) = \frac{1}{M_0^2} f(x) f(y). \quad (\text{A5})$$

675 Drawing random samples $\{X_i\}_{i=1}^N$ and $\{Y_i\}_{i=1}^N$ from $p1$, $I_{k,\text{coal}}$ can be estimated
 676 by

$$\hat{I}_{N,k,\text{coal}} = \frac{M_0^2}{2N} \sum_{i=1}^N ((X_i + Y_i)^k - X_i^k - Y_i^k) C(X_i, Y_i) \quad (\text{A6})$$

For the computation of the breakup integral $I_{k,\text{breakup}}$ the source and sink term are treated separately:

$$\begin{aligned} I_{k,\text{breakup}} &= \frac{1}{2} \int_0^{m'+m''} m^k P(m; m', m'') dm \int_0^\infty f(m') dm' \int_0^\infty f(m'') B(m', m'') dm'' \\ &\quad - \int_0^\infty m^k f(m) dm \int_0^\infty \frac{f(m'') B(m, m'')}{m + m''} dm'' \int_0^{m+m''} m' P(m'; m, m'') dm' \\ &=: I_{k,\text{breakup}}^{\text{source}} - I_{k,\text{breakup}}^{\text{sink}} \end{aligned} \quad (\text{A7})$$

677 Both terms ($I_{k,\text{breakup}}^{\text{source}}$ and $I_{k,\text{breakup}}^{\text{sink}}$) allow for the analytical integration of one
 678 of the variables, such that the Monte Carlo integration is only applied to the remaining
 679 double integrals. To achieve this, we make use of the fact that parts of the fragment dis-
 680 tribution, $P(m; m', m'') = \beta^2 (m' + m'') \exp(-\beta m)$ (Eq. 3), can be combined with m^k
 681 (for $I_{k,\text{breakup}}^{\text{source}}$) and m' (for $I_{k,\text{breakup}}^{\text{sink}}$) into Gamma distributions with known cumu-
 682 lative density functions (cdf). We start with the breakup source integral, where intro-
 683 ducing a normalization factor $\gamma = \frac{\beta^{\alpha_k - 2}}{\Gamma(\alpha_k)}$ allows for the construction of $\text{Gamma}(\alpha_k, \beta)$
 684 distributions with density

$$p_{2,k}(m) = \frac{\beta^{\alpha_k}}{\Gamma(\alpha_k)} m^{\alpha_k - 1} \exp(-\beta m), \quad (\text{A8})$$

685 with $\alpha_k = k+1$. Outside of $p_{2,k}(m)$, there is no dependence on m left in $I_{k,\text{breakup}}^{\text{source}}$,
 686 and so the integral $\int_0^{m'+m''} p_{2,k}(m) \frac{1}{\gamma} dm = \frac{1}{\gamma}$ can be integrated separately. Its solution
 687 is the cdf of a $\text{Gamma}(\alpha_k, \beta)$ distribution, which for positive integers α_k is given by

$$F(x; \alpha_k, \beta) = 1 - \sum_{i=0}^{\alpha_k - 1} \frac{(\beta x)^i}{i!} \exp(-\beta x), \quad (\text{A9})$$

688 evaluated at $x = m' + m''$. The remaining double integral is then computed us-
 689 ing the same technique as for the coalescence integral: Drawing random samples $\{Y_i\}_{i=1}^N$
 690 and $\{Z_i\}_{i=1}^N$ from $p1(x, y) = \frac{1}{M_0^2} f(x) f(y)$, the Monte Carlo estimate of the breakup source
 691 term is

$$\hat{I}_{N,k,\text{breakup}}^{\text{source}} = \frac{M_0^2}{2\gamma N} \sum_{i=1}^N (Y_i + Z_i) B(Y_i, Z_i) F(Y_i + Z_i; \alpha_k, \beta) \quad (\text{A10})$$

692 A similar approach is applied to the breakup sink integral, which is simplified to
 693 a double integral by the analytic integration of the density of a $\text{Gamma}(2, \beta)$ distribu-
 694 tion (resulting from the product of m' and $\beta^2 \exp(-\beta m)$). Drawing random samples $\{X_i\}_{i=1}^N$
 695 and $\{Z_i\}_{i=1}^N$ from $p1$, the breakup sink term is approximated by

$$\hat{I}_{N,k,\text{breakup}}^{\text{sink}} = \frac{M_0^2}{N} \sum_{i=1}^N X_i^k B(X_i, Z_i) F(X_i + Z_i; 2, \beta). \quad (\text{A11})$$

696 Putting everything together, the time rate of change of the k -th PMD moment due
 697 to collision-coalescence and collisional breakup (Eq. 7) is approximated by

$$\begin{aligned}
 \frac{\partial M_k}{\partial t} &= I_{k,\text{coal}} + I_{k,\text{breakup}} \\
 &\approx \hat{I}_{N,k,\text{coal}} + \hat{I}_{N,k,\text{breakup}} \\
 &= \hat{I}_{N,k,\text{coal}} + \hat{I}_{N,k,\text{breakup}}^{\text{source}} - \hat{I}_{N,k,\text{breakup}}^{\text{sink}}.
 \end{aligned}
 \tag{A12}$$

698 The results shown in this paper were obtained using $N = 200$ Monte Carlo sam-
 699 ples.

700 Appendix B Superdroplet Breakup

701 The superdroplet method to model stochastic collision-coalescence of droplets is
 702 extended in this work to include a breakup-like process for numerical experiments. As
 703 breakup introduces a competing process for coalescence, this stochastic breakup imple-
 704 mentation provides a more realistic set of dynamics to demonstrate the CES algorithm's
 705 ability to learn from data. Here we will briefly describe the modifications made to the
 706 existing package PySDM (Bartman & Arabas, 2021) in order to introduce a breakup-
 707 like process.

708 Maintaining the notation of (Shima et al., 2009), each superdroplet with label in-
 709 dex i for these box model simulations is assigned a multiplicity ξ_i and a mass m_i . A pair
 710 α of superdroplets collides and coalesces with scaled probability p_α , and a random num-
 711 ber ϕ is generated to determine the number of coalescences that occur, γ_α . In the im-
 712 plementation of PySDM, p_α is computed based on a collision rate from kinetic theory
 713 and a coalescence efficiency E_c , which combined are referred to as the coalescence ker-
 714 nel. In this new implementation, every time step also includes the potential for collisional
 715 breakup of a given superdroplet pair. Like p_α , the probability of a breakup is computed
 716 based on the collision rate multiplied with $(1 - E_c)$, and whether the breakup occurs
 717 is determined based on the generation of a new random number.

For a collision-coalescence event that occurs for the SD pair α with γ_α coalescences,
 the multiplicities of the superdroplets are updated as follows:

$$\begin{aligned}
 \xi_j &\leftarrow \xi_j - \gamma_\alpha \xi_k \\
 \xi_k &\leftarrow \xi_k \\
 M_j &\leftarrow M_j \\
 M_k &\leftarrow M_k + \gamma_\alpha M_j
 \end{aligned}
 \tag{B1}$$

718 In this process, the superdroplet j maintains its mass but loses multiplicity to coales-
 719 cence, while droplets k grow due to coalescence with droplets j .

As a substitute for collisional breakup, we treat the process as a collisional coales-
 cence of two superdroplets followed by spontaneous breakup into n_f uniform fragments.
 Thus if a breakup is determined to occur, the same quantity γ_α is computed to deter-
 mine the number of pre-coalescences that occur according to the same dynamics described
 above. Subsequently, the newly coalesced superdroplet k spontaneously fragments:

$$\begin{aligned}
 \xi_k &\leftarrow n_f \xi_k \\
 M_k &\leftarrow M_k / n_f
 \end{aligned}
 \tag{B2}$$

720 For the purposes of these simple test cases, n_f is set to 2 such that the resulting
 721 SD has the average mass of the two colliding SDs. However, we note that this choice is
 722 fundamentally unrealistic, as it would drive the system toward uniformly sized droplets
 723 that are the average size of the initial distribution. For the numerical experiments pre-
 724 sented in this paper, the simulation time is chosen to be far shorter than the time re-
 725 quired for this nonstochastic behavior to become apparent.

Appendix C Open Research

The code for the Calibrate-Emulate-Sample algorithm is available at <https://github.com/CliMA/CalibrateEmulateSample.jl>. The Cloudy repository can be found at <https://github.com/CliMA/Cloudy.jl>.

Acknowledgments

This work was supported by the generosity of Eric and Wendy Schmidt by recommendation of the Schmidt Futures program, by the Hopewell Fund, and the Paul G. Allen Family Foundation.

References

- Andrejczuk, M., Grabowski, W. W., Reisner, J., & Gadian, A. (2010). Cloud-aerosol interactions for boundary layer stratocumulus in the Lagrangian cloud model. *J. Geophys. Res. Atmos.*, *115*(D22). doi: 10.1029/2010JD014248
- Bartman, P., & Arabas, S. (2021). *Super-droplet kernels – backend-level routines for Monte-Carlo particle coagulation solvers: API proposal with CPU and GPU implementation in Python*.
- Bartman, P., Arabas, S., Górski, K., Jaruga, A., Lazarski, G., Olesik, M., ... Talar, A. (2021). *PySDM v1: particle-based cloud modelling package for warm-rain microphysics and aqueous chemistry*.
- Beard, K. V., & Ochs, H. T. (1995). Collisions between small precipitation drops. Part II: Formulas for coalescence, temporary coalescence, and satellites. *J. Atmos. Sci.*, *52*(22), 3977–3996. doi: 10.1175/1520-0469(1995)052<3977:CBSPDP>2.0.CO;2
- Beheng, K. D. (2010). The evolution of raindrop spectra: A review of microphysical essentials. In *Rainfall: State of the science* (pp. 29–48). American Geophysical Union (AGU). doi: 10.1029/2010GM000957
- Best, D. J. (1983). A note on Gamma variate generators with shape parameter less than unity. *Computing*, *30*(2), 185–188. doi: 10.1007/BF02280789
- Borowka, S., Heinrich, G., Jahn, S., Jones, S., Kerner, M., & Schlenk, J. (2019). A GPU compatible quasi-Monte Carlo integrator interfaced to pySecDec. *Comp. Phys. Comm.*, *240*, 120–137. doi: 10.1016/j.cpc.2019.02.015
- Brockwell, P. J., & Davis, R. A. (1996). *Introduction to time series and forecasting*. Springer New York.
- Brunton, S., Proctor, J., & Kutz, J. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl. Acad. Sci.*, *113*, 3932–3937.
- Cleary, E., Garbuno-Inigo, A., Lan, S., Schneider, T., & Stuart, A. M. (2021). Calibrate, emulate, sample. *J. Comp. Phys.*, *424*. doi: 10.1016/j.jcp.2020.109716
- Dunbar, O. R. A., Garbuno-Inigo, A., Schneider, T., & Stuart, A. M. (2021). Calibration and uncertainty quantification of convective parameters in an idealized GCM. *J. Adv. Model. Earth Syst.*, *13*(9), e2020MS002454. doi: <https://doi.org/10.1029/2020MS002454>
- Duncan, A. B., Stuart, A. M., & Wolfram, M.-T. (2021). *Ensemble inference methods for models with noisy and expensive likelihoods*.
- Dziekan, P., & Pawlowska, H. (2017). Stochastic coalescence in Lagrangian cloud microphysics. *Atmos. Chem. Phys.*, *17*(22), 13509–13520. doi: 10.5194/acp-17-13509-2017,2017
- Feingold, G., Tzivion, S., & Leviv, Z. (1988). Evolution of raindrop spectra. Part I: Solution to the stochastic collection/breakup equation using the method of moments. *J. Atmos. Sci.*, *45*(22), 3387–3399. doi: 10.1175/1520-0469(1988)045<3387:EORSPI>2.0.CO;2

- 776 Garbuno-Inigo, A., Hoffmann, F., Li, W., & Stuart, A. M. (2020). Interacting
777 Langevin diffusions: Gradient structure and ensemble Kalman sampler. *SIAM*
778 *J. Appl. Dyn. Syst.*, *19*(1), 412–441. doi: 10.1137/19M1251655
- 779 Garbuno-Inigo, A., Nüsken, N., & Reich, S. (2020). Affine invariant interacting
780 Langevin dynamics for Bayesian inference. *SIAM J. Appl. Dyn. Syst.*, *19*(3),
781 1633–1658. doi: 10.1137/19M1304891
- 782 Golovin, A. M. (1963). On the kinetic equation for coagulating cloud droplets with
783 allowance for condensation. *Izv. Acad. Sci. USSR, Ser. Geophys.*, *10*, 949–953.
- 784 Howland, M. F., Dunbar, O., & Schneider, T. (2021). Parameter uncertainty quan-
785 tification in an idealized GCM with a seasonal cycle. *ESSOAr*, *26*.
- 786 Iglesias, M. A., Law, K. J. H., & Stuart, A. M. (2013). Ensemble Kalman methods
787 for inverse problems. *Inverse Probl.*, *29*(4), 045001. doi: 10.1088/0266-5611/
788 29/4/045001
- 789 Kanzaki, J. (2011). Monte Carlo integration on GPU. *Eur. Phys. J. C*, *71*(2).
- 790 Kennedy, M. C., & O’Hagan, A. (2000). Predicting the output from a complex com-
791 puter code when fast approximations are available. *Biometrika*, *87*(1), 1–13.
- 792 Kennedy, M. C., & O’Hagan, A. (2001). Bayesian calibration of computer mod-
793 els. *J. R. Stat. Soc.: Series B (Statistical Methodology)*, *63*(3), 425–464. doi:
794 10.1111/1467-9868.00294
- 795 Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated
796 annealing. *Science*, *220*(4598), 671–680.
- 797 Kleijnen, A. A. N., Jack P. C. and Ridder, & Rubinstein, R. Y. (2013). Variance re-
798 duction techniques in Monte Carlo methods. In S. I. Gass & M. C. Fu (Eds.),
799 *Encyclopedia of operations research and management science* (pp. 1598–1610).
800 Boston, MA: Springer. doi: 10.1007/978-1-4419-1153-7_638
- 801 Liu, C., Martin, R., & Syring, N. (2017). Efficient simulation from a gamma distri-
802 bution with small shape parameter. *Comp. Stat.*, *32*(4), 1767–1775. doi: 10
803 .1007/s00180-016-0692-0
- 804 Long, A. B. (1974). Solutions to the droplet collection equation for polynomial ker-
805 nels. *J. Atmos. Sci.*, *31*(4), 1040–1052. doi: 10.1175/1520-0469(1974)031<1040:
806 STTDCE>2.0.CO;2
- 807 Low, T. B., & List, R. (1982). Collision, coalescence and breakup of rain-
808 drops. Part I: Experimentally established coalescence efficiencies and frag-
809 ment size distributions in breakup. *J. Atmos. Sci.*, *39*(7), 1591–1606. doi:
810 10.1175/1520-0469(1982)039<1591:CCABOR>2.0.CO;2
- 811 Morrison, H., Kumjian, M. R., Martinkus, C. P., Prat, O. P., & van Lier-Walqui, M.
812 (2019). A general n-moment normalization method for deriving raindrop size
813 distribution scaling relationships. *J. Appl. Meteorol. Climatol.*, *58*(2), 247–267.
814 doi: 10.1175/JAMC-D-18-0060.1
- 815 Morrison, H., van Lier-Walqui, M., Fridlind, A. M., Grabowski, W. W., Harrington,
816 J. Y., Hoose, C., . . . Xue, L. (2020). Confronting the challenge of modeling
817 cloud and precipitation microphysics. *J. Adv. Model. Earth Syst.*, *n/a*(n/a),
818 e2019MS001689. doi: 10.1029/2019MS001689
- 819 Morrison, H., van Lier-Walqui, M., Kumjian, M. R., & Prat, O. P. (2020). A
820 Bayesian approach for statistical-physical bulk parameterization of rain micro-
821 physics. Part I: Scheme description. *J. Atmos. Sci.*, *77*(3), 1019–1041. doi:
822 10.1175/JAS-D-19-0070.1
- 823 Posselt, D. J. (2016). A Bayesian Examination of Deep Convective Squall-Line Sen-
824 sitivity to Changes in Cloud Microphysical Parameters. *J. Atmos. Sci.*, *73*(2),
825 637–665. doi: 10.1175/JAS-D-15-0159.1
- 826 Posselt, D. J., & Vukicevic, T. (2010). Robust characterization of model physics
827 uncertainty for simulations of deep moist convection. *Mon. Wea. Rev.*, *138*(5),
828 1513–1535. doi: 10.1175/2009MWR3094.1
- 829 Pruppacher, H., & Klett, J. (1978). *Microphysics of clouds and precipitation* (1st
830 ed.). Reidel.

- 831 Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine*
832 *learning*. Cambridge, MA, USA: MIT Press.
- 833 Riechelmann, T., Noh, Y., & Raasch, S. (2012). A new method for large-eddy sim-
834 ulations of clouds with Lagrangian droplets including the effects of turbulent
835 collision. *New J. Phys.*, *14*(6), 065008. doi: 10.1088/1367-2630/14/6/065008
- 836 Robert, C. P., & Casella, G. (2005). *Monte Carlo statistical methods*. Berlin, Heidel-
837 berg: Springer Verlag.
- 838 Schneider, T., Lan, S., Stuart, A., & Teixeira, J. (2017). Earth system model-
839 ing 2.0: A blueprint for models that learn from observations and targeted
840 high-resolution simulations. *Geophys. Res. Lett.*, *44*(24), 12,396–12,417. doi:
841 10.1002/2017GL076101
- 842 Schneider, T., Stuart, A. M., & Wu, J. (2021). Ensemble Kalman inversion for
843 sparse learning of dynamical systems from time-averaged data. *J. Comp.*
844 *Phys.*.
- 845 Shima, S., Kusano, K., Kawano, A., Sugiyama, T., & Kawahara, S. (2009). The
846 super-droplet method for the numerical simulation of clouds and precipita-
847 tion: a particle-based and probabilistic microphysics model coupled with a
848 non-hydrostatic model. *Q. J. R. Meteorol. Soc.*, *135*(642), 1307–1320. doi:
849 10.1002/qj.441
- 850 Shima, S., Sato, Y., Hashimoto, A., & Misumi, R. (2019). Predicting the mor-
851 phology of ice particles in deep convection using the super-droplet method:
852 development and evaluation of SCALE-SDM 0.2.5-2.2.0/2.2.1. *Geosci. Model*
853 *Dev. Discuss.*, *2019*, 1–83. doi: 10.5194/gmd-13-4107-2020
- 854 Smoluchowski, M. V. (1916). Drei Vorträge über Diffusion, Brownsche Bewegung
855 und Koagulation von Kolloidteilchen. *Z. Phys.*, *17*, 557–585.
- 856 van Lier-Walqui, M., Morrison, H., Kumjian, M. R., Reimel, K. J., Prat, O. P.,
857 Lunderman, S., & Morzfeld, M. (2020). A Bayesian approach for statisti-
858 cal–physical bulk parameterization of rain microphysics. Part II: Idealized
859 Markov chain Monte Carlo experiments. *J. Atmos. Sci.*, *77*(3), 1043–1064. doi:
860 10.1175/JAS-D-19-0071.1
- 861 van Lier-Walqui, M., Vukicevic, T., & Posselt, D. J. (2014). Linearization
862 of microphysical parameterization uncertainty using multiplicative pro-
863 cess perturbation parameters. *Mon. Wea. Rev.*, *142*(1), 401–413. doi:
864 10.1175/MWR-D-13-00076.1
- 865 Weisberg, S. (2014). *Applied linear regression* (Fourth ed.). Hoboken NJ: Wiley.