Observation-Dependent Posterior Inflation for the Ensemble Kalman Filter

DANIEL HODYSS AND WILLIAM F. CAMPBELL

Naval Research Laboratory, Monterey, California

JEFFREY S. WHITAKER

NOAA/Earth System Research Laboratory/Physical Sciences Division, Boulder, Colorado

(Manuscript received 21 September 2015, in final form 28 March 2016)

ABSTRACT

Ensemble-based Kalman filter (EBKF) algorithms are known to produce posterior ensembles whose variance is incorrect for a variety of reasons (e.g., nonlinearity and sampling error). It is shown here that the presence of sampling error implies that the true posterior error variance is a function of the latest observation, as opposed to the standard EBKF, whose posterior variance is independent of observations. In addition, it is shown that the traditional ensemble validation tool known as the ''binned spread-skill'' diagram does not correctly identify this issue in the ensemble generation step of the EBKF, leading to an overly optimistic impression of the relationship between posterior variance and squared error. An updated ensemble validation tool is described that reveals the incorrect relationship between mean squared error (MSE) and ensemble variance, and gives an unbiased evaluation of the posterior variances from EBKF algorithms. Last, a new inflation method is derived that accounts for sampling error and correctly yields posterior error variances that depend on the latest observation. The new method has very little computational overhead, does not require access to the observations, and is simple to use in any serial or global EBKF.

1. Introduction

Inflation of the prior and/or posterior ensemble variances in ensemble-based Kalman filtering (EBKF) is commonly used to help prevent filter divergence (e.g., Anderson and Anderson 1999; Zhang et al. 2004; Furrer and Bengtsson 2007; Sacher and Bartello 2008; Anderson 2007, 2009; Whitaker and Hamill 2012; Menetrier and Auligne 2015). The basic idea is that a small ensemble size leads to a biased estimate of the posterior ensemble variance [Furrer and Bengtsson (2007); Sacher and Bartello (2008); also section 3 of this paper will discuss this in detail]. This bias in the variance manifests as an ensemble whose variance is, on average, smaller than the mean squared error (MSE) of the ensemble mean. Hence, ''inflating'' the ensemble variance (by multiplying it by a number greater than 1) offers a crude way to reduce this bias in the variance and form a more accurate ensemble; one whose ensemble variance is closer to being equal to the MSE of the ensemble mean.

Because the amount by which the ensemble variance is biased low is state and regionally dependent significant research has been performed to determine adaptive methods that will appropriately inflate when and where this bias is greatest. We will gather this previous work into two groups: 1) inflation methods aimed at getting the variance right on average (e.g., Anderson and Anderson 1999; Zhang et al. 2004; Furrer and Bengtsson 2007; Sacher and Bartello 2008; Whitaker and Hamill 2012; Menetrier and Auligne 2015) and 2) inflation methods aimed at getting the variance consistent with a Bayesian posterior (e.g., Anderson 2007, 2009). While the distinction between these two views of the literature is subtle, it is nevertheless important to realize that they are distinctly different. A careful discussion of why these two groups are different will be presented in section 2. Here, we will simply state that we will be deriving a method that fits into the second group by showing that a Bayesian posterior has an observation-dependent posterior variance but that the EBKF ensemble variance does not have an observation dependence. This we will

Corresponding author address: Dr. Daniel Hodyss, Naval Research Laboratory, Marine Meteorology Division, 7 Grace Hopper Ave., Stop 2, Monterey, CA 93943.
E-mail: daniel.hodyss@nrlmry.navy.mil

DOI: 10.1175/MWR-D-15-0329.1

show immediately implies a posterior inflation algorithm. While other observation-dependent inflation algorithms have been proposed (e.g., Anderson 2007, 2009), this manuscript will illustrate a new technique as well as clearly frame why those past methods were successful. We show specifically how sampling error and non-Gaussianity affect the data assimilation problem from the perspective of the Bayesian posterior and an EBKF, and then illustrate why including information from observations proves successful.

This manuscript is organized as follows. In section 2 we describe our approach to understand the issues with the EBKF that require inflation and illustrate a new issue that appears when comparing the EBKF to a Bayesian posterior. In section 3 we discuss in detail how a sampling error from a limited ensemble size implies that the posterior error variance should be a function of the observation. In section 4 we illustrate the posterior inflation methods discussed in section 3 as applied to the three-variable Lorenz equations and a two-level primitive equation model of the atmosphere. We have strived to render section 4 self-contained such that readers not interested in the detailed derivations of section 3 can apply the theory through an examination only of section 4. Section 5 closes the manuscript with a recapitulation of the most important results and a discussion of the application of these ideas to common EBKF algorithms.

2. Preliminaries and approach

a. Theory

Suppose we have an observing system simulation experiment (OSSE), and conduct a nature run for some length of time, generating $N_{\text{total}} = M \times N$ true states x_t, where M and N will be defined shortly. Assume that x_t is a scalar, and the state variable is observed directly at every analysis time. We then run an N_e-member EBKF, generating N_{total} ensemble mean analyses each with an associated N_e-member sample ensemble variance P_a. Construct the binned spread-skill diagram by partitioning the N_e-member sample ensemble variances P_a at analysis time into N bins, sorting from smallest P_a to greatest. In each of these bins there is an analysis corresponding to each of the M ensemble variances in that bin. Calculate the MSE for the ith bin using

$$S^i = \frac{1}{M}\sum_{j=1}^{M} e_j^2, \tag{2.1}$$

where $e_j = x_t - \overline{x}_a$. In practice the errors e_j would be calculated against an unassimilated set of observations and in this case an accounting for measurement errors must be made. In the limit as the number of bins and the number of samples in each bin approach infinity, the MSE is

$$S(P_a) = \int_{-\infty}^{\infty} e^2 \rho(e \mid P_a)\, de, \tag{2.2}$$

where $\rho(e \mid P_a)$ denotes the probability density function (pdf) of the error conditioned on the posterior ensemble variance. By design, the binned spread-skill diagram (e.g., Whitaker and Loughe 1998; Wang and Bishop 2003) obtains the MSE (skill) with respect to a particular ensemble variance (spread). Therefore, the EBKF user is content when the MSE is a linear function of P_a [i.e., (2.2) is very close to a straight line with a slope of 1].

Note that we could have chosen to bin on something other than ensemble variance. This is useful because if we bin on a variable that is not ensemble variance we may see how the MSE varies with respect to changes in that variable. One particularly illuminating way to see how MSE varies is to bin on innovations, $v_i = y - H\overline{x}_f$, where $\overline{x}_f$ is the prior mean and H is the observation operator, such that we again obtain N bins from the sorted innovations with the quantity in the ith bin being estimated as

$$P_v^i = \frac{1}{M}\sum_{j=1}^{M} P_a^j, \tag{2.3}$$

which after applying the same limiting procedure above obtains

$$P_v(v) = \int_{-\infty}^{\infty} P_a \rho(P_a \mid v)\, dP_a. \tag{2.4}$$

In contrast to (2.1), the observations in v_i are not an independent set for verification purposes, but the actual assimilated observations. Obviously, if we bin in this way there can no longer be any expectation that the result $P_v(v)$ (i.e., binned ensemble variances) should follow a line with slope one. However, we know that (2.4) is still supposed to be a prediction of the MSE in each corresponding bin, that is,

$$S(v) = \int_{-\infty}^{\infty} e^2 \rho(e \mid v)\, de. \tag{2.5}$$

The quantity (2.5) defines how the average error variance of the analysis depends on the innovation. It was shown in Hodyss (2011), Hodyss and Campbell (2013), and Posselt et al. (2014) that a Bayesian posterior is such that $S(v) = P_v(v)$. Therefore, an alternative ensemble diagnostic to the binned spread-skill diagram is the

relationship between (2.4) and (2.5) such that one may plot (2.5) versus (2.4) and expect a one-to-one relationship. We will show later in this manuscript that even when the traditional binned spread-skill diagram reveals a perfect one-to-one relationship between the MSE (skill) and the ensemble variance at analysis time (spread) that the relationship between (2.4) and (2.5) may be completely wrong for an EBKF. The relationship between (2.4) and (2.5) is not always one to one because the ensemble generation schemes used by the EBKF assume that the posterior ensemble should have a variance that satisfies

$$\overline{P}_a = (I - KH)P_f, \tag{2.6}$$

where K is the Kalman gain and P_f is the prior error variance. As shown by Hodyss (2011), Hodyss and Campbell (2013), and Posselt et al. (2014), even if we have an infinite ensemble, a perfect model, and therefore a perfect prior error variance, (2.6) does *not* result in the correct variance for nonlinear systems with non-Gaussian pdfs. Even in Gaussian situations, a limited ensemble size leads to sampling error, rendering (2.6) again incorrect, in that the true posterior error variance is a function of the latest set of observations and (2.6) is not. We believe this paper is the first attempt at delivering a posterior ensemble with a better relationship between (2.4) and (2.5) while retaining a good relationship between skill and spread. This will require a correction for sampling error that we will show to be easy to calculate, easy to use, and generally applicable to serial and global EBKFs.

Note that (2.6) is determined by solving for the *expected* variance of the posterior (EVP):

$$\overline{P}_a = \int_{-\infty}^{\infty} P_a(\upsilon)\rho(\upsilon)\, d\upsilon, \tag{2.7}$$

where $P_a(\upsilon)$ is the variance of the posterior (VP):

$$P_a(\upsilon) = \int_{-\infty}^{\infty} (x - \overline{x}_K)^2 \rho(x \,|\, \upsilon)\, dx, \tag{2.8}$$

and $\rho(\upsilon)$ is the marginal distribution of innovations υ, $\rho(x \,|\, \upsilon)$ is the posterior distribution, and the analysis in (2.8) is only an approximation to the true posterior mean $\overline{x}$:

$$\overline{x} \approx \overline{x}_K = \overline{x}_f + K[y - H\overline{x}_f], \tag{2.9}$$

where the subscript K makes clear that the Kalman estimate of the posterior mean is not necessarily equal to the true posterior mean. We emphasize here that we will compare the EVP, in (2.7), to the VP of the EBKF estimate in (2.8) throughout this manuscript and urge the reader to make note of this notation.

As alluded to in the introduction, we view the inflation literature as reasonably being broken into two groups. On the one hand, there are methods that aim to make the variance correct on average (e.g., Anderson and Anderson 1999; Zhang et al. 2004; Furrer and Bengtsson 2007; Sacher and Bartello 2008; Whitaker and Hamill 2012; Menetrier and Auligne 2015), which implies that they attempt to set the ensemble variance to (2.7). We will show in section 3 that sampling error leads to a systematic underestimation of (2.7) and that inflation can help to reduce this bias in the ensemble variance. On the other hand, there are methods that attempt to give an observation dependence to the ensemble variance (e.g., Anderson 2007, 2009), which we interpret as meaning that they attempt to set the ensemble variance to (2.8) and we further interpret this as aiming to construct a more Bayesian ensemble from an EBKF. Note that by getting the VP ensemble to have the variance in (2.8) that one obtains (2.7) by default. However, if one sets the VP ensemble to (2.7) the ensemble will not properly satisfy (2.8) and this fact motivates our goal in this paper to construct a method that tries to set the EBKF ensemble variance to (2.8).

Because (2.9) is an approximation to the true posterior mean, (2.8) differs from the VP about the true posterior mean:

$$P_t(\upsilon) = \int_{-\infty}^{\infty} (x - \overline{x})^2 \rho(x \,|\, \upsilon)\, dx \le P_a(\upsilon). \tag{2.10}$$

Equation (2.8) is in fact the correct VP for the errors about the flawed state estimate delivered by an EBKF algorithm (2.9). Equation (2.8) may be written in terms of the true VP, $P_t(\upsilon)$, plus the square of an observation-dependent bias:

$$P_a(\upsilon) = P_t(\upsilon) + (\overline{x} - \overline{x}_K)^2. \tag{2.11}$$

Equation (2.11) is derived from (2.8) by adding and subtracting the true posterior mean inside the squared parentheses. Equation (2.11) states that the VP for an EBKF should be the true VP inflated by the observation-dependent bias in the estimate of the posterior mean from (2.9). Note that the state estimate $\overline{x}_K$ may differ from the true posterior mean $\overline{x}$ because of 1) non-Gaussianity of the prior and/or observation likelihood, 2) sampling error, and 3) model error. While this manuscript will focus on the ways that sampling error induces differences between $\overline{x}_K$ and $\overline{x}$, we provide next an example of error introduced by non-Gaussianity, which will prove useful in section 4 where we study these differences in two different nonlinear models.

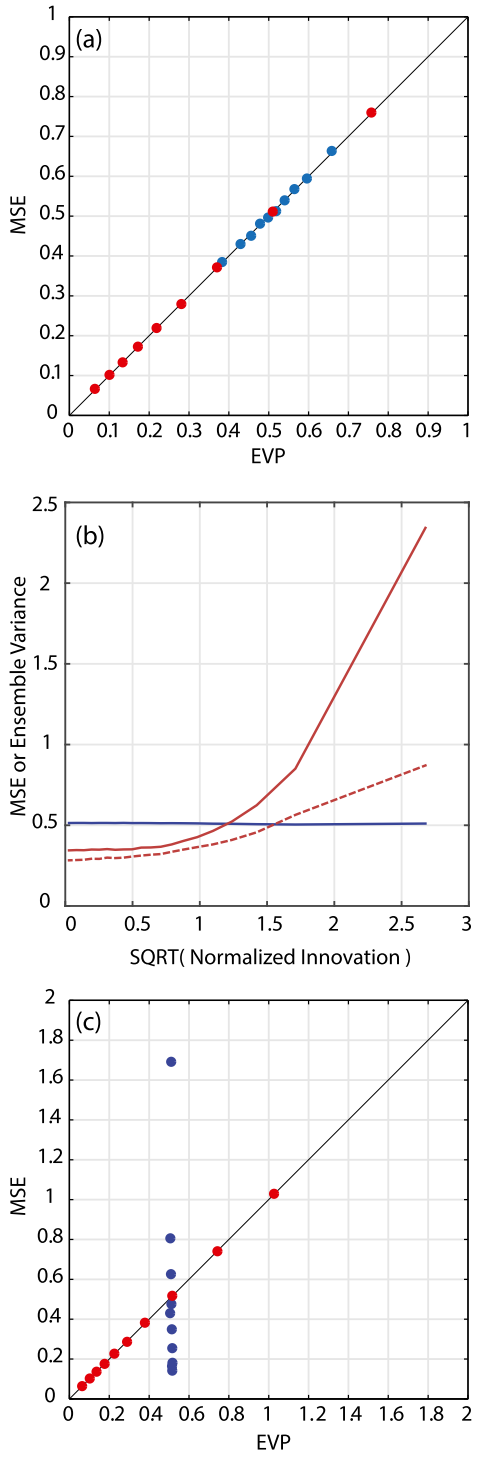

FIG. 1. Data assimilation diagnostics for a non-Gaussian problem. (a) The traditional binned spread-skill diagnostic for the Kalman filter (blue) is shown as well as the Bayesian solution (red). (b) The MSE (red solid) of the posterior mean and EVP (blue) from a Kalman filter binned as a function of normalized innovation. The red dashed line is the MSE as a function of normalized innovation for the Bayesian posterior. The Bayesian ensemble variance is identical to the MSE and cannot be seen under the red dashed line. (c) The new binned spread-skill diagnostic for the Kalman filter (blue) as well as the Bayesian solution (red).

b. *Non-Gaussian example*

Suppose we wish to estimate the state for a scalar in the case where the prior distribution is a non-Gaussian random variable x represented as

$$x = \overline{x}_f + \sqrt{\frac{P_f}{2}}(\chi^2 - 1), \tag{2.12}$$

where χ is a random draw from a Gaussian with $\mathcal{N}(0, 1)$ such that χ^2 is a random variable distributed according to the standard ''chi-square'' distribution, the prior mean $\overline{x}_f$ is equal to 1, and P_f is the true variance of the prior. We chose a chi-square-distributed variable here simply because it is a well-known and easily constructed model of a non-Gaussian process. Assume a Gaussian observation likelihood with an observation error variance R equal to 1. For each trial of this experiment, a true state will be created using (2.12) and random noise with variance R will be added to create an observation. In addition, for each trial, the parameter P_f will be drawn from an *inverse* Gamma distribution with shape and scale (k–θ) parameters equal to 10, which results in a mean prior variance of 1.11. We draw the prior variance from a distribution (rather than setting it to a single number) because we require that there be a distribution of EVP values across the trials to properly produce a binned spread-skill diagram.

Assume also that we have access to an infinite ensemble and simply make use of the known true variance of the prior P_f and the known true prior mean $\overline{x}_f$ in (2.9) to estimate the state, and similarly in (2.6) to estimate the posterior variance. Conduct 10^6 trials, sort the 10^6 EVPs we obtained from (2.6) from smallest to largest, and construct $N = 10$ equally populated bins ($M = 10^5$). Corresponding to the EVPs in each of these bins are state estimates from (2.9) for which we apply (2.1), where we simply validate against the known truth. Figure 1a reveals a perfect relationship between the EVP and the MSEs. For this simple example, we can calculate the complete Bayesian solution (done here using particle filtering methods) and compare (see Fig. 1a). The Bayesian solution must also produce a perfect relationship between the ensemble variance and the MSEs. Because the Bayesian solution does not make assumptions on the shape of the prior distribution while the Kalman filter does, the Bayesian solution will generally have a lower MSE, which can be seen by noting that the equally populated bins are generally shifted toward smaller values of MSE.

As discussed above, we may alternatively bin by the size of the innovations. To compare innovations

drawn from different priors (with different values of P_f), we need to normalize them. In Fig. 1b we normalize the innovations by their standard deviation [$v_N = \sqrt{v^2/(P_f + R)}$], sort the results into 10 bins, and then compute the analysis error in each bin [see (2.1)]. One can see that consistent with (2.6), the EVP is a horizontal line, implying that it does not vary with innovation. Equation (2.8) reveals how the MSE of the Kalman analysis varies with innovation. Figure 1b shows that the MSE of the Kalman estimate is in fact a strong function of the innovation, because the prior is non-Gaussian. It is important to realize that if the prior were Gaussian, then the posterior MSE curve as a function of innovation would be flat, which means that all innovations, whether small or large, result in the same expected squared error. This is a property of Gaussian posteriors and explains why the Kalman state estimate is correct for Gaussian posteriors even when its ensemble variance does not vary with innovation.

From Fig. 1b it is clear that the EVP does not predict the MSE in the state estimate, as the blue curve and the red curve are far from matching. By contrast, the MSE and the VP for the Bayesian result are identical, showing that the Bayesian result properly predicts the MSE as a function of innovation. In Fig. 1c we plot (2.5) versus (2.4). This new type of binned spread-skill plot is constructed by taking the data from the MSE (red curve) and the EVP (blue curve) in Fig. 1b and plotting them against each other. In other words, we plot a blue circle at the location $(x, y) = (\text{EVP}, \text{MSE})$. The blue circles show that the MSE is completely independent of EVP; by contrast, the Bayesian results (red circles) fall exactly along the one-to-one line.

Figure 1a gives the impression that during the data assimilation cycles when the EVP is large, the MSE is large, and vice versa. However, Fig. 1b shows that the EVP does not actually track the MSE. This is because during each of these data assimilation cycles, there is actually a different innovation, but the EVP does not vary with innovation. So, if the EVPs are not predicting the cycle-to-cycle bin-averaged MSE, then what does Fig. 1a actually show? The traditional binned spread-skill diagram in Fig. 1a shows that the EVP is an unbiased estimate of the true VP, such that an average of the VP over innovation [i.e. (2.7)] results in the correct *average* (or expected) VP, but *not* the correct VP for each cycle. In other words, for each of these cycles during the evaluation period the VP is equally distributed about the EVP but in fact is rarely if ever equal to it. Therefore, the binned spread-skill diagram in Fig. 1a does not imply that the EVP is correctly predicting the MSE for each cycle, or equivalently that the ensemble variance correctly changes with the size of the actual squared error. Indeed, Fig. 1b shows that the EVP is usually incorrectly predicting the MSE. We emphasize again that the Bayesian result does not suffer from this problem and, therefore, an ensemble that is more consistent with the Bayesian posterior (and therefore better) would result from accounting for this discrepancy. In the next section, we will illustrate how EBKFs incorrectly predict the MSEs in the presence of sampling error and discuss an inflation technique to address this issue.

3. Sampling error

Furrer and Bengtsson (2007) and Sacher and Bartello (2008) clearly describe the impact of sampling error on the EBKF state estimates and ensemble variances from the perspective of the EVP. Here, we extend Furrer and Bengtsson (2007) and Sacher and Bartello (2008) by showing that the impact of sampling error on the MSE of the EBKF is in fact a function of the observations when viewed from the perspective of the VP. In this way we will see the issues discussed in the previous sections emerge as a natural consequence of sampling error in the prior estimates.

a. Prior

We begin by writing the formulas describing the uncertainty in a sample estimate of the prior mean and variance, and we emphasize here that we are not assuming that the prior, $\rho(x)$, is necessarily Gaussian. We imagine the particular random variable of interest is a random draw from $\rho(x)$ such that

$$x_i = \overline{x}_f + \sqrt{P_f}\varepsilon_i, \tag{3.1}$$

where $\overline{x}_f$ is the true mean of $\rho(x)$, P_f is the true variance of $\rho(x)$, and ε_i has mean zero, variance one, and ε_i is distributed according to the standardized score:

$$\varepsilon_i = \frac{x_i - \overline{x}_f}{\sqrt{P_f}}. \tag{3.2}$$

The subscript i in (3.1) and (3.2) denote the ith random draw from $\rho(x)$ and refer to the members of our ensemble. We emphasize here that because we are not assuming that $\rho(x)$ is necessarily Gaussian, this implies that ε_i may be drawn from a non-Gaussian distribution. Therefore, while ε_i has mean zero and variance one its higher moments are arbitrary and not constrained to satisfy the properties of a Gaussian.

The sample mean and sample variance are

$$\overline{x}^s = \frac{1}{N_e}\sum_{i=1}^{N_e} x_i, \tag{3.3a}$$

$$P^s = \frac{1}{N_e - 1}\sum_{i=1}^{N_e} (x_i - \overline{x}^s)^2, \tag{3.3b}$$

where N_e is the number of samples (ensemble size) drawn from $\rho(x)$. We imagine that we calculate (3.3a) and (3.3b) many times with different random draws from $\rho(x)$ in order to build statistics. We denote these different calculations of (3.3a) and (3.3b) with the index j and by using (3.1) in (3.3a) and (3.3b) we may show that

$$\overline{x}_j^s = \overline{x}_f + \sqrt{P_f}\,\frac{1}{N_e}\sum_{i=1}^{N_e} \varepsilon_{ij}, \tag{3.4a}$$

$$P_j^s = a_j P_f, \tag{3.4b}$$

$$a_j = \frac{1}{N_e - 1}\left[\sum_{i=1}^{N_e} \varepsilon_{ij}^2 - \frac{1}{N_e}\left(\sum_{i=1}^{N_e} \varepsilon_{ij}\right)^2\right]. \tag{3.4c}$$

Hence, the statistics of the sample mean and variance depend on the characteristics of the ε_{ij}. To reduce the complexity of the presentation we will make use of the following notation for the expected value with respect to sampling:

$$\langle u_j \rangle = \lim_{J\to\infty} \frac{1}{J}\sum_{j=1}^{J} u_j, \tag{3.5}$$

where the convergence of this limit is to be understood in the sense of probability.

As is well known the sample mean and variance definitions in (3.4a)–(3.4c) are unbiased, which implies that their mean with respect to sampling error is the true value:

$$\langle \overline{x}_j^s \rangle = \overline{x}_f, \tag{3.6a}$$

$$\langle P_j^s \rangle = P_f, \tag{3.6b}$$

where we have used the fact that $\langle a_j \rangle = 1$. The variance due to sampling of the sample mean in (3.3a) is simply the standard central limit theorem (CLT) result

$$\langle (\overline{x}_j^s - \overline{x}_f)^2 \rangle = \frac{P_f}{N_e}, \tag{3.7}$$

while the variance due to sampling in the sample variance is

$$\langle (P_j^s - P_f)^2 \rangle = \langle (a_j - 1)^2 \rangle P_f^2 = \frac{1}{N_e}\left(F_f - \frac{N_e - 3}{N_e - 1} P_f^2\right), \tag{3.8a}$$

where we have used the fact that

$$\langle (a_j - 1)^2 \rangle = \frac{1}{N_e}\left(\frac{F_f}{P_f^2} - \frac{N_e - 3}{N_e - 1}\right), \tag{3.8b}$$

and F_f is the central fourth moment of $\rho(x)$. Note that if $\rho(x)$ were to be Gaussian such that $F_f = 3P_f^2$ then (3.8a) would reduce to the standard formula:

$$\langle (P_j^s - P_f)^2 \rangle = \frac{2}{N_e - 1} P_f^2. \tag{3.9}$$

We will use (3.7) and (3.8) below to quantify our uncertainty in sample estimates of the mean and variance of the prior distribution that we use in the Kalman formula.

b. Posterior

To understand the impact of sampling error on the ensemble Kalman filter estimate of the posterior mean we include perturbations from sampling:

$$\overline{x}_j^a = \overline{x}_j^s + \frac{P_j^s}{P_j^s + R}(y - \overline{x}_j^s), \tag{3.10a}$$

$$\overline{x}_j^s = \overline{x}_f + \delta_j^m, \tag{3.10b}$$

$$P_j^s = P_f + \delta_j^v, \tag{3.10c}$$

where the perturbations δ_j^m and δ_j^v are the errors in the estimate of the mean and variance, respectively, owing to sampling and are from (3.4a) and (3.4b):

$$\delta_j^m = \sqrt{P_f}\,\frac{1}{N_e}\sum_{i=1}^{N_e} \varepsilon_{ij}, \tag{3.11a}$$

$$\delta_j^v = (a_j - 1)P_f. \tag{3.11b}$$

Without loss of generality, we will throughout this section not make use of an observation operator in the innovation as this is a scalar system and we are observing the variable of interest. We emphasize, however, that including the observation operator in the analysis is trivial.

Using a Taylor-series expansion in terms of the perturbations δ_j^m and δ_j^v we may write (3.10a) as

$$\overline{x}_j^a = \overline{x}_K + \frac{\overline{P}_a}{P_f}\left(\delta_j^m + [\overline{x}_K - \overline{x}_f]\frac{\delta_j^v}{P_f}\right) + \cdots, \tag{3.12}$$

where the true Kalman mean without sampling error is

$$\overline{x}_K = \overline{x}_f + \frac{P_f}{P_f + R}(y - \overline{x}_f), \tag{3.13}$$

and $\overline{P}_a$ is the true EVP [(2.6)], evaluated with the true prior variance.

We may neglect quadratic and higher terms in (3.12) and find the variance with respect to sampling error within our Kalman estimate about the true posterior mean as

$$\langle(\overline{x}-\overline{x}_j^a)^2\rangle \approx \underbrace{(\overline{x}-\overline{x}_K)^2}_{\text{Non-Gaussianity}} + \underbrace{\left[\frac{\overline{P}_a}{P_f}\right]^2\left(\langle(\delta_j^m)^2\rangle + [\overline{x}_K-\overline{x}_f]^2\frac{\langle(\delta_j^v)^2\rangle}{P_f^2}\right)}_{\text{Sampling}}, \tag{3.14}$$

where $\overline{x}$ is the true posterior mean and differs from the true Kalman mean when the posterior is not Gaussian. Note that the quantities $\langle(\delta_j^m)^2\rangle$ and $\langle(\delta_j^v)^2\rangle$ are known from (3.7) and (3.8a), respectively. In the derivation of (3.14) we have used the fact that the correlation between the sampling errors in each statistic, $\langle\delta_j^m\delta_j^v\rangle$, can readily be calculated by using the same techniques as in section 3a, and is found to vanish (i.e. $\langle\delta_j^m\delta_j^v\rangle = 0$). Equation (3.14) shows that there are two components of the expected squared difference between the sample mean estimate and the true posterior mean: one owing to non-Gaussianity and the other due to sampling error.

Equation (3.14) can be used to calculate the MSE in our posterior mean estimate $\overline{x}_j^a$ as

$$S(y) = \int_{-\infty}^{\infty}\langle(x-\overline{x}_j^a)^2\rangle\rho(x\,|\,y)\,dx = P_t(y) + \langle(\overline{x}-\overline{x}_j^a)^2\rangle$$
$$\approx P_t(y) + (\overline{x}-\overline{x}_K)^2 + \left[\frac{\overline{P}_a}{P_f}\right]^2\left[\frac{P_f}{N_e} + \frac{1}{N_e}\left(\frac{F_f}{P_f^2} - \frac{N_e-3}{N_e-1}\right)(\overline{x}_K-\overline{x}_f)^2\right] \tag{3.15}$$

and $P_t(y)$ is the VP about the true posterior mean. The MSE (S) of $\overline{x}_j^a$ has three terms. The first term, $P_t(y)$, measures the natural variability about the true posterior mean. The second term measures the inflation of the MSE by the bias owing to the difference between the true Kalman estimate and the true posterior mean, which is invoked in non-Gaussian situations where the true Kalman state estimate (3.13) is not accurate. Last, the third term is the inflation of the MSE owing to sampling error in the Kalman filter equation and should be noted to be a function of the observations through $\overline{x}_K$. This third term is composed of two separate terms: the first term is the inflation owing to sample error in the prior mean estimate (3.7) and the second term is the inflation owing to sampling error in the prior variance (3.8a) used to calculate the Kalman gain.

It is interesting to compare the result in (3.15) to the EVP from a typical ensemble generation scheme as the expectation is that the EVP is attempting to predict S. We, therefore, may evaluate (2.6) using (3.10c) and the same Taylor expansion as (3.12), but this time keeping quadratic terms, to obtain

$$\overline{P}_{aj}^s \approx \overline{P}_a + \left(\frac{R}{P_f+R}\right)^2\delta_j^v - \frac{R^2}{(P_f+R)^3}(\delta_j^v)^2. \tag{3.16}$$

By applying (3.5) and making use of (3.11b) we obtain

$$\langle\overline{P}_{aj}^s\rangle \approx \overline{P}_a - \frac{1}{P_f}\langle(a_j-1)^2\rangle\overline{P}_a^2 + \frac{1}{P_f^2}\langle(a_j-1)^2\rangle\overline{P}_a^3. \tag{3.17}$$

Equation (3.17) shows that, except for very large EVP, the EVP including sampling error is less than the true EVP, which can be seen by the negative quadratic term.

As an example, note that if $\rho(x)$ is Gaussian then $\overline{x} = \overline{x}_K$ and $P_a(y) = \overline{P}_a$ such that

$$S_G \approx \overline{P}_a + \left[\frac{\overline{P}_a}{P_f}\right]^2\left[\frac{P_f}{N_e} + \frac{2}{N_e-1}(\overline{x}_K-\overline{x}_f)^2\right] \tag{3.18a}$$

and similarly

$$\langle\overline{P}_{aj}^s\rangle \approx \overline{P}_a - \frac{2}{N_e-1}\frac{\overline{P}_a^2}{P_f} + \frac{2}{N_e-1}\frac{\overline{P}_a^3}{P_f^2}. \tag{3.18b}$$

Again, we hope that the EVP in (3.18b) will predict the MSE of the ensemble mean in (3.18a) (i.e., if our ensemble was properly constructed we would have $S_G = \langle\overline{P}_{aj}^s\rangle$). Note, however, that there are two obvious ways that (3.18a) and (3.18b) are different: 1) the second-order terms in (3.18a) and (3.18b) have opposite signs and 2) (3.18a) is a function of the latest observation through $\overline{x}_K$ and (3.18b) is not. In section 4 we describe a posterior inflation algorithm that attempts to account for these two issues in the posterior ensemble generation step of an EBKF algorithm.

c. *Gaussian example*

We will illustrate the theory of section 3b using a simple Gaussian model:

$$x = \overline{x}_f + \sqrt{P_f}\chi. \tag{3.19}$$

Here, χ is a random draw from a normal distribution with mean 0 and variance 1, prior mean, $\overline{x}_f$ is equal to 1, and P_f is the true variance of the prior. As in the previous section, we will assume a Gaussian observation likelihood with an observation error variance R equal to 1. Here we employ 10^7 trials. For each trial of this experiment, a sample truth will be created using (3.19) and random noise with variance R added to create an observation. Again, for each trial, the parameter P_f will be drawn from an *inverse* Gamma distribution with shape and scale parameters equal to 10, which results in a mean prior variance of 1.11. The pdf for P_f is plotted in Fig. 2a. In contrast to the experiment in section 2, here we use a small ensemble of $N_e = 8$ members to estimate the prior mean and variance using samples from (3.19). The pdf of eight-member estimates of P_f from (3.3b) is plotted in Fig. 2a. One can see that the pdf of P_j^s is wider, indicating a larger variance, than the pdf of P_f, consistent with the definition of the marginal of the sampling distribution:

$$\rho(P^s) = \int_0^\infty \rho(P^s \mid P_f)\rho(P_f)\, dP_f. \tag{3.20}$$

In (3.20) $\rho(P_f)$ is the above-mentioned inverse Gamma distribution and $\rho(P^s \mid P_f)$ is the conditional pdf denoting the process in (3.10c). This excess variability from sampling errors in the prior variances results in excess variability in the EVP through calculations of (2.6), which are also shown in Fig. 2a.

To understand a little bit more about the relationship between the true EVP from (2.6) and its sample estimates we divide our samples of the true EVP into 10 equally populated bins and plot the mean sample EVP against the mean true EVP in Fig. 2b. Additionally, we evaluate (3.18b) in Fig. 2b, which shows that truncating the expansion as a cubic polynomial does an excellent job of predicting this relationship. This, both empirically and theoretically, confirms the well-known result that the sample estimates of the EVP underestimates the true EVP because the red curve and filled circles are below the one-to-one line.

As emphasized in section 2, the EVP differs from the true VP as is shown in Fig. 2c. Recall in section 2 that we showed that a truly Bayesian VP, unlike the EVP, must be equal to the MSE of the state estimate. To understand

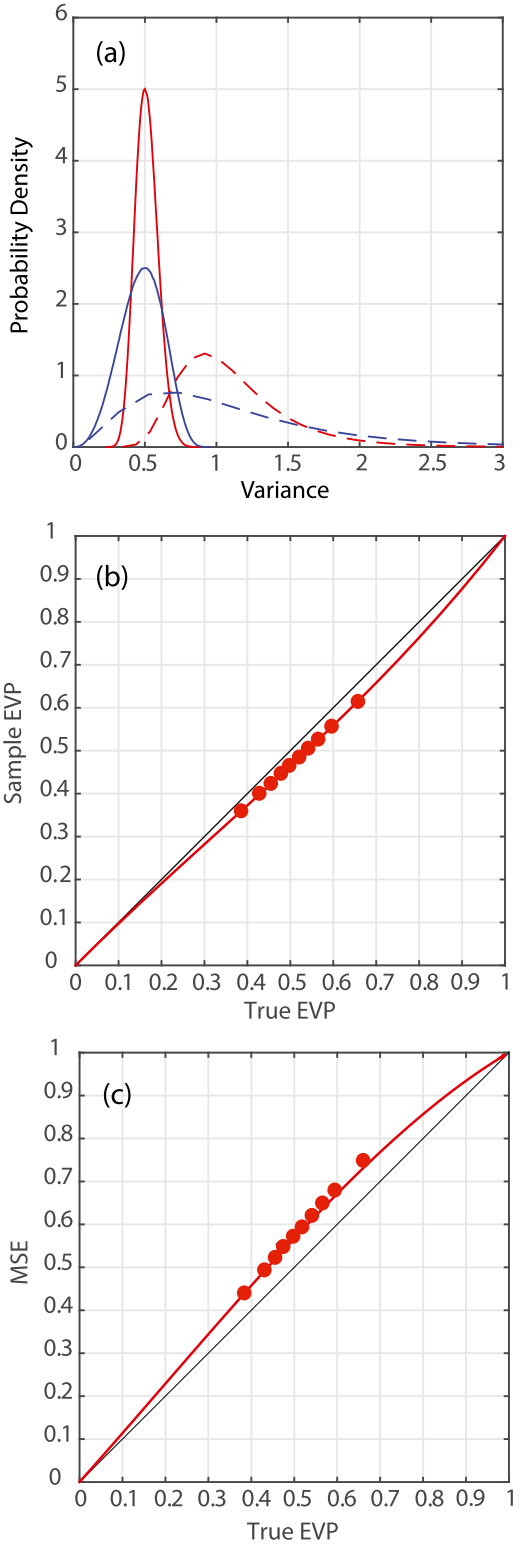

FIG. 2. Sampling statistics for the EVP. (a) The marginal distributions for the variances for the prior (dashed) and the posterior (solid) for the true distributions (red) and the sampling distributions (blue). (b) The binned sample EVP plotted as a function of true EVP (red closed circles). The red curve is (3.18b). (c) The binned MSE plotted as a function of the true EVP (red closed circles). The red curve is (3.21).

this better we again divide our samples of the true EVP into 10 equally populated bins and plot the MSE of the sample mean estimate in each bin [(2.2)]. To compare this to our estimate of the MSE in (3.18a) we must account for the fact that each bin is also equivalent to an average (integration) over innovation. We do this by integrating with respect to the marginal distribution of innovations:

$$\overline{S}_G = \int_{-\infty}^{\infty} S_G(v)\rho(v)\,dv \approx \overline{P}_a + \left[\frac{1}{N_e} + \frac{2}{N_e - 1}\right]\frac{\overline{P}_a^2}{P_f} - \frac{2}{N_e - 1}\frac{\overline{P}_a^3}{P_f^2}, \tag{3.21}$$

and subsequently plot this curve in Fig. 2c. Note that (3.21) provides an excellent prediction of the MSE as a function of the true EVP $\overline{P}_a$. As discussed in section 2, one can see from Fig. 2c that the relationship between MSE and the true EVP does not fall on the one-to-one line, and, therefore, that the true EVP does not predict the MSE of the sample posterior mean. To understand this better, it is interesting to compare (3.21) to (3.18b). While both (3.21) and (3.18b) are cubic polynomials in the EVP, the second and third terms are opposite in sign and (3.21) includes an inflation of the MSE owing to sampling error in the prior mean (first term in brackets in the second term on the right-hand side) and (3.18b) does not include this term. This term is missing from (3.18b) because (3.18b) is calculated from (2.6), which does not know about the prior mean.

This averaging with respect to innovation in (3.21) is unsatisfactory because it is our goal to understand and include observation-dependent variance information in the ensemble. Hence, in Fig. 3 we present the result of binning the MSE about the Kalman state estimate with respect to innovation. In Fig. 3a we plot the MSE as a function of the normalized innovation. As in section 2, one can see that small, normalized innovations typically result in smaller posterior MSE than large innovations. Because the prior and posterior are Gaussian in this case (and Gaussianity does not lead to observation-dependent variances) we know then that this observation-dependent error variance is a direct result of the sampling error.

In Fig. 3a we also plot the EVP for both the eight-member estimate as well as the true value of the EVP. Note that both are below the MSE but that the sample estimate is further below as was predicted by (3.18b). The "true" EVP is less than the actual MSE because the true EVP is measuring the expected error variance about a Kalman estimate with no sampling error. Note that while (2.6) cannot predict the MSE, (3.18a) can and is plotted in Fig. 3a. Equation (3.18a) provides a very good estimate of the MSE as a function of innovation albeit it is slightly too small. This underestimation of the MSE by (3.18a) is a direct result of the truncation of the Taylor expansion.

Comparing the MSE of the Kalman state estimate in Fig. 1b to that in Fig. 3a reveals a stronger degree of curvature in the non-Gaussian example of Fig. 1b. The sampling error theory of this section reveals that the lowest-order impact of sampling error is that the MSE picks up a quadratic dependency on innovation. In the non-Gaussian example of Fig. 1b, where we remind the reader that there was no sampling error, the strong innovation dependency arose from the observation-dependent bias term in (2.5). This same term can be seen in (3.15) and, using the theory of quadratic nonlinear regression as applied to Kalman filtering (Hodyss 2011), can be shown to lead to a quartic dependency on the innovation. This quartic dependency arises because the lowest-order difference between the true posterior mean and that estimated by a Kalman filter is a quadratic function of the innovation. We will see in the next section that in the nonlinear models of that section, where there exists both sampling error and non-Gaussianity, the MSE is a stronger function of the innovation than predicted by (3.18a), which is consistent with non-Gaussianity leading to a quartic dependency on innovation in the MSE.

In Fig. 3b we plot our new version of the MSE as a function of ensemble variance diagram. This new type of binned spread-skill plot is constructed by taking the data from the MSE (red curve) in Fig. 3a and the sample EVP (blue curve) in Fig. 3a and plotting a line through all the locations (x, y) = (EVP, MSE). Because the EVP does not predict the MSE the blue curve is a vertical line, which implies an inability to predict the MSE. Similarly, we also plot in Fig. 3b the same relationship using MSE as a function of the true EVP and this shows the same inability to predict the MSE. Last, we plot MSE as a function of the predicted MSE [(3.18a); black curve in Fig. 3a] and we see a very good prediction of the MSE as this black line is now very close to the one-to-one line.

In the next section, which provides example applications for two different nonlinear evolution equations, we will not have available the true prior variances to calculate the normalized innovations and we will not have the true EVP and posterior mean. This will lead us to normalize the innovations with the sample prior variances and to evaluate (3.18a) with sample estimates rather than the true values. In Figs. 3c and 3d we show how the results change when we improperly normalize the innovations. The only difference between Figs. 3a, b and 3c,d is that the normalization in Figs. 3a and 3b used the true prior variances and the normalization in Figs. 3c and 3d used the sample prior variances.

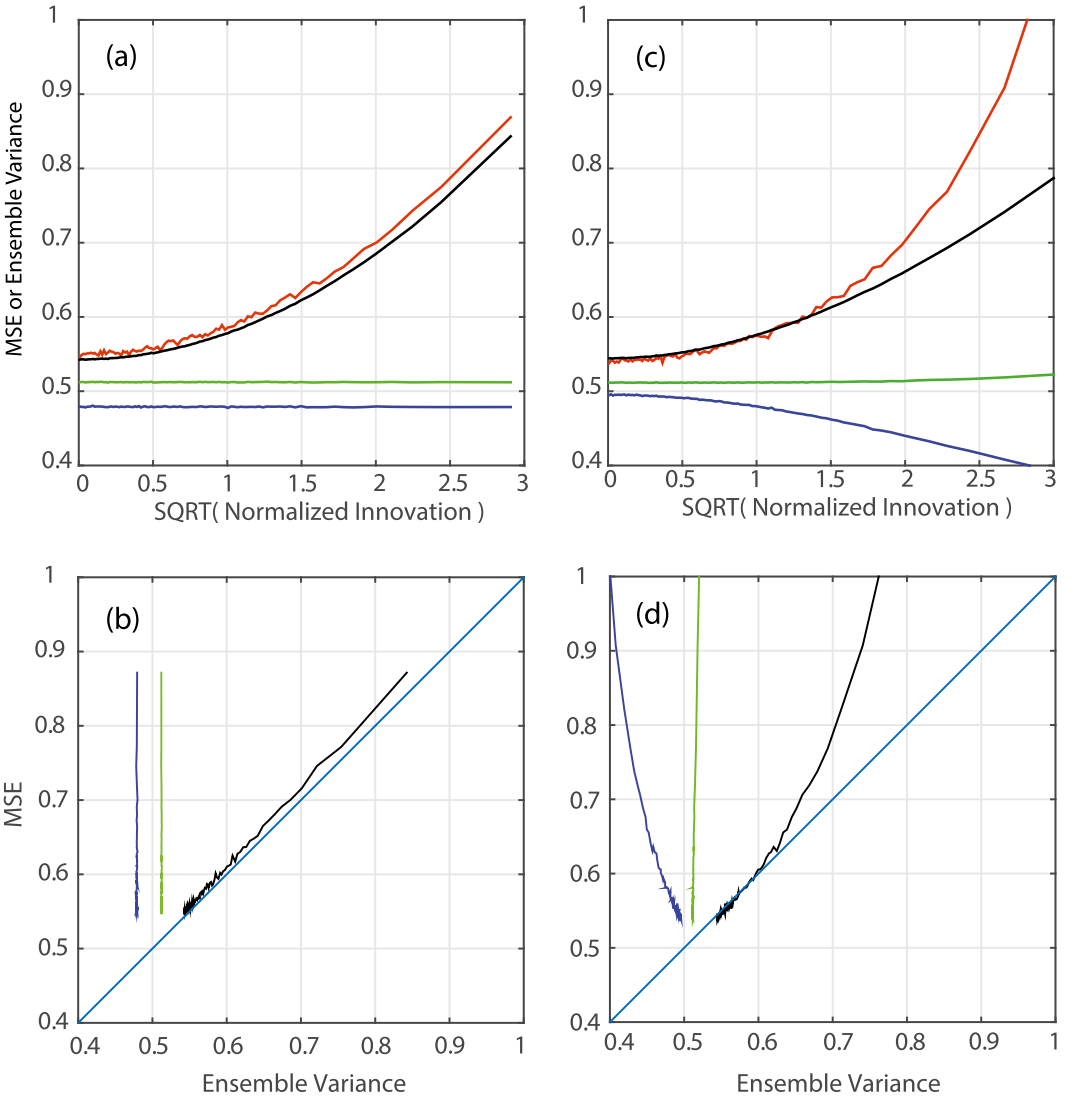

FIG. 3. Ensemble diagnostics as a function of normalized innovation. (a),(c) The MSE (red), sample EVP (blue), true EVP (green), and the prediction of the MSE from (3.18a) (black) are plotted. (b),(d) The new binned spread-skill diagnostic for MSE vs sample EVP (blue), MSE vs true EVP (green), and MSE vs the prediction of the MSE from (3.18a) (black) are plotted. In (a),(b) the normalized innovation uses the true prior variance and prior mean. In (c),(d) the normalized innovation uses the sample estimate of the prior variance and prior mean and (3.18a) uses sample estimates.

Therefore, Figs. 3c and 3d show us what we need to be aware of when applying this diagnostic to real-world scenarios. Figure 3c shows that it will appear that the EVP bends to smaller values for large normalized innovations and that the MSE has a larger dependence on the normalized innovation than it should be when the innovation is normalized using a sample estimate. Furthermore, these new relationships will lead to additional curvature in Fig. 3d. These features must be noted when using these diagnostics in real-world situations as will be illustrated next.

4. Applications

Our goal in this section is to illustrate how to apply this theory to a state vector of arbitrary length and to discuss how to apply the previous theory to cycling data assimilation in nonlinear systems.

a. Lorenz-63

In this subsection we perform data assimilation using the ensemble Kalman filter (EnKF) with the ensemble generation scheme referred to as perturbed

observations (Evensen 1994, 2003). The equation for the EnKF posterior update for the ith prior ensemble member $\mathbf{x}_f^i$ can be written as

$$\mathbf{x}_a^i = \mathbf{x}_f^i + \mathbf{P}_f\mathbf{H}^{\mathrm{T}}[\mathbf{H}\mathbf{P}_f\mathbf{H}^{\mathrm{T}} + \mathbf{R}]^{-1}(\mathbf{y} + \boldsymbol{\varepsilon}_o^i - \mathbf{H}\mathbf{x}_f^i), \quad (4.1)$$

where $\mathbf{y}$ are the observations and the observation operator $\mathbf{H}$ will be chosen here to simply observe the state variables x and z, both with an observation error variance of 0.1; $\boldsymbol{\varepsilon}_o^i$ will be a vector of length 2 consisting of random draws from a normal distribution with mean zero and variance equal to $\mathbf{R}$. Last, $\mathbf{P}_f$ is the (3×3) forecast error covariance matrix, which will be estimated using ensemble statistics in the standard way.

We show results for three separate types of inflation:

1) We may employ a constant prior inflation, which is simply applied by defining a matrix $\mathbf{Z}$ whose columns are the prior perturbations, $\boldsymbol{\varepsilon}_f^i = \mathbf{x}_f^i - \overline{\mathbf{x}}_f$, and subsequently transforming as

$$\mathbf{Z} \rightarrow \alpha\mathbf{Z}, \quad (4.2)$$

where α is a tunable parameter greater than 1 and $\overline{\mathbf{x}}_f$ is the prior mean. The resulting covariance matrix is then used in (4.1).

2) We may employ a constant posterior inflation, which is also applied by defining a new matrix $\mathbf{X}$, whose columns are the posterior perturbations, $\boldsymbol{\varepsilon}_a^i = \mathbf{x}_a^i - \overline{\mathbf{x}}_a$, and subsequently transforming as

$$\mathbf{X} \rightarrow \beta\mathbf{X}, \quad (4.3)$$

where β is a tunable parameter greater than 1 and $\overline{\mathbf{x}}_a$ is the posterior mean. The new perturbations in (4.3) are added to $\overline{\mathbf{x}}_a$ and then used as the posterior ensemble.

3) Last, we may employ the posterior inflation method of section 3 by extending (3.18a) to a vector as

$$\mathbf{S}_G = a\mathbf{p}_a + \left[\frac{\mathbf{p}_a}{\mathbf{p}_f}\right]^2\frac{\mathbf{p}_f}{K} + b\left[\frac{\mathbf{p}_a}{\mathbf{p}_f}\right]^2\left[\frac{2}{K-1}\mathbf{d}^2\right], \quad (4.4)$$

where

$$\mathbf{p}_a = \frac{\mathrm{diag}(\mathbf{X}\mathbf{X}^{\mathrm{T}})}{K-1}, \quad (4.5)$$

$$\mathbf{p}_f = \frac{\mathrm{diag}(\mathbf{Z}\mathbf{Z}^{\mathrm{T}})}{K-1}, \quad (4.6)$$

$$\mathbf{d} = \overline{\mathbf{x}}_a - \overline{\mathbf{x}}_f, \quad (4.7)$$

and the operation "diag" refers to forming the diagonal elements of a matrix into a vector. In (4.4), *all mathematical operations involving vectors are to be interpreted as element-wise operations*. Obviously, (4.4) is not strictly correct, as we have not accounted for sampling error in the covariance between state variables because we simply perform the calculation in (3.18a), which ignores the sampling error in the covariance for each variable in the state vector. This approximation has been made in order to result in a formula that may simply be evaluated using vector manipulations and, most importantly, is a function of the observations *without requiring access to the observations*. In other words, (4.4) only requires access to the vector difference between the analysis and the prior mean (commonly referred to as the "correction").

The parameters a and b are tunable parameters that are required because we will now apply the theory of section 3 to non-Gaussian distributions and because we will evaluate (4.4) using sample statistics rather than the true EVP. Each of these parameters is there to account for a specific effect missing from the theory of section 3. First, (3.15) shows that in non-Gaussian situations we need to know the quantity $P_t(y) + (\overline{x} - \overline{x}_K)^2$, which we do not know. As shown in Fig. 1b the MSE, which this quantity represents, is actually smaller than the EVP for small values of normalized innovation. This implies that we will tune the parameter $a < 1$ because $\mathbf{p}_a$ will tend to overestimate this quantity. Note, however, that in the presence of model error it may prove that tuning the parameter a less than 1 is undesirable. Nevertheless, because $a < 1$ this will lead to times when the inflation algorithm will in fact deflate the ensemble variance. Second, the quadratic dependency on the innovation in (4.4) will typically underestimate the actual relationship between the MSE and the ensemble variance in non-Gaussian situations because we are not explicitly accounting for the term $(\overline{x} - \overline{x}_K)^2$, which as discussed at the end of section 3 is likely to be a stronger function of the innovation than quadratic. This implies that we will tune the parameter $b > 1$. We will find in the applications to be described next that b needs to be about the same size as the ensemble size and that this value for b is a reasonable approach to start the tuning process. Alternatively, we could include a quartic term, rather than simply neglect $(\overline{x} - \overline{x}_K)^2$, following the theory of Hodyss (2011) and as discussed at the end of section 3. This we feel is outside the scope of the present study and will be left to future work.

Given (4.4) we may calculate the implied inflation as

$$\mathbf{g} = \sqrt{\frac{\mathbf{S}_G}{\mathbf{p}_a}}, \quad (4.8)$$

where again all mathematical operations are to be interpreted as element-wise. We use (4.8) in a way similar to (4.3):

$$\mathbf{X} \rightarrow \mathbf{g} \circ \mathbf{X}, \tag{4.9}$$

where $\circ$ represents the element-wise product.

In this section we will use the Lorenz (1963) equations as the first nonlinear dynamical system to be studied:

$$\frac{dx}{dt} = \sigma(y - x), \tag{4.10a}$$

$$\frac{dy}{dt} = x(r - z) - y, \tag{4.10b}$$

$$\frac{dz}{dt} = xy - \beta z, \tag{4.10c}$$

where the parameters will be $(\sigma, r, \beta) = (10, 28, 8/3)$. The system (4.10a)–(4.10c) will be solved using the explicit Runge–Kutta (4, 5) solver (ode45) found in recent versions of MATLAB.

We assimilate observations for 10^5 cycles with a time between observations of 0.15. If we did this with a small ensemble and without any sort of inflation of the prior or posterior, the EnKF described in (4.2) fails to track the truth for the entire 10^5 cycles (not shown). This sort of failure is common when using small ensemble sizes, and motivates the application of inflation to track the truth.

Each of the above three methods are individually tuned for minimum MSE averaged over the three state-variables for ensemble sizes of 8, 16, and 32 members and over the entire 10^5 cycles. This tuning for the constant prior and posterior inflation algorithm simply refers to the tuning of the α and β in (4.2) and (4.3) to obtain the minimum MSE. For the observation-dependent method, this required tuning a and b in (4.4) until the minimum MSE was found. For the $N_e = 8$ ensemble, which we will discuss in detail below, the constant prior inflation technique required $\alpha = 1.18$ and the constant posterior inflation technique required $\beta = 1.2$ to find the minimum MSE. The observation-dependent method found $a = 0.92$ and $b = 4$ to obtain the minimum MSE. The resulting MSE is shown in Fig. 4. This figure shows that the new posterior inflation algorithm results in the minimum MSE over the testing period for each ensemble size tested. In addition, a comparison between constant posterior and prior inflation does not appear to reveal a clear winner.

To understand why the new technique was superior we plot in Fig. 5 the binned MSE and ensemble variances for the $N_e = 8$ ensemble as a function of normalized innovation, which we define here as

$$\upsilon_N = \sqrt{(\mathbf{y} - \mathbf{H})^{\mathrm{T}}[\mathbf{H}\mathbf{P}_f\mathbf{H}^{\mathrm{T}} + \mathbf{R}]^{-1}(\mathbf{y} - \mathbf{H})}. \tag{4.11}$$

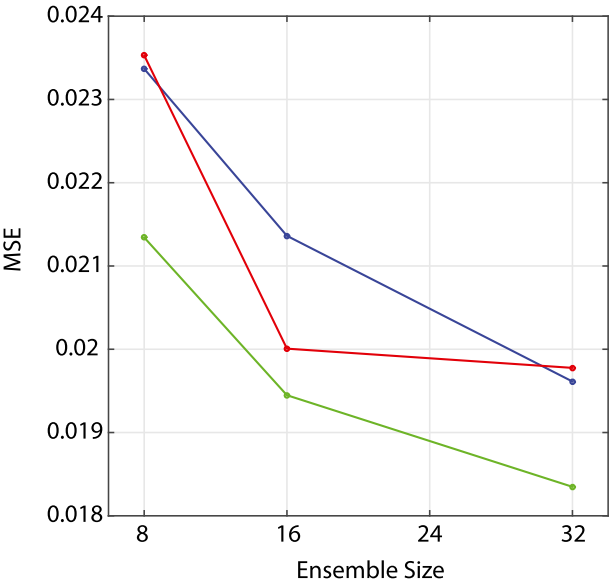

FIG. 4. Average MSE for the Lorenz-63 experiments. Constant prior (posterior) inflation as a function of ensemble size is shown in blue (red). In green the observation-dependent inflation technique in (4.4) is shown.

We note here that the $\mathbf{P}_f$ used in this calculation is the sample estimate. Recall that we saw in section 3 that sample estimates used to normalize the innovations should result in the MSE prediction in (4.5) to appear to underestimate the actual MSE and that the EVP from the EnKF should erroneously bend toward low values for large normalized innovation (see Fig. 3c).

In Fig. 5a we show the eight-member constant prior inflation methods MSE and ensemble variance for each of the three variables in Lorenz-63. The curves for the constant posterior inflation method are very similar (not shown). Here, it is seen that the MSE is a strong function of the normalized innovation but the ensemble variance is not. Also, consistent with Fig. 3c the ensemble variance bends slightly to smaller values for large innovation for the x variable. Compare this to Fig. 5b where we show the eight-member observation-dependent posterior inflation methods MSE and ensemble variance for each of the three variables in Lorenz-63. Here, it is seen that the ensemble variance now behaves more like the VP as it is a function of innovation and is generally correctly predicting the MSE. Further, note that the largest MSE in the observation-dependent method is substantially smaller than the largest MSE in the constant prior method. This result also holds for the constant posterior method (not shown). Hence, a substantial benefit of this new posterior inflation method is that the largest errors are smaller, and for those particular cycles that the ensemble does see large errors, the

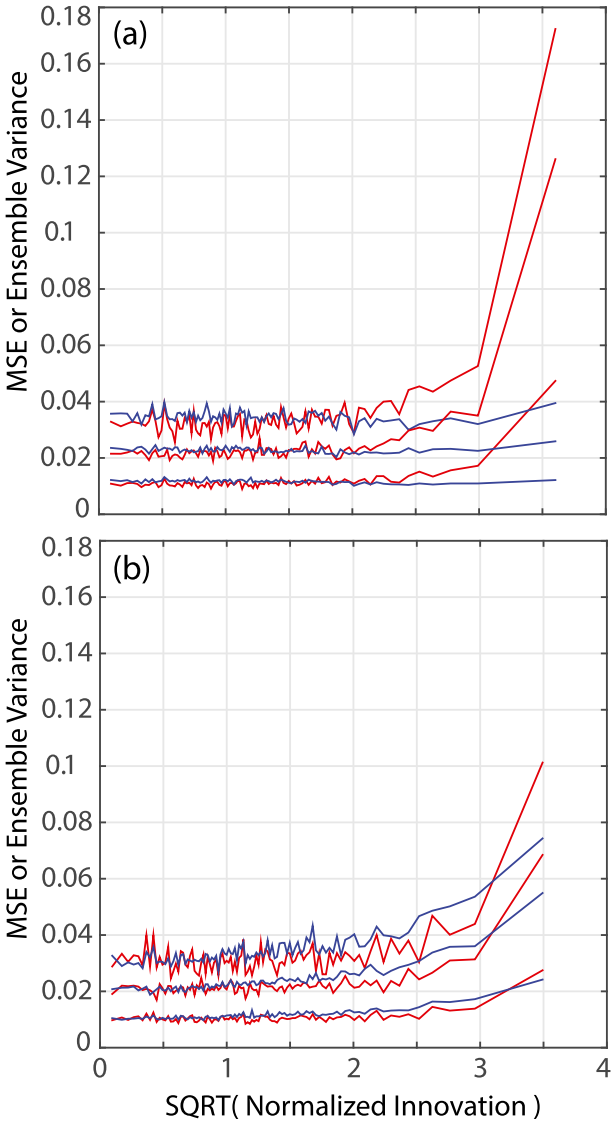

FIG. 5. Ensemble diagnostics as a function of normalized innovation for the eight-member ensemble. (a) The MSE (red) and EVP (blue) for the constant prior inflation method. (b) The MSE (red) and EVP (blue) for the observation-dependent inflation method. In both plots the curves are associated with the x, z, and then y variables arranged from lowest to highest, respectively.

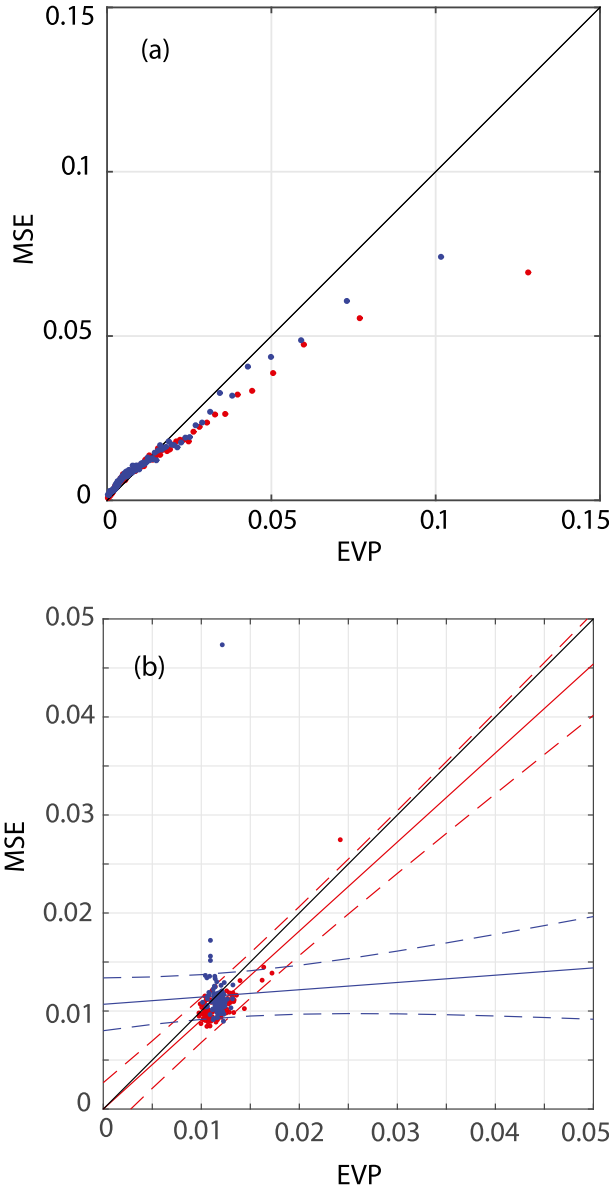

FIG. 6. Binned spread-skill plots for the Lorenz-63 experiments for the eight-member ensemble. (a) A traditional binned spread-skill plot and (b) the new binned spread-skill technique are shown. Blue closed circles represent the constant prior inflation method and red closed circles represent the observation-dependent method. The solid blue (red) line is the linear regression fit for the constant prior (observation dependent) method. Dashed lines are the confidence intervals for the linear regression fits for each technique at 95%.

ensemble is properly inflated to accurately predict these large MSE values.

We feel that it is possible that the skeptical reader will view Fig. 5 as showing that we should simply perform an "innovation check," whereby we simply discard all innovations larger than some threshold value, to eliminate this issue with large MSE from large innovations. To test this hypothesis we reran the same experiments described above (using constant prior inflation) but this time for every cycle we tested each individual observation for the size of its normalized innovation. When an observation had a normalized innovation that satisfied $v_N > c$, (where we tuned a parameter c for minimum posterior MSE) that observation was discarded. In all experiments performed, this tuning procedure resulted in the optimal c value being large enough that all innovations were kept. Hence, throwing out the largest

innovations was always suboptimal. We believe this shows that large innovations have considerable information about the way in which the data assimilation system is performing and that these innovations should be kept. In operational data assimilation, the common practice of performing an innovation check is relevant there because large innovations often imply something systematically wrong about the observing instrument. In our experiments, the observation statistics are precisely correct in the sense that the characteristics of the observational instrument (i.e., it is unbiased and its variance is $\mathbf{R}$) are known exactly.

In Fig. 6 we present the two versions of the binned spread-skill plot discussed in this manuscript. In Fig. 6a we show the traditional binned spread-skill plot for the x variable using the eight-member ensemble and for both the constant prior and observation-dependent technique. The shapes of these curves for the other variables were qualitatively similar. This shows that both techniques do a fairly good job for EVP values below 0.025. However, both techniques appear to overdo the EVP for values of EVP greater than 0.025. In Fig. 6b we show the new binned spread-skill plot for the x variable using the eight-member ensemble and for both the constant prior and observation-dependent technique. This binned spread-skill plot was constructed by taking the data from the MSE (red curve) and EVP (blue curve) in Fig. 5 and plotting a filled circle at the location $(x, y) = (\text{EVP}, \text{MSE})$ in Fig. 6b. This plot clearly shows that the constant prior inflation technique is not properly predicting the MSE as the scattering of data points is not aligned along the one-to-one line. By contrast, the observation-dependent technique has the scattering of data points more arranged along the one-to-one line revealing a better relationship between MSE and EVP. Moreover, statistical significance testing of the linear fit through the data reveals that the observation-dependent inflation technique has a statistically significant relationship to the true MSE, while the constant prior inflation method does not. This is seen in Fig. 6b, whereby the linear regression line and its statistical significant region, do not include the one-to-one line for the constant prior inflation technique. The relationships shown in Fig. 6b extend to the other variables of the Lorenz model as well as to other ensemble sizes and to the constant posterior inflation method.

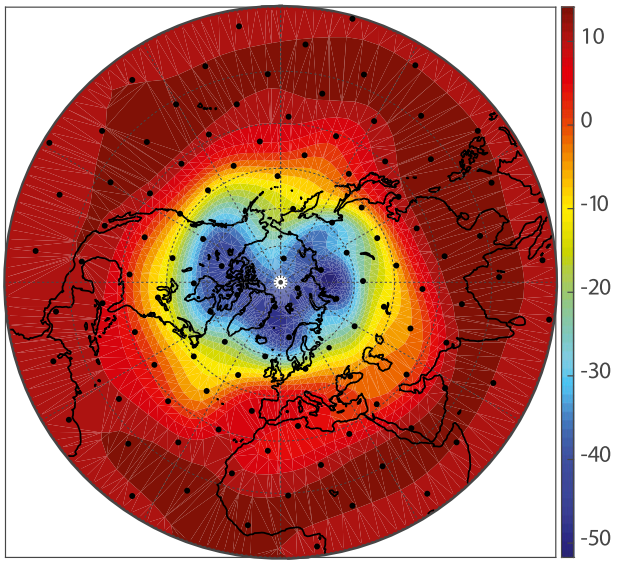

FIG. 7. Potential temperature and observing network. Shading reveals the potential temperature field on the last day of the validation period. Filled circles show the observing network.

b. *Two-level primitive equations*

Here we use the two-level primitive equation model of Whitaker and Hamill (2012). This is a triangular truncation at total wavenumber-32 (T32) spectral model very similar to that of Lee and Held (1993), which grossly simulates the general circulation of the atmosphere [details are found in Whitaker and Hamill (2012) and references therein]. The serial ensemble square root filter algorithm of Whitaker and Hamill (2002) will be used with a 10-member ensemble. We will assimilate 255 temperature observations (drawn from a T32 nature run) that are nearly equally spaced around the globe (i.e., at the Fibonacci points on the sphere) with an observation error variance of 1 K (see Fig. 7). We assimilate this observational network every 12 h for 2200 cycles, discard the first 200 cycles and calculate validation statistics on the last 2000 cycles (equivalent to 1000 days).

In this subsection, we compare the posterior inflation algorithm in (4.4) to the technique of Whitaker and Hamill (2012) referred to as relaxation to prior spread (RTPS):

$$\mathbf{g}_{\text{RTPS}} = \alpha \frac{\sqrt{\mathbf{p}_f} - \sqrt{\mathbf{p}_a}}{\sqrt{\mathbf{p}_a}} + \mathbf{1}, \tag{4.12}$$

TABLE 1. Posterior RMSE.

Method	U (lower)	U (upper)	V (lower)	V (upper)	W	θ
RTPS	1.1123	1.7137	1.0715	1.738	0.0737	0.457
Observation dependent	1.0836	1.6864	1.0475	1.7086	0.0732	0.449

where all mathematical operations are to be interpreted as element-wise, **1** is the one vector whose length is the same as $\mathbf{p}_a$, and α is a tunable parameter controlling the amount of inflation. We use (4.12) to inflate the posterior as

$$\mathbf{X} \rightarrow \mathbf{g}_{\mathrm{RTPS}} \circ \mathbf{X}. \tag{4.13}$$

It is important to realize that, while $\mathbf{g}_{\mathrm{RTPS}}$ knows the type and location of any observations that went into the calculation of $\mathbf{p}_a$, $\mathbf{g}_{\mathrm{RTPS}}$ does not know the value of the observation and its distance from the prior mean. Similarly, those posterior inflation algorithms that aim to make the variance correct on average (e.g., Zhang et al. 2004; Furrer and Bengtsson 2007; Sacher and Bartello 2008) also have this trait that they know the type and location of any observations, but they do not know the value of the observation and its distance from the prior mean. Therefore, in this subsection we will be comparing the performance of posterior inflation methods that know the value of the observation and its distance from the prior mean to methods that know only the type and location of the observation.

We tuned the RTPS method and the observation-dependent method in (4.4) for smallest root-mean-square error (RMSE) over the 2000-cycle validation period. For both techniques, we tuned the localization length scale and the particular tunable parameter for that technique. We found that both techniques minimized the RMSE with a localization length scale of 6000 km. For the RTPS method we found $\alpha = 0.4$ produced the lowest RMSE and for the observation-dependent method we found $a = 0.89$ and $b = 10.25$ produced the lowest RMSE. The RMSE for these settings for both techniques is presented in Table 1.

In Fig. 8 we present the traditional binned spread-skill diagram as well as the result of binning the MSE and ensemble variance as a function of normalized innovation. In Fig. 8a, the binned spread-skill diagram for potential temperature at the observation locations reveals that both methods produce a reasonable *average* relationship between MSE and ensemble variance for ensemble variances less than about 0.5. Both methods produce far too much ensemble variance for ensemble variances greater than about 0.5. This pattern, in which the binned spread-skill diagram reveals an overdispersive ensemble for large ensemble variances, was also seen in the Lorenz-63 experiments for similarly small ensemble sizes. We speculate that this curvature in the binned spread-skill diagram for small ensemble size is due to the relative lack of sensitivity of the Kalman gain to sampling error when the true prior variance is substantially larger than the observation error variance. In this case, the sampling error that leads to a substantially larger estimate of the prior variance than the true prior variance leads to a Kalman gain that is only a small amount closer to one, and therefore its impact on the state estimate is minimal. This leads to the MSE not increasing at the same rate as the ensemble variance,

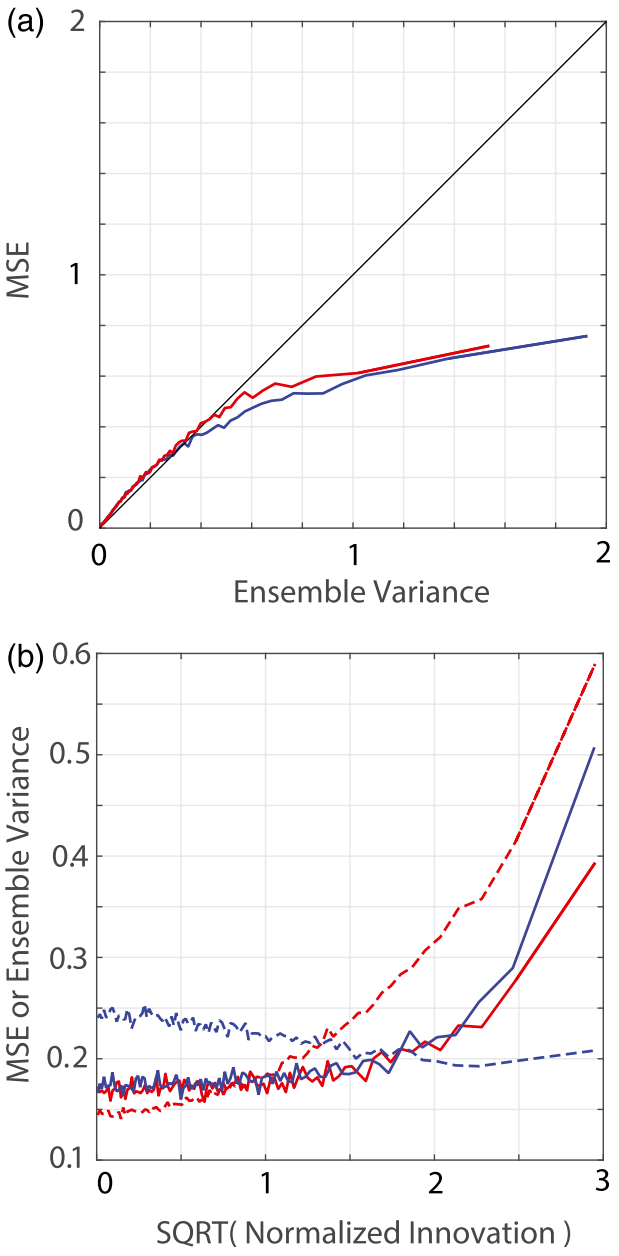

FIG. 8. Binned spread-skill diagram and MSE and ensemble variance as a function of innovation for potential temperature. (a) The binned spread-skill diagram for RTPS (blue) and the observation-dependent method (red). (b) The MSE (solid) and ensemble variance (dashed) for the RTPS (blue) and the observation-dependent method (red).

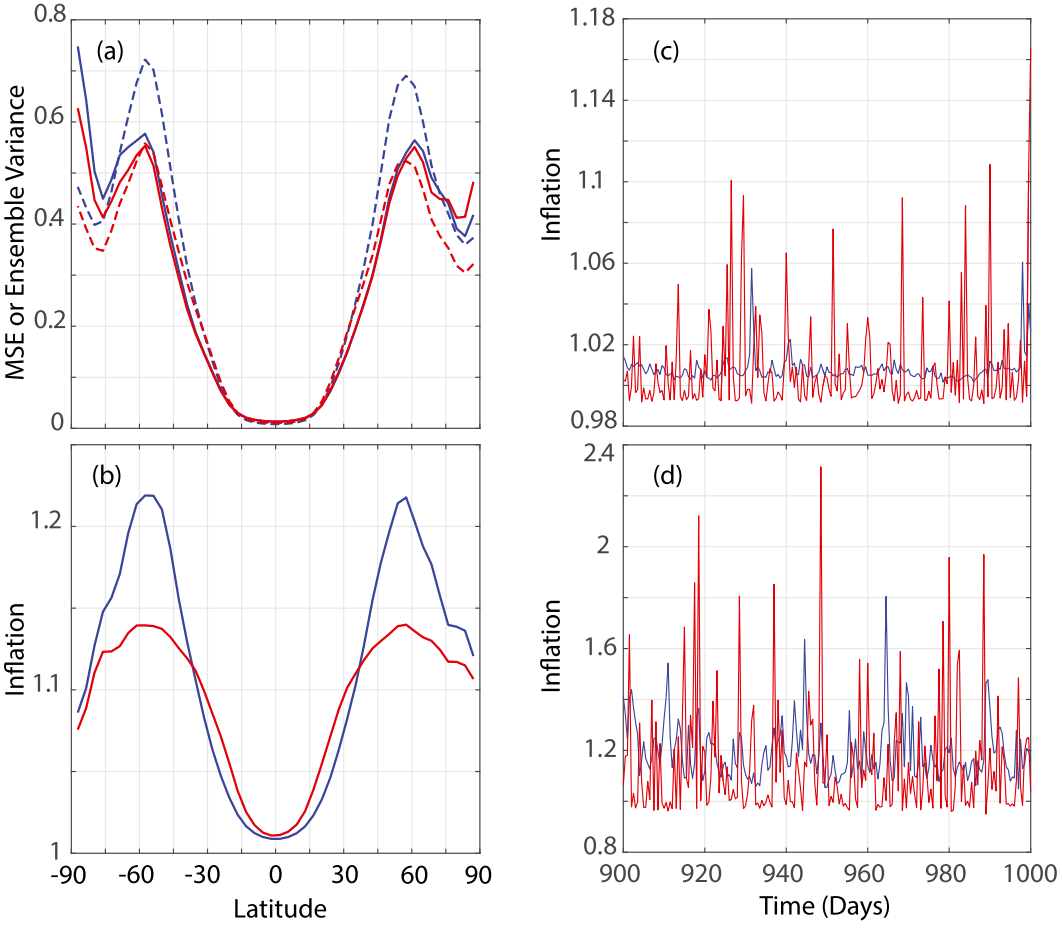

FIG. 9. Diagnostics for potential temperature. (a) The time and zonally averaged MSE (solid) and ensemble variance (dashed). (b) The time and zonally averaged inflation. (c),(d) The time series of the inflation for the grid points at 180° and 2°N and 50°N, respectively, and for the last 100 days of the validation period. In all frames, the RTPS method is blue and the observation-dependent method is red.

and hence a binned spread-skill diagram that curves to the right.

In any event, to understand this improvement we provide in Fig. 9 a comparison of the inflation for both techniques and the MSE for the potential temperature variable. In Fig. 9a the zonally averaged ensemble variance and MSE is plotted. The other variables (e.g., zonal, meridional, and vertical wind) have very similar properties. Here, we see that the major difference between the observation-dependent method and RTPS occurs near 55° latitude, where the RTPS method appears to create an ensemble variance that is far too large. In Fig. 9b we present the zonally averaged inflation for both techniques. Here, we see why the RTPS method has too much ensemble variance near 55° latitude as its inflation is typically much larger than the observation-dependent method. In Figs. 9c and 9d we present time series of the inflation for the grid points at 180° and 2°N and 50°N, respectively. We show these time series plots to show the difference between the methods, but also to show that the desirable properties discussed in section 4a can be seen here. In section 4a, we noted that we desired for the inflation to sometimes "deflate" the ensemble variance, but then at other times, to inflate depending on the size of the innovations. This pattern can be clearly seen in Figs. 9c and 9d, and by comparison we see that the RTPS method produces a relatively smoother and more temporally correlated inflation. The relatively noisy temporal pattern to the inflation from the observation-dependent method does not, however, result in a noisy spatial pattern. The spatial pattern for the inflation for both methods can be seen in Fig. 10. Here, we see that both methods produce a spatial pattern with similar spatial scales. Note, however, that the observation-dependent method produces a spatial pattern with large regions less than 1. We tested the

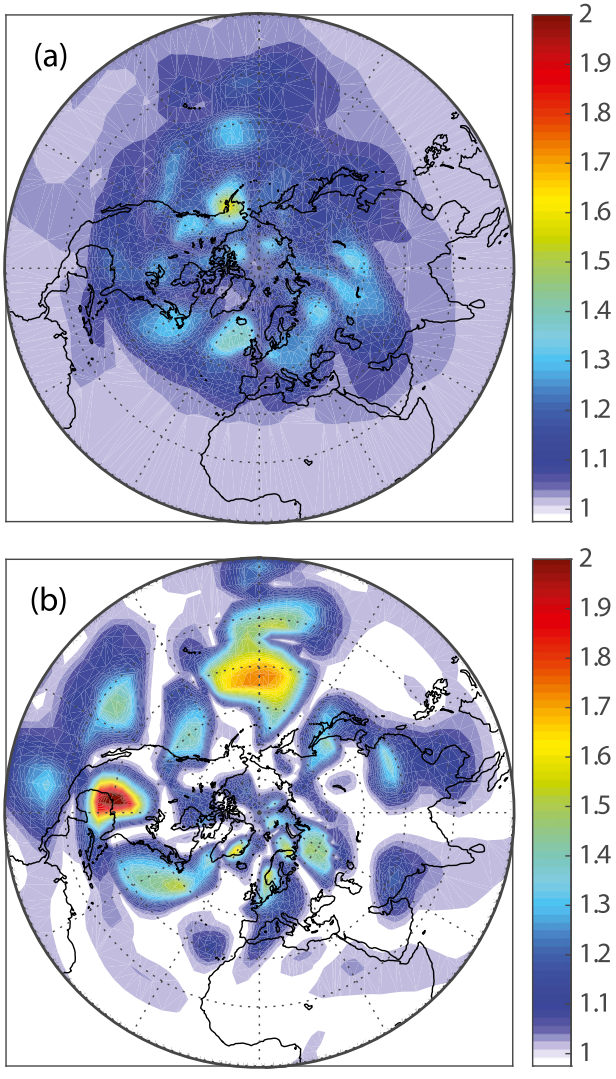

FIG. 10. (a) A snapshot of the inflation from the RTPS method on the last day of the validation period and for the potential temperature. (b) A snapshot of the inflation from the observation-dependent method on the last day of the validation period and for the potential temperature. Note that for both plots the color scales are identical and white denotes values less than 1.

beneficial aspects of these regions by increasing the parameter a until it was equal to 1 and reran this same experiment. The result was a severely degraded performance in all regions. This we believe shows that the proper ''deflation'' of the posterior variances is beneficial and should be included in posterior inflation algorithms that attempt to produce a Bayesian ensemble.

5. Summary and conclusions

This article has sought to illustrate the differences between the EVP [(2.1)] and the VP [(2.2)]. The reason this is of paramount importance to ensemble generation is because a Bayesian posterior has the property that the posterior mean squared error is equal to the VP [(2.2)] and not the EVP [(2.1)], but all EBKF ensemble generation methods attempt to deliver EVP instead of VP. The critical property that leads to differences between the VP and EVP is the fact that the EVP is independent of observations. We showed that there are two ways that the VP and EVP will typically differ; non-Gaussianity in the posterior leads to a VP that is quartic in the innovation, while the sampling error in the prior leads through the Kalman formula to a VP that is quadratic in the innovation. This observation-dependent VP led us to introduce both an algorithm to predict this dependency as well as a verification tool to measure this dependency.

We showed in a hierarchy of models that accounting for this discrepancy between the VP and EVP results in a better ensemble in the sense that the posterior ensemble mean has smaller MSE and the posterior ensemble variance better tracks the MSE. We have provided an algorithm in section 4 that can be applied to state vectors of any length. The main computational burden for the algorithm is the tuning of two parameters, which we would suggest performing regionally over a training set. Last, much of the experimental work presented here used quite small ensemble sizes (32 members or less). We would suggest that for ensembles sizes larger than this, which is likely to have little sampling error and, therefore, a much larger contribution from non-Gaussianity, this inflation parameterization may be better replaced by a scalar evaluation of the quadratic polynomial filter of Hodyss (2011). This would require employing the quadratic polynomial filter equation to predict the difference between the true posterior mean and the true Kalman filter mean, and using this result in the second term of (3.15). The major issue with this suggestion is that it would require having the observations available when calculating the inflation factor, which substantially increases the complexity of the algorithm from its present form.

This research has suggested that future work understanding the relationship of the ensemble moments from an EBKF as compared to the Bayesian solution should provide fruitful performance improvements. Work in this direction is under way.

Acknowledgments. We thank Jeff Anderson for helpful and insightful comments on this work. This research is supported by the Chief of Naval Research through the NRL Base Program, PE 0601153N.

REFERENCES

Anderson, J. L., 2007: An adaptive covariance inflation error correction algorithm for ensemble filters. *Tellus*, **59A**, 210–224, doi:10.1111/j.1600-0870.2006.00216.x.

——, 2009: Spatially and temporally varying adaptive covariance inflation for ensemble filters. *Tellus*, **61A**, 72–83, doi:10.1111/j.1600-0870.2008.00361.x.

——, and S. L. Anderson, 1999: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Wea. Rev.*, **127**, 2741–2758, doi:10.1175/1520-0493(1999)127<2741:AMCIOT>2.0.CO;2.

Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99**, 10 143–10 162, doi:10.1029/94JC00572.

——, 2003: The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dyn.*, **53**, 343–367, doi:10.1007/s10236-003-0036-9.

Furrer, R., and T. Bengtsson, 2007: Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *J. Multivar. Anal.*, **98**, 227–255, doi:10.1016/j.jmva.2006.08.003.

Hodyss, D., 2011: Ensemble state estimation for nonlinear systems using polynomial expansions in the innovation. *Mon. Wea. Rev.*, **139**, 3571–3588, doi:10.1175/2011MWR3558.1.

——, and W. F. Campbell, 2013: Square root and perturbed observation ensemble generation techniques in Kalman and quadratic ensemble filtering algorithms. *Mon. Wea. Rev.*, **141**, 2561–2573, doi:10.1175/MWR-D-12-00117.1.

Lee, S., and I. M. Held, 1993: Baroclinic wave packets in models and observations. *J. Atmos. Sci.*, **50**, 1413–1428, doi:10.1175/1520-0469(1993)050<1413:BWPIMA>2.0.CO;2.

Lorenz, E. N., 1963: Deterministic non-periodic flow. *J. Atmos. Sci.*, **20**, 130–141, doi:10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2.

Menetrier, B., and T. Auligne, 2015: Optimized localization and hybridization to filter ensemble-based covariances. *Mon. Wea. Rev.*, **143**, 3931–3947, doi:10.1175/MWR-D-15-0057.1.

Posselt, D. J., D. Hodyss, and C. H. Bishop, 2014: Errors in ensemble Kalman smoother estimates of cloud microphysical parameters. *Mon. Wea. Rev.*, **142**, 1631–1654, doi:10.1175/MWR-D-13-00290.1.

Sacher, W., and P. Bartello, 2008: Sampling errors in ensemble Kalman filtering. Part I: Theory. *Mon. Wea. Rev.*, **136**, 3035–3049, doi:10.1175/2007MWR2323.1.

Wang, X., and C. H. Bishop, 2003: A comparison of breeding and ensemble transform Kalman filter ensemble forecast schemes. *J. Atmos. Sci.*, **60**, 1140–1158, doi:10.1175/1520-0469(2003)060<1140:ACOBAE>2.0.CO;2.

Whitaker, J. S., and A. F. Loughe, 1998: The relationship between ensemble spread and ensemble mean skill. *Mon. Wea. Rev.*, **126**, 3292–3302, doi:10.1175/1520-0493(1998)126<3292:TRBESA>2.0.CO;2.

——, and T. H. Hamill, 2002: Ensemble data assimilation without perturbed observations. *Mon. Wea. Rev.*, **130**, 1913–1924, doi:10.1175/1520-0493(2002)130<1913:EDAWPO>2.0.CO;2.

——, and ——, 2012: Evaluating methods to account for system errors in ensemble data assimilation. *Mon. Wea. Rev.*, **140**, 3078–3089, doi:10.1175/MWR-D-11-00276.1.

Zhang, F., C. Snyder, and J. Sun, 2004: Impacts of initial estimate and observation availability on convective-scale data assimilation with an ensemble Kalman filter. *Mon. Wea. Rev.*, **132**, 1238–1253, doi:10.1175/1520-0493(2004)132<1238:IOIEAO>2.0.CO;2.