Pacific Northwest
NATIONAL LABORATORY

The BOSS microphysics framework

March 20, 2023

Sean Patrick Santos
with Marcus van-Lier Walqui and Hugh Morrison

U.S. DEPARTMENT OF ENERGY BATTELLE

PNNL is operated by Battelle for the U.S. Department of Energy

Overview

- What is BOSS?
  - Bayesian methodology
  - Physical/mathematical constraints
- Results from emulating the TAU bin model in idealized 1-D driver:
  - Evaluation contexts: direct vs. time-evolving
  - What information is most useful for predicting autoconversion?
- Further steps:
  - Single-category liquid microphysics
  - Emulation in LES
  - BOSS in EAM

What is BOSS?

- BOSS is the “Bayesian Observationally constrained Statistical-physical Scheme”, actually a family of microphysics schemes developed with similar methodologies.
- BOSS schemes are data-driven (supervised machine learning).
  - Input data theoretically can be any output of a model implementing BOSS.
    - ✓ Microphysical process rates themselves
    - ✓ Precipitation and liquid water content
    - ✓ Averaged outputs from climate models (e.g. annual-mean radiative forcings)
  - How to best combine/weight data from very different sources is not settled.
- BOSS schemes are not “black-box”.
  - The process rate equations in a BOSS scheme can be written down succinctly and look like those of traditional parameterizations.

Basic anatomy of a bulk microphysics scheme

- We start with a representation of the drop size distribution (DSD).
  - In a bulk scheme we typically use the distribution's moments to describe it.
    - ✓ For a DSD $N(D)$ that is a function of diameter D, the n-th moment is $M_n = \int_0^\infty D^n N(D) dD$.
    - ✓ 0th moment is total number, 3rd proportional to total liquid mass, 6th is radar reflectivity factor, etc.
- The goal of the microphysics scheme is to evolve these moments in time:
  - $\frac{dM_n}{dt} = \left(\frac{dM_n}{dt}\right)_{proc1} + \left(\frac{dM_n}{dt}\right)_{proc2} + \ldots$
  - Proc1, proc2, etc. denote the processes that affect a given moment, e.g. condensation, autoconversion, accretion, sedimentation…
  - For most processes, BOSS uses sums of power laws, e.g. $\left(\frac{dM_n}{dt}\right)_{coll} = \sum_i a_i \prod_j M_j^{b_{ij}}$
- For sedimentation we predict fall speed.
  - $\left(\frac{dM_n}{dt}\right)_{sed} = \frac{d}{dz}(v_n M_n)$, $v_n = \sum_i a_i \prod_j M_j^{b_{ij}}$, (with some corrections e.g. for air density)

How BOSS schemes are developed

1. Get data to use for training the scheme (and evaluation).
2. Choose prognostic moments for the microphysics scheme.
3. Choose an ansatz of process rate formulas with a moderately large number of adjustable parameters (~25-50 currently).
4. Choose an "error model", an assumed probability distribution for model-data discrepancy.
5. Use Markov chain Monte Carlo (MCMC) to find probability distribution of parameter values via Bayesian inference.
   - Need some (typically weak) prior distribution dictating plausible parameter values.
   - We typically use the maximum likelihood estimator (MLE) as estimated "best" values, but also get the whole posterior probability distribution.

Flow of data in MCMC used to train BOSS

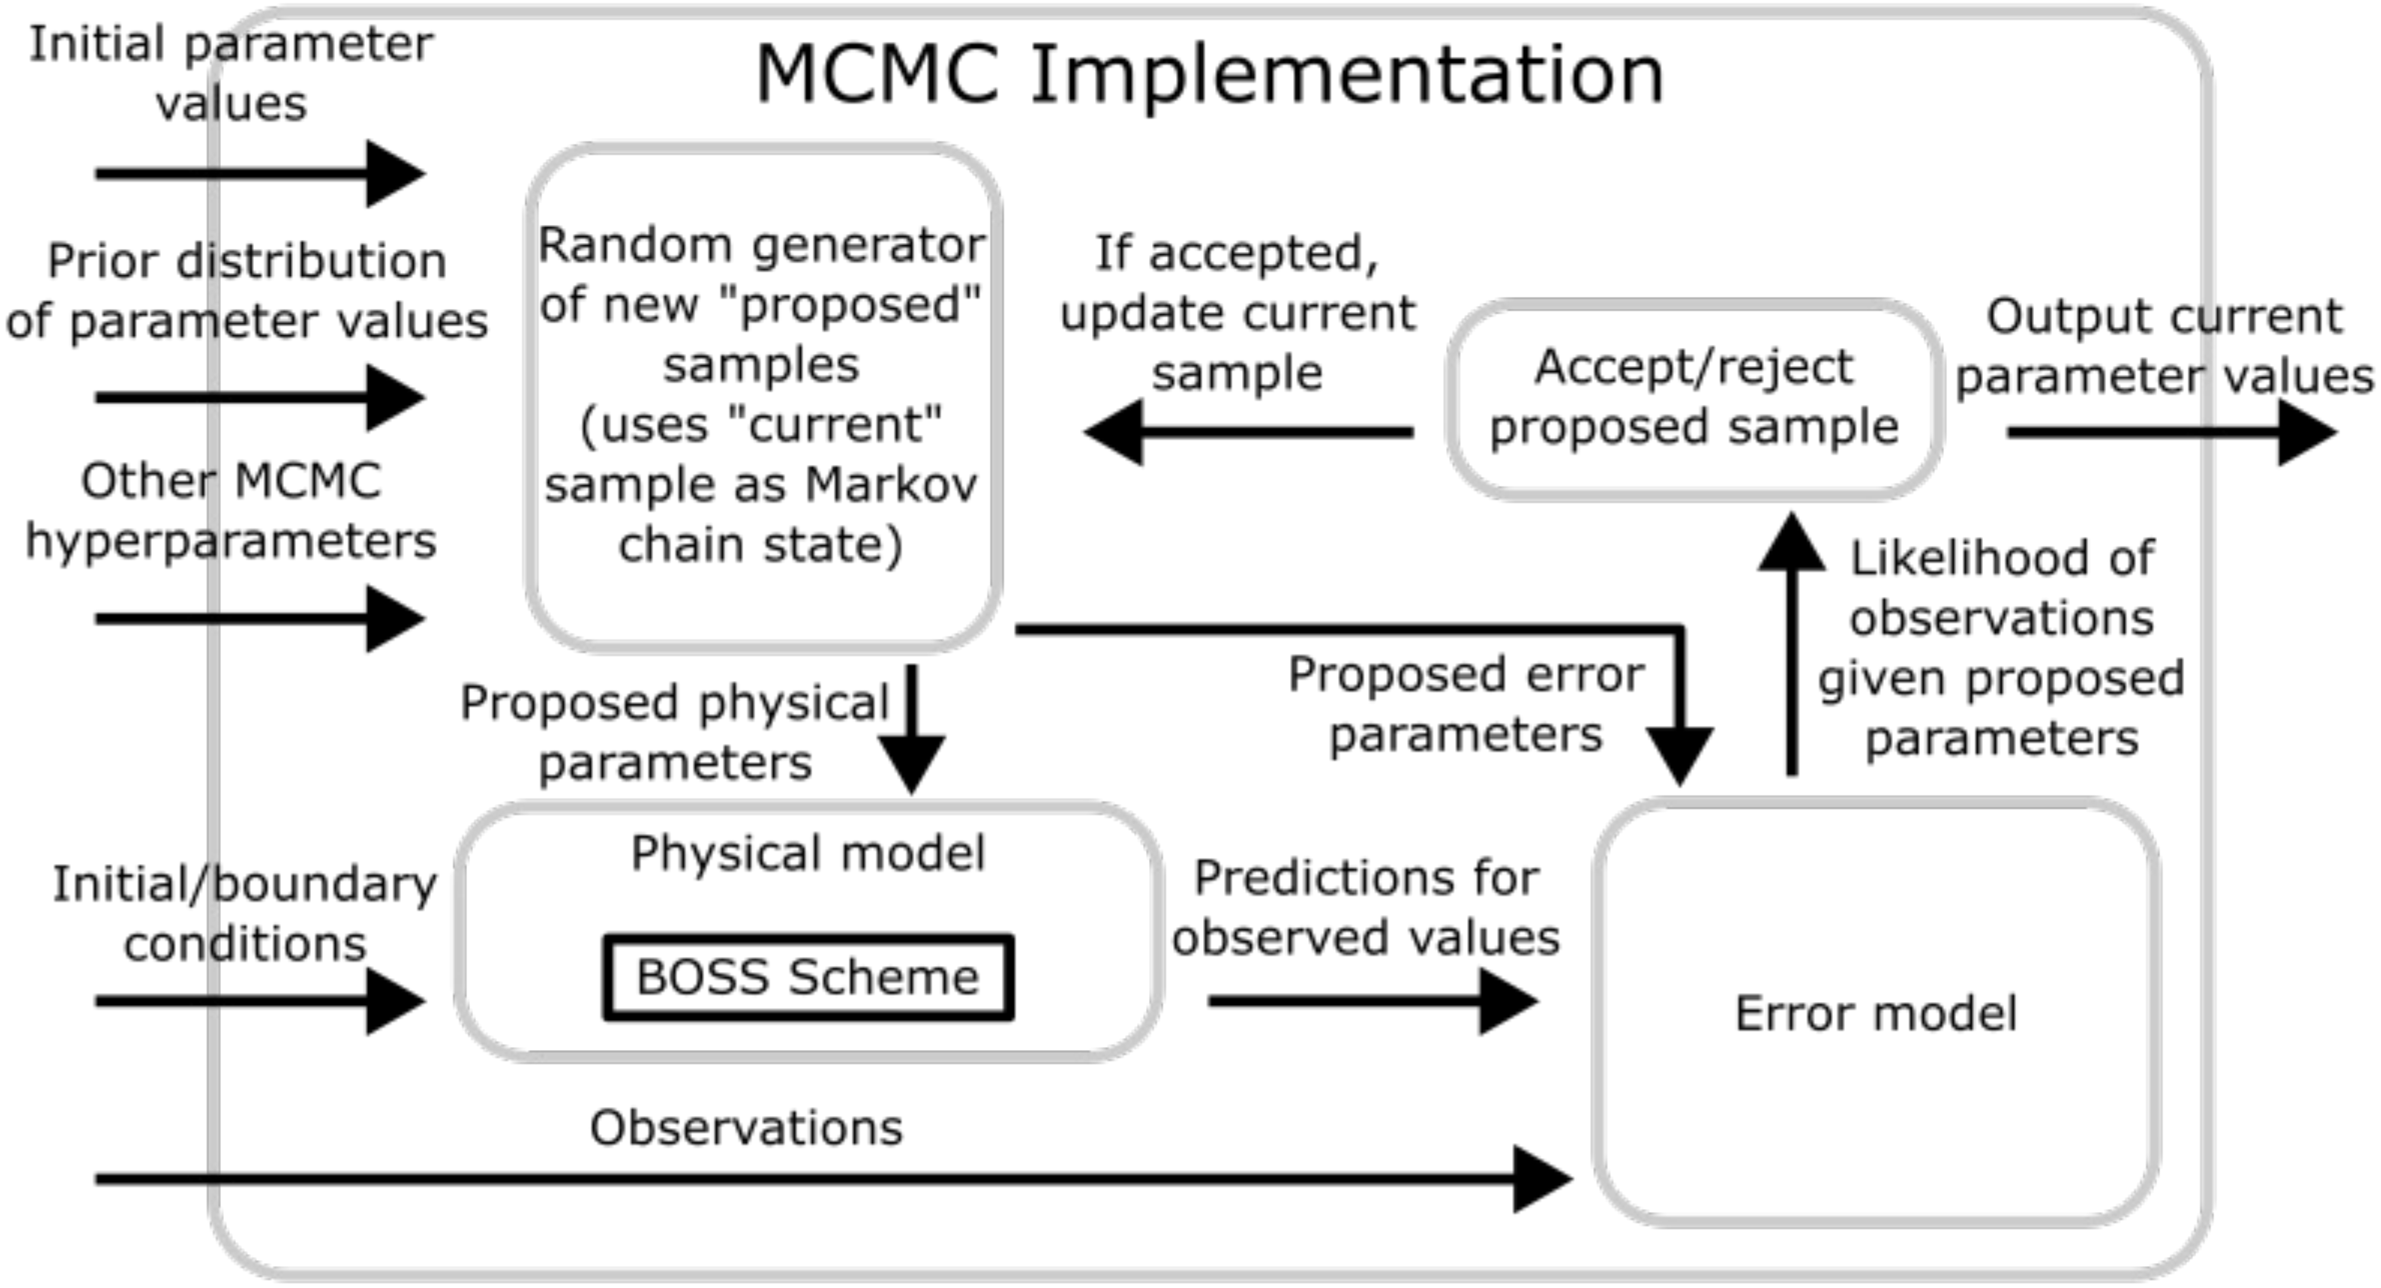

Using physical/mathematical constraints

- Power law series are more general than Taylor series, so we can approximate any analytic function (and some discontinuous ones) by using a large enough number of terms.
- However, we prefer to keep things as simple as possible:
  - Fewer power laws is computationally cheaper.
  - Fewer parameters means MCMC performs better.
  - Easier for humans to understand.
- We therefore rely on a few physical/mathematical constraints to specialize the power laws.
  - All microphysical processes exhibit certain scaling behaviors (e.g. average fall speeds are intrinsic quantities, scaling symmetry of the stochastic collection equation).
  - Moments must correspond to a valid distribution, e.g. for a three-moment scheme, the variance in particle size must be positive.

Emulating the TAU bin model

- As a proof-of-concept to see how well BOSS handles complex microphysics problems, we developed BOSS emulators of the Tel Aviv University's bin model (TAU).
  - I.e. the TAU bin model output was used as a source of "observations".

- Emulation was of the liquid microphysics only.

- Focused on emulation of the non-precipitating and drizzling stratocumulus regimes in a 1-D kinematic driver.

Bin model results in simulated driver

- Shown is cloud and rain mass over time for the strongest drizzling case.
- CCN = 10/cm^3 produces peak surface flux of ~0.25 mm/hr, so relatively “heavy” for drizzle.

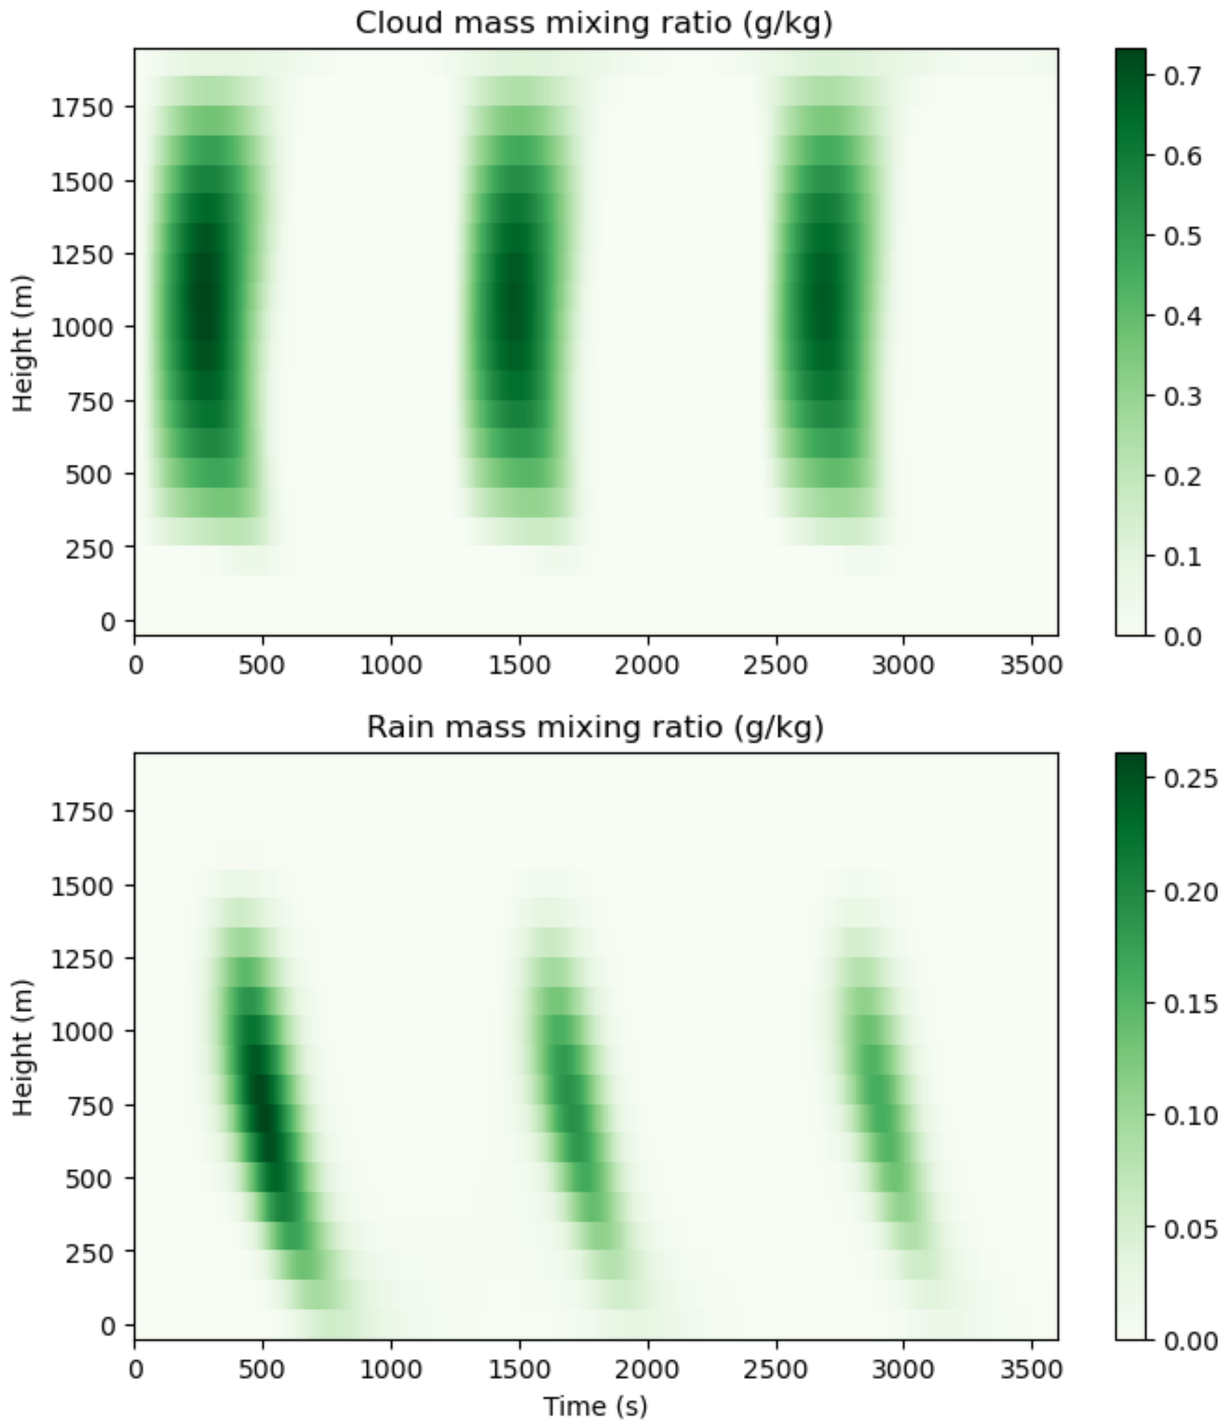

BOSS configurations

- Produced 6 different BOSS emulators.
- Half were standard two-moment schemes, like MG2 and P3 (labeled “2M”).
- The other half used three moments for cloud, two moments for rain (“3M”).
- For each number of moments, three different autoconversion formulas:
  - “NR1T” – Single power law term, depends on cloud moments only.
  - “NR2T” – Two power law terms, depends on cloud moments only
  - “RD” – Two power law terms, one “initiation” or “triggering” term depending only on cloud moments, and one “continuation” or “tail” term depending also on rain.
- First versions of these schemes were tuned using a traditional “offline” method, which ignores time evolution of the system.
  - Treat moment values as “inputs” and bin model process rates as “outputs”.
  - Simple regression problem: fit the process rate formulas to the data.

Offline tuning struggles with rain

- All schemes perform fine for non-precipitating cloud, but struggle with rain.
- All tend to underestimate rain production.
- Surface rates are OK for 2M-RD, 3M-RD, and 3M-NR2T.
  - However, this is partly due to compensating errors; less drizzle produced, but also less evaporation than bin.

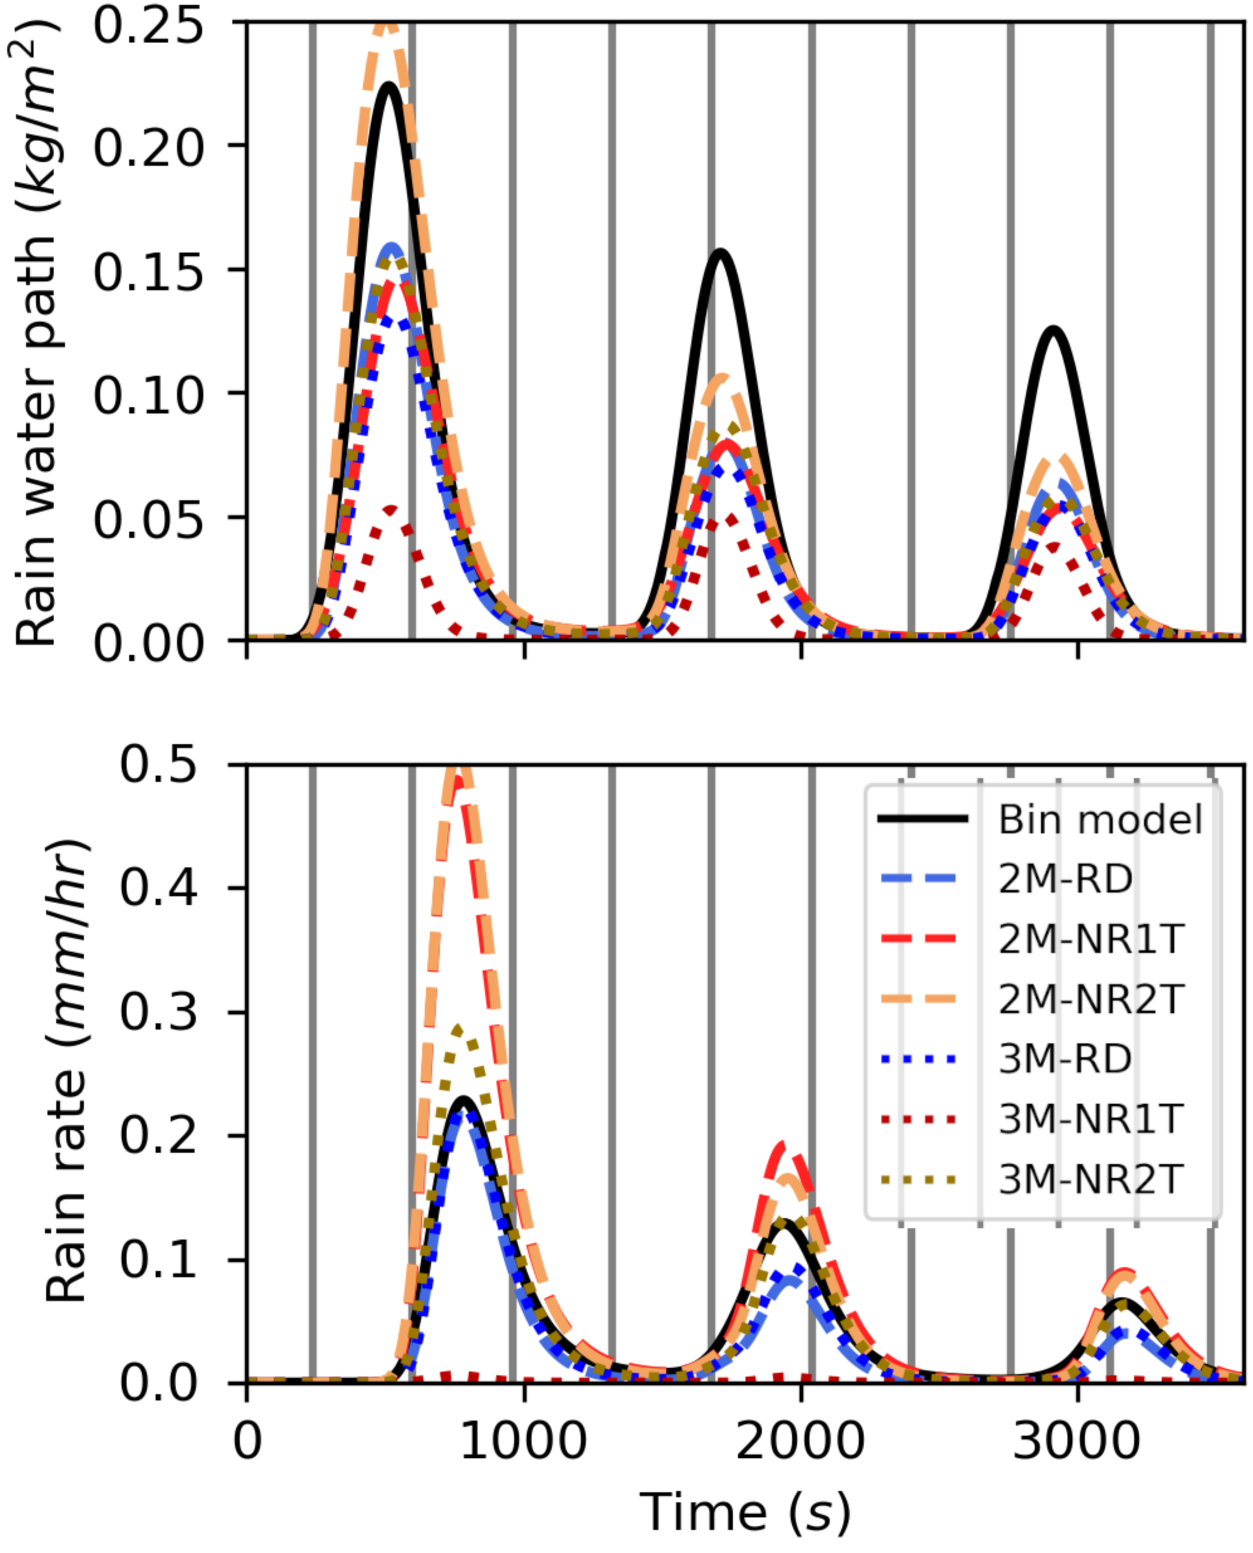

Pacific Northwest
NATIONAL LABORATORY

Trouble with autoconversion

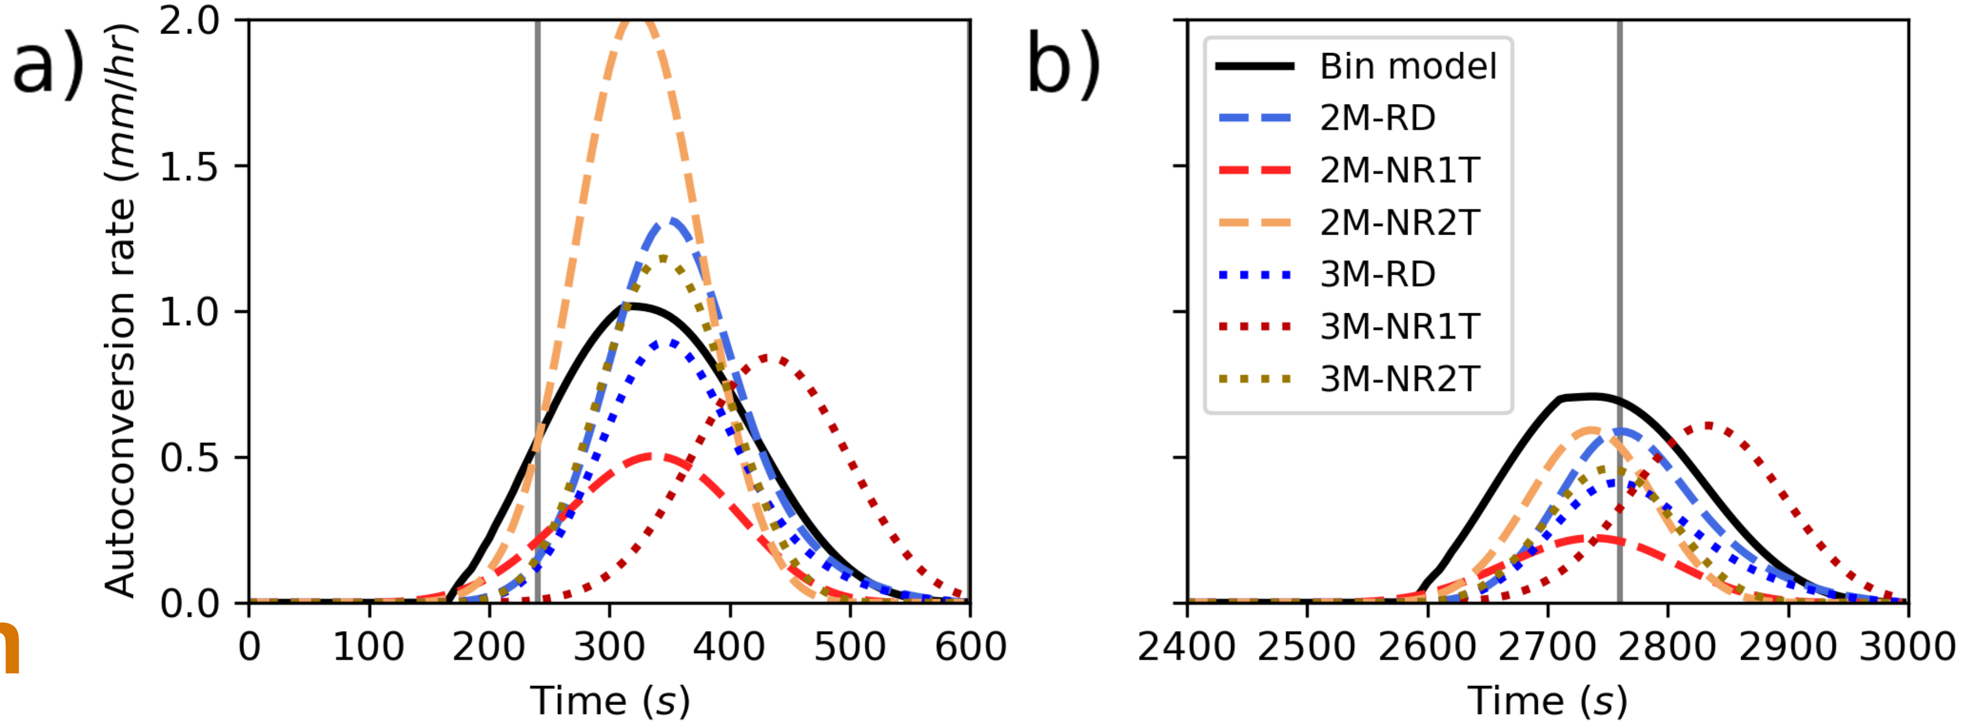

- Main problem for these schemes is capturing autoconversion.
  - Rates shown for (a) first oscillation, and (b) third oscillation of same run.
- Most schemes have peak autoconversion too late, comes on too slow.
  - Reduces overall rain production because accretion starts later.
  - Accretion/autoconversion ratio too high, so rain drops are too large.
- Schemes may be too sensitive to size of cloud drops.
  - Comparing oscillations with slightly larger drops (a) to slightly smaller drops (b), most BOSS schemes have a much larger change in autoconversion rate than bin model.

Inference in a time-evolving context

- We know that fitting instantaneous process rates “offline” is insufficient.
  - Seifert and Rasp (2020): Neural networks that are highly accurate offline emulators of a superdroplet scheme are still worse than a two-decade-old “traditional” scheme at capturing timing of precipitation.
- In the real world, we can’t directly observe most process rates for most moments anyway.
- What if we train BOSS using its performance in the 1-D driver instead?

Observations for time-evolving inference

- We chose four vertical levels (evenly spaced 500m apart).
- We chose ten times (evenly spaced 6 min apart).
- At each time, we "observe":
  - All modeled cloud and rain moments at each of the chosen vertical levels.
  - Column-integrated liquid water path.
  - Total (cloud+rain) surface fluxes of mass and number.
- These observations are used to train new variants of the BOSS schemes (labeled "TE" schemes).
- Instantaneous process rates not used except to help define a (very weak) prior on parameter values.

Time-evolving fits are better for rain

- Improvement here is unsurprising; these schemes are directly optimized for these observations.
  - However, the overall accuracy of most schemes shows that structural uncertainty of BOSS is not too high.
- Pink curve is 2M-RD with uniform prior for parameters.
  - Mainly shows that choice of prior is not that important.

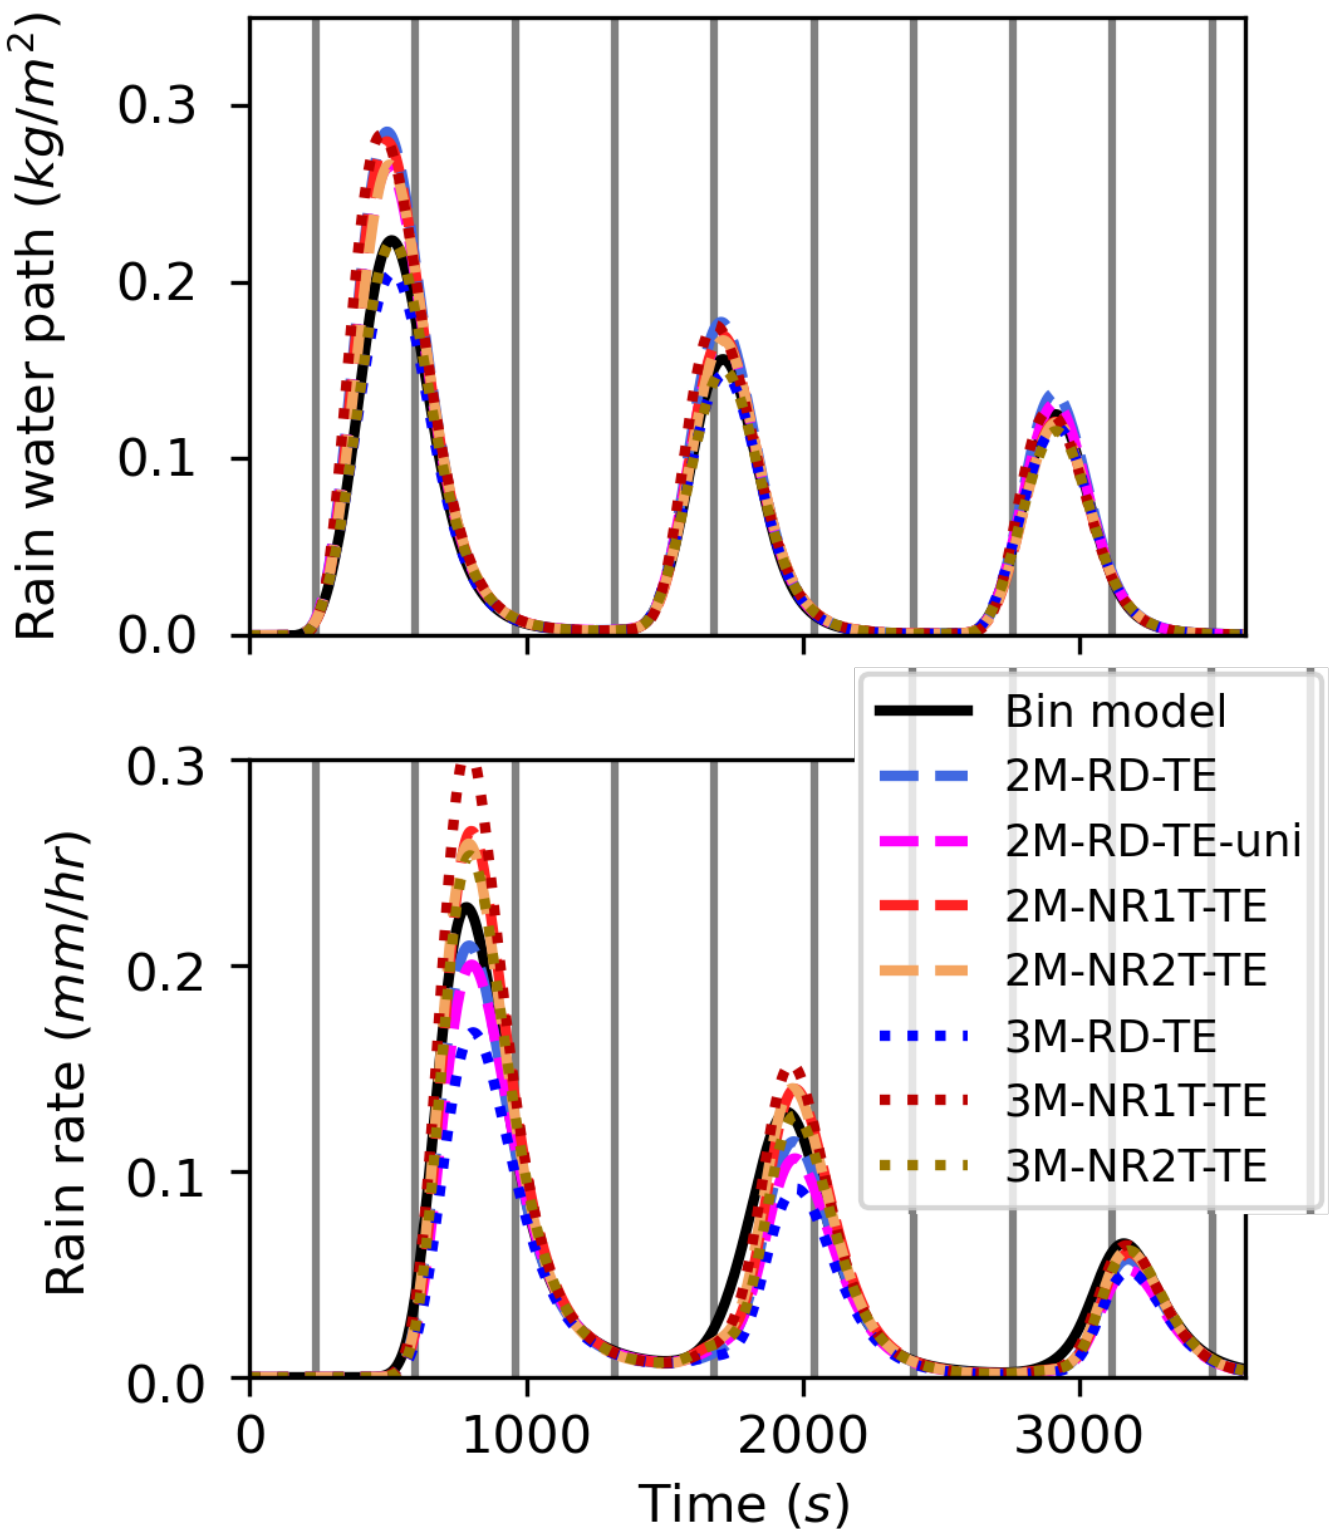

Autoconversion from BOSS TE schemes

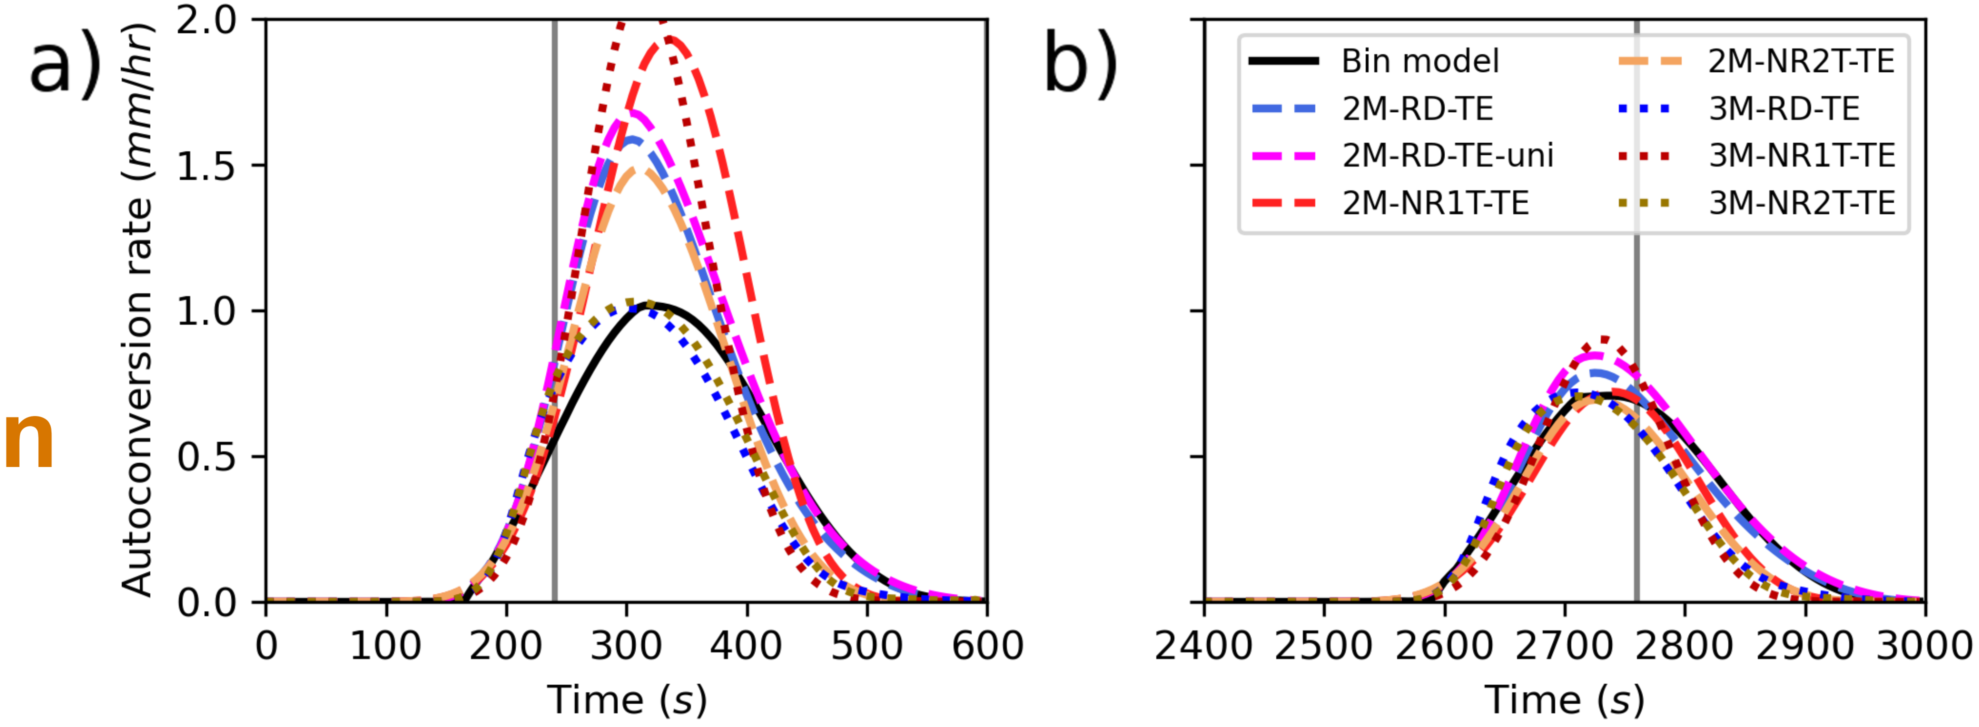

- The time-evolving regression has corrected the timing of autoconversion.
- 2M-*-TE schemes and 3M-NR1T-TE are not flexible enough to perfectly represent every regime in our data set.
  - Some excessive drop size sensitivity remains, and causes too much autoconversion for the largest drop sizes.
- The 3M-2T-TE and 3M-RD-TE schemes match autoconversion very well.
  - These schemes were not tuned on the autoconversion rate!
  - Inferred correct process-level physics from observations of rain number/mass.
  - Observations are pretty sparse: grey lines show observation times.

Rain process rates for TAU bin model and selected TE schemes

Rain mass

Rain mass autoconversion rate

Rain mass accretion rate

Rain mass evaporation rate

Rain mass-weighted fall speed

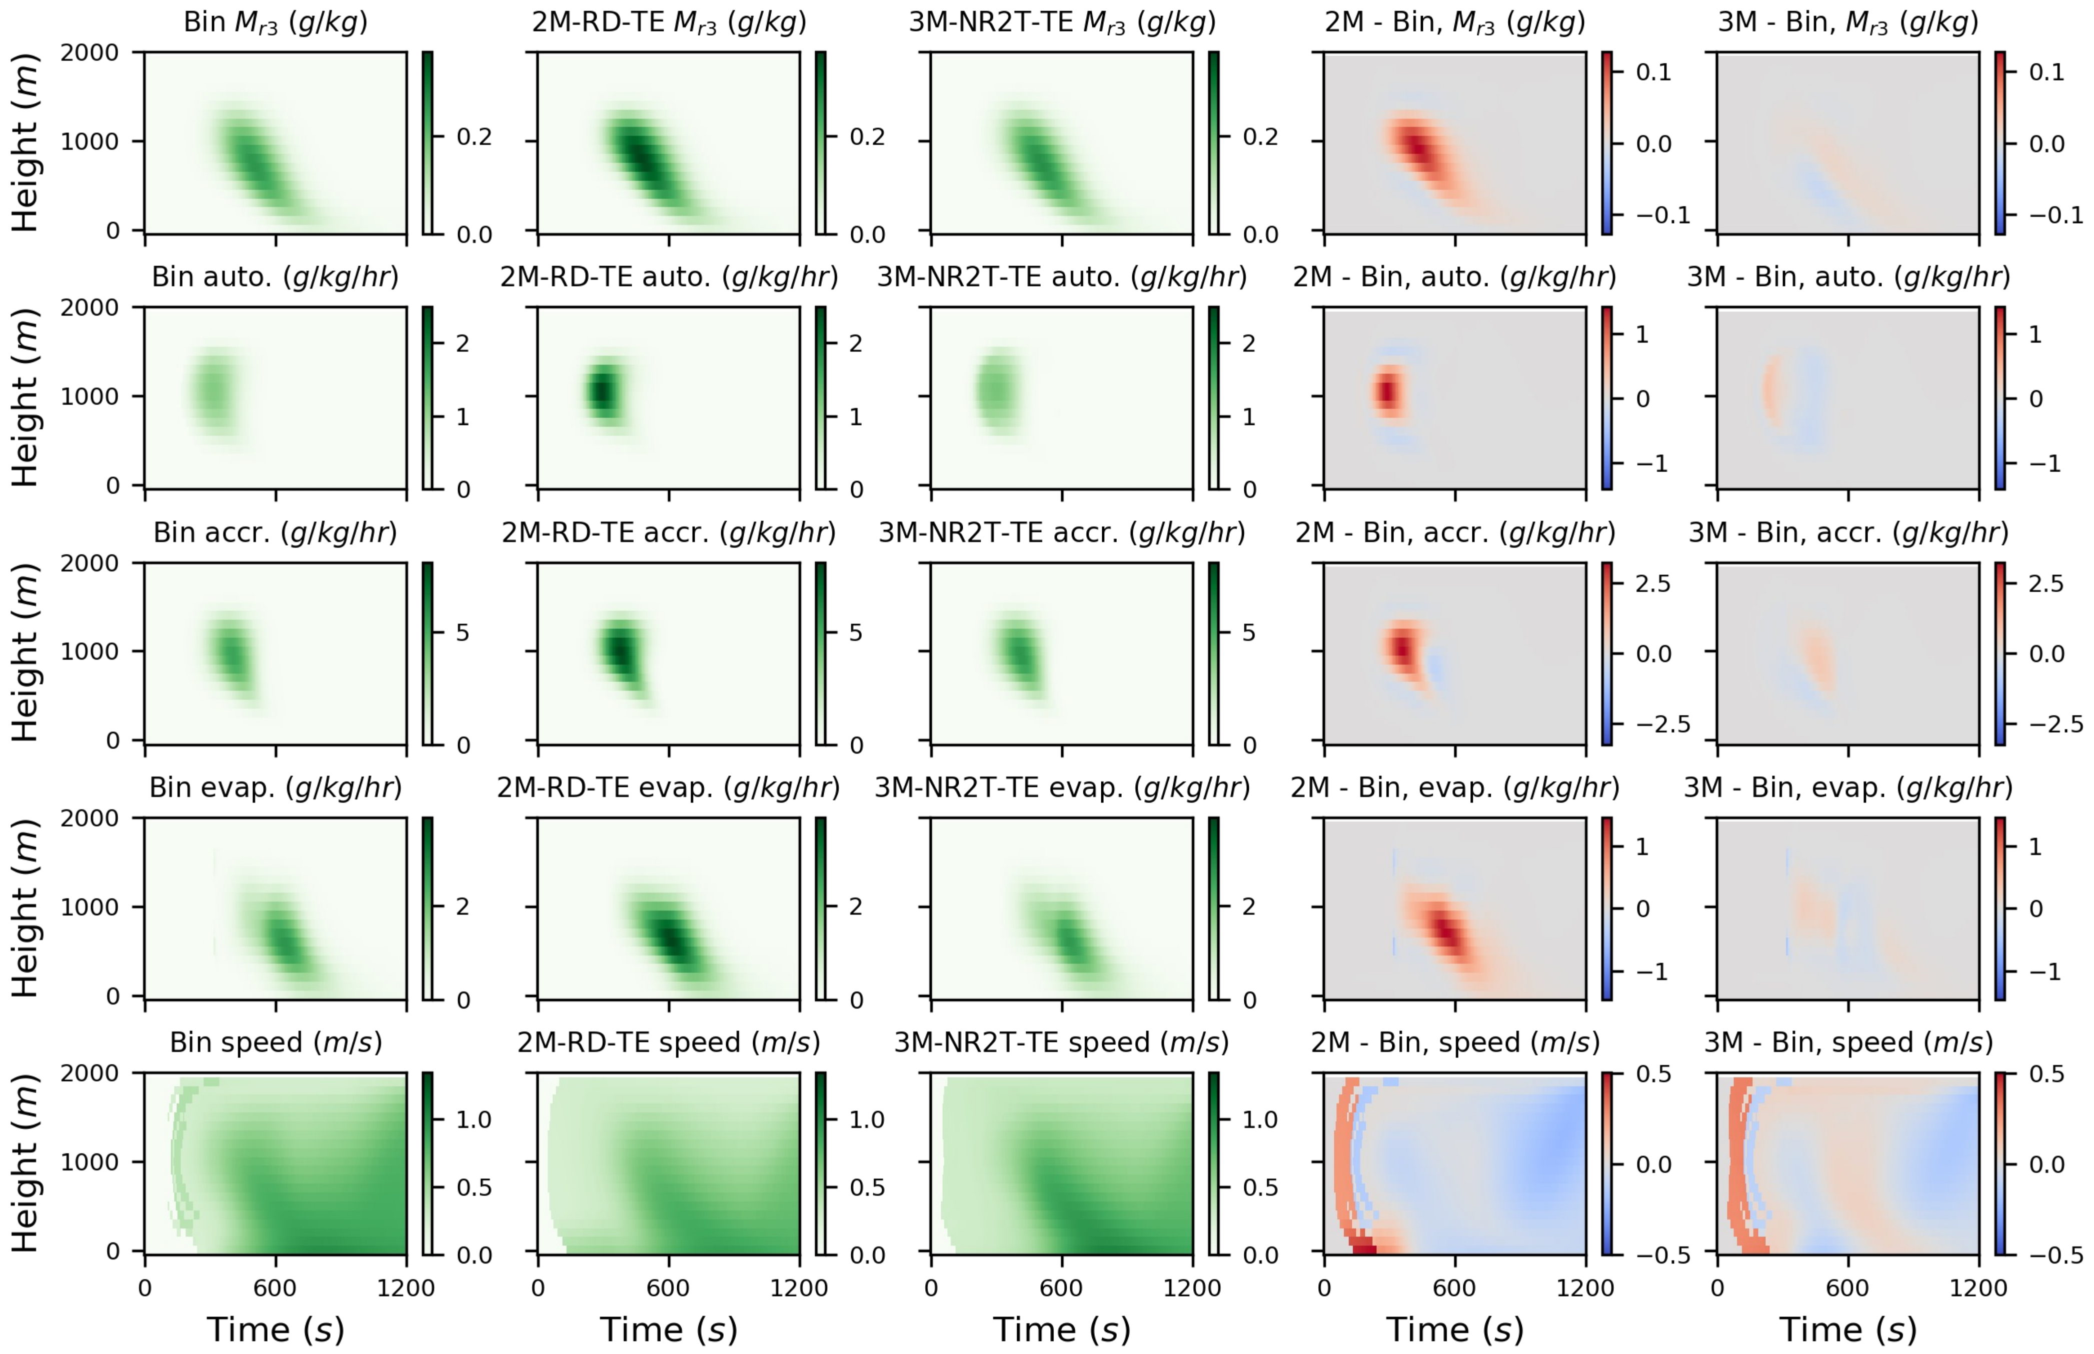

Remaining deficiencies of TE schemes

- Two-moment BOSS is still not perfect at handling autoconversion.
  - Seifert and Rasp (2020): Predicting timing of autoconversion from two cloud moments may be inherently ill-posed in some regimes.
  - Igel et. al. (2022): Dividing the drop spectrum into disjoint “cloud” and “rain” categories makes it difficult to accurately capture activity at the cloud-rain boundary. So two-category schemes are inherently disadvantaged for modelling collision-coalescence.
- Cloud number evaporation is difficult to model accurately.
  - This process is ignored by GCMs, though.
  - Hard because large “drizzle-like” cloud drops can coexist with small cloud drops.
    - ✓ Can be better represented with 3M schemes.
  - Occurs at very short time scales, so infrequent observations don’t capture it well.
- Two-moment rain schemes cause excessive size sorting in sedimentation.
  - Need three rain moments, or at least more complex sedimentation formula.

What's next for BOSS

- Development of "single-category" BOSS (with no cloud/rain distinction)
  - First prototype has been created.
  - ESMD project for EAM implementation.
- Testing/developing BOSS in LES models
  - LES implemented in CM1 for single-category and two-category models.
- GCM adaptation
  - How to deal with
  - I.e., how do we handle GCM spatial/temporal resolutions?
- Adding ice!

Single-category BOSS

- A single category scheme is expected to handle collision-coalescence better.
- Have developed prototype 4-moment scheme ("4M") using extended data set including heavy rain.
  - Also retrained 2M-RD-TE for comparison.
- 4M works better for heavier rain (top), but not a clear improvement when trained on original drizzling cases (bottom).
  - Probably possible to improve 4M further, but may also be some new issues, e.g. due to number evaporation difficulties.

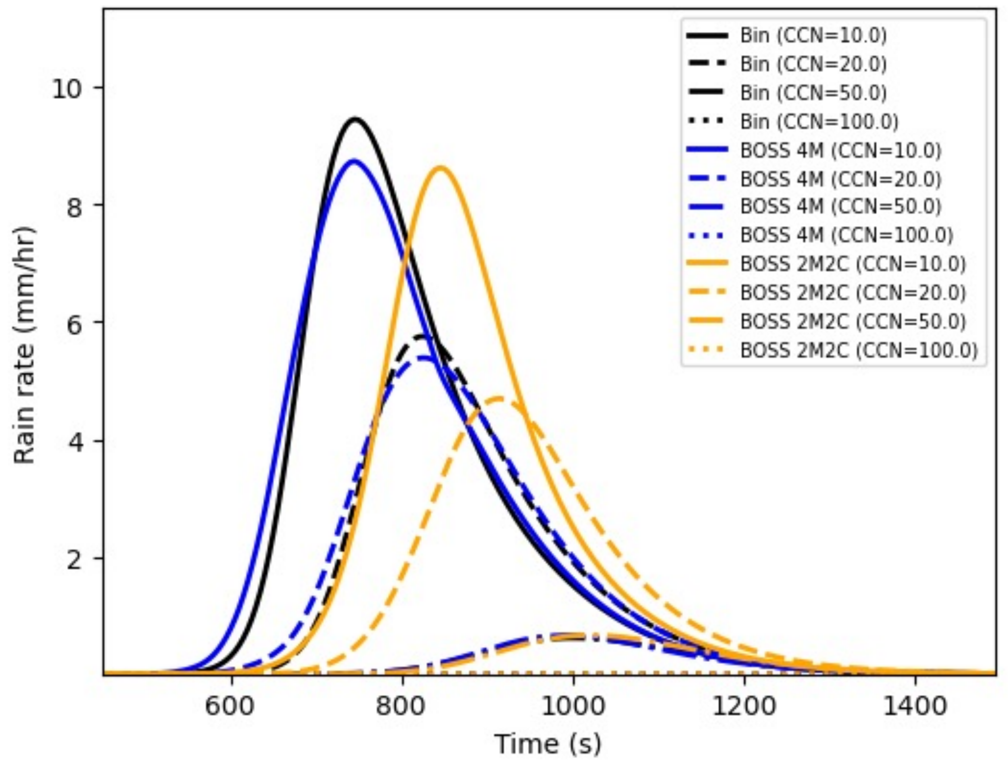

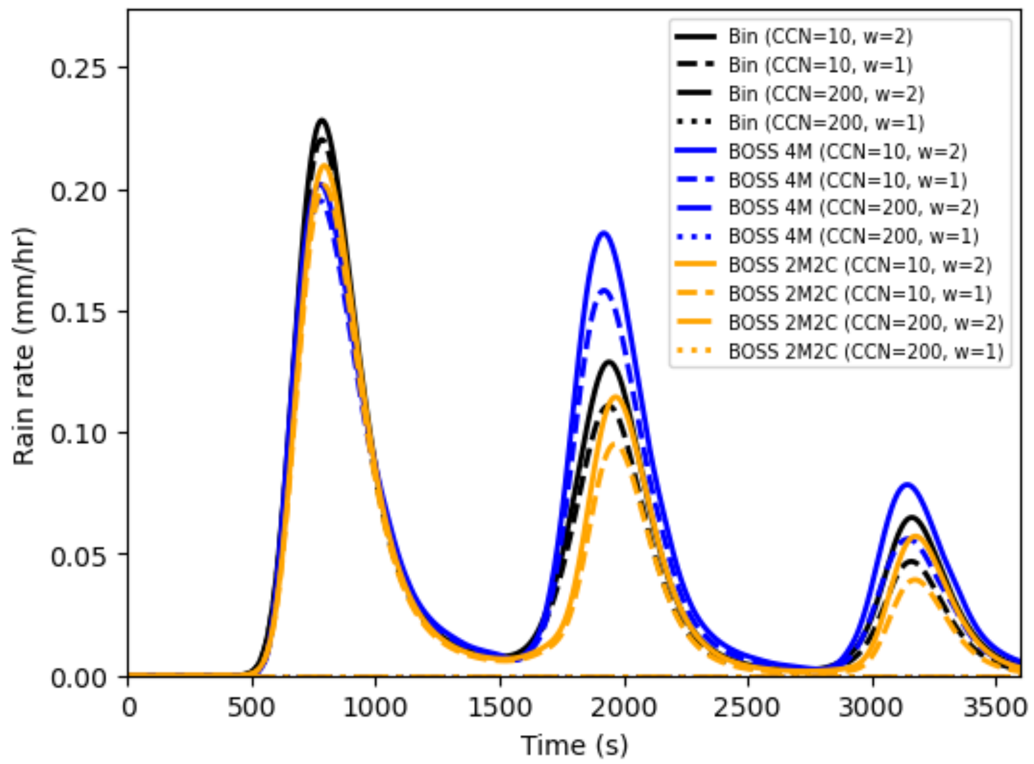

LES tests

- Two-category (two-moment) and single-category BOSS implemented in the CM1 LES model.
  - Rapid recent progress from Kaitlyn Loftus (@ Columbia), Hugh, and Marcus.
- Current ESMD effort to train BOSS in an LES context.
  - LES is too expensive to run MCMC on directly.
  - Plan to produce an emulator using machine learning, as a proxy for the full LES.
- Goal is to produce an emulator in a context with more realistic drop size distribution and spatial variability than the 1-D model.
- LES data sets may also be useful for analyzing sub-grid-scale variability.
  - Necessary to account for this when adapting to climate model resolutions.

GCM integration

- It is not trivial to either increase the number of cloud moments, or move to single-category, in GCMs.

- Main difficulty: other cloud physics parameterizations (e.g. CLUBB) only deal with cloud mass.
  - They don’t predict other moments!

- Open questions: Can we diagnose additional moments in a way that leads to improved accuracy in the microphysics? How?
  - Can we *prognose* more moments in CLUBB? What would that look like?

BOSS in EAM single column model

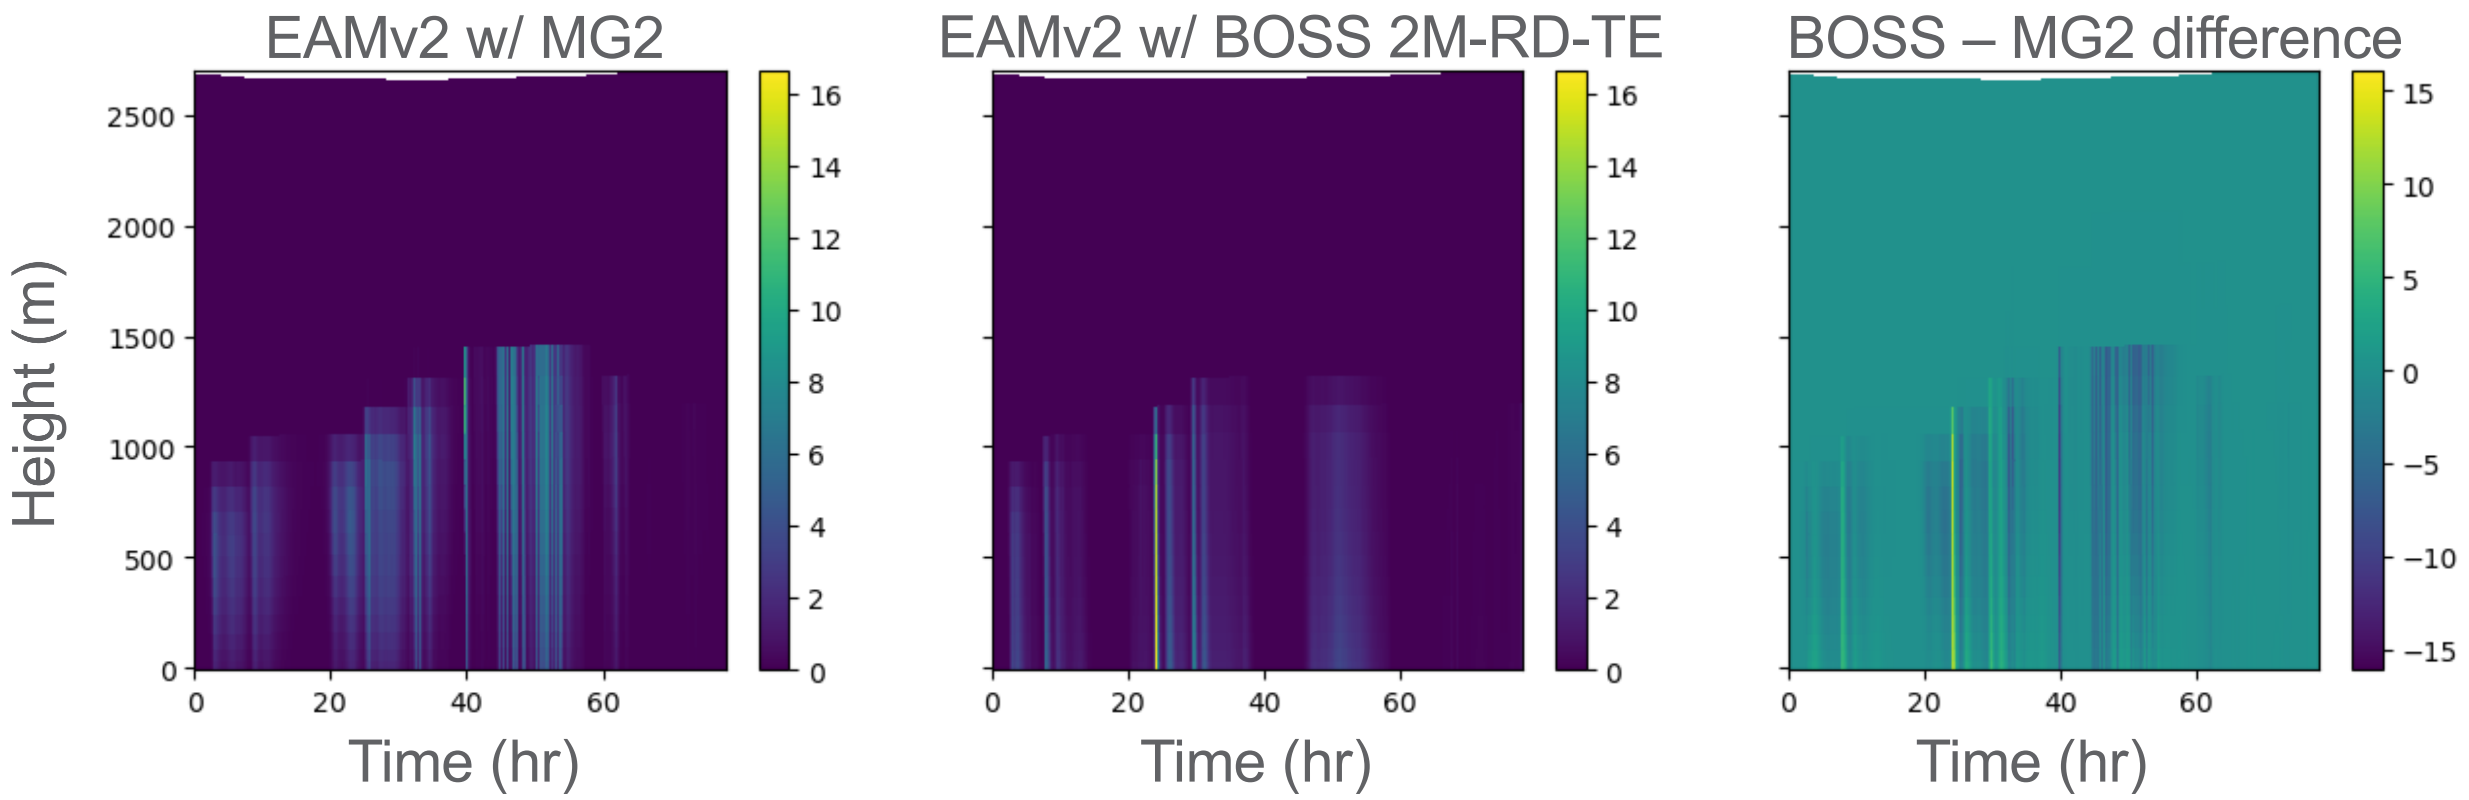

- Comparison of rain mass flux in single column case for ARM MAGIC leg 15A.
  - Using non-standard options: 5s time step, mass gradient precipitation fraction.
- Rain much more intermittent using BOSS than MG2. (Probably good!)
- Rain drops are too big with no virga when using BOSS. (Very bad!)
- BOSS not expected to be realistic without accounting for SGS variability.

Adjusting to longer time steps

- BOSS tests shown so far have used a 5 second time step.
  - On par with the time step used for sedimentation of rain in EAM.
  - Two orders of magnitude smaller than time step used for most other processes!
  - BOSS is probably cheaper than EAM's current microphysics, but 5s time step would still be expensive.
- The time resolution issue is not specific to BOSS.
  - Time integration of MG2 and P3 is not accurate for >1 minute time steps.
- New SciDAC-5 project "PAESCAL" is revising time integration and process coupling in E3SM.
  - Planning new microphysics coupler/integrator in C++.
  - Initial goal is to improve accuracy of time integration for P3.
  - However, software will allow new microphysics schemes to be "swapped in".

Thank you

The bulk of this work was funded by the U.S. Department of Energy under grant DE-SC0021270, and performed at NASA GISS.
Some data analysis and tests with E3SM coupling performed at PNNL as part of the PAESCAL SciDAC-5 project.

Extra slides

Overview

- What is BOSS?
- Theoretical constraints on process rates
  - Scaling symmetry and "normalization"
  - Preserving valid moment combinations
- Results from emulating the TAU bin model:
  - Evaluation contexts: direct vs. time-evolving
  - What information is most useful for predicting autoconversion?
- Further steps:
  - Single-category liquid microphysics
  - BOSS in EAM

Using physical/mathematical constraints

- Power law series are more general than Taylor series, so we can approximate any analytic function (and some discontinuous ones) by using a large enough number of terms.
- However, we prefer to keep things as simple as possible:
  - Fewer power laws is computationally cheaper.
  - Fewer parameters means MCMC performs better.
  - Easier for humans to understand.
- We therefore rely on a few physical/mathematical constraints to specialize the power laws.

Constraint I: Scaling symmetries and normalization

- Take a volume of cloudy air and remove half the particles.
  - The shape of the DSD is the same, but with half the total number.
  - Average particle properties (e.g. fall speeds) will not change.
  - Interactions with environment (e.g. evaporation) will halve (linear scaling).
  - Collisional processes (e.g. accretion) will be reduced by ¼ (quadratic scaling).
- Can use this to express processes as power laws of “normalized” moments.
  - Reduces the number of BOSS parameters.

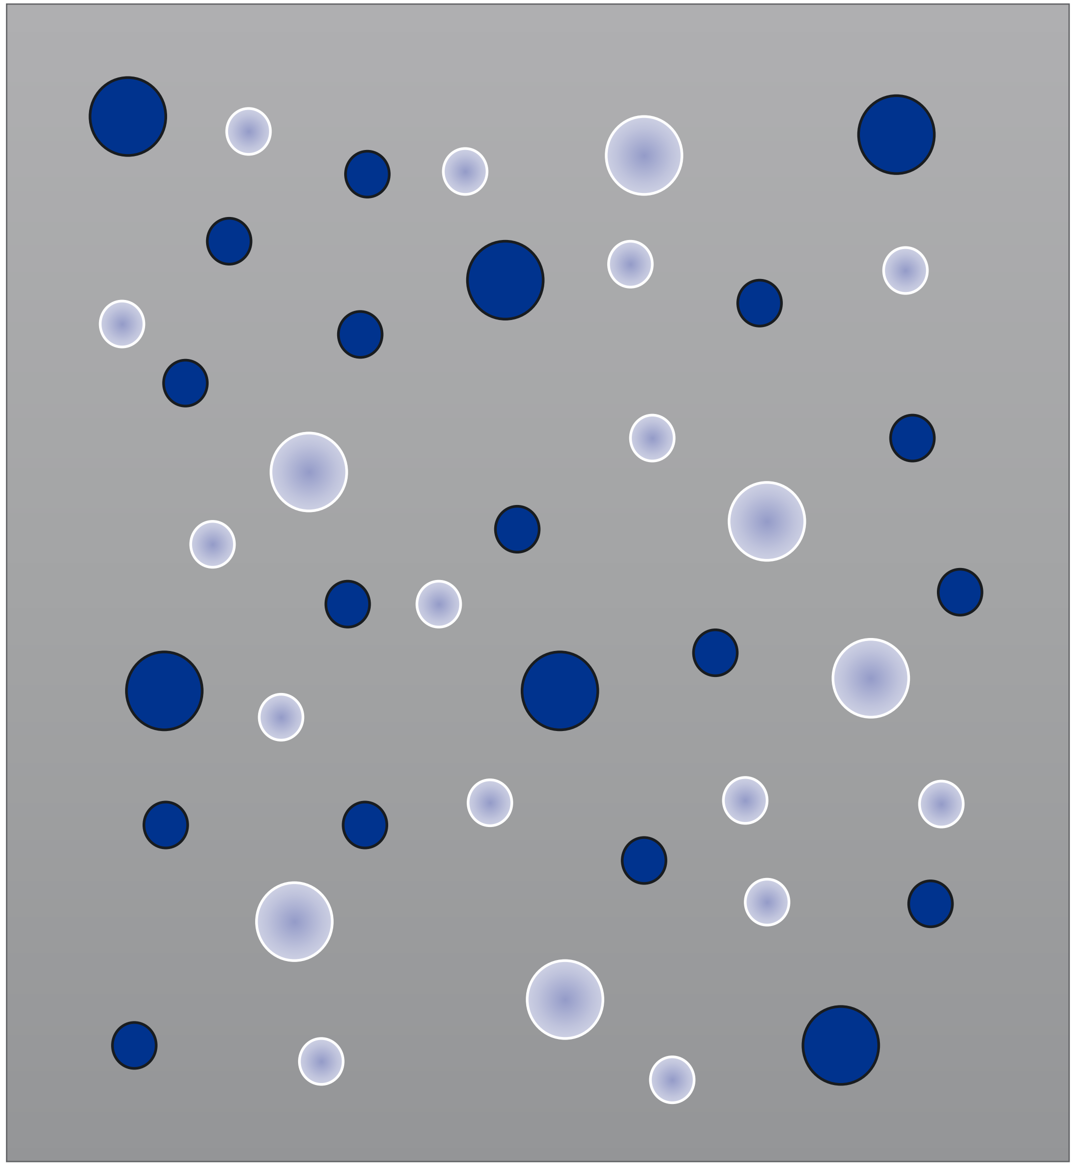

Constraint II: Valid moment combinations

- Since the DSD is positive, its moments cannot take on arbitrary values, but are subject to certain conditions.
- First condition: All moments must be positive ($M_n > 0$).
  - Virtually all microphysics parameterizations obey this for sufficiently small time step.
  - At longer time steps, enforced by conservation limiters.
- Second condition: Monotonicity of moment ratios ($\frac{M_{n+1}}{M_n} > \frac{M_n}{M_{n-1}}$)
  - Only applicable for schemes with three or more moments.
  - Equivalent to requiring measures of drop size to have positive standard deviation.
  - Multi-moment parameterizations sometimes ignore this condition or only enforce it using limiters.
  - For multi-moment BOSS, parameterizations chosen to respect this condition.
- (For a scheme with N_{mom} moments, there are $\left\lfloor \frac{N_{mom}+1}{2} \right\rfloor$ applicable conditions.)

Emulating the TAU bin model: A 1-D kinematic driver

- A simple 1-D kinematic driver produces non-precipitating or drizzling stratocumulus based on input forcing parameters.
  - Vertical velocity prescribed as simple sinusoidal oscillations over time.
  - An constant “latent heat flux” forcing was applied to all levels.
  - Cloud condensation nuclei (CCN) held constant, with Twomey droplet activation.
  - No other parameterizations active (i.e. no radiation, turbulence, etc.).
- Uniform 20-level vertical grid with model top of 2 km (i.e. Δz=100m).
- Eight different runs produced:
  - Simulations run for 1 hour with 20 minute oscillation period.
  - Oscillation amplitude, latent heat flux, and CCN concentration were varied.
- Moments and process rates written out at every vertical level and time step.