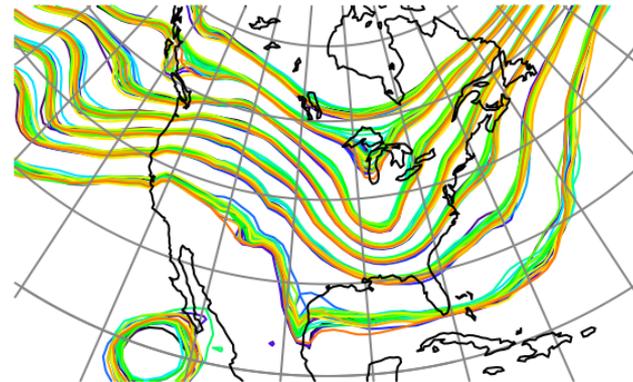


Data
Assimilation
Research
Testbed



DART_LAB Tutorial Section 5: Adaptive Inflation



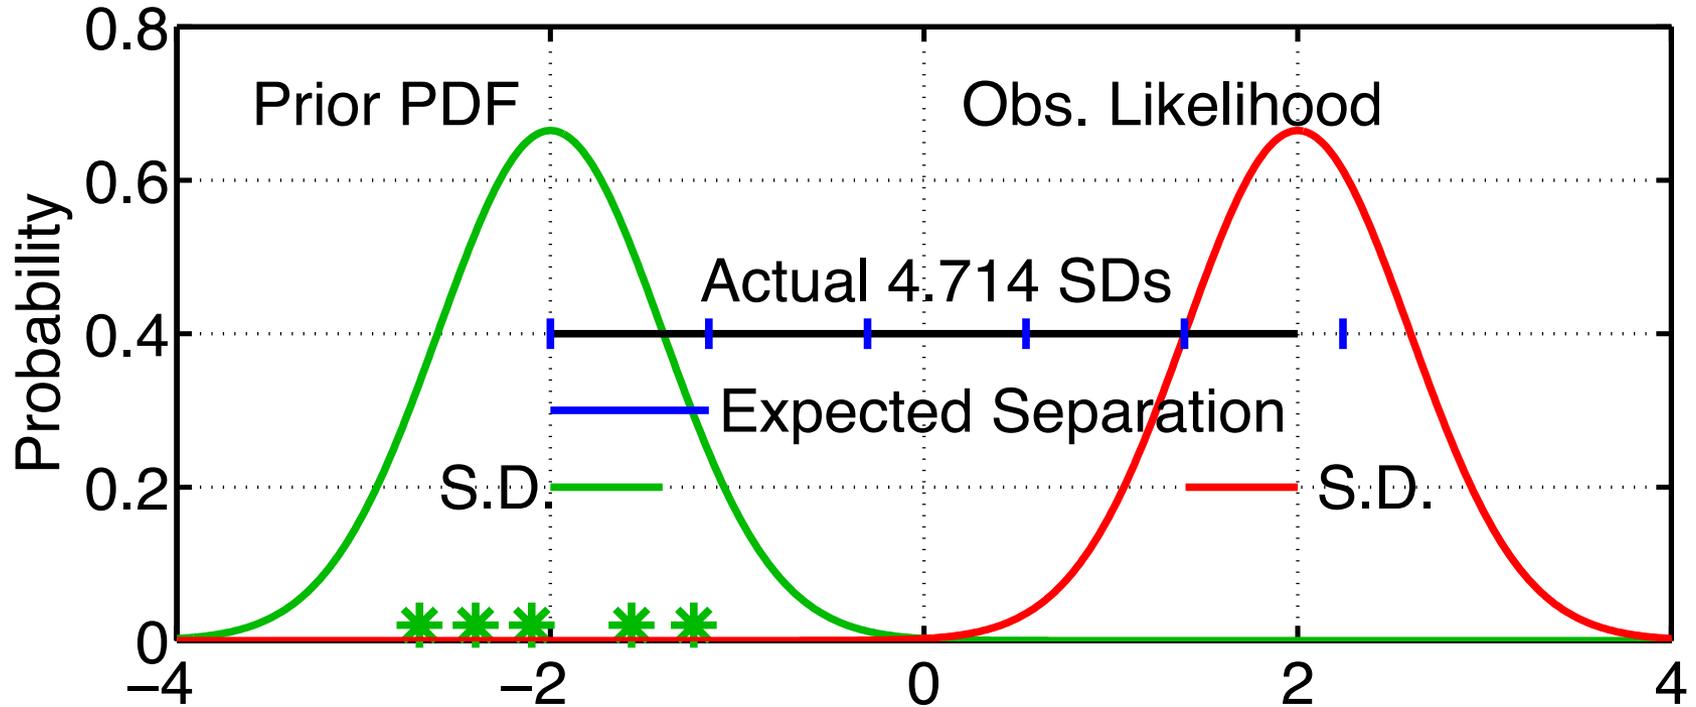
©UCAR 2014



The National Center for Atmospheric Research is sponsored by the National Science Foundation. Any opinions, findings and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

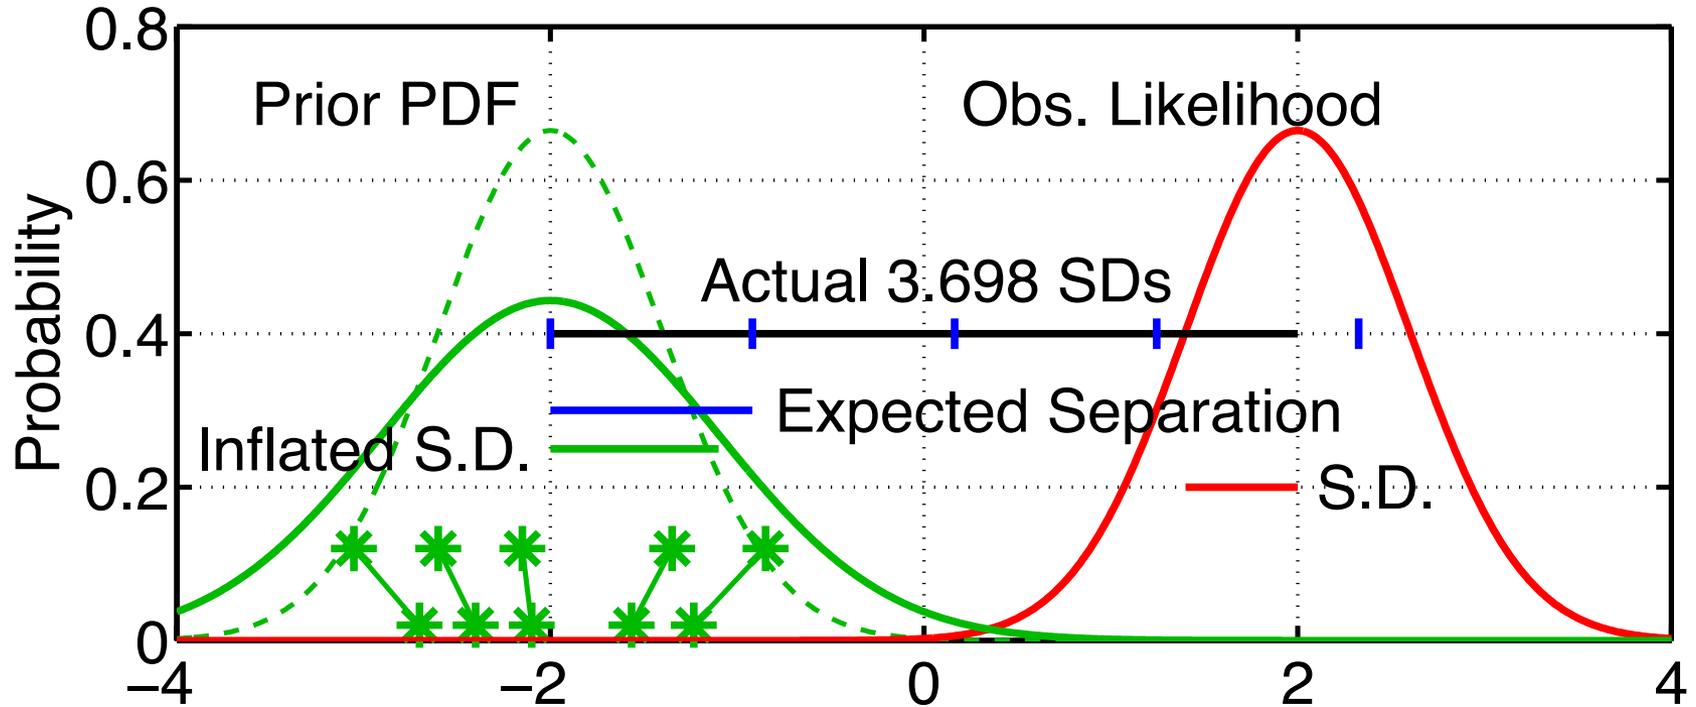
NCAR | National Center for
UCAR | Atmospheric Research

Variance inflation for observations: An adaptive error tolerant filter



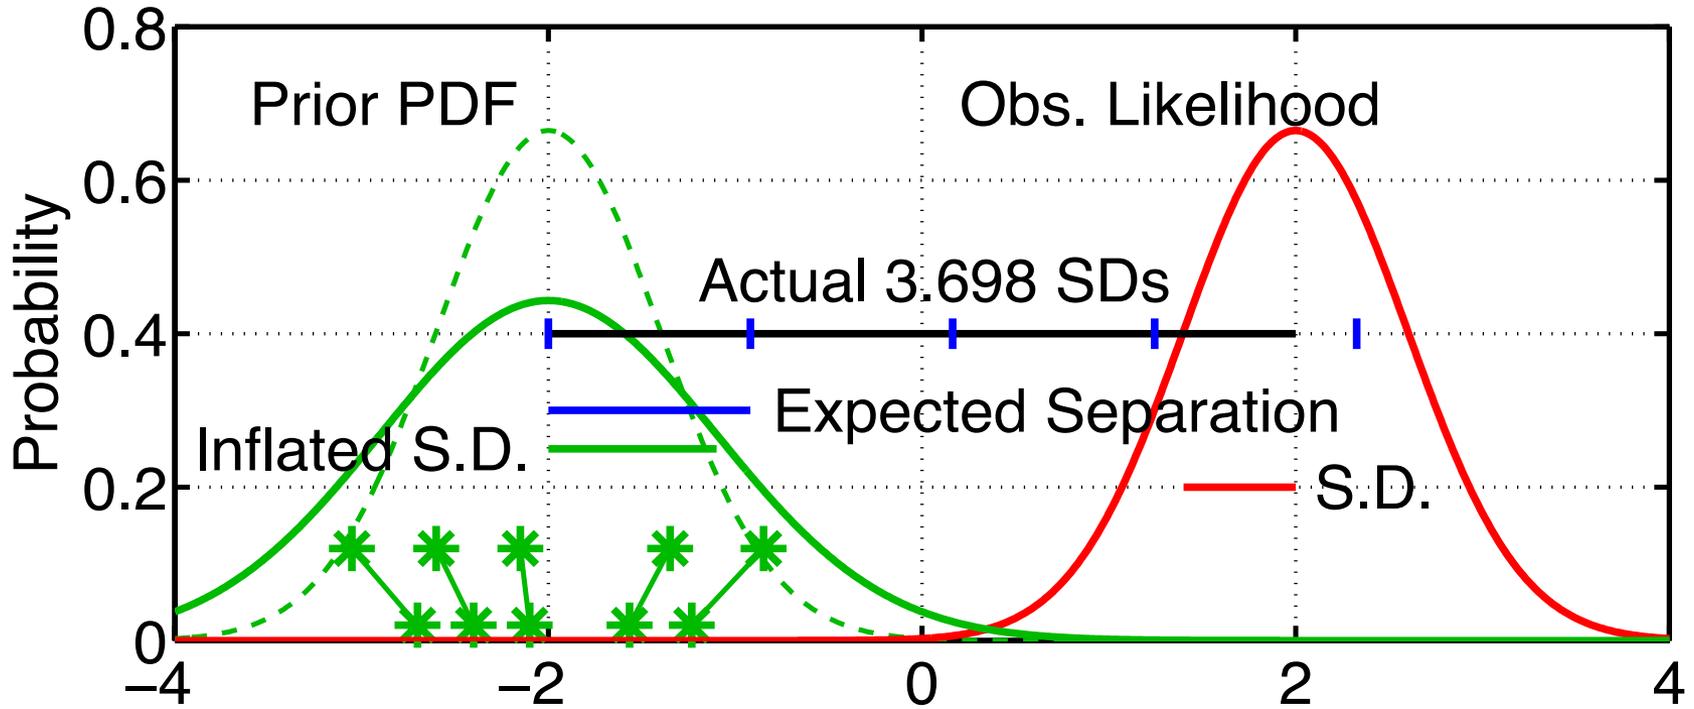
1. For observed variable, have estimate of prior-observed inconsistency.
2. Expected (prior_mean – observation) = $\sqrt{\sigma_{prior}^2 + \sigma_{obs}^2}$
Assumes that prior and observation are supposed to be unbiased.
Is it model error or random chance?

Variance inflation for observations: An adaptive error tolerant filter



1. For observed variable, have estimate of prior-observed inconsistency.
2. Expected (prior_mean – observation) = $\sqrt{\sigma_{prior}^2 + \sigma_{obs}^2}$
3. Inflating increases expected separation
Increases 'apparent' consistency between prior and observation.

Variance inflation for observations: An adaptive error tolerant filter

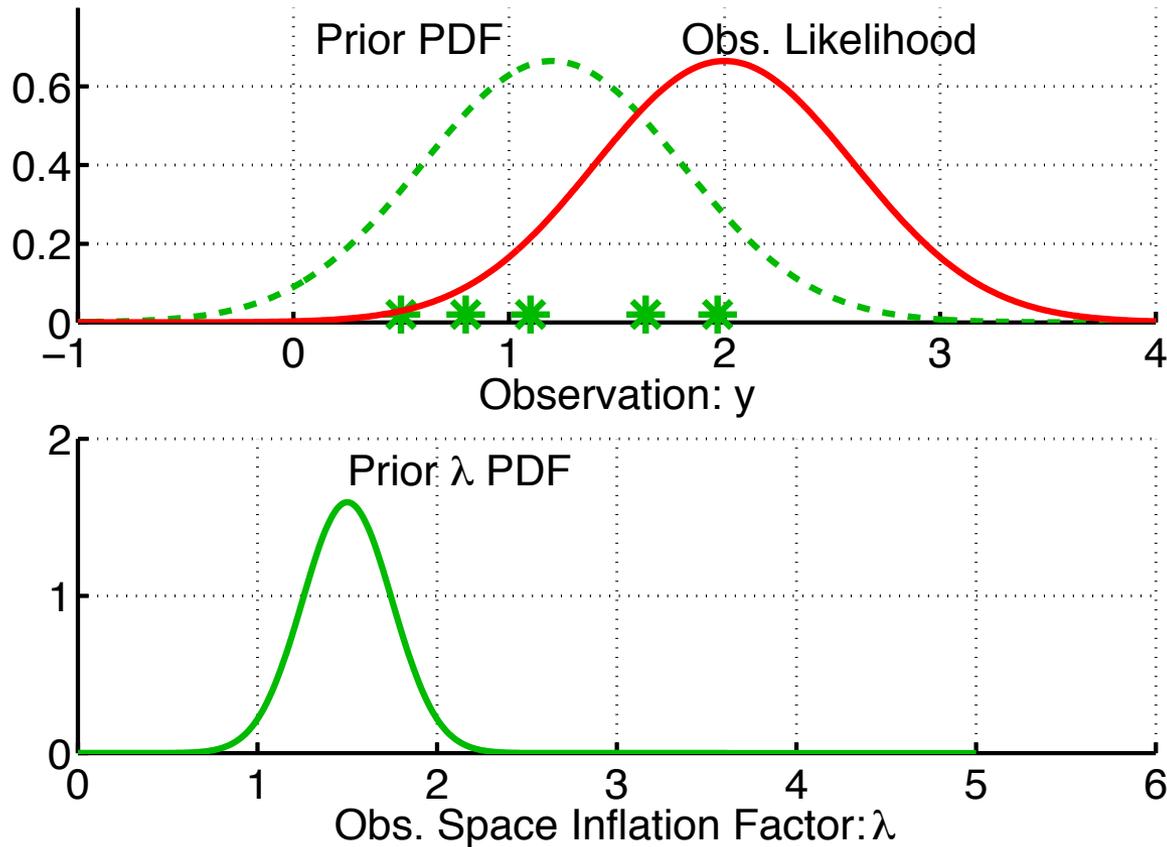


Distance D from prior mean y to obs is $N\left(0, \sqrt{\lambda\sigma_{prior}^2 + \sigma_{obs}^2}\right) = N(0, \theta)$

Prob y_0 is observed given λ : $p(y_o | \lambda) = (2\pi\theta^2)^{-1/2} \exp(-D^2/2\theta^2)$

Variance inflation for observations: An adaptive error tolerant filter

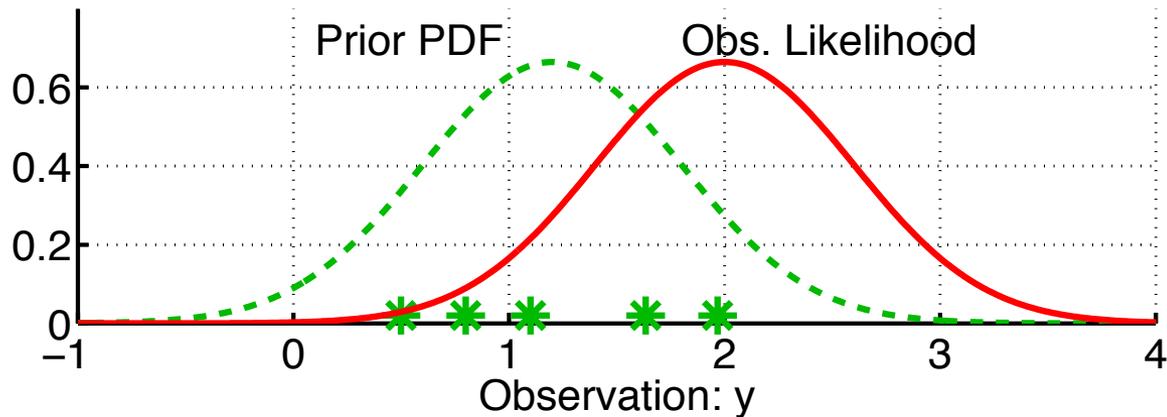
Use Bayesian statistics to get estimate of inflation factor λ .



Assume prior is gaussian: $p(\lambda, t_k | Y_{t_{k-1}}) = N(\bar{\lambda}_p, \sigma_{\lambda, p}^2)$

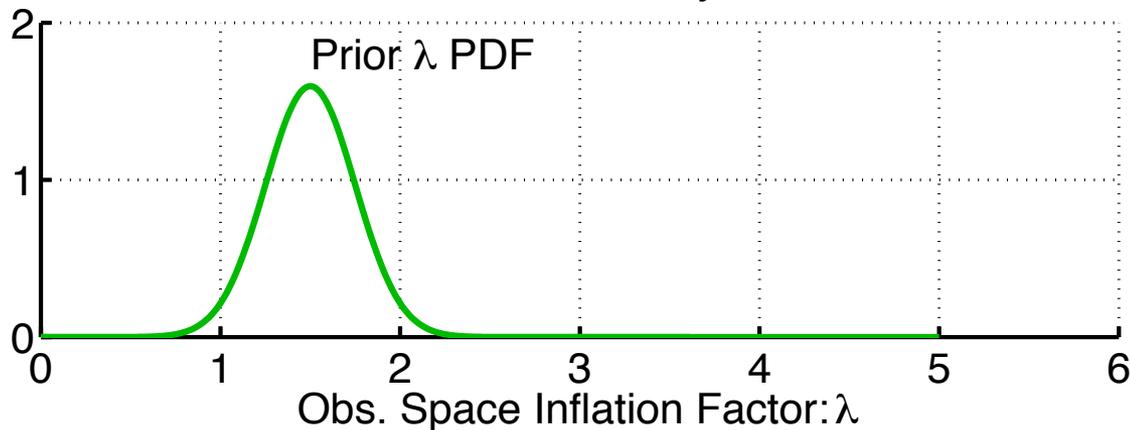
Variance inflation for observations: An adaptive error tolerant filter

Use Bayesian statistics to get estimate of inflation factor λ .



We've assumed a gaussian for prior.

$$p(\lambda, t_k | Y_{t_{k-1}})$$

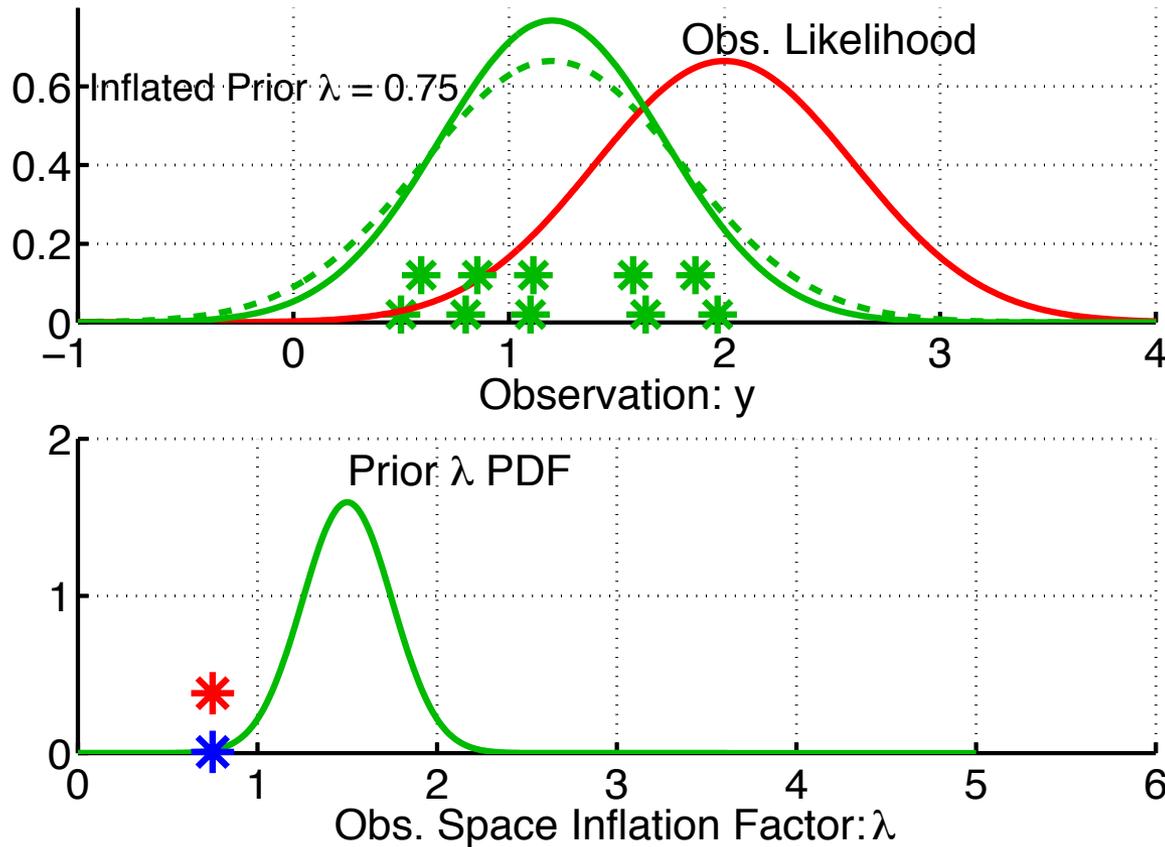


Recall that $p(y_k | \lambda)$ can be evaluated From normal PDF.

$$p(\lambda, t_k | Y_{t_k}) = p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}}) / \textit{normalization}$$

Variance inflation for observations: An adaptive error tolerant filter

Use Bayesian statistics to get estimate of inflation factor λ .



Get $p(y_k | \lambda = 0.75)$
from normal PDF.

Multiply by

$$p(\lambda = 0.75, t_k | Y_{t_{k-1}})$$

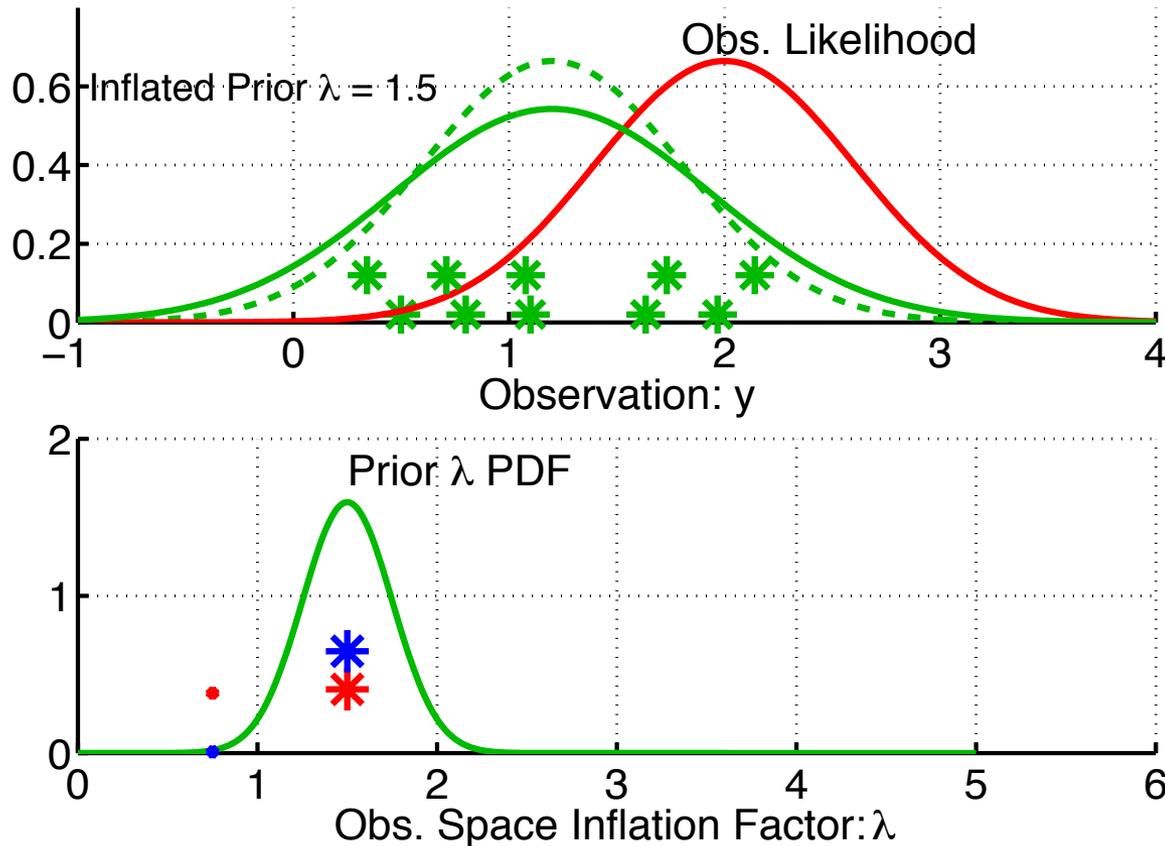
to get

$$p(\lambda = 0.75, t_k | Y_{t_k})$$

$$p(\lambda, t_k | Y_{t_k}) = p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}}) / \textit{normalization}$$

Variance inflation for observations: An adaptive error tolerant filter

Use Bayesian statistics to get estimate of inflation factor λ .



Get $p(y_k | \lambda = 1.50)$
from normal PDF.

Multiply by

$$p(\lambda = 1.50, t_k | Y_{t_{k-1}})$$

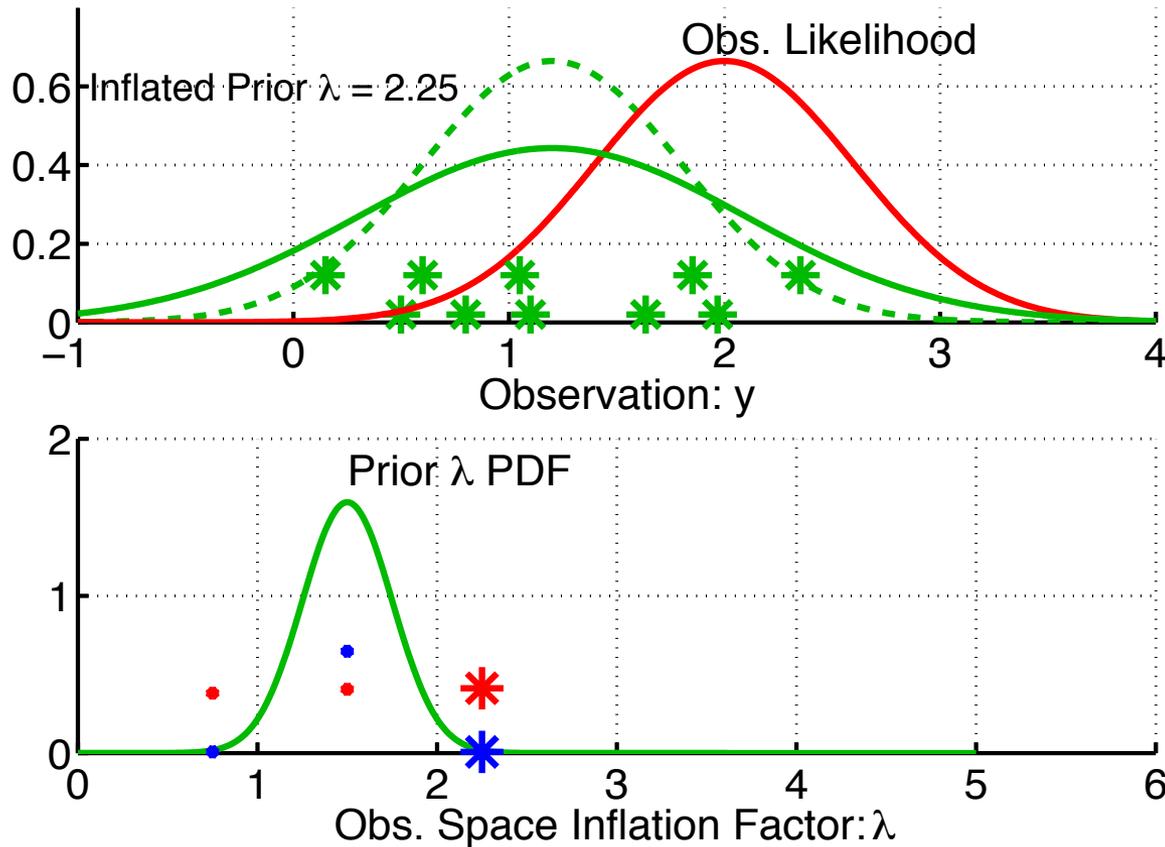
to get

$$p(\lambda = 1.50, t_k | Y_{t_k})$$

$$p(\lambda, t_k | Y_{t_k}) = p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}}) / \textit{normalization}$$

Variance inflation for observations: An adaptive error tolerant filter

Use Bayesian statistics to get estimate of inflation factor λ .



Get $p(y_k | \lambda = 2.25)$
from normal PDF.

Multiply by

$$p(\lambda = 2.25, t_k | Y_{t_{k-1}})$$

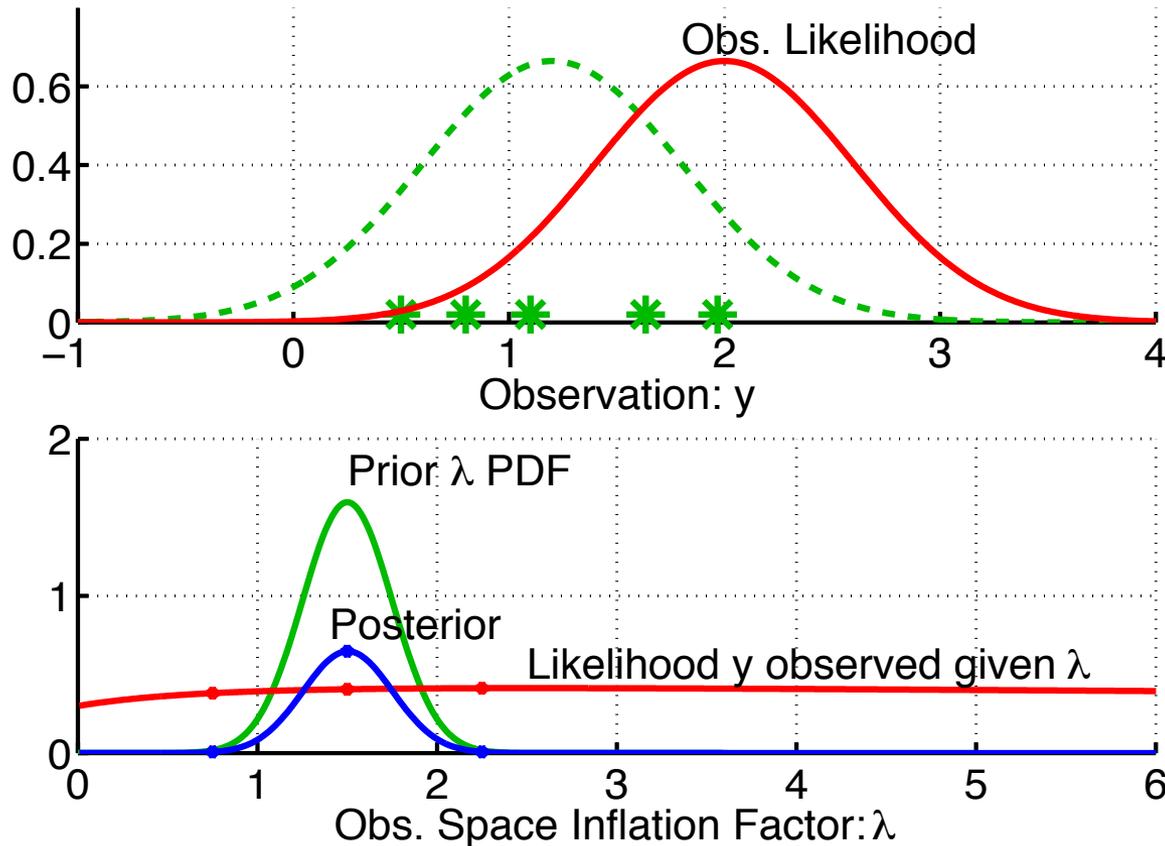
to get

$$p(\lambda = 2.25, t_k | Y_{t_k})$$

$$p(\lambda, t_k | Y_{t_k}) = p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}}) / \text{normalization}$$

Variance inflation for observations: An adaptive error tolerant filter

Use Bayesian statistics to get estimate of inflation factor λ .



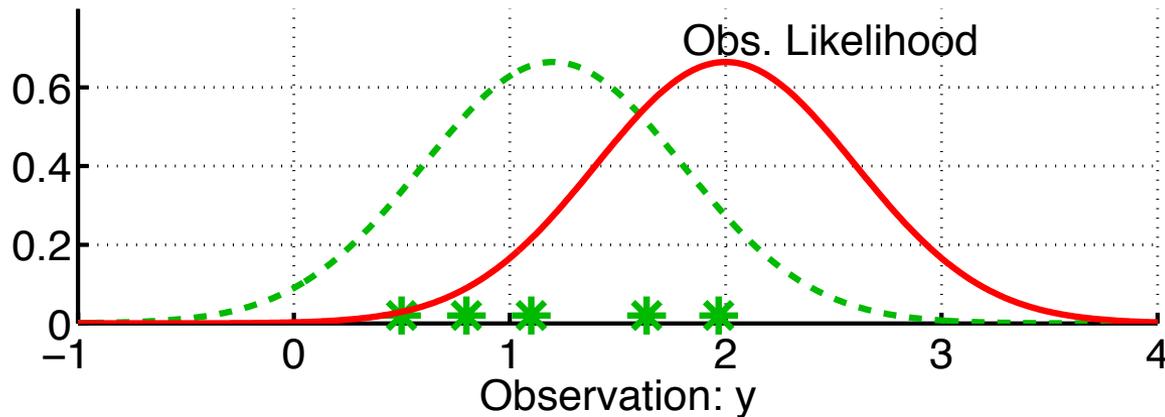
Repeat for a range of values of λ .

Now must get posterior in same form as prior (gaussian).

$$p(\lambda, t_k | Y_{t_k}) = p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}}) / \textit{normalization}$$

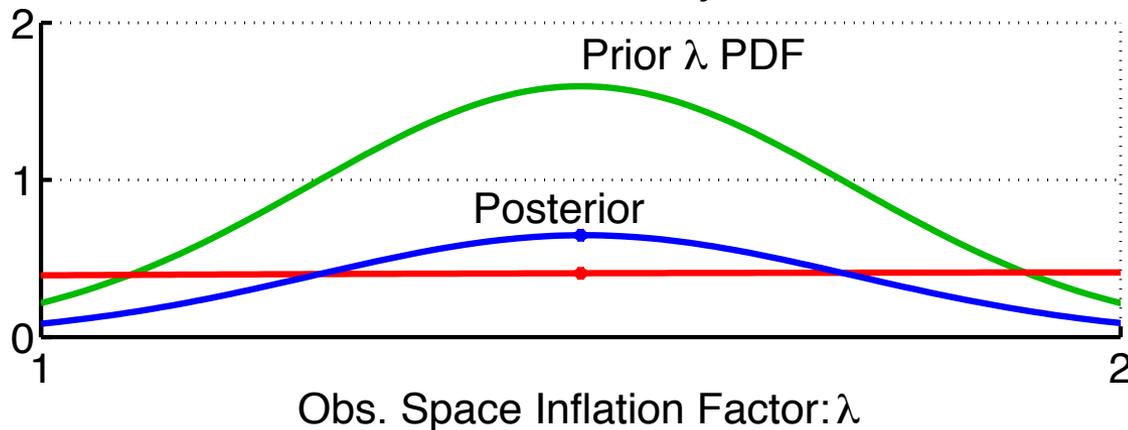
Variance inflation for observations: An adaptive error tolerant filter

Use Bayesian statistics to get estimate of inflation factor λ .



Very little information about λ in a single observation.

Posterior and prior are very similar.

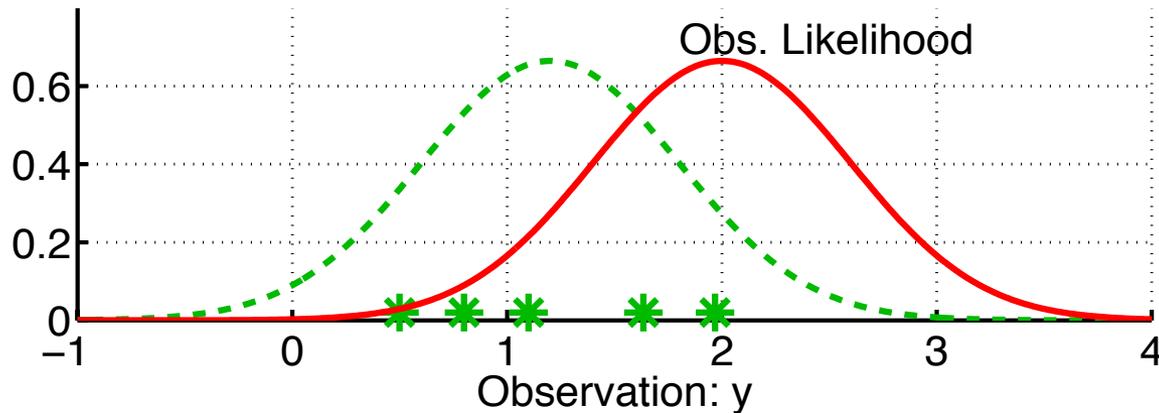


Normalized posterior indistinguishable from prior.

$$p(\lambda, t_k | Y_{t_k}) = p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}}) / \textit{normalization}$$

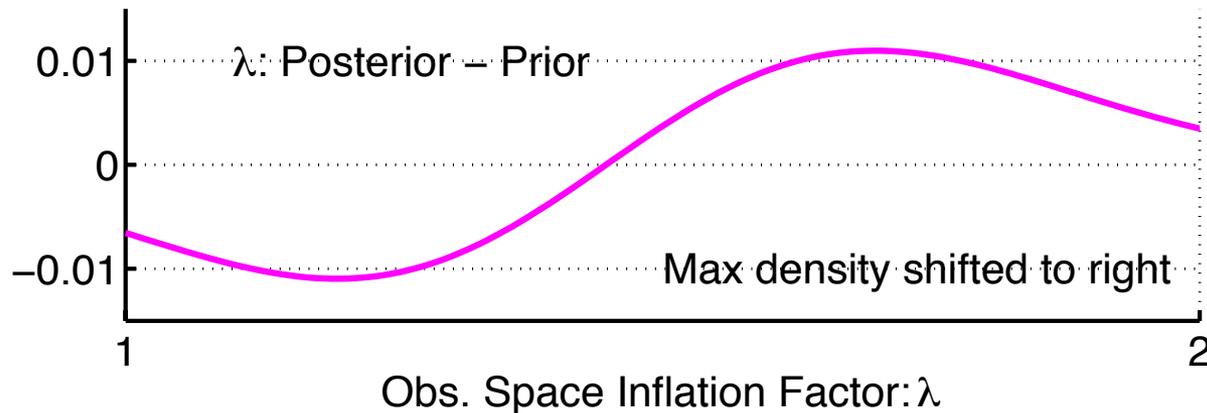
Variance inflation for observations: An adaptive error tolerant filter

Use Bayesian statistics to get estimate of inflation factor λ .



Very little information about λ in a single observation.

Posterior and prior are very similar.

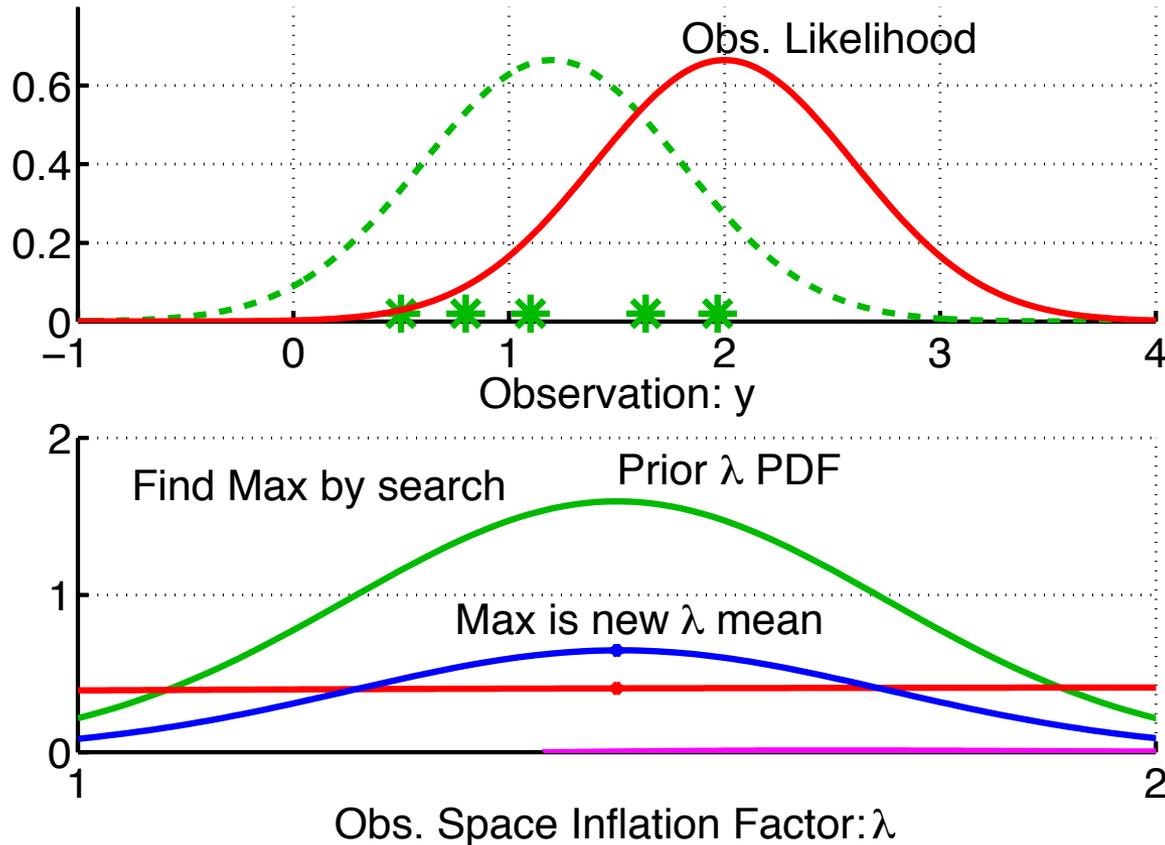


Difference shows slight shift to larger values of λ .

$$p(\lambda, t_k | Y_{t_k}) = p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}}) / \textit{normalization}$$

Variance inflation for observations: An adaptive error tolerant filter

Use Bayesian statistics to get estimate of inflation factor λ .



One option is to use Gaussian prior for λ .

Select max (mode) of posterior as mean of updated Gaussian.

Do a fit for updated standard deviation.

$$p(\lambda, t_k | Y_{t_k}) = p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}}) / \textit{normalization}$$

Variance inflation for observations: An adaptive error tolerant filter

A. Computing updated inflation mean, $\bar{\lambda}_u$.

Mode of $p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}})$ can be found analytically!

Solving $\partial [p(y_k | \lambda) p(\lambda, t_k | Y_{t_{k-1}})] / \partial \lambda = 0$ leads to 6th order poly in λ .

This can be reduced to a cubic equation and solved to give mode.

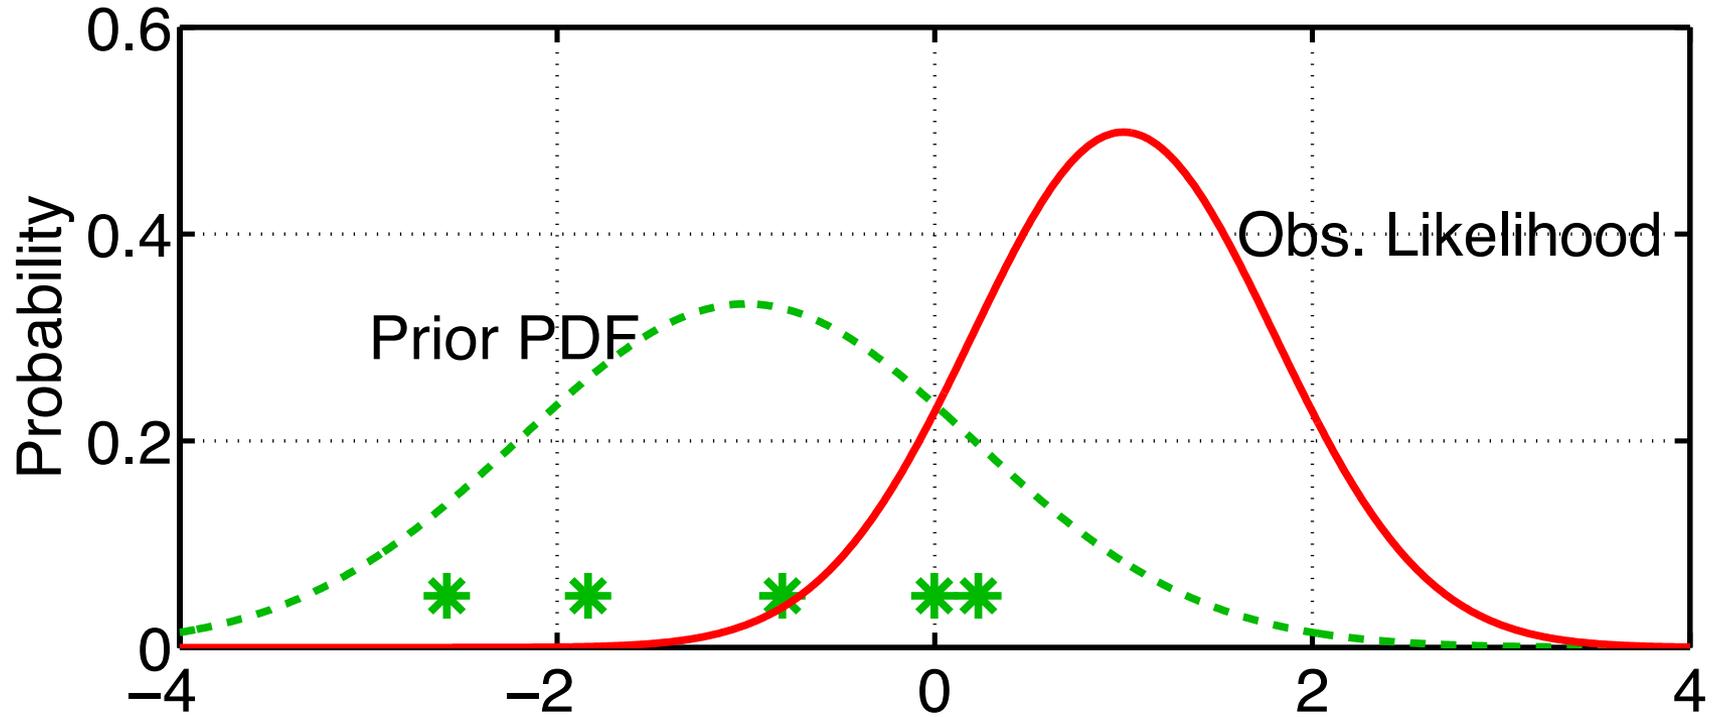
New $\bar{\lambda}_u$ is set to the mode.

This is relatively cheap compared to computing regressions .

B. Computing updated inflation variance, $\sigma_{\lambda,u}^2$.

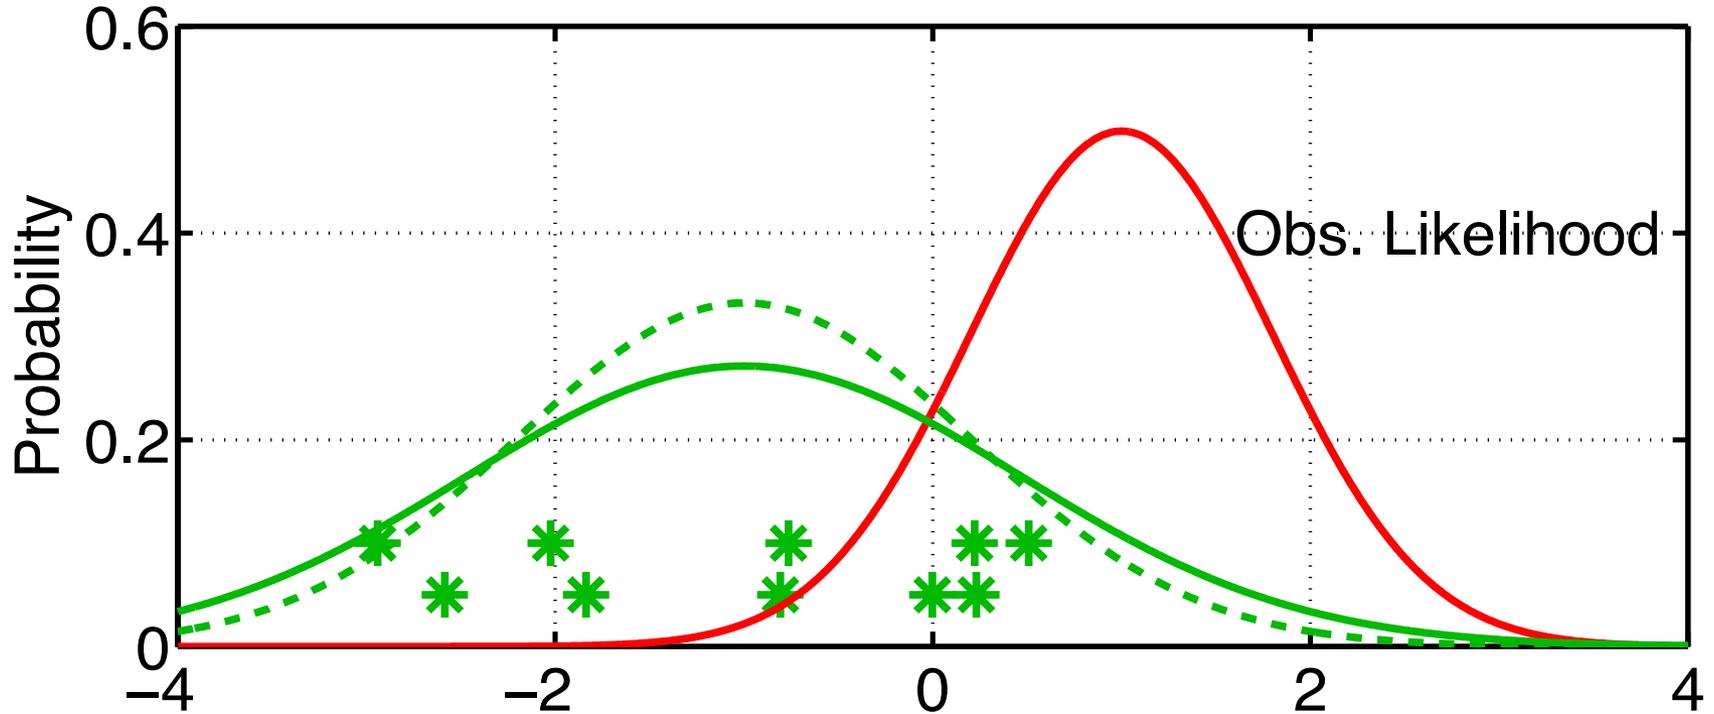
1. Evaluate numerator at mean $\bar{\lambda}_u$ and second point, e.g. $\bar{\lambda}_u + \sigma_{\lambda,p}$
2. Find $\sigma_{\lambda,u}^2$ so $N(\bar{\lambda}_u, \sigma_{\lambda,u}^2)$ goes through $p(\bar{\lambda}_u)$ and $p(\bar{\lambda}_u + \sigma_{\lambda,p})$
3. Compute as $\sigma_{\lambda,u}^2 = -\sigma_{\lambda,p}^2 / 2 \ln r$ where $r = p(\bar{\lambda}_u + \sigma_{\lambda,p}) / p(\bar{\lambda}_u)$

Single Variable Computations with Adaptive Error Correction



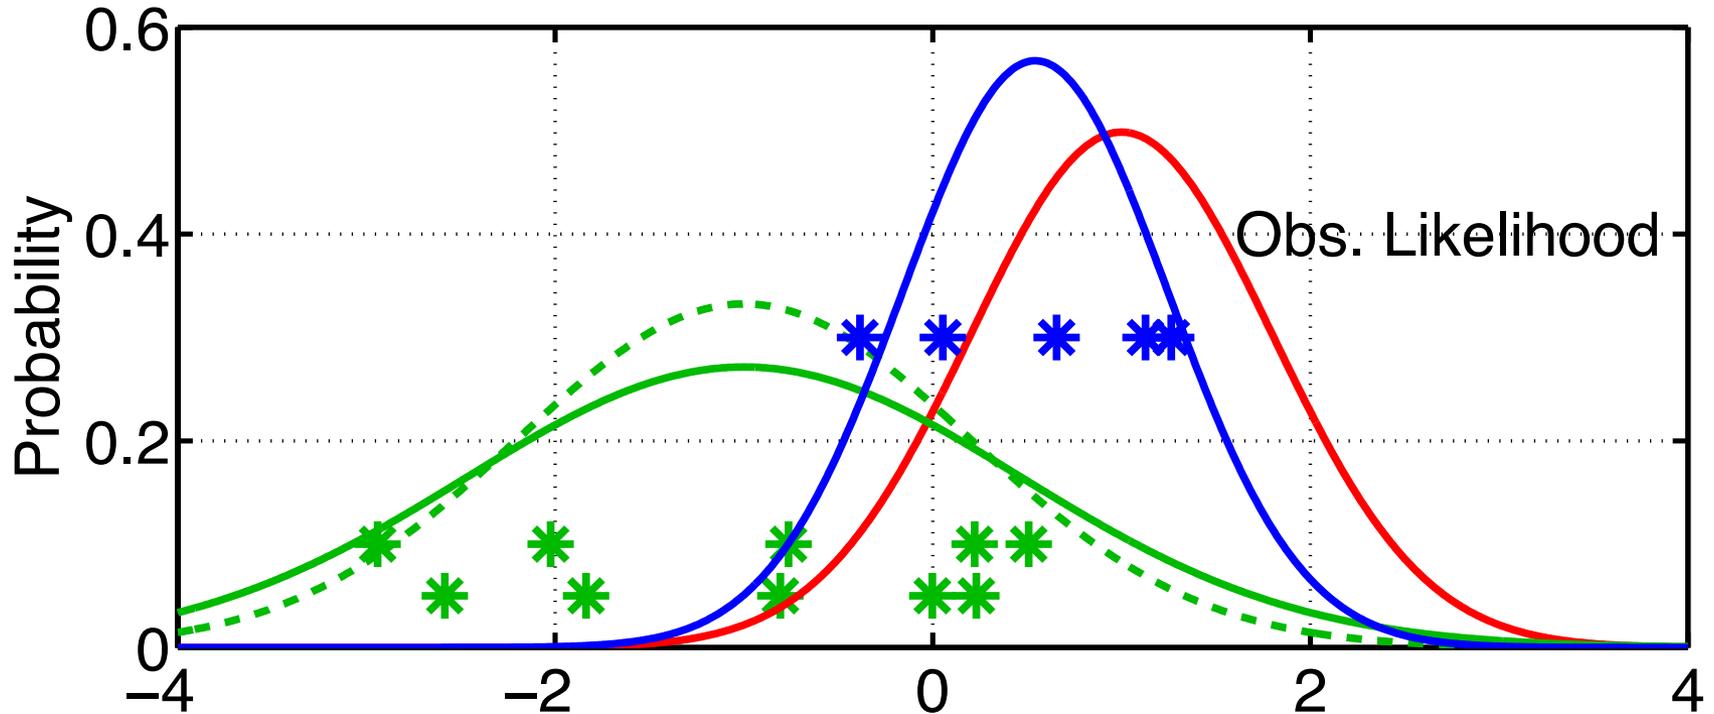
1. Compute updated inflation distribution, $p(\lambda, t_k | Y_{t_k})$.

Single Variable Computations with Adaptive Error Correction



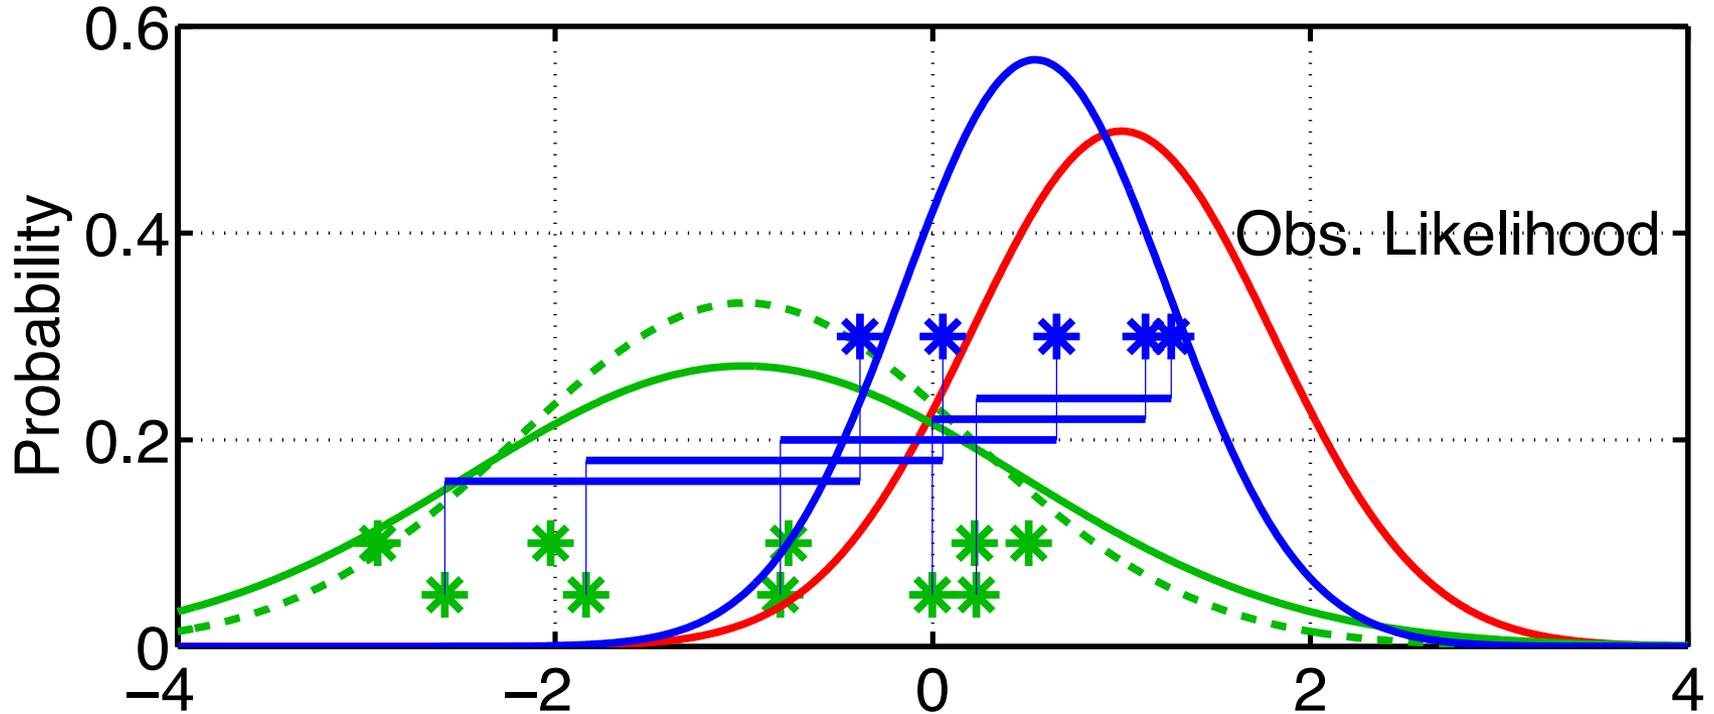
1. Compute updated inflation distribution, $p(\lambda, t_k | Y_{t_k})$.
2. Inflate ensemble using mean of updated λ distribution.

Single Variable Computations with Adaptive Error Correction



1. Compute updated inflation distribution, $p(\lambda, t_k | Y_{t_k})$.
2. Inflate ensemble using mean of updated λ distribution.
3. Compute posterior for y using inflated prior.

Single Variable Computations with Adaptive Error Correction



1. Compute updated inflation distribution, $p(\lambda, t_k | Y_{t_k})$.
2. Inflate ensemble using mean of updated λ distribution.
3. Compute posterior for y using inflated prior.
4. Compute increments from ORIGINAL prior ensemble.

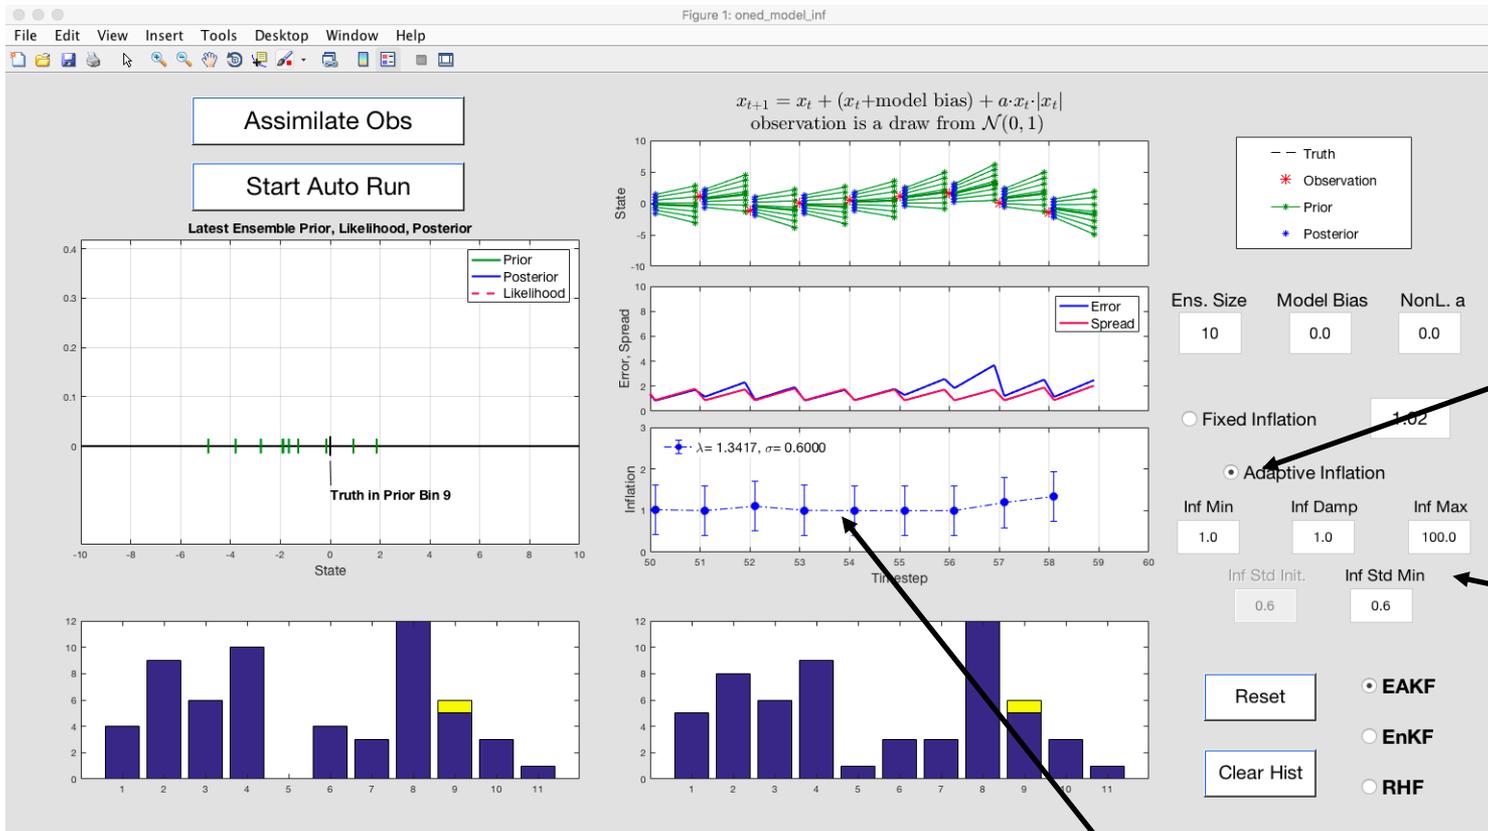
Single Variable Computations with Adaptive Error Correction

Adaptive inflation can be tested with matlab script *oned_model_inf.m*

Can explore 5 different values that control adaptive inflation:

- Minimum value of inflation, often set to 1 (no deflation).
- Inflation damping, more on this later. Value of 1.0 turns it off.
- Maximum value of inflation.
- The initial value of the inflation standard deviation.
- A lower bound on inflation standard deviation (it will asymptote to zero if allowed).

Single Variable Computations with Adaptive Error Correction



Turn on adaptive inflation.

Inflation settings.

Inflation value and its standard deviation.

Single Variable Computations with Adaptive Error Correction

Try adaptive inflation.

Pick a lower value for standard deviation.

Initial lower bound on inflation of 1.0, upper bound large (100).

Try introducing a model bias.

What happens if lower bound is less than 1?

Inflation Damping

Inflation mean damped towards 1 every assimilation time.

- *inf_damping* 0.9: 90% of the inflation difference from 1.0 is retained.

Can be useful in models with heterogeneous observations in time.

For instance, a well-observed hurricane crosses a model domain.

Adaptive inflation increases along hurricane trace.

After hurricane, fewer observations, no longer need so much inflation.

For large earth system models, following values may work:

inf_sd_initial = 0.6,

inf_damping = 0.9,

inf_sd_lower_bound = 0.6.

Adaptive Multivariate Inflation Algorithm

Suppose we want a global multivariate inflation, λ_s , instead.

Make same least squares assumption that is used in ensemble filter.

Inflation of λ_s for state variables inflates obs. priors by same amount.

Get same likelihood as before: $p(y_o | \lambda) = (2\pi\theta^2)^{-1/2} \exp(-D^2/2\theta^2)$

$$\theta = \sqrt{\lambda_s \sigma_{prior}^2 + \sigma_{obs}^2}$$

Compute updated distribution for λ_s exactly as for single observed variable.

Implementation of Adaptive Multivariate Inflation Algorithm

1. Apply inflation to state variables with mean of λ_s distribution.
2. Do following for observations at given time sequentially:
 - a. Compute forward operator to get prior ensemble.
 - b. Compute updated estimate for $\hat{\lambda}_s$ mean and variance.
 - c. Compute increments for prior ensemble.
 - d. Regress increments onto state variables.

Spatially varying adaptive inflation algorithm

Have a distribution for λ for each state variable, $\lambda_{s,i}$.

Use prior correlation from ensemble to determine impact of $\lambda_{s,i}$ on prior variance for given observation.

If γ is correlation between state variable i and observation then

$$\theta = \sqrt{\left[1 + \gamma \left(\sqrt{\lambda_{s,i}} - 1\right)\right]^2 \sigma_{prior}^2 + \sigma_{obs}^2}$$

Equation for finding mode of posterior is now full 12th order:

Analytic solution appears unlikely.

Can do Taylor expansion of θ around $\lambda_{s,i}$.

Retaining linear term is normally quite accurate.

There is an analytic solution to find mode of product in this case!

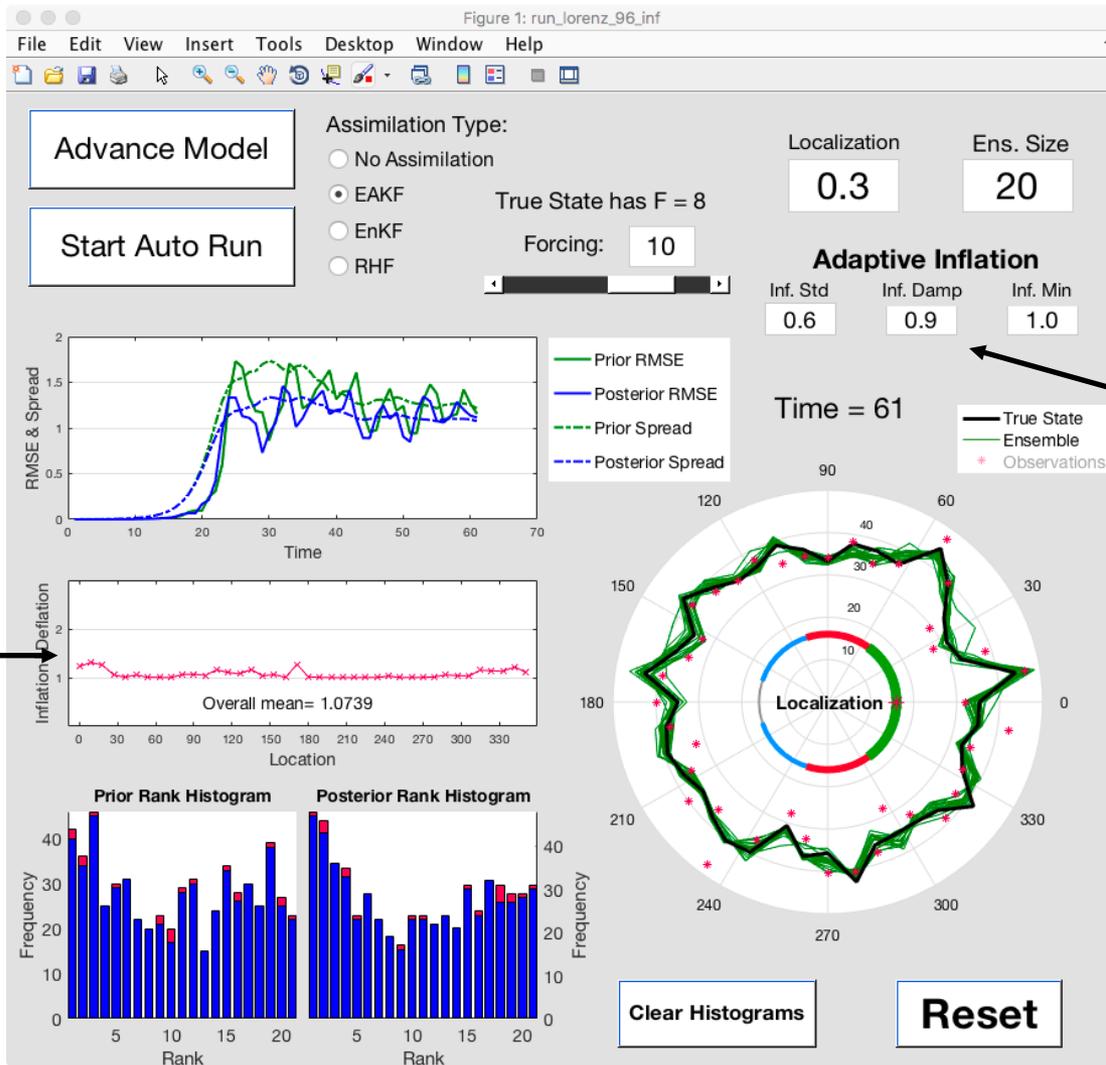
Spatially Varying Adaptive Inflation

Spatially Varying Adaptive inflation can be tested with matlab script *run_lorenz96_inf.m*

Can explore 3 different values that control adaptive inflation:

- Minimum value of inflation, often set to 1 (no deflation).
- Inflation damping. Value of 1.0 turns it off.
- The value of the inflation standard deviation.
 - Lower bound on standard deviation is set to same value.
 - In this case, standard deviation just stays fixed at selected value.

Spatially Varying Adaptive Inflation



Adaptive inflation spatial plot.

Adaptive inflation controls.

Spatially Varying Adaptive Inflation

Explore some of the following:

How does adaptive inflation change as localization is changed?

How does adaptive inflation change for different values of inflation standard deviation?

If the lower bound is smaller than 1, does deflation (inflation < 1) happen?

Simulating Model Error in 40-Variable Lorenz-96 Model

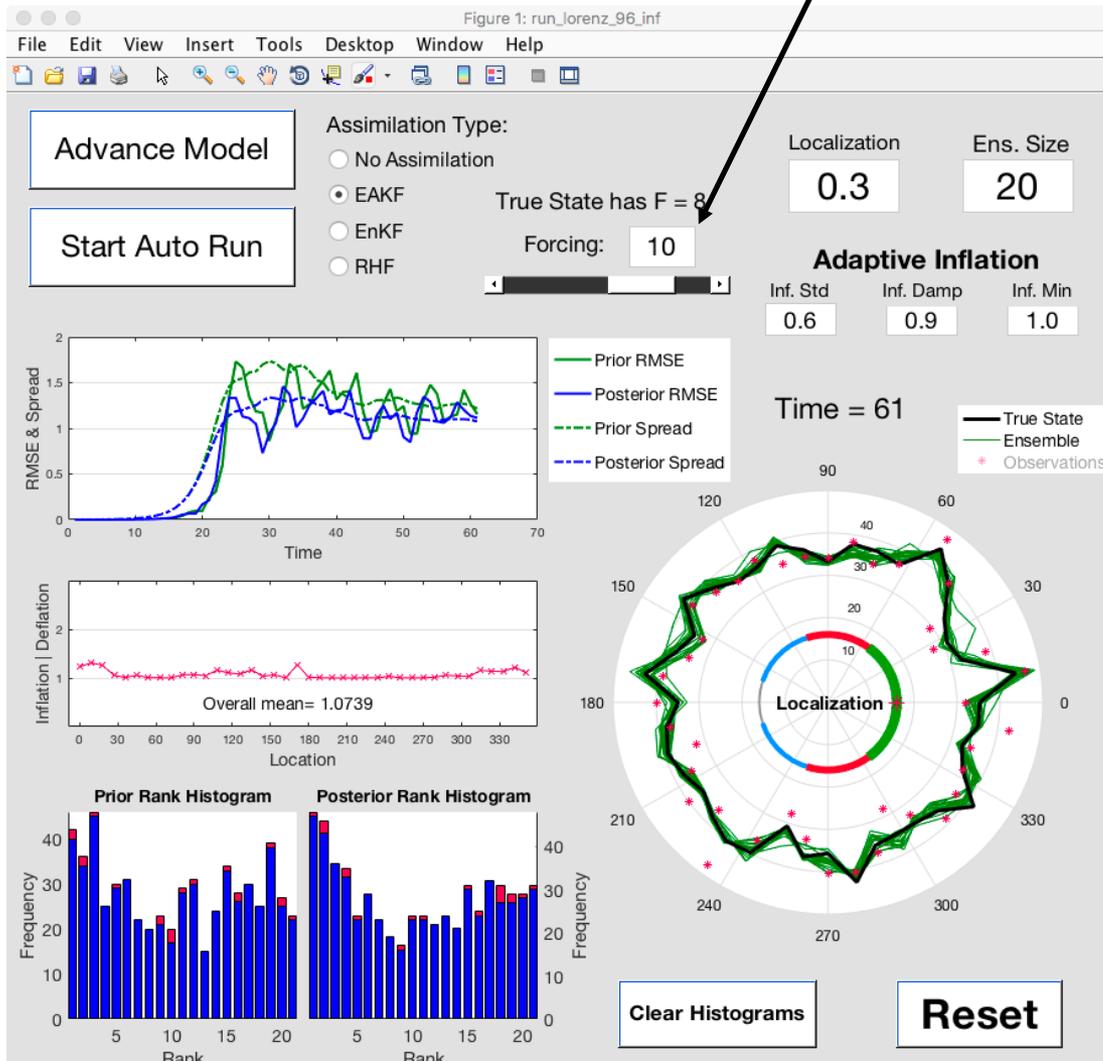
Inflation can deal with all sorts of errors, including model error.

Can simulate model error in Lorenz96 by changing forcing.
Synthetic observations are from model with forcing = 8.0.

Both *run_lorenz_96* and *run_lorenz_96_inf* allow model error.

Spatially Varying Adaptive Inflation

Model error: Change forcing for assimilating model here.



Spatially Varying Adaptive Inflation with Model Error

Explore some of the following:

Change the model forcing to a larger or smaller value (say 6 or 10).

How does adaptive inflation respond to model error?

Do good values of localization change as model error increases?

